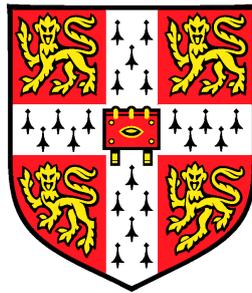


The use of spin-pure and non-orthogonal Hilbert spaces in Full Configuration Interaction Quantum Monte–Carlo

Simon Smart

Trinity College



This dissertation is submitted for the degree of Doctor of Philosophy at the
University of Cambridge, December 2013

For my mother

Diana Jean Smart

1956-2013

The use of spin-pure and non-orthogonal Hilbert spaces in Full Configuration Interaction Quantum Monte–Carlo

Simon Smart

Abstract

Full Configuration Interaction Quantum Monte–Carlo (FCIQMC) allows for exact results to be obtained for the ground state of a system within a finite-basis approximation of the Schrödinger equation. Working within imposed symmetry constraints permits dramatic reductions in the size and internal connectivity of the Hilbert space considered, with associated reductions in the computational cost involved, as well as permitting exclusion of the natural ground state to extract a series of excited states of the system. As all converged solutions are eigenfunctions of the square of the total spin operator, \hat{S}^2 , as well as the Hamiltonian and the projected spin, imposing *spin-purity* as an additional ‘symmetry’ is a natural extension.

In this thesis, the use of various spin-pure spaces is compared to the previously used determinantal spaces. Variations on the FCIQMC algorithm which work in non-orthogonal (and non-normalised) basis sets, and with the arbitrary discretisation of imaginary time removed, are considered along with the implications of the differences to the normal FCIQMC algorithm.

Preface

This dissertation contains an account of research carried out in the period between October 2009 and December 2013 in the Department of Chemistry at the University of Cambridge, under the supervision of Professor Ali Alavi. The following sections of this thesis are included in work that has either been published or is to be published:

Chapter 2 and parts of chapter 5 *Linear-scaling and parallelizable algorithms for stochastic quantum chemistry.* George H. Booth, Simon D. Smart and Ali Alavi. arXiv:2305.6981 and submitted.

Parts of chapter 2 *Full Configuration Interaction Excitations of Ethene and Butadiene: Resolution of an Ancient Question* Csaba Daday, Simon Smart, George H. Booth, Ali Alavi and Claudia Filippi. *J. Chem. Theory Comput.*, **8**, 4441 (2012).

This dissertation is a result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text. **Section 4.4** has been previously included in work submitted for the Certificate of Postgraduate Study. Other than this exception, this dissertation has not been submitted in whole or in part for any other degree or diploma at this or any other university. This dissertation does not exceed 60,000 words in length.

Simon D. Smart
December 2013

Acknowledgements

I would like to thank Ali, my supervisor, both for his ideas and (sometimes uncontrollable) enthusiasm, but also for his almost magical ability to make logistical problems go away when they rear their head. Without his support I certainly wouldn't have made it to the end of this PhD — let alone had so much fun along the way.

I am greatly indebted to George Booth for the huge amount of knowledge and advice he has (tried) to pass on to me, and to Alex Thom who has made a habit of turning up just when my attempts to solve problems appear to be turning into bashing my head on the desk — and then nudging me in the right direction.

Thanks, also, to James Spencer who has been an excellent sounding board and source of ideas for implementation related concerns, and generally stimulating my computer-geekery in an enjoyable fashion. To James Shepherd, thanks you for your insight into statistics and errors which always illuminates my weaknesses. To Michael Moeckel, I am glad that you always think about scientific problems in a very different way to me. Thanks Catherine for playing around with the real-coefficient code.

Thanks Deidre for always bringing a cheerful antipodean presence to tea, and Andreas and Rob for being generally helpful. Jenny, thank you for putting up with my idiosyncrasies in the office, and letting me get involved in what you were doing. Nick, thanks for your excellent proofreading.

I would also like to thank Trinity college for not just funding, housing and feeding me, but for providing an environment in which I have felt so at home. I am very grateful to Dr. Khalfa and everyone in the Graduate office who have created the space both temporally and financially for me to pick up the pieces when everything fell apart. Twice. To all my friends who have made my time here such a pleasure, you know who you are. Thanks especially to Kathy for your excellent proofreading (with a good sense of humour). I would also like to thank my dad, who has put a roof over my head, and somehow held everything together over the last year.

Finally I would like to thank my wife, Hilary Hall who has tirelessly read and re-read everything I have written and without whose support the last few years would have been much more difficult.

Contents

Abstract	i
Preface	iii
Acknowledgements	v
1 Introduction	1
1.1 Non-technical introduction	1
1.2 Thesis overview	3
1.3 Electronic structure problems	4
1.4 Hartree–Fock theory	6
1.4.1 Electron correlation	7
1.5 Moving beyond Hartree–Fock theory	8
1.5.1 Perturbative approaches, and Møller–Plesset perturbation theory	9
1.5.2 (Full) Configuration Interaction	11
1.5.3 Coupled cluster approach	13
1.5.4 Monte–Carlo approaches	14
1.6 Aims for further development	17
2 Full Configuration Quantum Monte–Carlo	19
2.1 What is FCIQMC?	19
2.2 Derivation of FCIQMC	21
2.2.1 Orthogonal FCIQMC (Booth et al.)	22
2.2.2 Non-orthogonal FCIQMC	23
2.3 Energy estimators	26
2.3.1 The projected energy, E_P	26
2.3.2 The shift, E_S	28
2.3.3 The weighted energy estimator	29
2.4 The ‘standard’ FCIQMC algorithm	31
2.4.1 Description, granularity and chance	31

2.4.2	Initialisation (trial wavefunction)	36
2.4.3	The ‘standard’ algorithm	37
2.4.4	The initiator approximation	39
2.4.5	A typical FCIQMC calculation	42
2.5	Issues for consideration	44
2.6	The imaginary time step, $\delta\tau$	47
2.7	Basis sets for FCIQMC	49
2.7.1	Symmetry	50
2.7.2	Time-reversal (and similar) symmetries	53
2.7.3	Mixed schemes	54
3	The role of spin in FCIQMC	57
3.1	Physical and practical views of spin	57
3.2	Evolution of spin in FCIQMC	60
3.2.1	Instantaneous spin for Slater determinants	60
3.2.2	Instantaneous spin for HPHF functions	62
3.2.3	Spin trajectories and the initiator approximation	63
3.3	Spin eigenfunctions	67
3.3.1	Size of the Hilbert Space	68
3.3.2	Spin paired (Rumer) spin eigenfunctions	68
3.3.3	Genealogical (Kotani-Yamanouchi) spin eigenfunctions	71
3.3.4	Serber-type spin eigenfunctions	73
3.3.5	Spin-projected (Slater) determinants	76
3.4	Implementation of spin eigenfunctions in FCIQMC	78
3.4.1	Representation of spin eigenfunctions	79
3.4.2	Non-local death and population control	81
3.4.3	Demonstration of non-local death	83
3.4.4	Algorithmic changes to suppress population explosion	86
3.4.5	Further issues with spin-projected determinants	87
4	Hamiltonian matrix element evaluation	89
4.1	Slater determinants and HPHF functions	89
4.1.1	Excitation generation integration	90
4.2	Configurational State Functions	91
4.2.1	Scaling of matrix element calculations	93
4.3	Rumer-type CSF	94

4.4	Kotani-Yamanouchi CSFs	95
4.4.1	X_μ and X_ν share the same spatial structure	95
4.4.2	X_μ and X_ν differ by one spatial orbital	100
4.4.3	X_μ and X_ν differ by two spatial orbitals	100
4.5	Serber-type CSFs	104
4.5.1	Line-up permutation	105
4.5.2	Generation of permutation representation matrices	105
4.5.3	Storage of permutation representation matrices	107
4.5.4	Evaluation of matrix elements	110
4.6	Spin-projected determinants	111
4.6.1	Alignment	112
4.6.2	Sanibel coefficient	113
4.6.3	Element classification	115
4.6.4	Integration with excitation generation	116
4.6.5	Matrix element generation	120
4.6.6	The overlap matrix	121
4.7	Mixed Hilbert spaces	122
4.8	Summary	122
5	Excitation generation using CSFs	125
5.1	Comparison to determinantal excitation generators	126
5.2	Generating an excitation	127
5.2.1	Spatial symmetry information	128
5.2.2	Choice of excitation type	130
5.2.3	Generation of (spatial) single excitations	131
5.2.4	Generation of (spatial) double excitations	132
5.2.5	Selection of spin structure or label	135
5.3	Determining type probabilities	135
5.4	Integration with matrix element calculation	136
5.5	Multiple-structure spawning	137
5.5.1	‘Blob-to-blob’ spawning	139
5.6	Spawning in mixed representation schemes	141
6	Nitrogen	145
6.1	The nitrogen atom	145
6.1.1	Different types of CSF	147

6.1.2	Convergence of errors	149
6.1.3	Summary	154
6.2	The nitrogen dimer	156
6.2.1	Hartree–Fock solutions	158
6.2.2	Binding curves for excited states	161
6.2.3	Summary	164
7	A brief aside — Continuous Time FCIQMC	165
7.1	Motivation	165
7.2	Derivation	165
7.3	Implementing continuous time FCIQMC	167
7.3.1	Performing annihilation	167
7.3.2	Generating new particles	168
7.3.3	Algorithm overview	171
7.3.4	The initiator approximation	172
7.4	Acceptance ratios	173
7.5	Comparison of FCIQMC and <i>ct</i> -FCIQMC	175
7.6	Summary	182
8	Concluding remarks	183
	Bibliography	185

List of Figures

1.1	Hierarchy of quantum chemical methods	2
2.1	Off-diagonal energy terms and the projected energy	30
2.2	Full Configuration Interaction Quantum Monte-Carlo	38
2.3	FCIQMC algorithm overview	40
2.4	Trajectories of <i>i</i> -FCIQMC with differing $n_{s,\min}$	43
3.1	Spin trajectories in FCIQMC and <i>i</i> -FCIQMC	64
3.2	Occupation of spin components	66
3.3	Extended Rumer diagrams	70
3.4	Branching diagram for Kotani-Yamanouchi spin functions	72
3.5	Branching diagram for Serber spin eigenfunctions	74
3.6	A particle population explosion caused by non-local death	81
3.7	Spawning connection strength for spin-projected determinants	82
3.8	Correct and incorrect death-like spawning	84
3.9	A model system with a duplicated site	85
3.10	Convergence failure in non-orthogonal FCIQMC	87
4.1	Scaling of representation components	93
4.2	Determinants differing by one spin orbital	101
4.3	Summation of matrix element components	102
4.4	Parity of matrix elements	104
4.5	Proportion of permutations required using Serber functions	108
4.6	On-the-fly generation of permutations	109
4.7	Spawning connection strengths for Serber functions	123
5.1	Determinantal and CSF excitation generators	127
5.2	Random excitation generation overview	129
5.3	Multiple site spawning	138
5.4	Time step variation with multiple-structure spawning	140
5.5	Histogram of p_{gen} values for stretched N_2	141
5.6	Degeneracy of determinant excitation generation	144
6.1	Convergence of FCIQMC with different CSFs	148
6.2	Wavefunction coefficients using CSFs	149
6.3	Initiator error in the Nitrogen atom	151

6.4	Convergence of energies for Nitrogen with basis set size	153
6.5	Basis set convergence of the excitation energies of Nitrogen	155
6.6	Reference Slater determinants for N_2	157
6.7	Hartree–Fock binding curves for N_2	160
6.8	FCIQMC Binding curves for N_2	163
7.1	Oversampling of spawning rates in <i>ct</i> -FCIQMC	170
7.2	Acceptance ratios in FCIQMC and <i>ct</i> -FCIQMC	174
7.3	Particle count trajectories in <i>ct</i> -FCIQMC	176
7.4	Dependence of plateau height and length on $\delta\tau$ in <i>ct</i> -FCIQMC	177
7.5	Computational efficiency in <i>ct</i> -FCIQMC	178
7.6	Energy convergence with $\delta\tau$ in <i>ct</i> -FCIQMC	181

List of Tables

4.1	The Slater–Condon Rules	90
4.2	Non-zero spin-projected Hamiltonian matrix elements	117
6.1	Experimental data for the Nitrogen atom	146
6.2	State and excitation energies for the Nitrogen atom	152
6.3	Hartree–Fock energies for N_2	159
6.4	FCIQMC energies for N_2	162

1 Introduction

1.1 Non-technical introduction

Chemistry can be accurately (if impractically) described as a subset of applied physics. In this mindset, we are solely concerned with describing molecules from the bottom-up — how do the constituent parts of a molecule (electrons, protons and neutrons) interact with each other and with the environment to give the behaviour we observe.

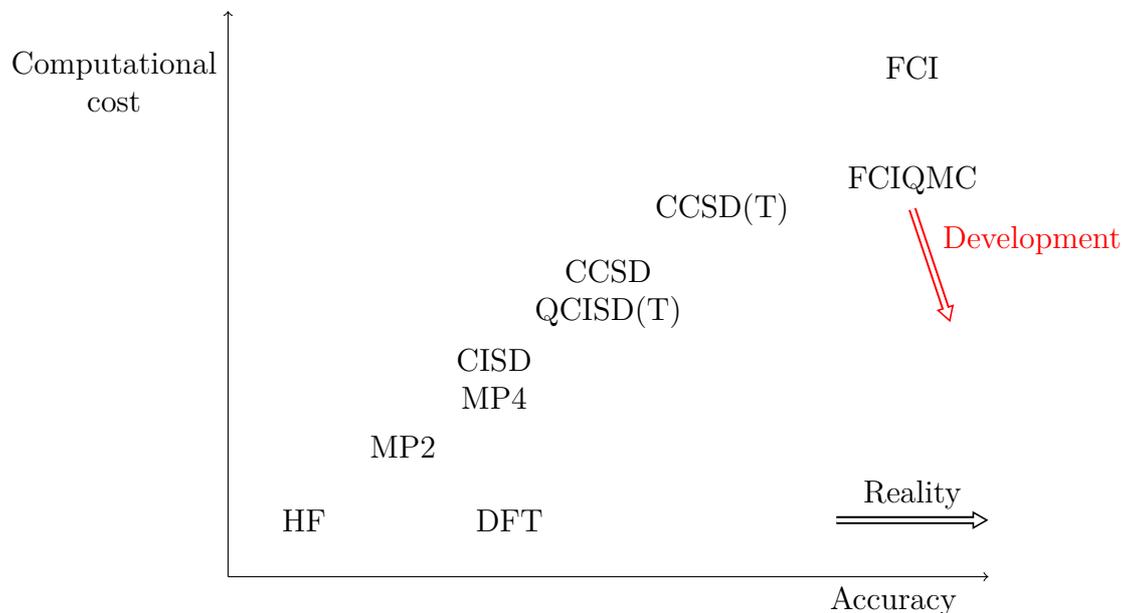
The devil is always in the details.

Unfortunately, it is incomprehensibly difficult to accurately model all of the parts of a molecule fully, and theoreticians are required to make approximations which compromise accuracy for tractability. Indeed, as Dirac famously stated in 1929:¹

The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only in that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed.

The most commonly used models, in particular Hartree–Fock theory, consider the components of a molecule one-by-one, subsuming all of the other parts into an averaged background. This works quite well, and explains a lot of chemistry, but is not quantitatively accurate. In practice, electrons repel each other like poles on a magnet, and will at any moment be found far apart from each other, even if their average position is similar. The dynamic effects of such movement, and similar instantaneous interactions, can never be described by a simple averaged model.

Improving on these approximations is necessary, but is a game of diminishing



HF	— Hartree–Fock theory ^{2–4}
DFT	— Density Functional Theory ^{5,6}
MP2	— Møller–Plesset perturbation theory to second order ^{7,8}
MP4	— Møller–Plesset perturbation theory the fourth order ⁹
CISD	— Configuration Interaction Singles and Doubles ¹⁰
QCISD(T)	— Quadratic Configuration Interaction Singles and Doubles with perturbative Triples ^{11,12}
FCI	— Full Configuration Interaction ^{13,14}
FCIQMC	— Full Configuration Interaction Quantum Monte–Carlo ^{15,16}
CCSD	— Coupled Cluster Singles and Doubles ¹⁷
CCSD(T)	— Coupled Cluster Singles and Doubles (with perturbative Triples) ¹²

Figure 1.1: Hierarchy of quantum chemical methods. All quantum chemical development is aimed at moving to the lower right of this diagram. FCIQMC is capable of achieving the accuracy associated with Full Configuration Interaction (FCI), using substantially less computational resources. This permits it to make use of larger basis sets than FCI, and hence achieve higher accuracy. Development is aimed at reducing the cost of this method further, whilst extending the range of systems it may be applied to, and the range of output values that can be calculated.

returns. Firstly we have to find a language to describe these dynamic effects. Then we have to analyse all of the components of this language, and the way they connect to each other, to find the best possible description of the dynamic effects. To do this systematically and fully for all molecules is nearly impossible — the problem is just too big for even the biggest computers.

Instead of a fully systematic approach, we choose to play a game of chance. We take a guess at a possible description of the molecule, and then we repeatedly make *random* changes to our description which may be accepted as good, or rejected as bad. It is our job to develop the rules of this game, so that our description will always tend to get better, and do so efficiently. We also want these rules to be able to incorporate anything that we already know about the molecule, so that the computational effort can be minimised.

1.2 Thesis overview

Full Configuration Interaction Quantum Monte-Carlo (FCIQMC) is a recent approach to obtaining the electronic correlation energy for chemical systems (molecular or otherwise). It is able to obtain Full Configuration Interaction results using an approach that builds on methods similar to those of Diffusion Monte-Carlo, but using a discrete antisymmetrised basis.

In this thesis, the position of FCIQMC within the ensemble of available quantum chemical techniques is laid out. Following this, a novel derivation of FCIQMC, applicable to systems that make use of non-orthogonal and non-normalised basis sets is presented along with all of the necessary material to make use of FCIQMC in an environment where the basis functions are eigenfunctions of the total spin squared operator, \hat{S}^2 . These methods are then applied to the nitrogen dimer.

For the remainder of this chapter the nature of the problem, i.e. electron correlation, is discussed. Many quantum chemical methods have been developed to approach the problem of electron correlation, and their strengths and weaknesses are explored to shed light on the choices made in developing FCIQMC.

In chapter 2 FCIQMC is derived, then the choices made in terms of representation of the evolving wavefunction, and the stochastic implementation of the method, are discussed along with its basis set requirements.

In chapter 3 the role of spin in FCIQMC is explored, before the general structure of Configurational State Functions (CSFs), i.e. functions each constructed from the product of a spatial wavefunction in an orbital basis and a total spin eigenfunction in turn constructed from primitive spin functions, is laid out and the specific construction of several sets of spin eigenfunctions which all map the relevant Hilbert space are discussed. The consequences of the use of CSFs are explored, along with any algorithmic modifications that are required.

In chapter 4 efficient schemes for calculating the Hamiltonian matrix elements between arbitrary CSFs of the specified types are presented. For these to operate efficiently they need careful integration with the excitation generation processes, which efficiently generate random CSFs connected to a currently occupied one, and are discussed in chapter 5. These techniques are then applied, in chapter 6, to the nitrogen atom and molecule.

In chapter 7 the consequences of the discretisation of imaginary time for the purposes of stochastic integration are discussed. Modifications to FCIQMC which eliminate this requirement, and therefore operate in ‘continuous’ imaginary time are considered.

Finally overall conclusions are discussed in chapter 8.

1.3 Electronic structure problems

The behaviour of chemical systems is normally well described by the Schrödinger equation,^{18,19 *}

$$i\hbar \frac{\partial}{\partial t} \Psi = \hat{H} \Psi. \quad (1.1)$$

The aim of (most) quantum chemical approaches is to find (a subset of) the approximate solutions to the time-independent Schrödinger equation,

$$\hat{H} \Psi = E \Psi, \quad (1.2)$$

*This formulation does not take into account relativistic effects. In principle the Dirac equation²⁰ can be solved to find *spinors* rather than wavefunctions. The drastically increased complexity of this problem is not generally needed to understand most chemistry. A substantial component of the relativistic effects observed may be recovered by application of perturbation theory to apply corrections to non-relativistic solutions²¹ or through the use of pseudopotentials.

for the wavefunction, Ψ , and the energy, E . Additionally, the Born–Oppenheimer approximation²² may be made, assuming that the motion of the electrons and the nuclei are separable. As such the nuclear positions, $\{\mathbf{R}_i\}$, are presumed do be known. Subject to the above, for molecular systems the Hamiltonian is given by

$$\hat{H} = -\frac{1}{2} \sum_i^N \nabla_i^2 - \sum_i^N \sum_j^{N_n} \frac{Z_j}{R_{ij}} + \sum_{i<j}^N \frac{1}{r_{ij}} + \sum_{i<j}^{N_n} \frac{Z_i Z_j}{\mathbf{R}_{ij}}, \quad (1.3)$$

in atomic units, where sums to N and N_n are over the electrons and nuclei in the system respectively, r_{ij} , R_{ij} and \mathbf{R}_{ij} give respectively the inter-electronic distances, the electron-nuclear distances and the nuclear-nuclear distances, and Z_i are the nuclear charges. By separating out the nuclear-nuclear terms, this can be reduced to a purely electronic view. The solutions, Ψ , must additionally satisfy the Pauli principle^{23,24}, in that they must be antisymmetric with respect to exchange of any pair of the N electrons.

Because the non-relativistic Hamiltonian, \hat{H} , is independent of spin, the physical properties associated with spin must be directly imposed. Also as a consequence of this independence, operators for the square of the total spin, \hat{S}^2 , and its z -axis projection, \hat{S}_z , commute with \hat{H} , and so it is possible to choose a complete set of eigenfunctions of \hat{H} that are also simultaneous eigenfunctions of these spin operators.

Due to the r_{ij} terms, which couple pairs of electrons, it is impossible to find analytic solutions to equation 1.2 for many-electron systems. Solutions are, however, still highly sought after. The field of electronic structure therefore involves finding approximations which best represent the solutions to equation 1.2.

Various schemes, of differing levels of complexity, have been proposed to help find these solutions. Here, Hartree–Fock theory will be discussed, followed by a discussion of why such solutions are insufficient when high accuracy is required. Perturbative corrections on top of Hartree–Fock theory are discussed followed by other methods building on top of Hartree–Fock theory for obtaining the correlation energy.

1.4 Hartree–Fock theory

Hartree–Fock theory^{2,3} is concerned with constructing the best possible approximate N -electron wavefunction out of only single-electron orbitals. In particular, the wavefunction is represented as a single Slater determinant²⁵, D_0 , which may be expressed as

$$\begin{aligned} |D_0\rangle &= \mathcal{A} \prod_i^N \phi_i^{\alpha_i} \\ &= \sum_{\mathbf{P}} (-1)^P \hat{\mathbf{P}} \prod_i^N \phi_i^{\alpha_i}(i). \end{aligned}$$

The method then attempts to optimise the set of single electron orbitals, $\{\phi_i\}$, to produce the optimal Slater determinant. In this context, ‘optimal’ means the orbitals that minimise the electronic energy,⁴ $E_0 = \langle D_0 | \hat{H} | D_0 \rangle$ (see table 4.1). See section 2.7 for a fuller discussion of Slater determinants. The orbitals, $\{\phi_i\}$, are constructed out of an underlying basis set (which is often constructed out of contracted gaussian functions²⁶), and the coefficients of these underlying terms are varied systematically subject to the constraint that $\langle \phi_i | \phi_j \rangle = \delta_{ij}$ to find the orbitals that give the minimal value of E_0 .

This problem may be recast into a different eigenfunction problem, known as the Hartree–Fock equation,*

$$\hat{f}_i |\phi_i\rangle = \epsilon_i |\phi_i\rangle,$$

where the Fock operator acting on the i^{th} electron,

$$\hat{f}_i = \hat{h}(i) + v_i^{\text{HF}},$$

is given by the one-electron kinetic energy term found in the Hamiltonian combined with an effective one-electron potential,

$$v_i^{\text{HF}} = \sum_{b \neq i} \int d\mathbf{x}_2 \phi_b(2) \frac{1 - \mathbf{P}_{12}}{r_{12}} \phi_b(2).$$

*In practice the Hartree–Fock equation is further recast into a matrix equation in an underlying basis set. This leads to the Roothaan equations,²⁷ or Pople–Nesbet²⁸ for unrestricted calculations, which are able to be tackled efficiently computationally. For the purposes of CSF development we restrict ourselves to restricted closed shell (RHF) or restricted open shell (ROHF) Hartree–Fock.

This potential essentially includes the averaged electron–electron interaction between an arbitrary electron and all of the electrons which are occupying the calculated orbitals.

This eigenvalue equation is complicated by the fact that the Fock operator, \hat{f}_i , itself depends on the orbitals, $\{\phi_i\}$. As such solutions must be found self-consistently, i.e. iteratively with a trial solution used in constructing the Fock operator, which is then used to obtain a new set of eigenfunctions, which are in turn used as the trial solutions until the trial and resultant functions converge.

1.4.1 Electron correlation

In reality, electrons do not interact with each other in an averaged fashion. The manner that their behaviour differs from averaged behaviour is generally split into two components, static and dynamic correlation — although the division between these is not entirely well defined.

Dynamic correlation

Electrons repel each other dynamically, and as a consequence are on average found further apart than is predicted if they are considered to move independently in an averaged repulsive field of other electrons. Their behaviour is also coupled in more subtle ways.

To describe this, terms dependent on the instantaneous distance between electrons must be added to the $3N$ -dimensional electronic wavefunction in all positions where electrons are close to each other²⁹.

Static correlation

It can be seen in a fully converged CI wavefunction (see section 1.5.2) that multiple determinants can achieve similar and significant weightings. This is especially true in systems with unpaired electrons in the ground state, systems which are nearly dissociated or systems with multiple localised spins. In these cases the wavefunction cannot be reasonably approximated by a single determinant, even with corrections made for the dynamic behaviour of electrons in each other's vicinity, and perturbation based approaches will fail.

The archetypal example of this is the failure of restricted Hartree–Fock the-

ory to effectively describe bond dissociation. A dissociated system is effectively made up of separate and non-interacting fragments. Unless all of the generated fragments are closed shell, attempting to describe these with a single determinant constructed from molecular orbitals over the whole system is bound to fail as these molecular orbitals necessarily include inter-atomic terms that should not appear. An example of this may be observed in figure 6.7, where the HF solution is not *size consistent*.

This problem may be approached using unrestricted solutions, permitting the breaking of spin-symmetry and localising some of the contributions, but this abandons any pretense that the wavefunction may represent an eigenfunction of \hat{S}_z or \hat{S}^2 , and does little to address static correlation in other types of system, such as localised spins in open-shell systems.

Slater determinants describe the behaviour of electrons according to an antisymmetrised product of one-electron functions. These can never include a description of dynamic electron-electron behaviour. Because the Hartree–Fock method can capture all of the electronic behaviour describable using single-electron orbitals in a given underlying basis set, it is used as a reference. The difference between the exact energy for the system, as described by the underlying basis set, and the Hartree–Fock energy may now be considered to describe only these coupled electron-electron effects, and is therefore labelled the *correlation energy*,

$$E_C = E_{\text{exact}} - E_{\text{HF}}. \quad (1.4)$$

The primary aim of post-Hartree–Fock electronic structure methods is to calculate this correlation energy, and thus capture the consequences of the dynamic behaviour of electrons.

1.5 Moving beyond Hartree–Fock theory

There have been many attempts made to move beyond Hartree–Fock solutions and improve the accuracy of quantum chemical calculations. Generally there is a trade-off between accuracy and computational cost as demonstrated in figure 1.1. Here several general categories of approach are considered.

1.5.1 Perturbative approaches, and Møller–Plesset perturbation theory

The simplest correction to the Hartree–Fock solution is to apply a perturbative correction using Rayleigh–Schrödinger perturbation theory^{7,30}. The Hamiltonian is written as

$$\hat{H} = \hat{H}_0 + \hat{H}_1,$$

and the energy and wavefunction written relative to a known solution to the model Hamiltonian, \hat{H}_0 . If the model Hamiltonian is taken as the Fock operator,

$$\begin{aligned} \hat{H}_0 &= \sum_i \hat{f}(i) \\ &= \sum_i [h(i) + v^{\text{HF}}(i)], \end{aligned}$$

and the zeroth order wavefunction $\Psi_0 = \Psi_{\text{HF}}$, then a perturbation may be used to ‘correct’ the overall Hamiltonian back to the true N -electron formulation, such that

$$\hat{H}_1 = \sum_{i<j}^N \frac{1}{r_{ij}} - \sum_i v^{\text{HF}}.$$

This formulation is known as Møller–Plesset Perturbation Theory (MPPT) or alternatively as Many-Body Perturbation Theory (MBPT), especially by physicists. The first and second order energy contributions can be extremely straightforwardly obtained, such that

$$\begin{aligned} E_0^{(1)} &= \langle \Psi_0 | \hat{H}_1 | \Psi_0 \rangle \\ &= -\frac{1}{2} \sum_{a,b} \left(\langle ab | \hat{h} | ab \rangle - \langle ab | \hat{h} | ba \rangle \right), \text{ and} \\ E_0^{(2)} &= \sum_{n \neq 0} \frac{|\langle 0 | \hat{H}_1 | n \rangle|^2}{E_0^{(0)} - E_n^{(0)}} \\ &= \sum_{\substack{a<b \\ r<s}} \frac{|\langle ab | \hat{h} | rs \rangle - \langle ab | \hat{h} | sr \rangle|^2}{\epsilon_a + \epsilon_b - \epsilon_r - \epsilon_s}, \end{aligned}$$

where $\{a, b, r, s\}$ refer to orbitals and orbital energies, and n refer to the available Slater determinants constructed out of these — the sum refers only to double excitations of the ground state as these are the only ones that contribute. Higher

order energy terms, and the expansion of changes in the coefficients, involve many more terms but can also be obtained³¹. Commonly the energy contributions of \hat{H}_1 are calculated to second (MP2⁸), third (MP3^{32,33}) or to fourth order (MP4⁹).

Møller–Plesset theory does an excellent job of effectively describing the local electron–electron behaviour throughout the system — it gets the electronic ‘cusp’ approximately right. It is also size-consistent in each order³⁴. However, as a form of perturbation theory, the final result can only be as good as the initial guess, the Hartree–Fock solution, permits. Systematic improvement on this result requires application of extremely high-order corrections, which rapidly becomes infeasible.

1.5.1.1 Other perturbative corrections

Perturbation theory is a powerful and general approach to obtaining more accurate results than would otherwise be available with a given level of theory. Corrections may be applied on top of any level of theory, and as such have wide applicability. A few examples are given here.

Electronic cusp conditions

The functional form of the wavefunction where two electrons approach each other can be obtained analytically³⁵, and is proportional to the interelectronic separation. R12 and F12 theory can include terms of this form perturbatively^{29,36}, with varying long-range effects.

Coupled cluster

Adding additional excitation levels to the cluster operator (see section 1.5.3) is both complicated and computationally expensive. CCSD(T)¹² uses perturbation theory to approximate the effect of including triple excitations in the cluster operator.

Relativistic corrections

By using the Schrödinger equation, it is implicit that the results will not be relativistically correct. For spectroscopically accurate results, correcting for relativistic effects may be required³⁷, and this is most easily done using perturbation theory.

Note that for all forms of perturbation theory, the quality of the final result is dependent on the perturbation being ‘small’. Ultimately this critically depends on

the quality of the underlying level of theory.

F12 and relativistic corrections can be applied to FCIQMC³⁸ results to improve convergence with respect to the dynamic correlation, and extend the theoretical reach beyond non-relativistic quantum mechanics respectively.

1.5.2 (Full) Configuration Interaction

Configuration Interaction (CI) methods^{13,14} attempt to construct eigenfunctions of the Hamiltonian within a basis set of $2M$ available spin orbitals. Wavefunctions are constructed as a linear combination of antisymmetrised basis functions, frequently Slater determinants, $\{D_{\mathbf{i}}\}^*$, which are in turn constructed as selections of N spin orbitals from the $2M$ available, such that

$$|\Psi_{\text{CI}}\rangle = \sum_{\mathbf{i}} c_{\mathbf{i}}^{\text{CI}} |D_{\mathbf{i}}\rangle,$$

where the coefficients $\{c_{\mathbf{i}}^{\text{CI}}\}$ remain to be found. The CI coefficients satisfy an eigenvector problem,

$$\sum_{\mathbf{j}} \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle c_{\mathbf{j}} = E_{\text{CI}} c_{\mathbf{i}},$$

where E_{CI} is the energy of the given wavefunction, and the solutions can be found by diagonalising the Hamiltonian matrix $H_{\mathbf{ij}} = \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle$ in the Slater determinant basis. The eigenvectors of this matrix give the coefficients of wavefunctions that are solutions to the Schrödinger equation within the given basis of Slater determinants, and the lowest energy eigenvalue corresponds to the ground state.

In general CI methods are not size consistent³⁹. If the complete set of Slater determinants that can be constructed out of the available $2M$ orbitals is used, then this calculation is known as a *Full* Configuration Interaction calculation. This form of calculation is size consistent¹⁰, and the obtained correlation energy is known as the *basis set correlation energy* as it is the maximum value obtainable for a wavefunction constructed out of the underlying basis set that the orbitals were constructed from.[†] Up to this limit, CI calculations are systematically improvable

*Throughout this thesis, indices into the Hilbert spaces used for CI are printed in bold, whereas indices to individual orbitals electrons *etc.* are printed in italic as normal.

[†]FCI calculations, and as a consequence FCIQMC calculations, are insensitive to the quality of the Hartree–Fock solutions used as input. Because a complete set of Slater determinants are used, all contributions are considered regardless, and only the statistical quality of the results in

by including a larger proportion of the available Slater determinants. The overall physical accuracy can then be systematically improved by making use of larger underlying basis sets to construct larger Hartree–Fock orbital bases. As the basis set size is increased, convergence of the electronic cusp behaviour (the dynamic behaviour of electrons in close proximity to each other) is slow^{40,41}. This is because trying to represent detailed local behaviour as a sum of long-range functions in a manner similar to a Taylor expansion is not efficient, whereas the larger scale structural components of the electron correlation are more easily converged.

Despite the exact nature of these results, the utility of the method is limited to rather small systems. The number of determinants in an N -electron system with $2M$ spin orbitals is given by*

$$N_{\text{dets}} = \binom{2M}{N} = \frac{2M!}{N!(2M - N)!},$$

increasing roughly exponentially with both M and N (see figure 4.1). Deterministic diagonalisation of the $N_{\text{dets}} \times N_{\text{dets}}$ matrix rapidly becomes impossible.

In practice several schemes exist to efficiently diagonalise these matrices, especially as a great deal is known about the nature of the solutions, including Lanczos⁴² and Davidson⁴³ diagonalisers. These have pushed the boundary of system size such that the largest FCI calculation known to date is of N_2 in a space of $\approx 9.7 \cdot 10^9$ determinants⁴⁴. This is many orders of magnitude smaller than the systems that may be approached with more approximate methods, or with FCIQMC. The application of these methods tends to be constrained by the memory requirements of needing to store at least two vectors of coefficients the size of the overall Hilbert space.

FCI results are extremely useful for providing both numerical and performance benchmarks for further development work.

FCIQMC are affected.

*This overestimates actual calculation size. Application of spatial symmetry, k-point symmetry spin and magnetic quantum number eigenfunctions radically cuts the size of the relevant Hilbert space. It still scales in the same way.

1.5.3 Coupled cluster approach

In order for Configuration Interaction calculations to be size consistent, all of the Slater determinants which can be constructed out of the Hartree–Fock orbitals (both occupied and virtual) must be considered. The Coupled Cluster Approach^{17,45} (CCA) is an attempt to provide an approximation to FCI that is size consistent.

The arbitrary n^{th} order excitation operator,

$$\hat{T}_n = \frac{1}{(n!)^2} \sum_{i_1 \dots i_n} \sum_{a_1 \dots a_n} c_{i_1 i_2 \dots i_n}^{a_1 a_2 \dots a_n} a_{a_1}^\dagger a_{a_2}^\dagger \dots a_{a_n}^\dagger a_{i_1} a_{i_2} \dots a_{i_n},$$

removes electrons from orbitals $\{i_n\}$, and places them instead into orbitals $\{a_n\}$, while combining this with an unknown associated, coefficient. Some selection of these operators are combined into the cluster operator

$$\hat{T} = \hat{T}_1 + \hat{T}_2 + \hat{T}_3 + \dots \quad (1.5)$$

which provides a range of excitations with associated coefficients relative to the reference determinant. The overall wavefunction is then constructed by application of the ansatz

$$\begin{aligned} |\Psi\rangle &= e^{\hat{T}} |D_0\rangle \\ &\approx (1 + \hat{T} + \hat{T}^2 + \dots) |D_0\rangle. \end{aligned}$$

The \hat{T}^2 and higher terms include excitations from combinations of multiple terms in the cluster operator, with multiple coefficients combined. As such, if the cluster operator in equation 1.5 is truncated to contain a subset of the relevant terms, the overall expression approximates contributions from higher-level excitations than are directly included through combinations of excitations with the associated products of their coefficients — i.e. the terms become *coupled*. As a consequence, the equations are now non-linear and cannot be solved by a straightforward diagonalisation.

Writing the Schrödinger equation using the Coupled Cluster wavefunction gives

$$\hat{H} |\Psi\rangle = \hat{H} e^{\hat{T}} |D_0\rangle = E e^{\hat{T}} |D_0\rangle.$$

Multiplying by $e^{-\hat{T}}$ and projecting onto the entire set of excited determinants generates a set of coupled equations,

$$\begin{aligned}\langle D_0 | e^{-\hat{T}} \hat{H} e^{\hat{T}} | D_0 \rangle &= E \\ \langle D_{i_1 \dots}^{a_1 \dots} | e^{-\hat{T}} \hat{H} e^{\hat{T}} | D_0 \rangle &= 0\end{aligned}$$

which can be solved for the coefficients $\{c\}$ by a variety of means.

Different levels of theory can be obtained by various truncations of the sum in the definition of the cluster operator (equation 1.5). If only the double excitations, \hat{T}_2 , are included, the theory is labelled CCD for ‘‘Coupled Cluster Doubles’’. Similarly, CCSD also includes single excitations in \hat{T}_1 , and CCSDT and CCSDTQ respectively include the triple- and quadruple-excitations although these are only realistically applicable to extremely small systems. CCSD(T) extends CCSD by approximating the effect of triples using perturbation theory, and is one of the most well known highly accurate methods available. In the limit where all terms up to \hat{T}_N are included in the cluster operator, then all terms in the FCI expansion explicitly appear in the ansatz above, and the method produces the same results as FCI, although less efficiently.

Coupled Cluster is an inherently single-reference method, in that the ansatz above is applied only to one reference determinant. Multiple-reference versions have been developed, although they increase the complexity substantially. Unlike FCI, Coupled Cluster is not variational, although modified methods to ensure this do exist⁴⁶. As a consequence of the non-linearity, a very large amount of information-mixing is present in any implementation, and as a result Coupled Cluster methods tend to be extremely difficult to parallelise over large computational hardware.

Recent work by Thom⁴⁷ has recast CCSD into a stochastic form, similar in both derivation and implementation to FCIQMC. This permits the memory efficiency of FCIQMC to be applied to the coupled cluster approach.

1.5.4 Monte–Carlo approaches

Most development in quantum chemical methods has focussed on increasingly clever (and correspondingly complicated) means to perform calculations deterministically in given basis sets, largely using extremely efficient matrix manipula-

tions. Monte–Carlo approaches, such as Variational Monte–Carlo⁴⁸, Path Integral Monte–Carlo⁴⁹, Coupled Cluster Monte–Carlo⁴⁷, Green’s Function Monte–Carlo⁵⁰, Møller–Plesset Monte–Carlo⁵¹, Auxilliary Field Monte–Carlo^{52–54}, Density Matrix Monte–Carlo⁵⁵ and others, represent a very different strand of thought, based on repeatedly considering the effects of stochastic changes and obtaining results through statistical analysis of output data. These approaches have various benefits.

Reduction in size of solution representation

In deterministic schemes, the size of the representation of the solution tends to be dependent on the size of the Hilbert space being considered. For Monte–Carlo representations, the size of the representation depends on the statistical accuracy desired, and in particular the number of ‘particles’ or ‘walkers’ that are going to be simulated in lieu of a wavefunction description.

Focussing of computational effort

In general, deterministic methods expend the same computational effort on relatively insignificant regions of the Hamiltonian matrix and the corresponding components of the wavefunction as they do on those that are critically important.

In Monte–Carlo approaches, computational effort is focussed on the regions that contribute the most, by concentrating either the representation (particle) density or the sampling in relevant locations. This causes computational effort to be used more sparingly.

Systematic (statistical) improvability

As a result of output data being obtained statistically, the accuracy may be improved systematically by either increasing the density of sampling (number of particles), or running the calculation for longer to improve the statistical accuracy of the results.

Parallelisability

In general Monte–Carlo methods represent systems with discrete particles that behave according to given rules such that the majority of the simulation effort for each of these particles is independent. Consequently, the algorithms tend to parallelise efficiently.

For the purposes of this thesis, the most illustrative commonly used Monte–Carlo method is Diffusion Monte–Carlo^{56–58} (DMC). In a similar way to FCIQMC, Diffusion Monte–Carlo attempts to find the long imaginary time limit of integrating the imaginary time Schrödinger equation (see section 2.2.1). This limit corresponds to the lowest eigenfunction of the Hamiltonian.

After substituting $t = i\tau$ into the Schrödinger equation (equation 1.1) and inserting a ‘trial’ energy offset term, E_T ,

$$\frac{\partial \Psi}{\partial \tau} = -(\hat{H} - E_T)\Psi$$

in atomic units. DMC treats this equation as a diffusion equation in imaginary time, and attempts to solve this using Monte–Carlo integration.

The amplitude of a trial wavefunction in real space is represented by the density of discrete ‘walkers’, each with its own position \mathbf{R} , distributed throughout the available space. These walkers move through real space by a series of random walks which are selected according to the kinetic energy component of the Hamiltonian, and created or destroyed by the potential component ($V - E_T$). E_T is adjusted throughout the simulation to control the number of walkers, and will equal the energy of the ground state when the distribution of walkers has converged to represent the ground state wavefunction.

The wavefunction converged on by DMC will be the lowest energy state available under the boundary conditions. As electrons are fermions, one of these conditions is the Pauli principle^{23,24} — this states that the wavefunction must be antisymmetric with respect to exchange of any two electrons. However, DMC has no means to represent negative values of the wavefunction or to ensure that a nodal-structure exists throughout the space to generate this antisymmetry. As such, the solution converged on will be the *bosonic* solution — one that is much lower in energy and does not correspond to physical behaviour of electrons. It is not possible to simply introduce negative walkers into the system as these will operate independently

leading to a second copy of the same wavefunction! This is known as the *fermion sign problem*.

The antisymmetry requirement may be dealt with using the fixed node approximation^{57,59}. This imposes a previously obtained nodal structure onto the evolving wavefunction, providing the additional boundary condition to enforce antisymmetry. The ground state may now be obtained if walkers are forbidden from moving between the regions demarcated by the fixed nodes. However, the accuracy of the energies obtained is highly dependent on the precise positioning of these nodes, which are difficult to predict *a priori*. Errors introduced due to inaccuracies in nodal surfaces generated using other forms of theory are difficult to systematically reduce.

The real strength of DMC is that, due to representing the wavefunction in real space, the local inter-electronic ‘cusp’ behaviour is well represented. As a consequence, DMC calculations have been some of the most accurate ones performed using established quantum chemical methods, and have been used to predict ground state energies, ionisation energies and similar properties to within ‘chemical accuracy’ (approx $1 \text{ kcal mol}^{-1} \approx 1.6 \text{ m}E_{\text{h}}$ ⁶⁰).

1.6 Aims for further development

Considerable effort has gone into development of a Monte–Carlo based approach to performing FCI calculations, known as Full Configuration Interaction Quantum Monte–Carlo. This method aims to take the beneficial aspects of Diffusion Monte–Carlo, namely its convergence after a stochastic integration of the imaginary time Schrödinger equation and the representation of the wavefunction out of many discrete particles, and to apply them in a method making use of the discrete antisymmetrised basis sets associated with FCI, recasting the nature of the fermion sign problem into something much more tractable.

This thesis is primarily concerned with the role of spin in FCIQMC calculations, considering how spin is involved in convergence and how it can be used to manipulate the behaviour of the method. This requires some modifications to the existing FCIQMC algorithms.

2 Full Configuration Quantum Monte–Carlo

In this section two independent derivations of the Full Configuration Interaction Quantum Monte–Carlo (FCIQMC) scheme are presented, along with the limitations of their applicability. Specifically, the derivation and algorithm published by Booth et al.¹⁵ is insufficient to generalise FCIQMC to cases where the Hilbert space is spanned by a non-orthogonal basis set.

The derivations of FCIQMC result in fairly general iterations schemes. The stochastic implementation of these schemes has a great deal of flexibility, and many different choices may be made. A discussion of the nature of discretisation of the representation of the wavefunction and the stochastic spawning process leads into an exposition of the ‘standard’ algorithm — which has been used in the majority of published work so far. Some alternative choices are useful in the following work and will be discussed as required.

An alternative FCIQMC scheme, re-interpreted to ask the question “when is the next particle to be spawned” rather than the usual question of “how many particles are spawned in this unit of imaginary time” is explored in its own chapter (chapter 7) later.

2.1 What is FCIQMC?

The *Full Configuration Interaction Quantum Monte–Carlo* (FCIQMC) method, recently developed by Booth, Thom and Alavi^{15,61}, stochastically integrates the imaginary time Schrödinger equation to obtain a representation of the FCI wavefunction. Although the FCI wavefunction is never instantaneously represented, it

is efficiently dynamically sampled to obtain the energy (or in principle the expectation values of any other operator^{62,63}) to arbitrary accuracy.

This approach shares much in common with Diffusion Monte–Carlo formulations (see section 1.5.4). In particular;

- The wavefunction is constructed out of particles, or walkers*, each of which carry a miniscule proportion of the overall magnitude.
- The wavefunction is evolved stochastically, according to rules which modulate the distribution of these particles.
- Answers are obtained to a statistical level of accuracy, which is dependent on the number of particles in the system, the length of time allowed for equilibration and the number of iterations that output data are collected over.

The substantive difference between DMC and FCIQMC is in the handling of the fermion sign problem. In DMC, this is approached by making a fundamentally uncontrolled approximation — the nodal structure is fixed according to a prior reference calculation[†] — which introduces a source of error into the final wavefunction that is not systematically improvable.

It is no surprise that simulations get stuck on this hurdle. Troyer et al.⁶⁵ have demonstrated that the fermion sign problem is in the category of non-deterministic polynomial complete decision problems (NP complete). As a consequence, any method which reduces the fermion sign problem to a methodology which scales polynomially would in principle provide a solution to all NP complete problems and imply that $NP = P$. This is generally believed not to be true.

FCIQMC represents the wavefunction as a linear sum of intrinsically antisymmetrised basis functions, and integrates the imaginary time Schrödinger equation in this space rather than real space. This prevents the collapse of the wave-

*The Monte–Carlo community is unable to agree on consistent terminology for the ‘particles’ used in simulations. These have often been labelled as ‘walkers’¹⁵ in a similar way to other similar Monte–Carlo methods, but this seems inappropriate for FCIQMC as they do not move. The term ‘psips’ has been suggested by some,^{57,64} to indicate ‘psi-particles’, but has gained little traction. ‘Spawners’ would be a fairly accurate description. I have chosen to use the term ‘particles’ to be both generic and avoid generating a preconception as to the dynamics within the system.

[†]Work has been carried out on developing DMC methods which work to improve this nodal surface. This does not, however, remove the fundamental issues associated with having a nodal surface.

function to the bosonic solution, and recasts the fermion sign problem into an attempt to determine the relative sign-structure of the coefficients of the given basis functions. This eliminates uncontrolled errors in the wavefunction at the cost of restoring exponential scaling of the computational cost with basis set size — a property of NP complete problems. Even when using the initiator approximation (see section 2.4.4) the overall computational scaling is exponential. The initiator approximation reduces the exponent, but it does not reduce the complexity to polynomial.

Once a series of propagation equations are developed for coefficients within this antisymmetrised Hilbert space, the coefficients are then represented by an ensemble of particles. The efficiency of the stochastic representation of the dynamics of these particles is a result of the *shape* of the Hilbert space — it is both very highly connected and extremely local. Each site in the Hilbert space is connected to only a few thousand others, and yet (in the same way as in the concept of “six degrees of separation” in human population dynamics) it is possible to move from a point in the space to any other point in $\frac{N}{2}$ steps, where N is the number of electrons in the system.

In considering FCIQMC, it is important to consider the consequences of manipulating the antisymmetrised basis set, the ways particles can be represented and the trade-off between memory usage, computational time and statistical noise in implementations.

2.2 Derivation of FCIQMC

Within FCIQMC an ensemble of particles is evolved stochastically according to a series of simple rules. These rules simulate the evolution of a set of coefficients, describing the wavefunction in terms of a basis set. Two different routes to obtaining functional rules are shown below. Note that the dynamics modelled in the two approaches are subtly different, to avoid requiring application of the inverted overlap matrix when using non-orthogonal basis functions.

2.2.1 Orthogonal FCIQMC (Booth et al.)

In common with Diffusion Monte–Carlo, FCIQMC is derived from the imaginary time Schrödinger equation,

$$\frac{\partial \Psi}{\partial \tau} = -(\hat{H} - E_S)\Psi, \quad (2.1)$$

where E_S is an additional energy offset. Integrating this with respect to imaginary time, τ , gives an iterable relationship,

$$\Psi(\tau + \delta\tau) \propto e^{-\delta\tau(\hat{H} - E_S)}\Psi(\tau). \quad (2.2)$$

Writing the wavefunction as a linear combination of the eigenfunctions of the Hamiltonian, ϕ_i , with eigenvalues E_i ,

$$\begin{aligned} \Psi(\tau) &\propto e^{-\tau(\hat{H} - E_S)}\Psi(\tau = 0) \\ &= \sum_i c_i e^{-\tau(E_i - E_S)}\phi_i. \end{aligned} \quad (2.3)$$

By noting that $E_i \geq E_0 \forall i$, where E_0 is the lowest energy eigenvalue, it is clear that the components of Ψ which do not correspond to the lowest eigenvalue decay exponentially with τ more rapidly than ϕ_0 . As such the long time limit of the projection,

$$\Psi_0 \propto \lim_{\tau \rightarrow \infty} e^{-\tau(\hat{H} - E_S)}\Psi(\tau = 0),$$

reveals the ground state.

It should be noted that if $E_S = E_0$, then the coefficient of ϕ_0 no longer decays, whilst those of the other eigenfunctions continue to do so, and a pure ground state wavefunction is obtained. If $E_S < E_0$ the overall amplitude will decline, and if $E_S > E_0$ it will grow, while the structure of the wavefunction continues to converge. This provides a great deal of control over the simulation and provides a measure of the energy, see section 2.3.2.

Generally the correlation energy, $E_C = E_0 - E_{\text{ref}}$, is desired. As a consequence, the substitution $E_S = E_{\text{ref}} + E_C$ is made.

If the time-step, $\delta\tau$, is small, equation 2.2 can be approximated by the first-order

Taylor expansion

$$\Psi(\tau + \delta\tau) \propto (1 - \delta\tau(\hat{H} - E_{\text{ref}} - E_C))\Psi(\tau). \quad (2.4)$$

At this point, the wavefunction is expressed as a linear combination of orthogonal basis functions $\{D_{\mathbf{j}}\}$,

$$\sum_{\mathbf{j}} c_{\mathbf{j},\tau+\delta\tau} |D_{\mathbf{j}}\rangle = (1 - \delta\tau(\hat{H} - E_{\text{ref}} - E_S)) \sum_{\mathbf{j}} c_{\mathbf{j},\tau} |D_{\mathbf{j}}\rangle, \quad (2.5)$$

from which the coefficient of each basis function can be projected out by integrating over the relevant $|D_{\mathbf{i}}\rangle$;

$$\begin{aligned} c_{\mathbf{i},\tau+\delta\tau} &= c_{\mathbf{i},\tau}(1 - \delta\tau(H_{\mathbf{ii}} - E_{\text{ref}} - E_S)) - \delta\tau \sum_{\substack{\mathbf{j} \neq \mathbf{i} \\ \mathbf{j} \leftarrow \mathbf{i}}} H_{\mathbf{ij}} c_{\mathbf{j},\tau} \\ \Delta c_{\mathbf{i},\tau \rightarrow \tau+\delta\tau} &= -\delta\tau(H_{\mathbf{ii}} - E_{\text{ref}} - E_S)c_{\mathbf{i},\tau} - \delta\tau \sum_{\substack{\mathbf{j} \neq \mathbf{i} \\ \mathbf{j} \leftarrow \mathbf{i}}} H_{\mathbf{ij}} c_{\mathbf{j},\tau}, \end{aligned}$$

where $H_{\mathbf{ij}} = \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle$.

For the purposes of FCIQMC, the above iterable step is considered as two separate processes,

$$\textit{Spawning} \quad c_{\mathbf{i},\tau} - \delta\tau \sum_{\substack{\mathbf{j} \leftarrow \mathbf{i} \\ \mathbf{j} \neq \mathbf{i}}} K_{\mathbf{ij}} c_{\mathbf{j},\tau} \longrightarrow c_{\mathbf{i},\tau+\delta\tau}, \quad \text{and} \quad (2.6a)$$

$$\textit{Death} \quad c_{\mathbf{i},\tau} - \delta\tau(K_{\mathbf{ii}} - E_S)c_{\mathbf{i},\tau} \longrightarrow c_{\mathbf{i},\tau+\delta\tau}, \quad (2.6b)$$

where $K_{\mathbf{ij}} = H_{\mathbf{ij}} - E_{\text{ref}}\delta_{\mathbf{ij}}$. $\mathbf{j} \leftarrow \mathbf{i}$ indicates the sites \mathbf{j} that are connected to the site \mathbf{i} , such that the Hamiltonian matrix element $H_{\mathbf{ij}} \neq 0$.

2.2.2 Non-orthogonal FCIQMC

The spin-projected determinants, introduced as a basis set in section 3.3.5 form an *over-complete* and therefore *non-orthogonal* basis. In addition, although they can be normalised, a derivation that does not assume this is preferable. The previous derivation of FCIQMC breaks down when being used with non-orthogonal basis functions, as projecting out the coefficients of each basis function (from equation 2.5) is no longer clean. As such, the derivation by Thom⁴⁷ intended for

Coupled Cluster Quantum Monte–Carlo (which was in turn inspired by FCIQMC), is adapted back to FCIQMC. Note that the dynamics derived here are subtly different to those for canonical FCIQMC. The same dynamics could be achieved by applying the inverse overlap matrix, which would avoid the consequences of the energy entering the off-diagonal terms and some of the observed pathological behaviour. The complexity of doing this is avoided in this derivation.

Starting from the Schrödinger equation for the converged CI wavefunction,

$$(\hat{H} - E) |\Psi_{\text{CI}}\rangle = 0,$$

it is clear that the operator $[1 - \delta\tau(\hat{H} - E)]$, where $\delta\tau$ is some positive real number, will project Ψ_{CI} onto itself,

$$(1 - \delta\tau(\hat{H} - E)) |\Psi_{\text{CI}}\rangle = |\Psi_{\text{CI}}\rangle,$$

and has the same eigenfunctions. This (converged) CI wavefunction is expressed as a linear combination of an arbitrary set of basis functions, $\{F_j\}$,

$$|\Psi_{\text{CI}}\rangle = \sum_{\mathbf{j}} C_{\mathbf{j}} |F_{\mathbf{j}}\rangle,$$

and integrated across an arbitrary basis function $F_{\mathbf{i}}$ giving

$$\sum_{\mathbf{j}} C_{\mathbf{j}} \langle F_{\mathbf{i}} | F_{\mathbf{j}} \rangle - \delta\tau \sum_{\mathbf{j}} C_{\mathbf{j}} \langle F_{\mathbf{i}} | \hat{H} - E | F_{\mathbf{j}} \rangle = \sum_{\mathbf{j}} C_{\mathbf{j}} \langle F_{\mathbf{i}} | F_{\mathbf{j}} \rangle \quad (2.7)$$

Subtracting all of the overlap terms, $C_{\mathbf{j}} \langle F_{\mathbf{i}} | F_{\mathbf{j}} \rangle$, where $\mathbf{i} \neq \mathbf{j}$, from both sides gives

$$C_{\mathbf{i}} \langle F_{\mathbf{i}} | F_{\mathbf{i}} \rangle - \delta\tau \langle F_{\mathbf{i}} | \hat{H} - E | F_{\mathbf{i}} \rangle C_{\mathbf{i}} - \delta\tau \sum_{\substack{\mathbf{j} \leftarrow \mathbf{i} \\ \mathbf{j} \neq \mathbf{i}}} \langle F_{\mathbf{i}} | \hat{H} - E | F_{\mathbf{j}} \rangle C_{\mathbf{j}} = C_{\mathbf{i}} \langle F_{\mathbf{i}} | F_{\mathbf{i}} \rangle. \quad (2.8)$$

$\mathbf{j} \leftarrow \mathbf{i}$ indicates the sites \mathbf{j} that are connected to the site \mathbf{i} , such that the Hamiltonian matrix element, $H_{\mathbf{ij}} - ES_{\mathbf{ij}} \neq 0$. This can be considered as an iterable process, such that the converged coefficients are iterated and regenerate themselves.

If instead the set of coefficients $\{c_{\mathbf{j}}\}$, which are not necessarily converged to the

FCI wavefunction, are considered then the iteration step

$$c_{\mathbf{i}} - \delta\tau \frac{(H_{\mathbf{ii}} - ES_{\mathbf{ii}})c_{\mathbf{i}} + \sum_{\substack{\mathbf{j} \leftarrow \mathbf{i} \\ \mathbf{j} \neq \mathbf{i}}} (H_{\mathbf{ij}} - ES_{\mathbf{ij}})c_{\mathbf{j}}}{S_{\mathbf{ii}}} \longrightarrow c_{\mathbf{i}} \quad (2.9)$$

defines a new process for evolution of the coefficients, where the Hamiltonian and overlap matrix elements are given by $H_{\mathbf{ij}} = \langle F_{\mathbf{i}} | \hat{H} | F_{\mathbf{j}} \rangle$ and $S_{\mathbf{ij}} = \langle F_{\mathbf{i}} | F_{\mathbf{j}} \rangle$. This process may be split into two component processes,

$$\textit{Spawning} \quad c_{\mathbf{i},\tau} - \delta\tau \frac{\sum_{\substack{\mathbf{j} \leftarrow \mathbf{i} \\ \mathbf{j} \neq \mathbf{i}}} (K_{\mathbf{ij}} - E_C S_{\mathbf{ij}})c_{\mathbf{j},\tau}}{S_{\mathbf{ii}}} \longrightarrow c_{\mathbf{i},\tau+\delta\tau}, \text{ and} \quad (2.10a)$$

$$\textit{Death} \quad c_{\mathbf{j},\tau} - \delta\tau \frac{(K_{\mathbf{ii}} - E_C S_{\mathbf{ii}})c_{\mathbf{i},\tau}}{S_{\mathbf{ii}}} \longrightarrow c_{\mathbf{i},\tau+\delta\tau}, \quad (2.10b)$$

where $K_{\mathbf{ij}} = H_{\mathbf{ij}} - E_{\text{ref}} S_{\mathbf{ij}}$, for implementation in a similar way to the original FCIQMC scheme. It is worth noting that the original FCIQMC scheme is obtained as a special case of this process where $S_{\mathbf{ii}} = \delta_{\mathbf{ij}}$. The overlap matrix terms enter the equations as a consequence of projecting the coefficients back onto the basis set being used. In cases where this is not done, such as in Auxilliary-Field Quantum Monte-Carlo⁵⁴, this term does not arise and the associated problems such as non-local death disappear.

Although equation 2.8 demonstrates that if this propagation is applied to the CI wavefunction, it will project out the same wavefunction, it does not demonstrate that the system will converge given a different trial wavefunction. If the converged CI wavefunction is not used, equation 2.7 does not represent the propagation step well, as the coefficients on the left would refer to imaginary time τ , and on the right $\tau + \delta\tau$, and the subtraction taken to reach equation 2.8 would not be valid.

As a consequence, if, and only if, simulations of this scheme converge, the obtained wavefunction is an eigenfunction of the Hamiltonian. By analogy with the orthogonal scheme outlined earlier, this will be the ground state wavefunction for the system. However, it is entirely possible for the scheme to be unstable even with a small value of $\delta\tau$. One of the consequences of this is explored in section 3.4.2.

2.3 Energy estimators

The primary expectation value of interest for an FCIQMC simulation is the energy of the resultant wavefunction. The expectation value of that is defined as

$$E(\tau) = \frac{\langle \Psi(\tau) | \hat{H} | \Psi(\tau) \rangle}{\langle \Psi(\tau) | \Psi(\tau) \rangle}$$

for any instantaneous wavefunction. For the ground state, this metric has the advantage of being variational, and so can provide an accurate upper bound on the energy. However, this estimator is biased due to correlations in the terms summed, especially in the denominator.⁵⁵ There is no straightforward way around this. This estimator also suffers from scaling problems similar to calculation of an accurate instantaneous value of the total spin expectation value (see section 3.2.1) in that the connectivity of every occupied site to every other occupied site through the Hamiltonian operator must be considered — and this scales extremely poorly.

A number of different metrics for the energy of the system are used in FCIQMC; the *projected energy*, E_P , the *shift*, E_S and weighted variants of those. These are discussed below. A blocking error analysis^{66,67} is performed on the output of these energy estimators to obtain the final output results.

Additionally, work by Cleland and Overy⁶⁸ permits stochastic evaluation of the one- and two-body density matrices during the output data collection phase of the calculation. It is possible to obtain the energy of the wavefunction from this metric. This is not discussed further in this report.

2.3.1 The projected energy, E_P

Consider the projection of the time-independent Schrödinger equation onto any arbitrary reference state;

$$E_P(\tau) = \frac{\langle \alpha | \hat{H} | \Psi(\tau) \rangle}{\langle \alpha | \Psi(\tau) \rangle}.$$

It is important to note that this metric is no longer variational. Its quality as a metric of the energy is strongly dependent on the choice of the projection state α — essentially the better the overlap between the wavefunction and the state α , i.e. the magnitude of the denominator of the expressed fraction, the better the estimate

and the more stable it will be to stochastic fluctuations in the wavefunction through a calculation.

In most cases the reference state, $F_{\mathbf{0}}$, is used for projection, such that

$$E_P = \frac{\langle F_{\mathbf{0}} | \hat{H} | \Psi \rangle}{\langle F_{\mathbf{0}} | \Psi \rangle} = \frac{\sum_{\mathbf{i} \leftarrow \mathbf{0}} H_{\mathbf{0i}} c_{\mathbf{i}}}{\sum_{\mathbf{i} \leftarrow \mathbf{0}} S_{\mathbf{0i}} c_{\mathbf{i}}}, \quad (2.11)$$

where $\mathbf{i} \leftarrow \mathbf{0}$ indicates that only those sites with non-zero Hamiltonian or overlap terms connecting the to the reference site need to be considered for the summation. In the case of orthonormal basis functions being used, the projected energy simplifies to

$$E_P = \frac{\sum_{\mathbf{i} \leftarrow \mathbf{0}} H_{\mathbf{0i}} c_{\mathbf{i}}}{c_{\mathbf{0}}}.$$

In the long imaginary time limit, an estimate of the ground state energy of the system (as opposed to the instantaneous energy of an ensemble of particles) may be obtained by averaging both the numerator and the denominator independently;

$$\langle E_P \rangle = \frac{\langle \langle \alpha | \hat{H} | \Psi(\tau) \rangle \rangle}{\langle \langle \alpha | \Psi(\tau) \rangle \rangle}.$$

It is important to note that the individual E_P values should not be averaged themselves, as this will introduce biases to the final value and cause issues determining statistical accuracy as inevitably there are correlations between the numerator and denominator in the expression. Additionally, it is worth noting that the shift energy estimator (see section 2.3.2) can introduce a bias⁶⁹ that affects the projected energy estimator if it is permitted to vary too rapidly

If the reference is well chosen for the problem under consideration, as is generally the case if the Hartree–Fock ground state is used as the reference, then it will carry a substantial proportion of the overall wavefunction and the denominator will be large, minimising statistical noise in this metric. Conversely, if the reference is poor, either due to the system being highly multi-configurational or the Hartree–Fock solution simply being poor, this metric for the energy will be poor. In some cases an alternative site becomes the most strongly weighted within a simulation, and in these cases the reference site should be changed on the fly.

It is possible to overcome these limitations by use of a more sophisticated trial wavefunction α . This has been of particular use in the semi-stochastic representation of FCIQMC⁷⁰, although this comes at additional computational cost.

It has been demonstrated by Booth⁷¹ that the average of the non-variational projected energy used within FCIQMC converges much more rapidly to the correct CI energy than the variational energy or any average of it. This is largely due to the variational energy requiring a much better instantaneous representation of the wavefunction to give a reasonable energies, and that it is skewed by being always higher than the true energy.

2.3.2 The shift, E_S

If the ensemble of coefficients, $\{c_i\}$, associated with a converged CI wavefunction is iterated according to the propagation equation defined in equation 2.9, the difference in the coefficient c_i between two steps is given by

$$\Delta c_i = -\delta\tau \frac{\sum_j c_j (H_{ij} - E S_{ij})}{S_{ii}}.$$

Considering the time-independent Schrödinger equation,

$$\begin{aligned} (\hat{H} - E_{\text{exact}}) |\Psi_{\text{CI}}\rangle &= 0, \\ \sum_j (\hat{H} - E_{\text{exact}}) |F_j\rangle c_j &= 0, \end{aligned}$$

then adding zero to the expression for Δc_i yields

$$\begin{aligned} \Delta c_i &= -\delta\tau \frac{\sum_j c_j (H_{ij} - E S_{ij})}{S_{ii}} + \overbrace{\delta\tau \frac{\sum_j c_j (H_{ij} - E_{\text{exact}} S_{ij})}{S_{ii}}}_{\text{zero}} \\ &= \frac{E - E_{\text{exact}}}{S_{ii}} \delta\tau \sum_j c_j S_{ij}. \end{aligned}$$

If the energy term appearing in this expression is considered separately for the diagonal and off-diagonal terms, then

$$\Delta c_i = \delta\tau (E_{\text{diag}} - E_{\text{exact}}) c_i + \frac{(E_{\text{off-diag}} - E_{\text{exact}})}{S_{ii}} \delta\tau \sum_{j \neq i} c_j S_{ij}. \quad (2.12)$$

If orthonormal basis functions are being used, such that $S_{ij} = \delta_{ij}$, then

$$\Delta c_i = \delta\tau(E_{\text{diag}} - E_{\text{exact}})c_i, \quad (2.13a)$$

$$|\Psi(\tau + \delta\tau)\rangle = e^{(E_{\text{diag}} - E_{\text{exact}})\delta\tau} |\Psi(\tau)\rangle. \quad (2.13b)$$

As a consequence, the value of E_{diag} does not influence the convergence of the wavefunction, it merely scales all of the coefficients by the factor $e^{(E_{\text{diag}} - E_{\text{exact}})\delta\tau}$ per iteration. This means that it may be used as a variable parameter to investigate the ground state energy of the system, and is labeled the ‘shift’, E_S ; $E_{\text{diag}} \equiv E_S$. Once there are enough particles in the simulation, the shift is (slowly) varied to maintain the L_1 -norm of the wavefunction, $N_w(\tau) = \sum_i |c_i(\tau)|$, by a weighted measure of the fractional change in the L1-norm between iterations,

$$E_S(\tau) = E_S(\tau - n_{\text{up}}\delta\tau) - \frac{\zeta}{n_{\text{up}}\delta\tau} \ln \frac{N_w(\tau)}{N_w(\tau - n_{\text{up}}\delta\tau)},$$

where n_{up} indicates the number of iterations between each occasion the ‘shift’ is updated, and ζ is a parameter indicating the strength of the adjustments and hence controlling the relationship between the smoothness of the variations (which is desirable because of its effects controlling the simulation dynamics) and its effectiveness at maintaining the L1-norm. Additionally, if this parameter is too large it may introduce a bias to the projected energy.⁶⁹

The value that E_S settles to provides a good estimate of the energy of the system, although it is worth noting that it has a potential bias,⁵⁵ especially in the low particle count limit. If the particle distribution attempts to explore a favourable area it will cause a (short-term) growth in particle count which this shift modifier will work to suppress, and similarly it will actively feed particle growth as the distribution moves into less favourable portions of the space.

2.3.3 The weighted energy estimator

If non-orthogonal basis functions are being used, i.e. $S_{ij} \neq \delta_{ij}$, the energy modifies the off-diagonal terms in the spawning as well as the diagonal ones. As demonstrated in figure 2.1, it is necessary for this value to be correct in order for the wavefunction to evolve correctly, and for the projected energy to tend to the correct value. As such, it is important that this value be independent of the shift

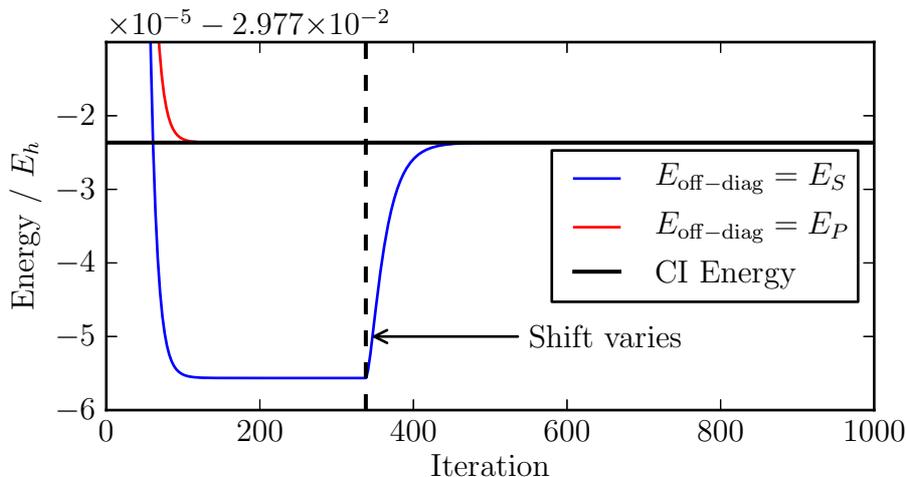


Figure 2.1: A plot of the projected energy for a model system (He_2 in a minimal basis set) where equation 2.9 is iterated exactly, relative to the exact CI energy obtained by diagonalisation. In blue is the projected energy if $E_{\text{off-diag}} \equiv E_S$ as described in equation 2.12. It can be seen that until the shift is allowed to vary (and converge to the correct energy) the projected energy fails to converge to the correct value. If a weighted average of the projected energy is used instead, $E_{\text{off-diag}} \equiv E_{\text{weighted}}$, as in the red line, then convergence occurs directly. This demonstrates that the off-diagonal terms require a good estimate of the energy of the system even outside of variable-shift mode, and that the off-diagonal terms and the shift should be decoupled.

— the structure of the wavefunction must converge correctly whether the shift is being used to encourage particle growth or to maintain particle number.

An estimator based on the projected energy is, conveniently, independent of the shift. However, the projected energy can vary extremely rapidly, especially early in the calculation when the denominator in equation 2.11 is small. It is important that the estimator for the off-diagonal elements varies smoothly and slowly to maintain sane behaviour in the particle dynamics whilst averaging to the correct value in the long time limit. This is more important than producing good statistics for analysing this particular value. As such, an exponentially weighted average of the projected energy is used in this term,

$$E_{\text{off-diag}}(\tau + \delta\tau) = E_{\text{weighted}}(\tau + \delta\tau) = \beta E_P(\tau) + (1 - \beta) E_{\text{weighted}}(\tau), \quad (2.14)$$

where β is a parameter which determines how strongly variations in E_P affect the

averaged value.

Once this estimator is used, the cross terms in equation 2.12 (on average) disappear as $E_{\text{weighted}} \approx E_{\text{exact}}$, and the scaling behaviour described in equation 2.13 is recovered. This enables the shift to be used in the same manner when using non-orthogonal basis sets.

2.4 The ‘standard’ FCIQMC algorithm

The propagation relations given in the above derivations (see equations 2.6 & 2.10) give a very general definition of the processes involved. A naïve implementation where the ensemble of coefficients $\{c_i\}$ are stored and iterated according to these relationships requires at least one full vector of the full size of the Hilbert space and the enumeration and calculation of all non-zero Hamiltonian matrix elements.

The benefits of FCIQMC are related to maintaining a sparse representation of the wavefunction and having an efficient stochastic scheme for sampling the Hamiltonian. There are a large number of different choices which can be made in the implementation of the processes described above. In the section below the range of choices made for the ‘standard’ implementation of the algorithm are discussed, along with some of the alternatives that have been used or considered.

2.4.1 Description, granularity and chance

The primary memory usage benefits of FCIQMC are related to the representation of the evolving wavefunction in memory – in particular the representation of the coefficients associated with the basis functions in use. The wavefunction is extremely sparse within the Hilbert space, and the majority of the non-zero coefficients are very small.

It is important that the representation, and the dynamics, are sufficient that the coefficients up to the double excitations of the reference determinant, i.e. those which contribute to the projected measure of the energy (see section 2.3.1), are represented accurately by their long term averages, while effectively reflecting the sparsity of the wavefunction.

There are a number of significant choices, and correspondingly a great deal of flexibility, in the ways that moves are made throughout the Hilbert space, and in how both spawned and stored particles are represented. These can be generally summarised by noting that increasing the granularity generally makes the calculation faster while decreasing the statistical accuracy of the answer — and that a balance needs to be struck between these competing priorities.

The two primary means of increasing the granularity of the calculation are;

Cutoffs

The sparsity of the wavefunction is represented by the existence of a minimum coefficient size, c_{\min} . Magnitudes smaller than this are stochastically rounded, such that

$$c \rightarrow \begin{cases} c & \text{if } |c| > c_{\min} \\ \text{sgn}(c) \times c_{\min} & \text{with probability } \frac{|c|}{c_{\min}} \text{ if } |c| < c_{\min} \\ 0 & \text{otherwise.} \end{cases}$$

The value of the coefficient is represented *on average* across many iterations. The computational and storage cost savings of representing a particle on a few iterations, rather than a smaller particle on every iteration, are substantial.

Discretisation

Similarly, for ease of representation, the coefficient may be discretised into a multiple of a granularity constant, ω , such that

$$c \rightarrow \begin{cases} \omega \times \left\lceil \frac{c}{\omega} \right\rceil & \text{with probability } \left(\frac{c}{\omega} - \left\lfloor \frac{c}{\omega} \right\rfloor \right) \\ \omega \times \left\lfloor \frac{c}{\omega} \right\rfloor & \text{otherwise.} \end{cases}$$

ω would normally be equal to one, such that the coefficients are restricted to integers and any fractional parts stochastically rounded up or down. If the representation of the coefficients is not being artificially discretised, i.e. $\omega = 0$, then the above process (with its implicit division by zero) should not be applied.

2.4.1.1 Discretisation of represented wavefunction

In an algorithm intending to represent the wavefunction with an ensemble of different particles, it would be extremely awkward to restrict the representation of the wavefunction such that $\langle \Psi | \Psi \rangle = 1$. In practice there are two stages to any FCIQMC calculation;

1. *Initial growth.* The overall amplitude of the wavefunction is allowed to grow to encourage the representation to spread out into the entire Hilbert space.
2. *Collection of data.* The overall L1-norm, $N_w = \sum_{\mathbf{i}} |c_{\mathbf{i}}|$, of the wavefunction is constrained by permitting the value of E_S to vary (see section 2.3.2). If the wavefunction is converged, then this (on average) maintains the wavefunction amplitude $\sum_{\mathbf{i}} |c_{\mathbf{i}}|^2$.

The primary output of an FCIQMC calculation is an estimator of the correlation energy of the system. Calculation of this using the projected energy (see section 2.3.1) is sensitive to the quality of the representation of the coefficients. The more granular the representation, the noisier the energy estimators will become as the distribution and the output values will fluctuate more in order to maintain the correct averages.

Conversely, a discrete representation of the coefficients is computationally more efficient, especially as a very large proportion of sites are correctly represented by amplitudes lower than the granularity will permit — and as a consequence only a small fraction will be occupied at any particular instant. This drastically reduces the computational cost per iteration, and the total memory requirements of the instantaneous representation of the wavefunction, at the cost of decreasing the statistical accuracy of the resultant energies.

The projected energy (see section 2.3.1) is dependent only on the coefficients of the reference site and the direct connections to it via the Hamiltonian. No sites which are more than double excitations away from the reference site contribute directly to this energy metric. As such, in order to maintain the best possible description of the wavefunction for the projected energy, whilst using memory efficient dynamics, it is helpful to use a split representation where the sites up to double or triple excitations from the reference site are represented using floating point coefficients, whereas the remainder of the Hilbert space is discretised. This can lead to up to 3 orders of magnitude of reduction in statistical noise⁷² compared

to a fully discrete representation.

2.4.1.2 Discretisation of spawning

There is a choice of how spawning is to be carried out. The step described in equation 2.10a is to be carried out stochastically for each of the sites in the occupied list. There are a number of choices to be made based on a tradeoff between the computational cost of each iteration and the associated stochastic noise.

Selection of target sites

It is necessary that equation 2.10a is faithfully reproduced *on average* throughout the calculation. For implementation, this expression is reversed,

$$-\delta\tau \frac{(K_{ij} - ES_{ij})c_j}{S_{ii}} \longrightarrow \Delta c_i \quad \forall i \leftarrow j, i \neq j,$$

such that for each occupied site, \mathbf{j} , the component of the value Δc_i corresponding to the connection $\mathbf{i} \leftarrow \mathbf{j}$ is calculated, and the overall sum in equation 2.10a is calculated when all of these components are combined.

It is not necessary for the full ensemble of connections to be considered each iteration. Any stochastically determined subset, $\{\mathbf{k}\}$, may be considered, provided that each of the terms is adjusted by the likelihood of it being selected in a given iteration, $p_{\text{gen}}(\mathbf{k}|\mathbf{j})$, such that

$$-\delta\tau \frac{(K_{jk} - ES_{jk})c_j}{p_{\text{gen}}(\mathbf{k}|\mathbf{j})} \longrightarrow \Delta c_k \quad \forall \mathbf{k} \in \{\mathbf{k}\} \in \{\mathbf{i} \leftarrow \mathbf{j}, \mathbf{i} \neq \mathbf{j}\}. \quad (2.15)$$

In the ‘standard’ case described below, the set $\{\mathbf{k}\}$ is reduced to containing only one element, i.e. by only choosing one target, \mathbf{k} , for each spawning event. In the case of same-spatial structure spawnings, this is not always the best choice (see section 5.5).

Subdivision of spawning

In equation 2.10a the magnitude of the connections made is multiplied by the coefficient c_j . Reducing the stochastic noise can be achieved by focussing more attention on the most significant (highly weighted) regions of the space. As such, the number of spawning attempts may be made in proportion with the magnitude of the coefficients, such that the expression in equation 2.15

is approximated to

$$\frac{|c_j|}{\gamma} \times \left[-\delta\tau\gamma \frac{(K_{\mathbf{j}\mathbf{k}} - ES_{\mathbf{j}\mathbf{k}}) \operatorname{sgn} c_j}{p_{\text{gen}}(\mathbf{k}|\mathbf{j})} \longrightarrow \Delta c_{\mathbf{k}} \quad \forall \mathbf{k} \in \{\mathbf{k}\} \in \{\mathbf{i} \leftarrow \mathbf{j}, \mathbf{i} \neq \mathbf{j}\} \right]. \quad (2.16)$$

where γ indicates the weight given to each spawning step. If γ does not divide c_j exactly, then $\frac{|c_j|}{\gamma}$ is rounded up to the nearest integer with probability equal to the fractional part, or down otherwise. This results in $\frac{|c_j|}{\gamma}$ independent spawning steps being performed. In the most frequently used implementation, $\gamma = 1$ and FCIQMC may be considered to use discrete, unit, signed particles as the source of spawns.

Discretisation of spawned particles

Discretisation and truncation of spawned particles serves the purpose of reducing the number of particles which have to be combined with the main particle list during annihilation (see section 2.4.3). This is significant as this list of spawned particles contributes the majority of data to communicate between computational nodes on a parallel machine. The parameter, c_{\min} , is labelled $n_{\text{s,min}}$ when applied to spawned particles.

The more aggressively cutoffs are applied as described above, the lower the computational cost — but at a cost of stochastic noise.

2.4.1.3 Sensible choices of granularisation and cutoffs

Canonical FCIQMC

The majority of FCIQMC calculations discussed in the literature make use of the original formulation of FCIQMC, such that all particles are considered to be the same, with integer weight. This can be considered as setting $c_{\min} = n_{\text{s,min}} = \omega = \gamma = 1$ for both storage and spawning. For each integer particle located on a site, one spawning attempt is made to a (non-uniformly) randomly selected connected site during each iteration.

Real-coefficients

As demonstrated by Overy⁷², FCIQMC may be implemented by relaxing c_{\min} and ω for both the stored and spawned particles. Generally c_{\min} is set extremely low for spawned particles, such that small incremental spawning effects occur slowly and only essentially zero sized spawns are trimmed to

reduce communication overhead, whereas c_{\min} is maintained for particle storage to maintain the sparsity of the representation. ω is set to zero, such that above the threshold all detail in the representation is retained. This results in a substantial improvement in the statistical quality of the results.

Partial real-coefficients

The vast majority of the statistical benefit of the real-coefficient scheme is obtained by representing the ‘core’ of the wavefunction — those close to the reference site which contribute directly to the projected energy, or strongly influence those sites that do. A computational saving may be made by only representing those sites with real coefficients, and using the canonical FCIQMC scheme for the remainder.

Semi-stochastic FCIQMC

Umrigar et al. have demonstrated⁷⁰ that treating the ‘core’ of the wavefunction deterministically is beneficial. The region of the space to consider in this way should be obtained through lower cost calculations, but is generally highly populated with particles — as such treating it exactly does not incur the substantial overhead it would in the sparsely occupied regions of the space. This region of exact integration of the imaginary time Schrödinger equation is then coupled with an FCIQMC scheme for the remainder of the space.

The use of CSFs will require some further modifications of the granularisation of FCIQMC calculations. In particular the subsets of connected sites spawned to in each spawning attempt are modified, as described in section 5.5. The majority of calculations in this thesis are performed using the partial real-coefficients scheme.

2.4.2 Initialisation (trial wavefunction)

The derivation of FCIQMC gives an algorithm that may be used to iterate the ensemble of particles. The quality of the initial distribution of particles, representing a trial-wavefunction, determines how far the simulation has to evolve to converge to a sensible result, and therefore how much computational effort is required.

There are two primary approaches used for initialising calculations

Single particle

A single particle (or a few particles, to reduce the likelihood of all the particles in the simulation dying) are placed on one site, and the shift is set so that the total particle population grows. The computational cost grows proportionally with the total number of particles in the system⁷³ and, especially if the initiator approximation is used, this encourages the population to grow coherently. This avoids computational time being unnecessarily spent on generating a coherence from an incoherent wavefunction through annihilation.

Approximate wavefunction

Alternatively, an approximate wavefunction calculated by a cheaper method,* may be used to initialise the wavefunction. It is assumed that the majority of the wavefunction weight described by these approximate solutions is sign-coherent and as such it can avoid the unnecessary computational cost associated with developing this sign structure through annihilation.

Previous FCIQMC calculation

The dumped output of a previous FCIQMC calculation may be loaded, and the calculation resumed either to collect additional data to improve the the statistical quality of the output, or to grow the number of particles.

2.4.3 The ‘standard’ algorithm

Once an initial distribution is established, the ensemble is iterated by repeatedly applying the following three processes for each timestep, $\delta\tau$. This is also demonstrated in figure 2.2.

1. Spawning

For each occupied determinant D_i , the stored coefficient is discretised into units of γ , with the remainder rounded to $\lceil c_i \rceil$ with probability equal to the fractional part of the coefficient if it is non integer, otherwise $\lfloor c_i \rfloor$. For each of these particles, a coupled site, D_j , is selected with a normalised probability $p_{\text{gen}}(\mathbf{j}|\mathbf{i})$ and a particle with weight n_s is created and stored in the spawned

*A variety of alternative methods may be used. In particular either a Complete Active Space SCF (CASSCF) calculation is used to populate a subset of the Hilbert space, or particles are distributed according to the MP2 wavefunction.

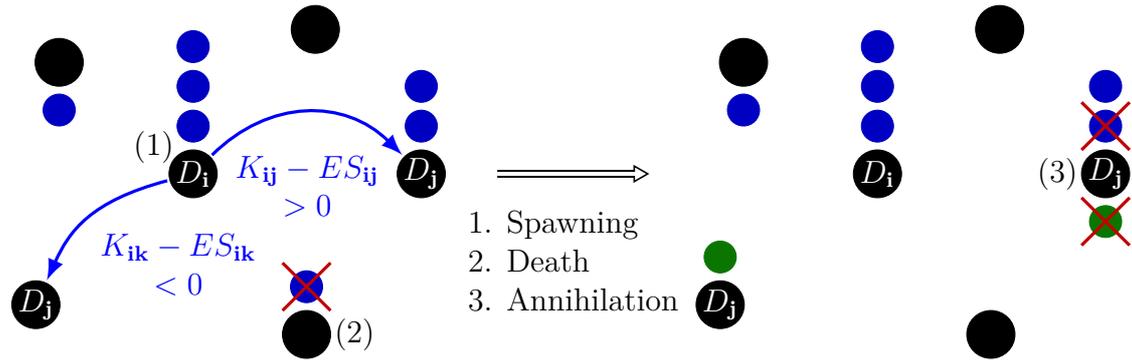


Figure 2.2: A diagrammatic representation of Full Configuration Interaction Quantum Monte-Carlo (FCIQMC). This shows the three components of the algorithm; (1) spawning, where the sign of the spawned particles depends on the sign of the connection, $K_{ij} - ES_{ij}$, (2) death, where particles on a site are destroyed depending on the diagonal matrix element, and (3) annihilation, where particles with opposite signs are removed from the simulation.

list, where

$$n_s(\mathbf{j}|\mathbf{i}) = -\delta\tau\gamma \frac{K_{ij} - E_{\text{weighted}}S_{ij}}{p_{\text{gen}}(\mathbf{j}|\mathbf{i})S_{jj}}. \quad (2.17)$$

If the magnitude of this is lower than a cutoff $n_{s,\min}$ then it is rounded to an amplitude of $n_{s,\min}$ with probability $\frac{|n_s|}{n_{s,\min}}$ or discarded otherwise. If purely integer coefficients are being used, n_s is stochastically rounded to an integer value as described above.

2. (Diagonal) death

Reduce the magnitude of the coefficient on each occupied site by

$$n_d(\mathbf{i}) = \delta\tau \frac{K_{ii} - E_S S_{ii}}{S_{ii}} c_i. \quad (2.18)$$

If integer coefficients are being used, then this value should be stochastically rounded to an integer value. If this value is negative, the coefficient is instead augmented (particles are cloned).

3. Annihilation

The list of newly-spawned particles is combined with the list of particles remaining after (diagonal) death. All pairs of particles which occupy the same site but with differing signs are removed from the simulation. After this, remaining particles with coefficients smaller than c_{\min} are rounded as described above. It is at this stage in the algorithm that various other control measures (see the initiator approximation and non-local death in

sections 2.4.4 and 3.4.2) are implemented.

Implementationally, this step is extremely important. When working with parallel computational resources, the communication of spawned particles from their origin to the processor containing any existing particles on the same site is the only step which must be performed synchronously on all processors — and as such its scaling is of critical importance. The use of hash functions to distribute particles efficiently across computational resources is discussed in a great deal more detail in other published work⁷³.

If $S_{ij} = \delta_{ij}$ this algorithm reduces to the orthogonal algorithm which has been previously published. A general overview of the implementation of the standard FCIQMC algorithm is found in figure 2.3.

2.4.4 The initiator approximation

A modification to the above proposed algorithm has been suggested by Cleland Booth and Alavi⁷⁴, known as *Initiator* Full Configuration Interaction Quantum Monte–Carlo (*i*-FCIQMC). This method attempts to approach the problem of low productive wavefunction development in the (very noisy) rapid particle growth at the start of a calculation, and avoid the formation of an annihilation plateau (see section 2.4.5).

The new method introduces an additional *survival criterion* for spawned particles. The occupied basis functions in the space are split into two sets; *initiators*, where the amplitude on that site is larger than a specified threshold such that $|c_i| > n_{\text{init}}$, and *non-initiators*, which includes all other sites. When a new particle is spawned, it survives only if

- its parent is located on an initiator site,
- the target site is already occupied, or
- more than one spawn, with non-initiator parent sites, occurs onto the same site with the same sign on the same iteration.

This modification to the algorithm is based on a single premise, that sites occupied by a large weight of particles are more likely to have the correct sign structure, and as such should be given precedence during spawning. In the many-particle limit, the approximation disappears.

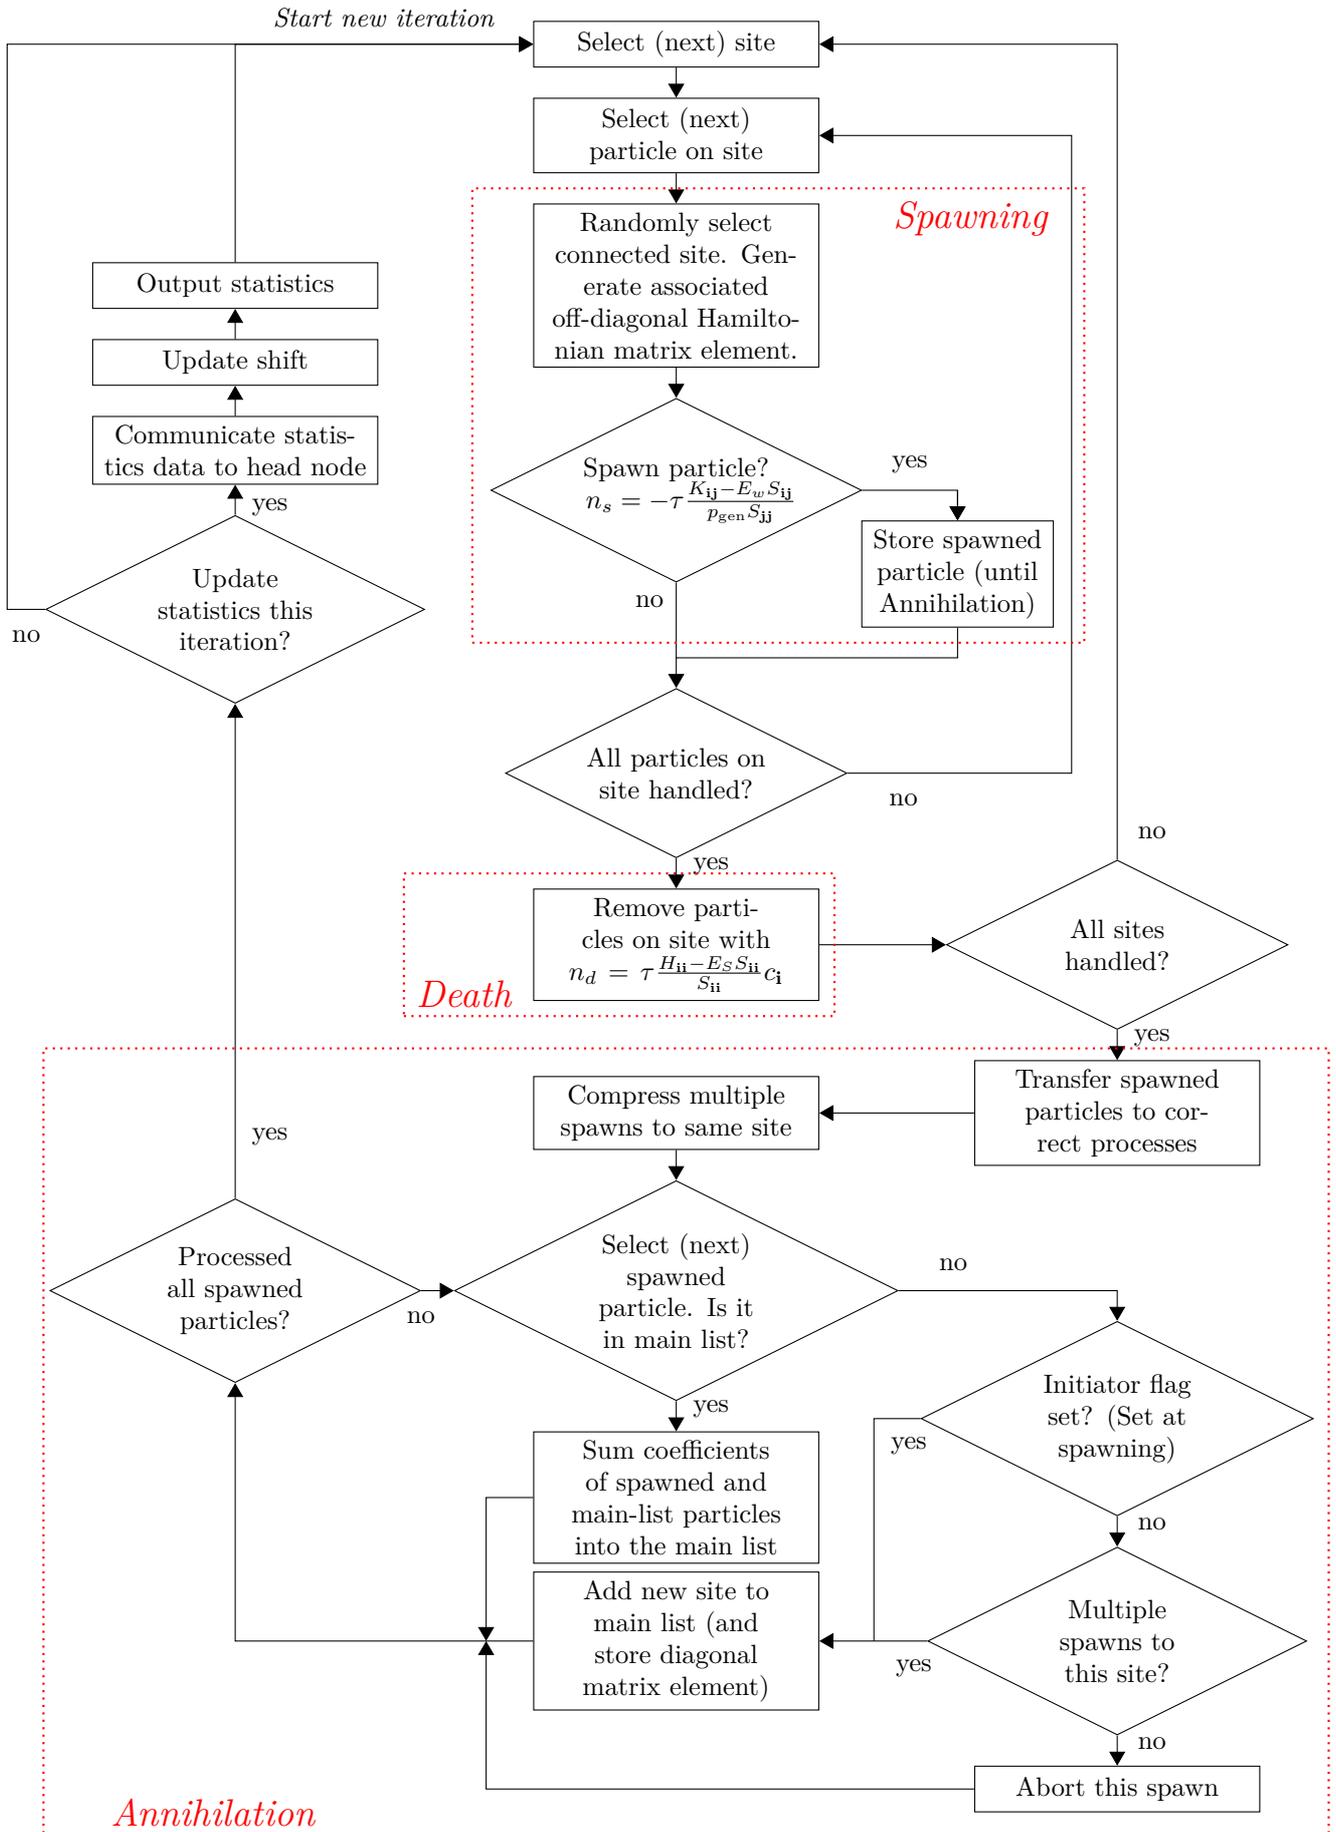


Figure 2.3: Overview of the FCIQMC algorithm, showing spawning, death and annihilation steps of the main iteration loop.

Figure 2.4 demonstrates the primary benefit of the initiator approximation — by favouring growth from sites that are more likely to be sign coherent, the initial incoherent growth of particles that results in the annihilation plateau is suppressed. This results in substantially faster convergence onto the correct wavefunction both in terms of particle count (the wavefunction becomes coherent with fewer particles), and computational time (the area between the annihilation plateau and the trajectory with the initiator approximation is essentially wasted). The growth of particles on the reference site is very similar, in line with the coherent growth of the resultant wavefunction. The number of particles on the reference site in *i*-FCIQMC represent a larger proportion of the total number of particles due to multitudinous low weight spawns into sparsely occupied regions of the Hilbert space being suppressed.

The benefits associated with the initiator approximation come at a cost. Restricting the spawning from insufficiently occupied sites effectively soft-truncates the Hilbert space as a function of the particle count. If there are insufficient particles in the system, the energy will not converge to the correct value, and the restriction on the statistical average of this value is not variational. This introduces an error, dependent on the particle number, known as the *initiator error*. Furthermore, due to the absence of the annihilation plateau, there is no straightforward means to determine that a simulation has sufficient particles, other than running it with many more particles and observing that the convergence does not change.

Prior analysis of the initiator approximation has been carried out in the canonical FCIQMC scheme, with unit particles — in this regime it is clear that if $n_{\text{init}} = 0$ the simulation is equivalent to canonical FCIQMC as all occupied sites will be initiators, and for $n_{\text{init}} = 1$ a substantial threshold has been introduced.

Using real coefficients permits a subtly closer analysis. It appears that the onset of initiator approximation induced behaviour occurs as a phase change — it is sudden. This is illustrated by the two choices in figure 2.4. In the first case, $n_{\text{init}} = 0.9$, which is smaller than the smallest permissible particles size — as such the behaviour is as for canonical FCIQMC. As soon as $n_{\text{init}} = 1.0$, and sites must have more than the minimum occupation by some amount, the annihilation plateau is suppressed. This is despite the minimum spawn size being 0.2 in this simulation. Beyond this value, *i*-FCIQMC is very insensitive to the value of n_{init} , with the choice of this primarily influencing the size of the initiator error.

This is a somewhat unexpected result, but it can be explained by consideration of the initial behaviour of the simulation. In a normal FCIQMC simulation, each of the particles may spawn progeny, which will spawn progeny and so forth. This results in a particle growth rate which is roughly proportional to the number of particles, and is eventually matched by an annihilation rate that instead grows quadratically.⁶⁴ Instead, in an *i*-FCIQMC simulation, as soon as the threshold is above the minimum particle size, the initial set of particles spawned from a site *cannot* successfully spawn progeny. They are only able to spawn once their magnitude increases (even infinitesimally), as a result of further spawning. As such, repeated pairs of spawns are required to generate new structure in the Hilbert space, resulting in a reduction of the spawning progress to more closely match the growth of annihilation.

From this it is reasonable to conclude that sign incoherent spawning in FCIQMC may be controlled by considering an entirely marginal preference for particles with some metric for being more sign coherent.

2.4.5 A typical FCIQMC calculation

A typical FCIQMC simulation contains two stages, particle growth and the constant amplitude stages. These have various characteristics, and differ substantially between normal and initiator FCIQMC calculations. Figure 2.4 shows a canonical FCIQMC simulation and an *i*-FCIQMC simulation for comparison.

Growth phase in FCIQMC

For a typical FCIQMC simulation, the growth phase contains three distinct sections. Firstly the simulation starts with rapid particle growth. As there is no developed sign structure, and few particles across the Hilbert space to annihilate with, the spawning patterns are dominated by the *amplitude* of the Hamiltonian matrix elements, and the simulation heads towards a pseudo-bosonic solution⁶⁴.

Once a sufficient number of particles have been scattered through the Hilbert space, spawning begins to cause annihilation. This phase brings the growth of the total number of particles to a halt in the *annihilation plateau*. During this stage the structure of the wavefunction develops, as noted by the improvements in the projected energy estimator and the growth of particle

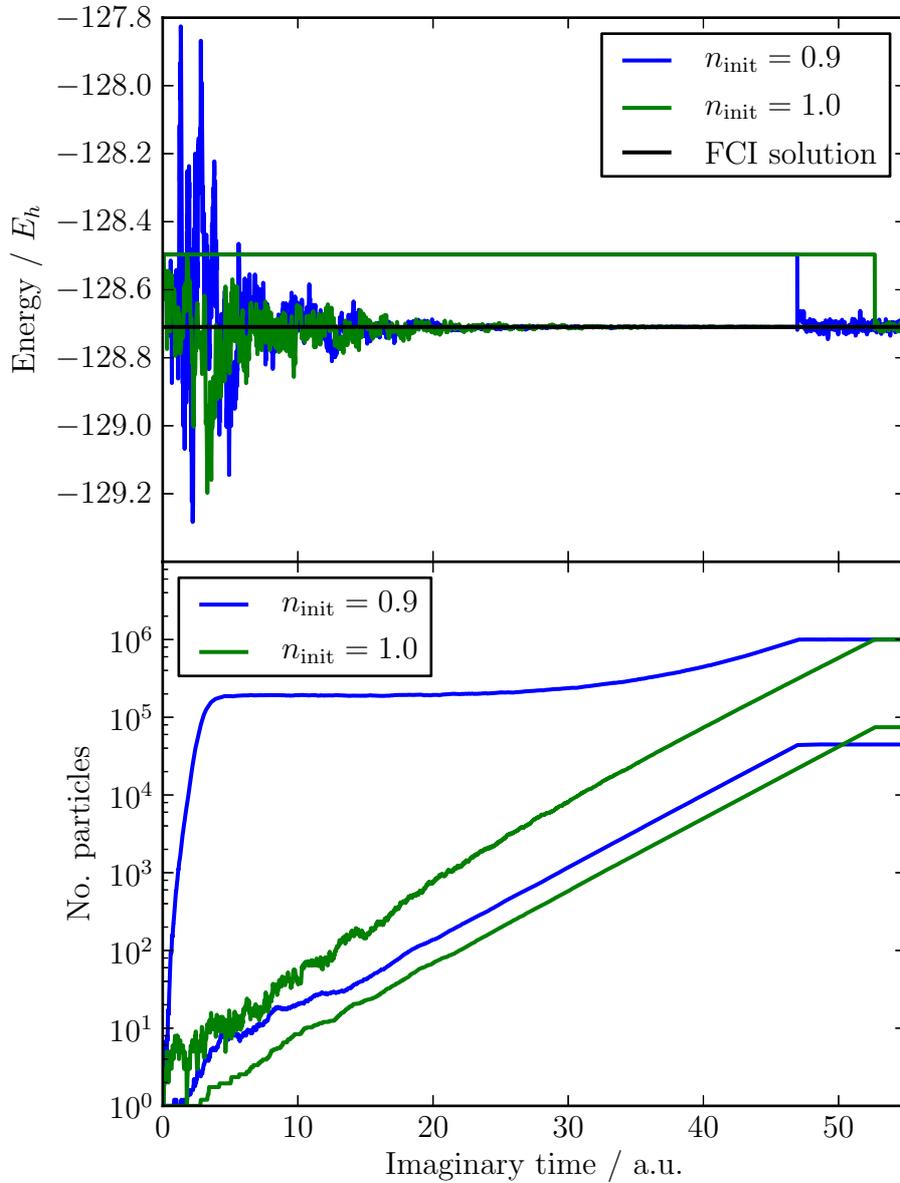


Figure 2.4: The trajectories of two *i*-FCIQMC simulations of Ne in an aug-cc-pVDZ basis set with different values of $n_{\text{init}} = 0.9, 1.0$. The simulations are run with $c_{\text{min}} = 1.0$ and $n_{\text{s,min}} = 0.2$. The top plot shows the projected energy estimator, which is joined by the shift energy estimator when the total number of particles in the simulation is constrained. The lower plot shows the total particle weight in the two simulations, along with the number of particles on the reference site for each simulation. Note that as 0.9 is smaller than the minimum particle size, the simulation with $n_{\text{init}} = 0.9$ is indistinguishable from non-initiator FCIQMC.

weight on the reference site.

At some point, sufficient structure develops in the wavefunction that new spawns are sign-coherent, and build overall particle weight. Once the growth of overall particles becomes roughly parallel to the growth on the reference site on a log plot, i.e. the whole wavefunction is scaling together, the simulation is converged.

Growth phase in *i*-FCIQMC

In *i*-FCIQMC the initial uncontrolled growth of particles is suppressed. Particle growth occurs more smoothly, and more coherently, with a gradual progression to a converged wavefunction with the growth of the total particle weight parallel to the growth on the reference site in the log plot. In keeping with the coherent growth, the projected energy estimator tends towards the correct value much more rapidly, with far less noise.

The extent to which the number of particles needs to be grown to converge the wavefunction well is not immediately obvious. See section 6.1.2 for examples of how to determine when this growth phase has generated enough particles.

Constant amplitude phase

Once the wavefunction is converged, and there are sufficient particles in the simulation, the shift estimator of the energy, E_S , is permitted to vary in a damped fashion to maintain the total number of particles in the system. If the overall and reference particle weights do not both track horizontally on the plot, then the wavefunction is not well converged.

2.5 Issues for consideration

When developing Monte–Carlo algorithms, one of the nice features is that, provided that the deterministic propagation equations are simulated correctly in a statistical limit, the correct answer will be resolved given sufficient particles to represent the solution and enough computational time, and in principle the many possible implementational decisions described above do not matter. However, the trick for algorithmic development is to ensure that “enough computational time” is in practice reasonably short. To that end some specific issues are discussed in

this section that have a bearing on efficient implementation — and on the areas that have not yet been satisfactorily resolved.

Convergence onto states other than the ground state

FCIQMC, by integration of the imaginary time Schrödinger equation, projects out the ground state wavefunction subject to the available basis set. A consequence of this is that FCIQMC is not generally suitable for finding excited states. Some headway can be made by utilising symmetry to exclude low lying excited states.

Other portions of the excitation spectrum can, in principle, be obtained by other related means. These include explicit orthogonalisation schemes, inverse Laplace transformation of the decay of the estimated energy through imaginary time and (real) time dependent modifications of FCIQMC. Work to approach these is ongoing.

Parallelisation

Other than annihilation, which only requires communication of the newly spawned particles between computational processes, FCIQMC processes particles individually and independently. As such, FCIQMC lends itself extremely well to efficient parallelisation.

Current schemes distribute the occupied sites uniformly across across the available MPI processes⁷³. In practice the non-uniformity of occupation of these sites determines the limit of efficient scaling with the number of computational processes. Work is ongoing in collaboration with James Spencer to overcome this limitation and distribute particles efficiently based on both the number of sites and the occupation level of these sites.

Number of particles required

In the canonical expression of FCIQMC, convergence is fairly straightforward to demonstrate — if a simulation has passed through the annihilation plateau, and exponential growth has resumed such that the entire wavefunction is scaling together, then the simulation is converged. This can be seen on a log-plot of particles against imaginary time as the gradient of the total number of particles becomes equal to the gradient of the number of particles on the reference site.

When making use of the initiator approximation, this is no longer so straight-

forward. It is necessary that the exponential scaling criterion as above is fulfilled, but there is no longer a clear metric for how many particles are required. In particular, it is necessary for there to be a certain number of particles in the system before the initiator approximation stops skewing the results by its restrictions on spawning. At present the only means to determine convergence by particle number is to run a series of calculations with different numbers of particles and observe the asymptotic behaviour.*

It is also worth noting that, independent of the convergence of the wavefunction, it is strongly advantageous to grow the simulation until there is a substantial weight of particles (1000-10000 particles) on the reference site. As this coefficient forms the denominator in the projected energy expression, this works to reduce the magnitude of the stochastic fluctuation in the projected energy. Some of the algorithmic trade-offs that influence the stochastic noise dramatically influence the number of particles that are realistically required.

Poor convergence of electronic cusp conditions

FCIQMC, in the same way as FCI, constructs wavefunctions by combining basis functions. In the same manner as a Taylor series, describing local detail in the wavefunction is much more difficult than describing the overall macroscopic ‘shape’ of the solution. As such, while the long-range correlation effects are well described, the convergence of the local electron-electron behaviour, i.e. the electronic ‘cusps’, with respect to the size of the basis set is very slow. DMC does a much better job at this.

Perturbation theory may be used to apply corrections to the FCIQMC wavefunction and inject the effect of explicit formulations of short-range inter-electronic behaviour. The formulations of R12 and F12 theory may be used^{29,36,38}, and these substantially improve the quantitative accuracy of FCIQMC calculations with smaller basis sets.

Low lying excited states

As demonstrated by equation 2.3, the rate of convergence of the wavefunction to the ground state in imaginary time is dependent on the size of the energy gap between the ground state and (especially) the first excited state of the

*It is possible to effectively perform this procedure within one calculation, by performing a calculation with the shift energy parameter set so as to grow the particle weight extremely slowly. This is equivalent to running multiple calculations, although it requires somewhat different error analysis⁷⁵.

system.

The use of symmetry subdivides the Hilbert space and block-diagonalises the Hamiltonian matrix, such that any individual calculation is only considering one block. If a basis set includes symmetry, then this permits selection of specific solutions, as above, but it also permits the exclusion of low-lying excited states that interfere with convergence and as such improve the efficiency of the calculation.

Cost per unit imaginary time

The computational cost of performing an FCIQMC simulation is ultimately dependent on three factors; *a*) the length of imaginary time required, *b*) the computational cost per iteration, and *c*) the imaginary time step, $\delta\tau$. Point (a) is influenced by consideration of excluding excited states, as well as earlier trade-offs between memory usage and stochastic noise. The noisier the simulation, the more data collection will be required to obtain the same statistical accuracy. Point (b) is critically influenced by the number of particles required, and is thus similarly influenced by statistical noise issues. The cost of each iteration is also heavily influenced by the efficiency of excitation generation (see chapter 5) and Hamiltonian matrix element calculation (see chapter 4). The choice of the imaginary timestep, (c), is also influenced by the basis set and the choices regarding dynamics and is discussed in section 2.6.

2.6 The imaginary time step, $\delta\tau$

In equation 2.4 it is assumed that the imaginary time step, $\delta\tau$, is ‘small’. This is similar to saying that the rate of particle production in the spawning step in section 2.4.3 must be low, normally less than 5% of the stored wavefunction magnitude per iteration. The choice of $\delta\tau$ is important as it has very significant consequences for the efficiency of a calculation. There are several restrictions on its value;

Death rate

Particle death is carried out on every iteration. The desired behaviour is a gradual exponential decline in coefficient magnitude on each site, with the fractional decrease given by $f_d = \delta\tau \frac{K_{ii} - E_S S_{ii}}{S_{ii}}$ (see equation 2.18).

If $f_d < 1$, then smooth exponential decay is observed. If $1 < f_d < 2$ then the coefficient on the site is “more than killed”, with its sign inverting and its magnitude decreasing by the fraction $2 - f_d$. If $f_d > 2$ then convergence becomes impossible, as the magnitude of the affected coefficients will increase while their sign will flip between each iteration!

It is important that $\delta\tau$ is chosen such that strictly $f_d < 2$, i.e. $\delta\tau < \frac{2S_{ii}}{K_{ii} - E_S S_{ii}}$, for all sites. For the simulation to be well behaved, $\delta\tau$ should be chosen such that $f_d < 1$, i.e. $\delta\tau < \frac{S_{ii}}{K_{ii} - E_S S_{ii}}$, for all sites with significant, non-fleeting, occupation.

Spawning step

In principle there is no restriction on the magnitude of spawns that occur. In order to minimise statistical noise, the number of particles spawned in any one step should be smaller than the average occupancy of the affected sites. Given that the majority of sites will only ever contain low single-digit occupancy, then the maximum spawning size should be similar.

If the initiator approximation is in use, this requirement becomes even more stringent. If individual spawns are of sufficient size that the occupancy of a target site becomes larger than the initiator threshold in one step, this causes the approximation to cease working correctly. Consider the case of an unoccupied site in a region of the Hilbert space that should contain little amplitude. If one of these sites becomes, inadvertently, an initiator, it is able to continually spawn particles into its surrounding region — and the initiator approximation has suppressed any other sites from being able to spawn in this region, so it takes a long time before annihilation brings this under control.

As a consequence, $\delta\tau$ is normally (dynamically) adjusted to maintain the maximum spawn size as $n_{\text{init}} - 1$ at a maximum. As the maximum spawn size depends on K_{ij} and $p_{\text{gen}}(\mathbf{j}|\mathbf{i})$ as described in equation 2.17, this means the largest spawn is roughly proportional to the largest Hamiltonian matrix element, K_{ij} , and inversely proportional to the smallest generation probability, $p_{\text{gen}}(\mathbf{j}|\mathbf{i})$. This is sometimes a slightly too aggressive criterion, as there can be a few very strong connections well inside the occupied region of the space. The value of $\delta\tau$ may be manually relaxed to permit only a few of these ‘too large’ spawns to occur.

There is one major problem with any method for choosing a globally unique time step in an FCIQMC calculation. In most cases the maximum value will be determined by the diagonal matrix elements corresponding to the highest energy basis functions available — a value which is strongly dependent on the size of the underlying Hartree–Fock basis set. However, very few particles ever occupy these extremely high-lying sites, resulting in the vast majority of computational effort being expended unnecessarily. In chapter 7 a novel approach is explored for minimising this cost, albeit with other implications.

2.7 Basis sets for FCIQMC

The FCIQMC algorithm will converge on a set of coefficients, $\{C_i\}$, such the ground state wavefunction is best represented by

$$|\Psi\rangle = \sum_i C_i |D_i\rangle$$

for a set of basis functions $\{D_i\}$. There are a large range of suitable basis sets available.

The Pauli principle²³ states that fermionic wavefunctions must be anti-symmetric with respect to exchange of any two electrons. The simplest general set of functions obeying this are the Slater determinants,²⁵ constructed from an antisymmetrised linear combination of spin-orbitals, $\{\phi_j^{\alpha_j}\}$,

$$|D_i\rangle = \mathcal{A} \prod_j^{n_{\text{elec}}} \phi_j^{\alpha_j},$$

where the antisymmetriser, $\mathcal{A} = (\sqrt{N!})^{-1} \sum_{\mathbf{p}} (-1)^P \mathbf{P}$, generates a signed sum over all permutations of the electronic (and spin) coordinates. Given a structure with g spatial orbitals doubly occupied (both α and β), and N_o orbitals occupied by one electron of which n_β have $m_s = -\frac{1}{2}$ and n_α have $m_s = \frac{1}{2}$, there are

$$n_{\text{det}}(N_o) = \binom{N_o}{n_\alpha} = \binom{N_o}{n_\beta}$$

ways of choosing the spins α, β to assign to a given choice of spatial orbitals $\{\phi\}$. More generally, a basis function may be written as the product of a spatial and

spin component of the wavefunction

$$|D_{\mathbf{i}}\rangle = \mathcal{A} \left(\prod_j^{n_{\text{elec}}} \phi_j \right) \theta_{\mathbf{i}}(n_{\text{elec}}, M_s), \quad (2.19)$$

where $\theta_{\mathbf{i}}(n_{\text{elec}}, M_s)$ is an eigenfunction of \hat{S}_z for n_{elec} electrons.

The dynamics of particles in FCIQMC are determined by the pattern of connectivity, and the magnitudes of the Hamiltonian matrix elements associated with these connections, throughout the Hilbert space. It is clear that these are strongly dependent on the choice of basis set used — even if exploring the same Hilbert space. As an example, the convergence of a system modelled using plane waves will be different to one using localised orbitals. Similarly, a basis which recognises the symmetry associated with a problem will behave much more happily than otherwise.

2.7.1 Symmetry

The solution to the Schrödinger equation must obey the symmetry of the problem. For each relevant symmetry, each of the eigenfunctions of the Hamiltonian will transform according to one of the irreducible representations of the point group corresponding to the symmetry of the problem.

If the elements of the basis set in use transform according to elements of the relevant symmetry group, then it is straightforward to partition the Hilbert space into those basis functions which transform according to the desired irreducible representation, and those that do not. The Hamiltonian matrix is block diagonal, with the blocks defined by the different irreducible representations, as

$$\left. \begin{aligned} \langle F_{\mathbf{i}} | \hat{H} | F_{\mathbf{j}} \rangle &= 0 \\ \langle F_{\mathbf{i}} | F_{\mathbf{j}} \rangle &= 0 \end{aligned} \right\} \quad \text{if } \Gamma_{\mathbf{i}} \neq \Gamma_{\mathbf{j}},$$

where $\Gamma_{\mathbf{j}}$ is the irreducible representation of basis function $F_{\mathbf{j}}$ under the relevant symmetry group. Where the symmetry group contains degenerate irreducible representations, multiple blocks will be equivalent to each other, e.g. being related to each other by a rotation, with the only impact being to reduce the size of the calculation. Otherwise blocks will correspond to physically distinct symmetries

with different solutions.

In order to make use of this symmetry effectively, the excitation generators used to generate random steps through the Hilbert space need to be written to explicitly take these symmetries into account. Doing so has two major benefits;

Reduced size of Hilbert Space

Reducing the overall size of the Hilbert space has several different advantages, all of which correspond to a general decrease in size of the computational problem;

- The connectivity around each site is reduced, increasing the generation probabilities for each step and increasing the timestep, $\delta\tau$, that may be used (see section 2.6).
- Reducing the size of the space reduces the number of particles needed to effectively explore it.
- The reduced size of the space and connectivity tends to increase the rate of convergence of the wavefunction.

Discrimination between eigenfunctions

FCIQMC converges the wavefunction onto the ground state of the Hamiltonian under the basis set in use. If the effective basis set is restricted by symmetry, the ground state under this symmetry will be found. As such, a series of excited states, distinguished by symmetry, can be calculated.

The basis sets utilised in FCIQMC are constructed as antisymmetrised products of orbitals. The symmetries of the basis functions, $\{F_i\}$, are determined from the properties of these orbitals, according to rules associated with the particular types of symmetries.

Spatial symmetries

For molecular systems, the basis functions in use should transform according to the irreducible representations of the point group in use. The number of irreducible representations contained in the point group, and thus the extent to which the Hilbert space can be subdivided, depends on the symmetry of the molecule. In general the more symmetric a molecule is the more that the space can be subdivided according to spatial symmetry.

If the orbitals, $\{\phi_i\}$, contained in a basis set transform according to the irreducible representations of the applicable symmetry groups, Γ_{ϕ_i} , then the

symmetry of the basis functions $\{F_j\}$ used in the calculation is given by

$$\Gamma_j = \bigotimes_i \Gamma_{\phi_i^{(j)}}.$$

If the point group of the system in question is Abelian then the \otimes operator is equal to its inverse, and implementationally may be performed using an XOR operation.

z-component of the total spin quantum number, M_s

If Slater determinants are used, or any other basis functions which can be expressed as a linear combination of Slater determinants, they have a well defined secondary spin quantum number

$$M_s = \sum_i m_s(i)$$

where $m_s(i)$ is the secondary spin quantum number corresponding to the i -th electron (the spin-projected on an arbitrary z -axis). The total M_s value of all basis functions in use is preserved by the excitation generators.

Magnetic quantum number, M_l

For atomic or linear systems, if the basis set in use specifies (projected) magnetic quantum numbers for each of the orbitals in use, then a well defined (projected) magnetic quantum number

$$M_l = \sum_i m_l(i)$$

defines independent blocks in the Hamiltonian. The excitation generators can ensure that this value is maintained.

Total spin quantum number, S

Although the square of the total spin operator, \hat{S}^2 , commutes with the Hamiltonian operator, and as a consequence any eigenfunction will also be an eigenfunction of \hat{S}^2 , determinantal basis functions are not eigenfunctions of \hat{S}^2 — they do not transform according to irreducible representations of the symmetric group under permutation of the spin indices (see section 3.3). This means that Slater determinants do not have a well defined symmetry label associated with total spin.

In order to block diagonalise the Hamiltonian matrix, linear combinations of \hat{S}_z eigenfunctions (i.e. determinants) must be used as the basis functions. These can be constructed in a number of ways, and are the subject of section 3.3. Further chapters 4 and 5 discuss implementational issues.

Total orbital angular momentum, L

For atomic or linear systems, total orbital angular momentum eigenfunctions can be constructed out of determinantal \hat{L}_z eigenfunctions in a similar way as \hat{S}^2 eigenfunctions. These eigenfunctions are not discussed further here.

Time-reversal symmetries

As discussed in the following section (section 2.7.2), time-reversal symmetry gives restrictions on the coefficients of related determinants within the system. These are not symmetries in the same sense as the others, but they do subdivide the space in a related way.

Other symmetries

The symmetries available ultimately depend on the properties of the system being examined. For example, if the system is modelling a solid using plane waves, the momentum of each basis function, described by the k -vector, needs to be maintained. These symmetries need to be included in the relevant excitation generator for efficient computation.

2.7.2 Time-reversal (and similar) symmetries

For systems with an even number of electrons, every spin state, S , contains a degenerate eigenfunction with an M_s value of zero. A time-reversal symmetry may be applied in this domain, as detailed in Ref. [16], relating the coefficients of spin-coupled pairs of determinants related by flipping the spin associated with each electron. The associated coefficients differ only by a sign-change,

$$C_{I_\alpha J_\beta} = (-1)^{S + \frac{N_\alpha}{2} + 1} C_{J_\alpha I_\beta},$$

where I and J represent the string of orbitals associated with α and β electrons respectively^{76,77}.*

*Note that this corresponds to flipping *all* of the spins. As a consequence, this will put all of the spin orbitals in the closed shell section of the representation out of order. The permutations involved in restoring the normal ordering (section 3.4.1) must be considered in obtaining a

total spin of the system and the number of unpaired electrons. If all orbitals are doubly occupied, there is no associated pair.

A basis of these spin-coupled pairs of determinants, alternatively known as Half-Projected Hartree-Fock (HPHF) functions⁷⁸, $|X_{IJ}\rangle$, may be used,

$$|X_{IJ}\rangle = \begin{cases} |D_{I\alpha J\beta}\rangle & \text{if } I = J \\ \frac{1}{\sqrt{2}} \left[|D_{I\alpha J\beta}\rangle + (-1)^{S + \frac{n_{\text{unpaired}}}{2} + 1} |D_{J\alpha I\beta}\rangle \right] & I > J \end{cases}, \quad (2.20)$$

containing roughly half the number of elements of the underlying determinantal basis. This may be represented implementationally by picking one of the paired determinants as the ‘standard’ representation, and excluding the other from the simulation.

The use of these HPHF functions results in an approximately 3-4 fold⁷³ saving in overall computational cost as a result of a reduction in the size and complexity of the space which *a)* roughly halves the number of particles required to converge the simulation, *b)* allows a roughly doubled timestep to be used, and *c)* reduces the duration of imaginary time required to converge the solution. The confluence of beneficial effects observed strongly indicates that these spin-reversal functions should be used whenever the system has an even number of electrons. It also strongly suggests that using full spin-eigenfunctions as a basis could be beneficial.

2.7.3 Mixed schemes

An FCIQMC simulation will tend towards representing the ground state wavefunction within the Hilbert space spanned by the basis set. As any state which can be represented with CSFs can be represented using the underlying determinantal basis, there is no issue with mixing the two representations in the same calculation. It is critically important to ensure that the entire Hilbert space remains well mapped, and that the boundary between the two representations is well defined.

All of the CSF regimes explored have some scaling issues associated with an increase in the number of unpaired electrons, beyond those generally experienced with Slater determinants. As such it is useful to consider the mixed scheme where sites are represented in CSFs (Kotani-Yamanouchi, Serber or spin-projected) for

sign relationship between the normal determinants involved. Thus care is needed to get the hamiltonian matrix elements in section 4.1.

all spatial structures with $N_{o,\max}$ or fewer unpaired electrons, and in Slater determinants for all other sites.

Making the description dependent on the *spatial* orbitals ensures that there is no overlap between the regions of the space represented in CSFs or in Slater determinants.

There are two elements which *must* be attended to in this mixed scheme. Firstly the Hamiltonian matrix elements between the Slater determinants and the CSFs in use must be easily calculable (section 4.7). In all of the cases considered here, the expansion of the CSFs into a linear combination of determinants is well known, and as such the Hamiltonian matrix elements across this boundary can be calculated in, at worst, $\mathcal{O}(n_{det}(N_{o,\max}))$ time.

Secondly, excitation generation across this boundary needs careful consideration (section 5.6). It is important to not excessively inhibit the operation of the excitation generators within their own regions of the space, where the majority of the excitations occur, but they must generate all connections with non-zero Hamiltonian matrix elements with the correct probability.

3 The role of spin in FCIQMC

The role of spin in FCIQMC simulations is the primary topic of this thesis. This first requires thinking about several different aspects of spin — what is spin physically, and how does this interact with the representation and implicit meanings ascribed to spin in the context of computer simulation. After discussing this, observations are made as to the evolution of the spin associated with the ensembles of particles in an FCIQMC simulation, and as a consequence of this evolution, why using Configurational State Functions (CSFs) might be a useful modification.

A number of different schemes exist for the construction of total spin eigenfunctions out of primitive spin functions, and several are presented, followed by a brief discussion of a few immediate implementational consequences.

Other schemes were considered for projecting out chosen spin components stochastically during FCIQMC calculations. These proved to be ineffectual within FCIQMC calculations and so have not been discussed here.

3.1 Physical and practical views of spin

Spin is a somewhat ephemeral property — although measurable⁷⁹ and well characterised, it is difficult to give a good explanation of what spin actually *is* given that humans habitually live in a very classical world.

Spin is an intrinsic property of matter associated with relativistic quantum mechanics, as presented by Dirac²⁰. Interestingly, it does not appear in non-relativistic quantum mechanics, where the Hamiltonian operator may be considered to be *spin-free*. The effects of spin are observed as an angular momentum, with its associated magnetic moment, whose total magnitude, S , is quantised according to specific rules and for which any measurement will only register the quantised projection, M_s , onto a given (experimental) axis. As the square of the total spin

operator, \hat{S}^2 , and the projected spin operator, \hat{S}_z , commute with the Hamiltonian operator, physical solutions of the Schrödinger equation are eigenfunctions of all three operators.

Slater determinants, as used in Hartree–Fock theory, are constructed as antisymmetrised products of single electron orbitals with defined spin z-projection values, m_s , associated with each electron. As a consequence the z-component of the total spin, $M_s = \sum m_s$, is well defined. These functions are not, however, eigenfunctions of the overall spin operator,⁸⁰ \hat{S}^2 — they are merely the best solutions available within the basis set paradigm. They may be expanded as linear sums of eigenfunctions of the square of the total spin operator, \hat{S}^2 , with differing eigenvalues, S , and the reverse is also true — spin eigenfunctions of chosen S may be constructed out of combinations of Slater determinants.

In this context, it is most important to consider spin from the perspective of the roles that it fulfils, and its consequences for the implementation of calculations.

Spin as a quantum number

The total spin is a quantum number, corresponding with measurable properties of a physical system. In particular, the total spin of a system directly dictates a substantial portion of its magnetic properties.

In addition to the magnetic properties, the different spin states available to a system correspond to physically distinct electronic states of molecules. Electronic transitions can be measured between these states spectroscopically. Understanding the relative properties of different spin states in a molecule, especially those involving complex elements such as transition metals, is important to unlocking their chemical behaviour.

Spin as a symmetry label

From the perspective of an FCI calculation, symmetry is the process of formally block-diagonalising the Hamiltonian matrix. If two particles in an FCIQMC simulation occupy sites with different symmetry labels, they belong to two entirely non-interacting subsets of the simulation that will evolve independently.

The total spin, S , is a property of a spin eigenfunction that behaves in this way. The Hamiltonian matrix element between any two spin eigenfunctions with differing total spin eigenvalues is zero — the Hamiltonian matrix is

block diagonalised. This has several consequences for FCIQMC calculations:

- The total spin, S may be specified as a parameter to a calculation. This permits convergence to states with a given spin even if they are not the global ground state (the ground state of the relevant block of the block-diagonal Hamiltonian is found).
- Consequently, the reference state for the projected energy and the trial wavefunction for initialisation of the calculation must be of the correct spin, or the ground state will not be found.
- The overall size of the Hilbert space is reduced, as only the sites with the correct symmetry need to be considered. This reduces the magnitude of the problem being attempted, and should lead to a reduction in the number of particles required, and the height of the plateau in non-initiator calculations.

Spin eigenfunctions as labels

For a molecular system, the non-relativistic Hamiltonian does not contain any spin dependent terms. Indeed, for an FCIQMC calculation the primary input is a dump file containing one- and two- electron integrals across spatial orbitals from another computational package. Although the Slater–Condon rules (table 4.1) do refer to the relative values of m_s for each electron, it is clear that the overall *specification* of the system is *spin-free* even if the wavefunction is not.

As a consequence, it is possible to formulate the quantum mechanical behaviour of such a system without reference to spin. Each concept in a conventional spin-based formulation is replaced with a corresponding spin-free concept⁸¹. The form of the available spin-free functions depends on the permutational symmetry of the Hamiltonian, and the state is labelled by a particular permutation in the same way that spin-states are constructed within the range permitted by the total spin, acting as a symmetry.

If both spin-free and standard spin-dependent formulations of quantum chemistry are constructed, they may be mapped directly on to each other. A consequence of this is that spin may be viewed as an indicator that modulates the form of functions constructed according to the symmetry of the symmetric group under the constraints of the problem being examined —

and nothing more. The different spin eigenfunctions that are available, and discussed later, are merely labels on the permutational behaviour which may be applied to the spatial components being considered.

This gives a strong indication that group-theoretical approaches to expressing and implementing the behaviour of spin eigenfunctions should be considered seriously. Further work may consider expressions of FCIQMC in an explicitly spin-free formulation at a later date.

3.2 Evolution of spin in FCIQMC

In all FCIQMC calculations the projected spin value, M_s , remains constant through the calculation even if it is not constrained to be so as all matrix elements between Slater determinants with an M_s value other than that of the reference determinant are zero. The same cannot be said of the total spin value, the eigenfunction of the operator \hat{S}^2 .

As the calculation progresses, the wavefunction tends towards an eigenfunction of the Hamiltonian and, as this commutes with the square of the total spin operator, it also tends towards an eigenfunction of \hat{S}^2 . The permitted eigenvalues of \hat{S}^2 are both discrete and well known, with permitted values of $S(S + 1)$ where S is half-integral and $|S| \leq \frac{N_o}{2}$ where N_o is the smallest number of unpaired electrons in any occupied basis function. As such the evolution of this value through the calculation gives a good insight into the progression of the calculation and the extent of convergence.

3.2.1 Calculation of instantaneous spin eigenvalues for Slater determinants

The instantaneous expectation value of \hat{S}^2 may be obtained by

$$\langle \hat{S}^2 \rangle = \frac{\langle \Psi | \hat{S}^2 | \Psi \rangle}{\langle \Psi | \Psi \rangle}.$$

The normalisation factor is already (trivially) calculated by summing the values of $|c_i^2|$ during the main computational loop. The operator can be expanded in terms

of the spin raising and lowering operators, \hat{S}_- , \hat{S}_+ and the z-component of the total spin operator, \hat{S}_z , such that $\hat{S}^2 = \hat{S}_- \hat{S}_+ + \hat{S}_z(\hat{S}_z + 1) = \hat{S}_+ \hat{S}_- + \hat{S}_z(\hat{S}_z - 1)$ ⁸². As a consequence,

$$\begin{aligned} \langle \Psi | \hat{S}^2 | \Psi \rangle &= \sum_{\mathbf{ij}} c_{\mathbf{i}}^* c_{\mathbf{j}} \langle D_{\mathbf{i}} | \hat{S}_- \hat{S}_+ + \hat{S}_z(\hat{S}_z + 1) | D_{\mathbf{j}} \rangle \\ &= \frac{\sum_{\mathbf{ij}} c_{\mathbf{i}}^* c_{\mathbf{j}} \langle D_{\mathbf{i}} | \hat{S}_- \hat{S}_+ | D_{\mathbf{j}} \rangle}{\langle \Psi | \Psi \rangle} + M_s(M_s + 1). \end{aligned} \quad (3.1)$$

The operators \hat{S}_- , \hat{S}_+ are further defined as

$$\begin{aligned} \hat{S}_- &= \sum_m \hat{s}_{-,m} & \hat{s}_{-,m} |\phi_m^\alpha\rangle &= |\phi_m^\beta\rangle & \hat{s}_{-,m} |\phi_m^\beta\rangle &= 0 \\ \hat{S}_+ &= \sum_m \hat{s}_{+,m} & \hat{s}_{+,m} |\phi_m^\alpha\rangle &= 0 & \hat{s}_{+,m} |\phi_m^\beta\rangle &= |\phi_m^\alpha\rangle. \end{aligned}$$

where $s_{+,m}$ and $s_{-,m}$ act on the m^{th} electron in the determinant. Application of the raising operator generates a linear combination of determinants each with one β electron raised to α . Application of the lowering operator then generates a sequence of $n_\beta(n_\alpha + 1)$ determinants. All of the terms where the raising or lowering operator returns zero are dropped. This results in a list of determinants where every possible pair of sites, one with an α spin and one with a β spin, have had their spins swapped. As such

$$\langle D_{\mathbf{i}} | \hat{S}_- \hat{S}_+ | D_{\mathbf{j}} \rangle = \begin{cases} 1 & \text{if } D_{\mathbf{i}} \text{ and } D_{\mathbf{j}} \text{ have the same spatial structure,} \\ & \text{and have either the same spin structure or} \\ & \text{differ only by swapping two spin labels} \\ 0 & \text{otherwise.} \end{cases}$$

To perform the sum in equation 3.1, each determinant in the occupied list is considered. For each determinant the list of connected determinants is generated and each of these is looked up in the occupied list and the relevant value of $c_{\mathbf{i}}^* c_{\mathbf{j}}$ added to the sum. When working on a multi-processor machine, with multiple MPI threads, there are two sensible options;

1. Transmit the determinants to all relevant processors to perform the calculations, or
2. Use a determinant distribution scheme such that the processor that a determinant is located on depends only on the spatial structure.

For small systems the poor load-balancing caused by scheme (2) causes the system to operate inefficiently, whereas for larger systems it is the optimum scheme to reduce communication overhead. It is worth noting that this calculation (clearly) scales badly. Work by Cleland and Overy permits the approximation of the total spin expectation value by calculation of the density matrix within FCIQMC⁶⁸. This method scales much better, and is appropriate for determining the spin of a converged wavefunction within FCIQMC, but is not effective for obtaining an instantaneous value for the total spin expectation value within a calculation, and therefore of little use for analysing the trajectories of FCIQMC simulations with an eye to development.

An instantaneous value of the spin expectation value for the ensemble of particles is not a good estimator for the spin of the converged wavefunction. There are correlations between the numerator and the denominator in the expression calculated, and these would need to be averaged separately to generate an actual estimator. Furthermore, there are correlations between the terms included in each of the numerator and denominator, leading to systematic bias in the same way as the variational energy⁵⁵ (see section 2.3). This is particularly notable if the state being considered has $S = 0$, as the expression defined is positive definite and any averaged values will be systematically skewed. Fortunately, the permitted values for the spin are discrete and well defined, and it is trivially clear which value a wavefunction is converging to. This metric is useful as a measure of the state of convergence, and as an indicator of particle dynamics.

3.2.2 Calculation of instantaneous spin eigenvalues for HPHF functions

Half-Projected Hartree–Fock (HPHF) functions are constructed from pairs of determinants, as described in section 2.7.2, such that

$$|X_i\rangle = \begin{cases} |D_i\rangle & \text{if no unpaired electrons} \\ \frac{1}{\sqrt{2}} [|D_i\rangle \pm |\overline{D}_i\rangle] & \text{otherwise,} \end{cases}$$

where \overline{D}_i is D_i with the spin of all the unpaired electrons flipped. The sign relationship between the two is known (see section 2.7.2). The relevant term in

the sum is now transformed with the change of basis, such that

$$\sum_{ij} c_i^* c_j \langle D_i | \hat{S}_- \hat{S}_+ | D_j \rangle \implies \sum_{ij} c_i^* c_j \langle X_i | \hat{S}_- \hat{S}_+ | X_j \rangle.$$

, where

$$\begin{aligned} \langle X_i | \hat{S}_- \hat{S}_+ | X_j \rangle &= \frac{1}{2} \langle D_i \pm \overline{D}_i | \hat{S}_- \hat{S}_+ | D_j \pm \overline{D}_j \rangle \\ &= \frac{1}{2} \left[\langle D_i | \hat{S}_- \hat{S}_+ | D_j \rangle + \langle \overline{D}_i | \hat{S}_- \hat{S}_+ | \overline{D}_j \rangle \right. \\ &\quad \left. \pm \left(\langle D_i | \hat{S}_- \hat{S}_+ | \overline{D}_j \rangle + \langle \overline{D}_i | \hat{S}_- \hat{S}_+ | D_j \rangle \right) \right] \end{aligned}$$

which by symmetry

$$= \langle D_i | \hat{S}_- \hat{S}_+ | D_j \rangle \pm \langle \overline{D}_i | \hat{S}_- \hat{S}_+ | D_j \rangle.$$

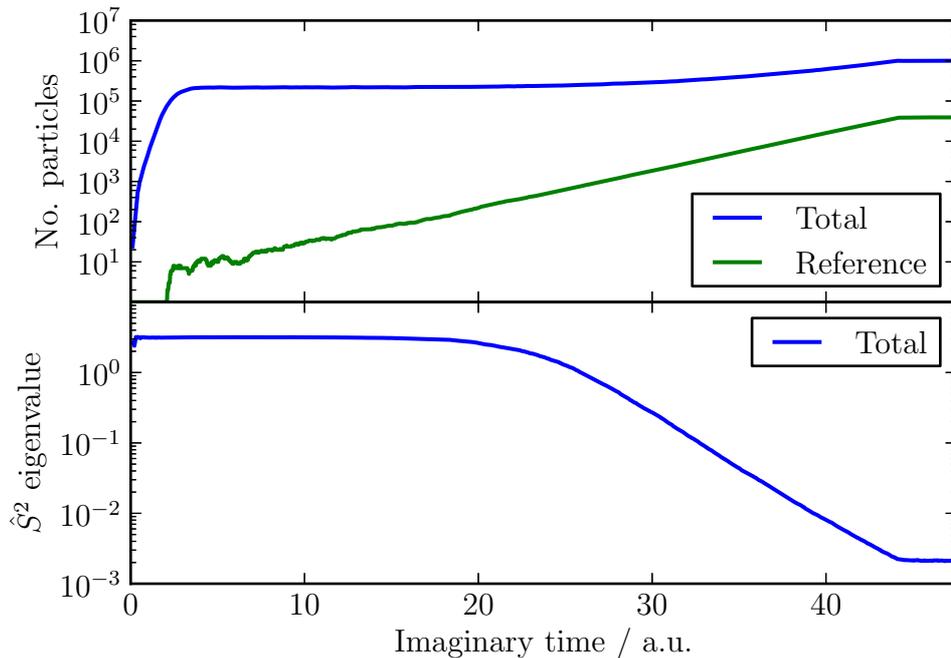
Each of these terms is equal to one or zero as previously discussed. Notably, for spatial structures with more than four unpaired electrons, only one of these terms can ever be non-zero, as application of $\hat{S}_- \hat{S}_+$ generates a string of determinants which differ by at most two electrons. These determinants are then all inverted by the inversion used in HPHF functions, resulting in $N_o - 2$ inversions.

For each HPHF function, X_j , consider the associated determinant D_j . Each of the determinants, D_i , which are connected to it by $\hat{S}_- \hat{S}_+$ are generated, and for each of these, if it is the determinant which is used as the canonical representation for the HPHF, then $c_i^* c_j$ should be included in the sum, otherwise $\pm c_i^* c_j$ should be used (depending on the sign associated with the HPHF).

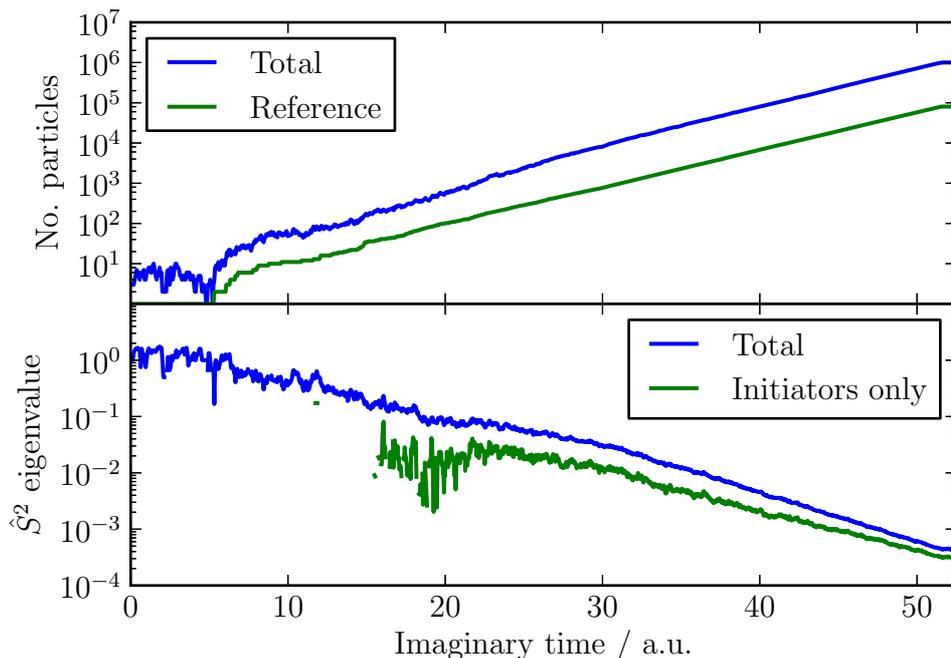
3.2.3 Spin trajectories and the initiator approximation

Examining the trajectories of the expectation value of the \hat{S}^2 operator in FCIQMC simulations makes it obvious that the role of spin ought to be of interest.

In figure 3.1a it can be observed that the sign-incoherent growth of particles at the start of a normal FCIQMC simulation, resulting in the annihilation plateau, also results in substantial disorder in the spin structure of the wavefunction with an expectation value of $\langle \hat{S}^2 \rangle \approx 3.17$. By contrast, the ground state of the Ne atom is known to have a spin of zero — making any non-zero component a good measure



(a) Normal FCIQMC calculation



(b) Initiator FCIQMC calculation

Figure 3.1: Spin trajectories in FCIQMC and *i*-FCIQMC simulations of Ne in an aug-cc-pVDZ basis set. The sign-incoherent growth, resulting in the annihilation plateau, in the FCIQMC simulation corresponds to a large value of the spin expectation value. This begins to fall before the overall number of particles begins to grow again. In *i*-FCIQMC, this unconstrained growth never occurs, and the overall structure is always more sign coherent, with overall much lower values obtained. The spin values associated with only the initiator sites are noticeably lower, corresponding to their better convergence.

of error. As the wavefunction coalesces during the annihilation plateau, indicated by the steady growth of particles on the reference site, the associated spin begins to fall. This occurs before a noticeable growth in the total number of particles occurs.

By contrast, in figure 3.1b *i*-FCIQMC behaves significantly differently. By suppressing the initial incoherent growth of particles, the corresponding growth in spin expectation value is also not observed. The additional coherence of the overall wavefunction may be observed in the fact that the final spin expectation value associated with 1 million particles is nearly an entire order of magnitude smaller than for normal FCIQMC. It is notable that when considered on their own, the ensemble of initiators have an even lower value, indicating that assumptions about their internal sign coherence are likely to be correct.

Figure 3.2 confirms the picture of the trajectory taken to convergence of the spin. These histograms from an *i*-FCIQMC calculation demonstrate the projection of the ensemble of determinants associated with the most highly occupied spatial structure with eight unpaired electrons onto all of the available Kotani-Yamanouchi CSFs (see section 3.3.3) with the given value of M_s . During the initial growth of particles (even with the initiator approximation in use) the growth of components in CSF space is fairly chaotic. These additional components die away as substantial order and symmetry exert themselves in determinantal space resulting in a final solution which scales up as the number of particles grows, containing essentially no components on the CSFs with values of $S \neq 0$, and a resulting structure on the determinantal coefficients.

This thesis attempts to answer the following question:

Does restricting the basis set to prevent spin-incoherent growth assist in the convergence of the wavefunction in the remaining space, or is spin primarily useful as a metric of convergence?

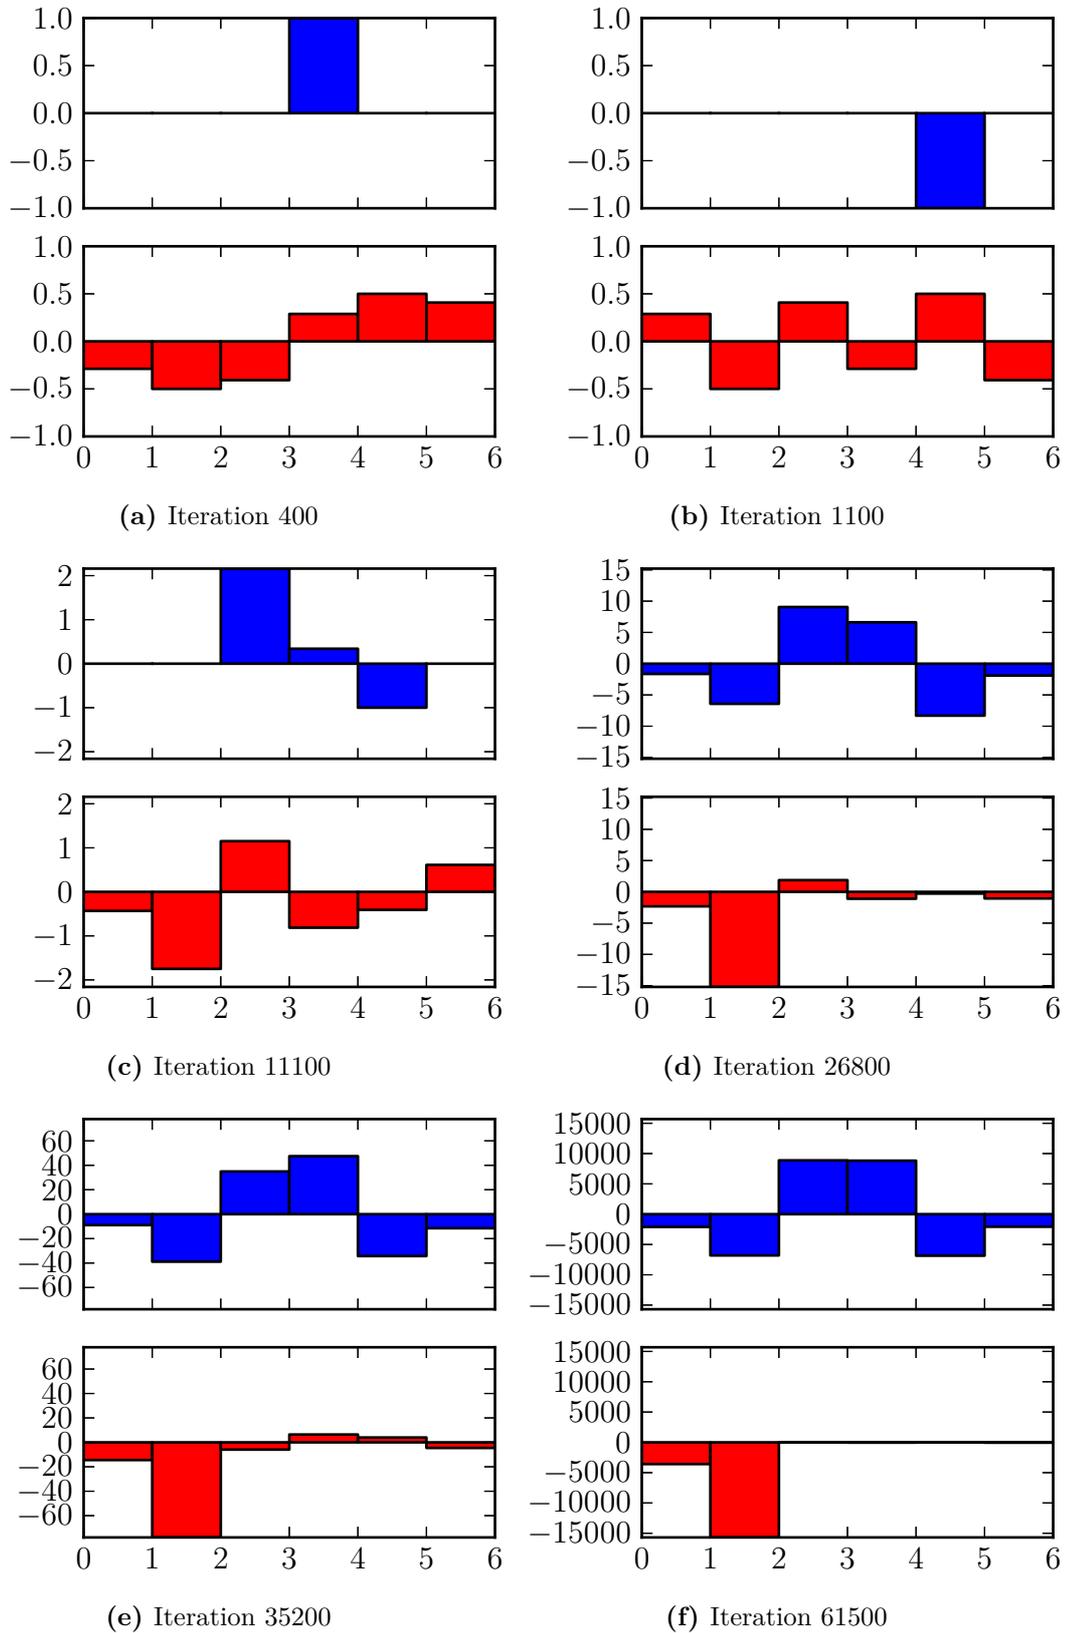


Figure 3.2: Occupation of the determinants in the most highly occupied spatial structure with eight unpaired electrons in an *i*-FCIQMC calculation for Ne in an aug-cc-pVDZ basis set. In blue are the instantaneous occupations of the determinants, and in red are the projections of these coefficients onto the Kotani-Yamanouchi CSFs (see section 3.3.3) with the same spatial structure. The two left-most CSFs are associated with $S = 0$, the remainder are associated with higher spin states. The initial growth of particles is spin-incoherent, and as the simulation progresses, the components projected onto the incorrect spin eigenfunctions tend to zero. Note the resultant symmetry in the distribution of the determinants.

3.3 Spin eigenfunctions

Although \hat{S}_z and \hat{S}^2 commute, eigenfunctions of \hat{S}_z are not generally eigenfunctions of \hat{S}^2 . As the Hamiltonian operator, \hat{H} , is spin-free it commutes with both the operators \hat{S}^2 and \hat{S}_z — thus eigenfunctions of the Hamiltonian are eigenfunctions of both of these spin operators too. As a consequence it is possible to write eigenfunctions of \hat{S}^2 as linear combinations of those of \hat{S}_z .

The elements of the basis set are required to transform under exchange of spin-labels in the same manner as \hat{S}^2 transforms — that is to say they must transform internally as linear combinations of degenerate eigenfunctions,

$$\begin{aligned}\hat{P}\Psi_\alpha &= \epsilon(P)\Psi_\alpha \\ \hat{P}_\sigma\Psi_\alpha &= \sum_\beta \Psi_\beta [P]_{\beta\alpha}^{NS},\end{aligned}$$

where $[P]_{\beta\alpha}^{NS}$ is the element of the permutation representation matrix of the permutation P , as defined in equation 4.8. The eigenfunctions now transform under permutations of the electronic *spin* coordinates according to an irreducible representation of the symmetric group. Under simultaneous permutation of the *spin* and *space* coordinates they still transform according to the antisymmetric representation.

θ_β^{NS} , defined in equation 2.19 obeys the spin symmetry constraint

$$P_\sigma\theta_\alpha^{NS} = \sum_\beta \theta_\beta^{NS} [P]_{\alpha\beta}^{NS}.$$

If this condition is satisfied by expanding the spin eigenfunction as a linear combination of primitive spin functions (eigenfunctions of \hat{S}_z), then these are known as Spin-Adapted Antisymmetrised Products (SAAPs)⁸³. The following sections demonstrate several different schemes for constructing SAAPs out of primitive spin functions. When these SAAPs are combined with a specific set of spatial orbitals in a given basis set, i.e. the overall wavefunction can be constructed out of Slater determinants, these are known as Configurational State Functions (CSFs).

All eigenfunctions with a given S are represented degenerately in all choices $-S \leq M_s \leq S$. The regime $S = M_s$ is used for all calculations, unless otherwise specified, as this simplifies the expansion and representation of the associated spin eigenfunc-

tions by minimising the number of primitive spin functions they are constructed from.

For use in FCIQMC, random excitation generation and Hamiltonian matrix element evaluation are also required. For discussions of these, see chapters 4 and 5. Implementation is made more straightforward by the fact that the structure of CSFs is strongly related to that for determinants. As the infrastructure for manipulating, storing and transmitting Slater determinants is already mature, it may be readily co-opted.

3.3.1 Size of the Hilbert Space

The number of CSFs is most easily obtained by considering primitive spin functions with n_α unpaired spin ‘up’, and n_β unpaired spin ‘down’ electrons such that $N_o = n_\alpha + n_\beta$ and $M_s = \frac{1}{2}(n_\alpha - n_\beta)$. From these eigenfunctions of \hat{S}_z , eigenfunctions of \hat{S}^2 may be constructed with $S = \frac{1}{2}N_o, \frac{1}{2}N_o - 1, \dots, M_s + 1, M_s$ with the dimension of the subspace being $\binom{N_o}{n_\alpha} \equiv \binom{N_o}{n_\beta}$.

All of the states in the subspace $M_s + 1$ have related functions with the same total spin in the subspace M_s , and as a consequence the number of unique functions with $S = M_s$ is given by

$$\begin{aligned} n_{\text{csf}}(N, S) &= \binom{N_o}{n_\alpha} - \binom{N_o}{n_\alpha + 1} = \binom{N_o}{n_\beta} - \binom{N_o}{n_\beta + 1} \\ &= \binom{N_o}{\frac{1}{2}N_o - S} - \binom{N_o}{\frac{1}{2}N_o - S - 1} \\ &= \frac{4S + 2}{N_o + 2S + 2} \binom{N_o}{\frac{N_o - 2S}{2}}. \end{aligned}$$

This formula can also be verified by induction by considering the genealogical construction (see section 3.3.3).

3.3.2 Spin paired (Rumer) spin eigenfunctions

For a two electron spin function, the singlet state for electrons i, j is given by

$$v(i, j) = \frac{1}{\sqrt{2}}[\alpha(i)\beta(j) - \beta(i)\alpha(j)].$$

An N -electron spin eigenfunction may be constructed by considering products of singlet-terms for g pairs of electrons, where $g = \frac{1}{2}(N - 2S)$, followed by $2S$ single-electron terms α , such that

$$V_{\mathbf{k}} = \prod_i^g \frac{1}{\sqrt{2}} v(e_{\mathbf{k}}(2i-1), e_{\mathbf{k}}(2i)) \prod_{i=2g+1}^N \alpha(e_{\mathbf{k}}(i)),$$

where $\{e\}$ specifies an ordered choice of electrons. Any doubly occupied spatial orbitals must be placed in the paired region. The number of spin paired eigenfunctions available is given by

$$\nu(N, S) = \binom{N}{2g} (2g-1)!! = \frac{N!}{2^g (N-2g)! g!},$$

which is significantly larger than the number of independent spin eigenfunctions, n_{csf} . A substantial number of these are related to each other by symmetry, but this still leaves a massively over-complete set. If n_{csf} linearly independent functions could be systematically generated, then these would form a basis which spans the entire Hilbert space.

Rumer et al.⁸⁴⁻⁸⁶ have described a procedure based on a series of diagrams known as *Extended Rumer Diagrams*. As shown in an example for a 5-electron system in figure 3.3, the numbers 1 to N are arranged on the circumference of a circle, along with an additional point, known as the pole, P . g arrows are drawn between the g pairs of coupled electrons, and all remaining uncoupled electrons are connected to the pole. If Rumer spin eigenfunctions are constructed with the linked pairs coupled, and the remaining electrons assigned to α spins, then the number of functions constructed is equal to n_{csf} , and they are all linearly independent.

If Schmidt orthogonalisation is applied to this series of functions, starting with the first Rumer function V_1 , then the Kotani-Yamanouchi branching functions (see section 3.3.3) are obtained⁸⁶. Similarly, the Serber functions (see section 3.3.4) may be obtained by Schmidt orthogonalisation while enforcing symmetry and anti-symmetry in pairs⁸⁷.

The expansion of these spin eigenfunctions into products of one-electron spin functions is particularly compact. While this benefits algorithms that depend on expansions into determinants, these algorithms are not really suitable for use deep within tight computational loops as they scale badly. As such these functions

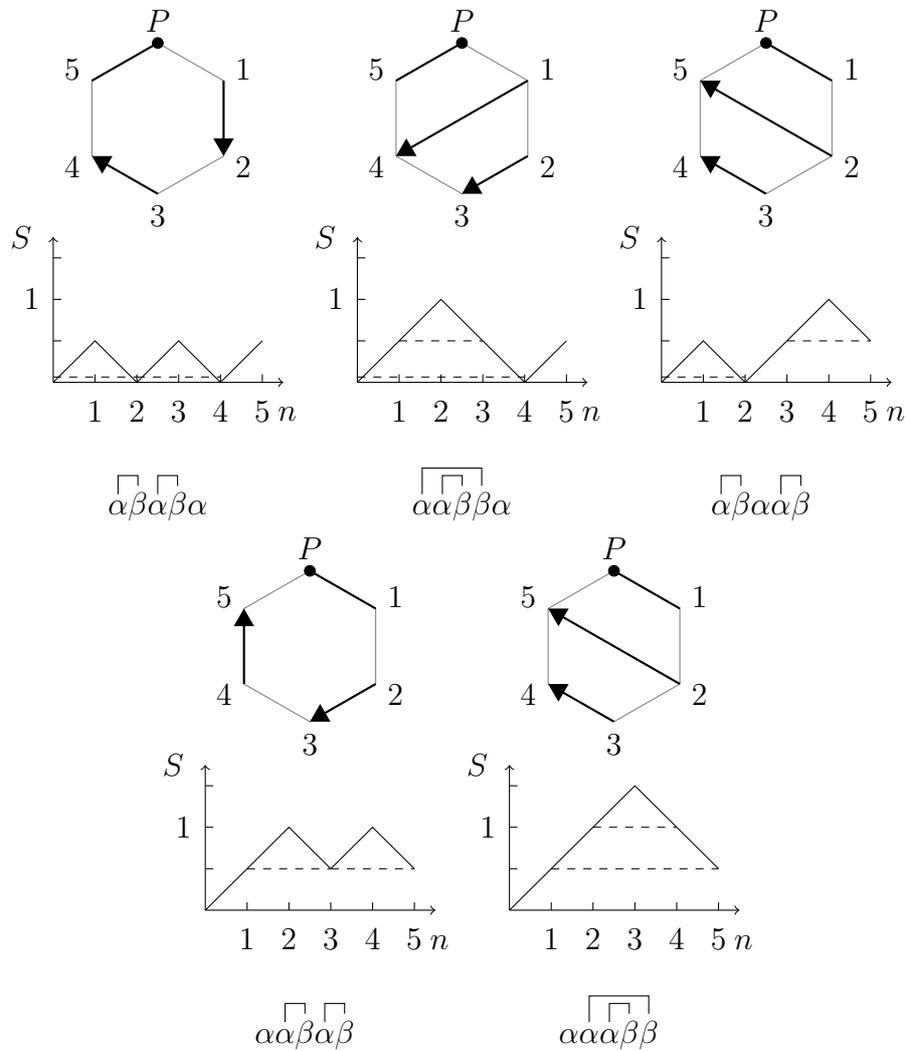


Figure 3.3: Extended Rumer diagrams, branching diagrams, and leading terms for $N = 5, S = \frac{1}{2}$. The paired electrons are indicated by arrows, dashed lines and square brackets respectively. Note that the Rumer functions indicated by the extended Rumer diagrams with no crossed arrows can be labelled with the same branching diagram labels as the Kotani-Yamanouchi spin eigenfunctions, which are produced by performing a Schmidt orthogonalisation on these functions.

are primarily useful for their relationship to the other types of spin eigenfunction. In particular, it is more straightforward to calculate the representation matrices for permutations in this space, and they are therefore useful for calculations in Kotani-Yamanouchi and Serber-type spaces.

3.3.3 Genealogical (Kotani-Yamanouchi) spin eigenfunctions

Kotani^{88,89} suggests a means to construct spin eigenfunctions in a stepwise fashion, constructing an eigenfunction with N electrons and given total and projected spins S , M_s from those available with $N - 1$ electrons, and corresponding spins $S \pm \frac{1}{2}$, $M_s \pm \frac{1}{2}$:

Addition

$$X(N, S, M; k) = [(S + M)^{\frac{1}{2}} X(N - 1, S - \frac{1}{2}, M - \frac{1}{2}; k') \alpha(N) + (S - M)^{\frac{1}{2}} X(N - 1, S - \frac{1}{2}, M + \frac{1}{2}; k') \beta(N)] (2S)^{-\frac{1}{2}}. \quad (3.2a)$$

Subtraction

$$X(N, S, M; k) = [-(S - M + 1)^{\frac{1}{2}} X(N - 1, S + \frac{1}{2}, M - \frac{1}{2}; k') \alpha(N) + (S + M + 1)^{\frac{1}{2}} X(N - 1, S + \frac{1}{2}, M + \frac{1}{2}; k') \beta(N)] \times (2S + 2)^{-\frac{1}{2}} \quad (3.2b)$$

States labelled k can be constructed from as many states k' with the properties of $S \pm \frac{1}{2}$ and $M_s \pm \frac{1}{2}$ as can be found. This continues until there is only one unpaired electron where the only primitive spin functions are $\alpha(1)$ and $\beta(1)$. Note that it is not physically possible to construct a spin function with total spin $S < 0$, thus the construction of states with spin S from $S - \frac{1}{2}$ is only possible for $S \geq \frac{1}{2}$. This limits the number of genealogical construction pathways. See figure 3.4 for a diagram showing the permitted branching pathways⁹⁰.

The construction of each spin function $X_{\mathbf{i}}$ can be labelled by a Yamanouchi symbol, $B_{\mathbf{i}}$. This is represented by a string of digits, either 1 or 2, representing addition and subtraction respectively. Each element, r , of $B_{\mathbf{i}}$ may equivalently be represented by the term $b_{\mathbf{i}r} = \pm \frac{1}{2}$ for addition and subtraction respectively (equation 3.2), as used below.

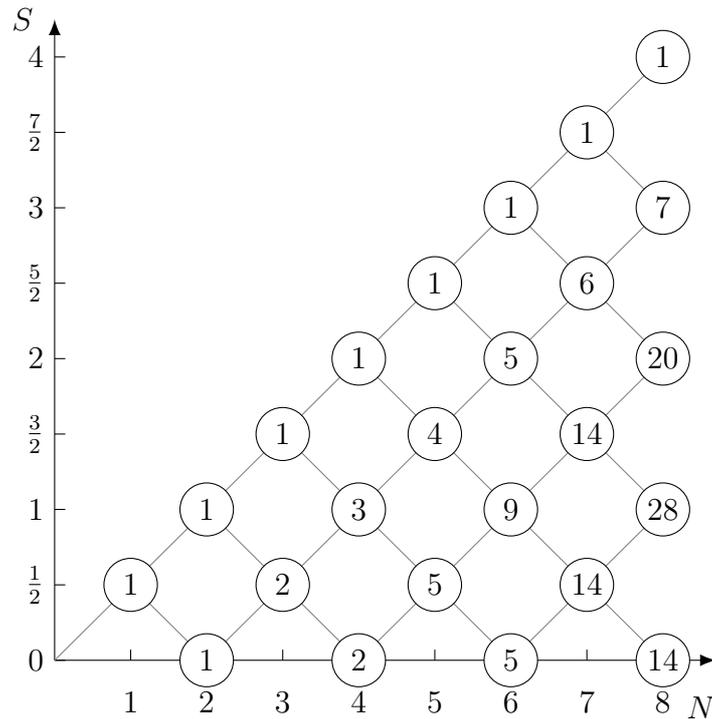


Figure 3.4: Branching diagram for Kotani-Yamanouchi spin functions. The circled numbers indicate the number of possible construction routes, and therefore the number of available eigenfunctions, with spin S for N electrons. These correspond to the number of routes from the origin to the given point, whilst only moving right, where ‘up’ and ‘down’ respectively correspond to 1 and 2 in the Yamanouchi symbol. Note that routes are not permitted to pass through any point with $S < 0$.

The coefficients of the primitive spin functions in each Kotani-Yamanouchi spin eigenfunction can be obtained more directly than by following the genealogical construction. It can be shown inductively that the coefficient of the primitive spin function θ_j in the spin eigenfunction $X_{\mathbf{i}}$, where $X_{\mathbf{i}} = \sum_j A_{\mathbf{i}j} \theta_j$, is given by

$$A_{\mathbf{i}j} = \prod_{r=1}^N C(b_{ir}, m_{jr}; S_{ir}, M_{jr}),$$

where $m_{jr} = \pm \frac{1}{2}$ is the m_s eigenvalue associated with the r^{th} element in the primitive spin function θ_j , $M_{jr} = \sum_{s=1}^r m_{js}$ is the partial resultant projected spin associated with the first r elements in the primitive spin function and $S_{ir} = \sum_{s=1}^r b_{is}$.

The relevant Clebsch–Gordan coefficients are given by⁸⁹

$$C(b, m; S, M) = \begin{cases} \sqrt{\frac{S+2mM}{2S}} & \text{if } b = \frac{1}{2} \\ -2m\sqrt{\frac{S+1-2mM}{2(S+1)}} & \text{if } b = -\frac{1}{2}. \end{cases}$$

3.3.4 Serber-type spin eigenfunctions

An alternative scheme for the construction of spin eigenfunctions may be obtained by considering adding *pairs* of electrons sequentially, rather than individually. States constructed in this manner are known as Serber functions.⁹¹ States with N electrons, total spin S and its z-projection M_s may be constructed from those with $N - 2$ electrons and spins $S, S \pm 1$ and $M_s, M_s \pm 1$ in a total of four ways.

Addition of singlet state

$$Z(N, S, M; k) = Z(N - 2, S, M, k')g_0(N - 1, N), \quad (3.3a)$$

(S - 1) → S by addition of a triplet state

$$\begin{aligned} Z(N, S, M; k) = & \{ [(S + M)(S + M - 1)]^{\frac{1}{2}} Z(N - 2, S - 1, M - 1; k')g_1(N - 1, N) \\ & + [2(S + M)(S - M)]^{\frac{1}{2}} Z(N - 2, S - 1, M; k')g_2(N - 1, N) \\ & + [(S - M)(S - M - 1)]^{\frac{1}{2}} Z(N - 2, S - 1, M + 1; k')g_3(N - 1, N) \} \\ & \times [2S(S - 1)]^{-\frac{1}{2}}, \end{aligned} \quad (3.3b)$$

S → S by addition of a triplet state

$$\begin{aligned} Z(N, S, M; k) = & \{ -[(S + M)(S - M + 1)]^{\frac{1}{2}} Z(N - 2, S, M - 1; k')g_1(N - 1, N) \\ & + 2^{\frac{1}{2}} M Z(N - 2, S, M; k')g_2(N - 1, N) \\ & + [(S - M)(S + M + 1)]^{\frac{1}{2}} Z(N - 2, S, M + 1; k')g_3(N - 1, N) \} \\ & \times [2S(S + 1)]^{-\frac{1}{2}}, \end{aligned} \quad (3.3c)$$

(S + 1) → S by addition of a triplet state

$$\begin{aligned} Z(N, S, M; k) = & \{ [(S - M + 2)(S - M + 1)]^{\frac{1}{2}} Z(N - 2, S + 1, M - 1, k')g_1(N - 1, N) \\ & - [2(S - M + 1)(S + M + 1)]^{\frac{1}{2}} Z(N - 2, S + 1, M; k')g_2(N - 1, N) \\ & + [(S + M + 1)(S + M + 2)]^{\frac{1}{2}} Z(N - 2, S + 1, M + 1; k')g_3(N - 1, N) \} \\ & \times [(2S + 2)(2S + 3)]^{-\frac{1}{2}}, \end{aligned} \quad (3.3d)$$

The construction of each Serber function, Z_i , can be labelled by a Serber branching diagram symbol, B_i . This is represented by a string of letters, $ABCD$ corresponding to singlet and triplet components in the listed order, with the additional term α and β for the last term if there are an odd number of electrons. Each element, r , of B_i may equivalently be represented by the term $b_{ir} = 0, 1, 2, 3$ respectively, as used below.

In a similar manner to the Kotani-Yamanouchi spin eigenfunctions, the coefficients of the contributing primitive spin functions can be directly obtained. In this case the spin eigenfunctions are no longer constructed from Slater determinants, but from Geminal Spin Product functions (GSPs). These GSPs are in turn constructed from the product of the geminal spin functions, $g_0 \cdots g_3$,

$$G(SP) = \prod_{r=1}^n g_{SP_r}(2r-1, 2r),$$

where n is the number of pairs of electrons, and SP is a vector of terms $sp_r = 0, 1, 2, 3$ indicating the sequence of geminal functions. The coefficient of each of these GSPs in the Serber function Z_i is given by,

$$A_{ij} = \prod_{r=1}^n C(b_{ir}, sp_{jr}; S_{ir}, M_{jr}),$$

where $Z_i = \sum_j A_{ij} G(SP_j)$. $M_{jr} = \sum_{s=1}^r m_{js}$ is the partial resultant projected spin associated with the first r terms in the geminal spin function where $m_{js} = -1, 0, 1$ correspondingly. $S_{ir} = \sum_{s=1}^r$ provides the same measure for the serber spin label. The relevant coefficients are given by

$$\begin{aligned} C(0, 0; S, M) &= 1 \\ C(1, 1; S, M) &= c(1)\sqrt{(S+M)(S+M-1)} \\ C(1, 2; S, M) &= c(1)\sqrt{2(S+M)(S-M)} \\ C(1, 3; S, M) &= c(1)\sqrt{(S-M)(S-M-1)} \\ C(2, 1; S, M) &= -c(2)\sqrt{(S+M)(S-M+1)} \\ C(2, 2; S, M) &= c(2)M\sqrt{2} \\ C(2, 3; S, M) &= c(2)\sqrt{(S-M)(S+M+1)} \\ C(3, 1; S, M) &= c(3)\sqrt{(S-M+2)(S+M+1)} \end{aligned}$$

$$C(3, 2; S, M) = -c(3)\sqrt{2(S - M + 1)(S + M + 1)}$$

$$C(3, 3; S, M) = c(3)i\sqrt{(S + M + 1)(S + M + 2)}$$

and

$$c(1) = \frac{1}{\sqrt{2S(2S - 1)}}$$

$$c(2) = \frac{1}{\sqrt{2S(2S + 1)}}$$

$$c(3) = \frac{1}{\sqrt{(2S + 2)(2S + 3)}}.$$

It is worth noting that for the Hamiltonian matrix element evaluation it is important that these spin functions are considered to be constructed out of N electrons, with the paired electrons being restricted to occupying singlet terms, rather than only considering the unpaired electrons as for Kotani-Yamanouchi spin functions. This is as a consequence of the matrix element evaluation depending on the full N -electron permutations between basis functions.

3.3.5 Spin-projected (Slater) determinants

In the previous two sections, spin eigenfunctions were constructed by building them up out of component fragments with fewer electrons. Here an alternative approach is examined, where the CSF corresponding to the components of the correct spin of an N electron function is obtained by a projection process on this function.

Löwdin⁹² defined a spin projection operator,

$$\hat{O}_S = \prod_{k \neq S} \frac{\hat{S}^2 - k(k + 1)}{S(S + 1) - k(k + 1)},$$

in which the components of the eigenfunctions of \hat{S}^2 (other than the one chosen) are subtracted out sequentially, and some normalisation is applied so that the spin component selected retains its magnitude*. In practice this operator can be

*It is worth noting that the same approach may be taken to obtain total angular momentum

applied by performing a resolution of the identity through Kotani-Yamanouchi spin functions (or Serber functions), X , of the given spin, S , such that

$$\hat{O}_S = \sum_X |X, S\rangle \langle X, S|.$$

If the spin-pure components of a spin eigenfunction are projected out, the eigenfunction is returned unchanged, thus \hat{O}_S is idempotent. If this operator is applied to a basis of Slater determinants, a new basis of ‘spin-projected’ determinants, $\{F_i\}$,

$$|F_i\rangle = \hat{O}_S |D_i\rangle,$$

is generated. The primary benefit of this basis set is that the Hamiltonian and overlap matrix elements are directly calculable, as laid out in section 4.6.

It should be clear that, unless the chosen value of M_s is maximal for the system, as the number of possible eigenstates of \hat{S}_z exceeds the number of eigenstates of \hat{S}^2 , then the newly formed basis must be over-complete. An orthonormal basis can be constructed from these new functions, but in doing so all other advantages over CSFs constructed using the Genealogical scheme are lost.

3.3.5.1 Normalisation

The spin-projected functions described above are not only non-orthogonal, they are not normalised. The self-overlap elements are given in section 4.6.6 as

$$S_{ii} = C_k(S, M_s, N_o),$$

and depend primarily on the number of unpaired electrons, N_o , as the spin is specified and constant for all sites in a given simulation. A consequence of this is that the amplitude of a coefficient, c_i , in the main list is of differing significance depending on its location within the Hilbert space.

If decisions regarding granularisations and coefficient cutoffs are going to be made in a meaningful sense for the ensemble of particles within a calculation, it is helpful to maintain a constant ‘weight’ for each site in the system. As such, the basis

eigenfunctions, generally starting with a projected angular momentum eigenfunction⁹³.

functions can be optionally redefined such that

$$F_{\mathbf{i}} \longrightarrow \frac{1}{\sqrt{S_{\mathbf{ii}}}} F_{\mathbf{i}}.$$

In practice this does not have a qualitatively large impact on the effectiveness of simulations. Quantitatively, using normalised basis functions reduces the total number of particles required, and thus reduces the overall computational cost.

3.3.5.2 Truncated spin-projected spaces

If a spin-projected space is constructed in the regime $S = M_s = 0$ the space may be simplified to remove the over-completeness, even though this does not eliminate the non-orthogonality.

Considering the two determinants, $D_{\mathbf{i}}, \overline{D}_{\mathbf{i}}$, associated with an HPHF function, where the spins corresponding to each electron have been flipped, the two spin-projected functions are related such that $F_{\mathbf{i}} = \pm \overline{F}_{\mathbf{i}}$, with the sign being the same as the sign linking the two determinants in the relevant HPHF function. As these two functions are representing exactly the same region of the Hilbert space an arbitrary choice can be made between them, and only one included in the simulation. This is done in the same way as for HPHF functions.

This simplification cannot be straightforwardly applied outside of the regime $S = M_s = 0$ as inverting all of the spins results in a determinant, and hence a spin-projected determinant, where $M_s(\overline{D}_{\mathbf{i}}) = -M_s(D_{\mathbf{i}})$ and as such is not in the space being considered.

3.4 Implementation of spin eigenfunctions in FCIQMC

The main structure of the FCIQMC algorithm is well defined by previous work. However, efficient usage of spin eigenfunctions requires modification to several portions of the algorithm and its implementation.

Hamiltonian matrix element evaluation

Various schemes can be used to calculate the diagonal and off-diagonal matrix elements used for the death and spawning steps respectively. In the case of Serber spin-functions there is a clear trade-off between computational cost and memory usage, and in all cases careful integration between the excitation generator and the matrix element evaluation routines is required to avoid unnecessary work. Matrix element evaluation schemes are outlined in chapter 4.

Excitation generation

The excitation generation scheme required for spin eigenfunctions is structurally similar to that for Slater determinants. The connectivity of the Hilbert space is noticeably different, requiring different acceptance rules, and substantially different generation probabilities. This is discussed in chapter 5.

Spawning and death

When using non-orthogonal or non-normalised basis functions, the terms included in the spawning and death steps need careful consideration. This is particularly true when considering the nature of spawning between multiple sites with the same spatial structure and a non-zero overlap term, where *spawning* could be considered to implement *non-local death*. This is discussed below in section 3.4.2.

Energy calculation

It is important to include all of the relevant terms in the projected energy calculation. The contributing terms differ from those in determinantal FCIQMC. Calculation of a stable energy estimate is extremely important for non-orthogonal FCIQMC, where this estimate enters into the off-diagonal terms used for spawning.

3.4.1 Representation of spin eigenfunctions

When performing FCIQMC calculations using Slater determinants or HPHF functions, two equivalent representations of the individual sites are used;

Expanded (natural) integer representation

For calculation of Hamiltonian matrix elements, and excitation generation, a list of the occupied spin-orbitals associated with a given determinant, or with

the canonical determinant in the HPHF function, is required. This is made up of a list of spin-orbitals in increasing numerical order. These spin orbitals are ordered such that all of the odd numbers represent β spin-orbitals, and the even ones α , with each *pair* of integers representing a spatial orbital.

Compacted bit-representation

A compact representation may be obtained as

$$\xi(D_\nu) = \sum_i 2^{\phi_i^{(\nu)} - 1},$$

where one bit is set for each occupied spin orbital. This representation requires $2M$ bits, one for each available spin-orbital in the basis set. This requires substantially less memory than the expanded form, and is used for storage of particles between iterations and especially for communication between computational nodes, as required in the annihilation step. The coefficient on a site and any associated flags may be stored and appended immediately following this compact orbital bit representation.

If spin eigenfunctions are being used, this representation needs to change a little.* The representations must now store the spatial structure of the spin eigenfunction and a label describing which spin structure is being used.

The numbering for the integer representation is adjusted, such that each pair of integers still represents one spatial orbital, but the odd integer now corresponds to the first electron occupying a spatial orbital, and the even integer is now only used to represent the second electron in a doubly occupied orbital. This relabelling also carries through to the bit-representation. For a spatial structure with g doubly occupied spatial orbitals and N_o unpaired electrons, the integer representation is now ordered beginning with

1. all $2g$ orbitals associated with the g pairs of doubly occupied orbitals, sorted in numerical order, followed by
2. N_o unpaired orbitals, in numerical order.

Both the expanded integer and the bit representations of the spatial structure are immediately appended with an integer containing a label (or an index to a label) for the spin structure of the relevant spin eigenfunction – either a Yamanouchi

*With the exception of spin-projected determinants, where the Slater determinants to be projected are stored as usual.

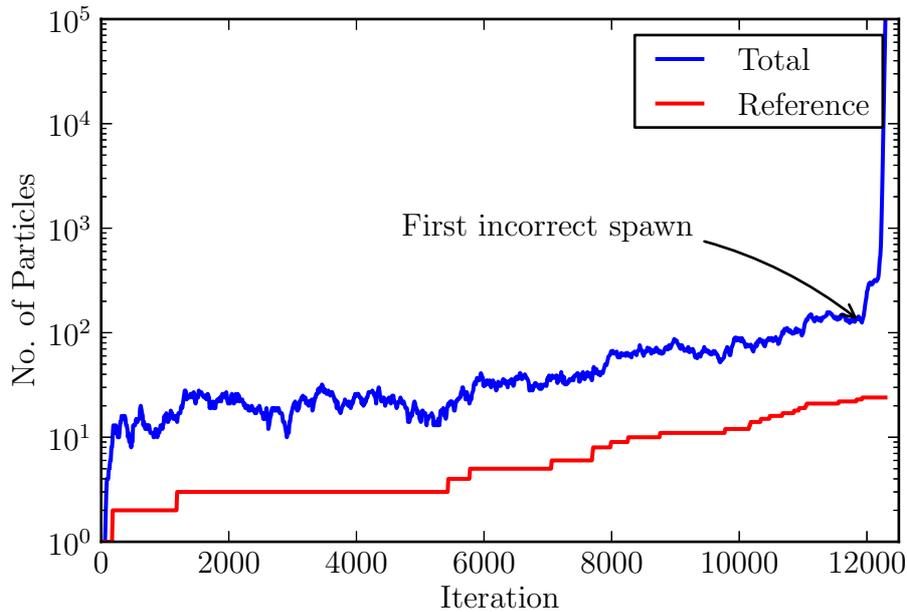


Figure 3.6: A particle population explosion caused by non-local death-like terms appearing in the off-diagonal spawning expressions. This plot shows the trajectory of the total particle population and the population on the reference site in the course of an FCIQMC calculation, showing that once an incorrect non-local death-like spawn has occurred, the positive feedback induced causes an explosion in the total particle population, whilst the reference is unaffected.

symbol, or a Serber eigenfunction label.

3.4.2 Non-local death and population control

Testing FCIQMC using spin-projected determinants, which are non-orthogonal, reveals extremely unexpected behaviour. As system size is increased, a dramatic instability is observed. As illustrated in figure 3.6, at some point during a calculation, a very large and dramatic increase in total particle population is observed — the total particle population grows by several orders of magnitude over the course of only a few iterations.

The onset of this is clearly triggered by a specific process, which is accessed stochastically, as it will occur at radically different times within a calculation with only a change of the random number seed. Similarly, decreasing the granularity of the representation or using much smaller values of the imaginary time step, $\delta\tau$, has

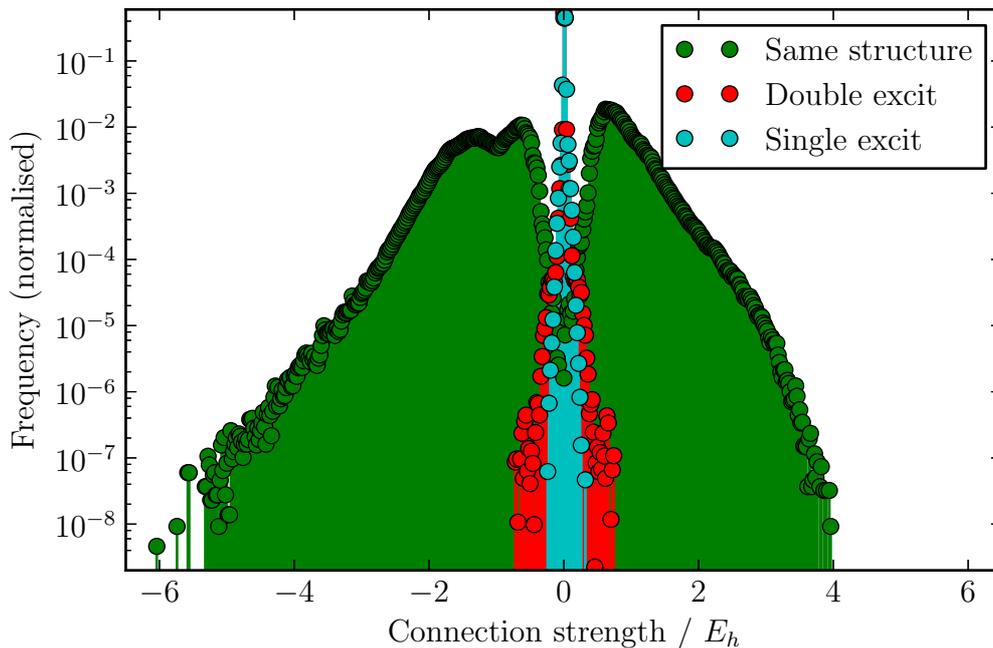


Figure 3.7: A histogram of the spawning connection strength, $\frac{K_{ij} - E_{\text{weighted}} S_{ij}}{S_{jj}}$, for the spawn $\mathbf{i} \rightarrow \mathbf{j}$, using spin-projected determinants for N_2 with a bond length of $6.00 a_0$ in a cc-pVDZ basis set. Note the much broader range of values generated for same spatial structure spawns, and that the distribution is somewhat asymmetric.

no impact on the emergence of this behaviour.

It is notable that this behaviour is not observed in two specific contexts;

- if spawning between sites with the same spatial structure is disabled, although this leads to convergence on a wavefunction with an incorrect energy, or
- if the correct wavefunction is converged on and there are many particles in the simulation, either through seeding with the results of a previous calculation or switching the same spatial structure spawning on slowly once the simulation has grown sufficiently.

This makes it clear that the spawning is related to the off-diagonal matrix elements between sites with the same spatial structure, and to some degree to (stochastic) deviations from the correct wavefunction. This is not entirely surprising; the matrix elements associated with same spatial structure spawns,

$$\begin{aligned}
\langle F_\mu | \hat{H} | F_\nu \rangle &= \sum_{\mathbf{i}, \mathbf{j}} c_{\mathbf{i}}^{(\mu)*} c_{\mathbf{j}}^{(\nu)} \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle \\
&= \sum_{\mathbf{i}} c_{\mathbf{i}}^{(\mu)*} c_{\mathbf{i}}^{(\nu)} H_{\mathbf{i}\mathbf{i}} + \sum_{\substack{\mathbf{i}, \mathbf{j} \\ \mathbf{i} \neq \mathbf{j}}} c_{\mathbf{i}}^{(\mu)*} c_{\mathbf{j}}^{(\nu)} H_{\mathbf{i}\mathbf{j}},
\end{aligned}$$

can be effectively approximated by noting that (according to table 4.1) $\langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle$ between two determinants with the same spatial structure is very close to $\langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{i}} \rangle$, and that the number of contributing terms to the remaining double sum is likely to be small (as the terms can differ by no more than two spatial orbitals). As a consequence,

$$\langle F_{\mathbf{i}} | \hat{H} | F_{\mathbf{j}} \rangle \approx S_{\mathbf{i}\mathbf{j}} \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{i}} \rangle + \text{small terms},$$

and the spawning matrix elements for same spatial structure spawns are largely dictated by the diagonal matrix elements of the relevant Slater determinants and the overlap matrix. This spawning could be considered to be equivalent to *non-local death* — i.e. the effect of death on one site causing changes in coefficients for other sites that represent the same region of the Hilbert space.

Because diagonal matrix elements for Slater determinants are much larger than off-diagonal ones, the same spatial structure Hamiltonian matrix elements are generally much larger than those associated with single or double excitations (see figure 3.7). They therefore have a large impact on the dynamics of simulations, place more restrictions on the choice of generation probabilities (see section 5.3) and the permissible imaginary time step, $\delta\tau$ (see section 2.6).

It can be observed that the explosive growth of particles observed in figure 3.6 begins with a pair of sites where positive feedback causes particles to be mutually spawned between them more rapidly than particles are killed by (local) death (see figure 3.8).

3.4.3 Demonstration of non-local death

In this section we consider a model two site system being expanded to three sites so as to generate non-orthogonality and overcompleteness. This is intended as an illustration of the nature of spawning associated with non-orthogonal basis functions, and it demonstrates that if the simulation converges, then the two cases are only equivalent if the spawning associated with the non-orthogonal connections

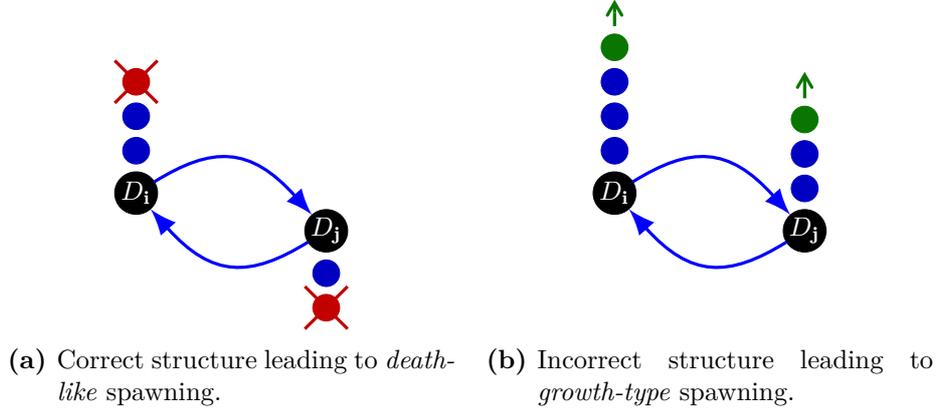


Figure 3.8: Correct and incorrect death-like spawning. In (a) spawns from D_i to D_j , and vice versa, create particles of opposite sign to those already present, and therefore reduce the amplitude on the target site (the red particles are ‘killed’). In (b) the reverse happens, and spawns in both directions cause amplitude to increase on both sites (the green particles are ‘born’). This suffers from positive feedback, as next iteration the number of particles will grow in proportion to those present.

is death-like. For the two site simulation given in figure 3.9a the change in the coefficients associated with the two sites per iteration is given by

$$\begin{aligned} S_{11}\Delta c_1 &= -\delta\tau c_0 H_{01} - \delta\tau c_1 [H_{11} - (E_{ref} + E_S)S_{11}] \\ S_{00}\Delta c_0 &= -\delta\tau c_1 H_{01} - \delta\tau c_0 [H_{00} - (E_{ref} + E_S)S_{00}]. \end{aligned}$$

Similarly, for the model system where the second site is duplicated to create a third with $|\overline{D_1}\rangle = -|D_1\rangle$, as in figure 3.9b, the expressions for the changes in the coefficients may be reduced to

$$\begin{aligned} S_{11}(\Delta c_1 - \Delta c_{\overline{1}}) &= -2\delta\tau c_0 H_{01} - \delta\tau(c_1 - c_{\overline{1}})[H_{11} - (E_{ref} + E_S)S_{11}] \\ &\quad - \delta\tau(c_1 - c_{\overline{1}})[H_{11} - (E_{ref} + E_C)S_{11}] \\ S_{00}\Delta c_0 &= -\delta\tau(c_1 - c_{\overline{1}})H_{01} - \delta\tau c_0 (H_{00} - (E_{ref} + E_S)S_{00}) \end{aligned}$$

by noting that $H_{11} = H_{\overline{1}\overline{1}} = -H_{1\overline{1}} = -H_{\overline{1}1}$ and $S_{11} = S_{\overline{1}\overline{1}} = -S_{1\overline{1}} = -S_{\overline{1}1}$.

It is clear that the trajectories followed by the two simulations are not the same — there is a factor of two that appears in the simulation with the duplicated site. However, if we consider the fully converged case, with the value of $E_S = E_C$, in the two-site system $\Delta c_0 = \Delta c_1 = 0$. In the three site system, if the coefficients are written such that $c_1 - c_{\overline{1}} = c_{1,\text{two-site}}$, then the same steady state behaviour is

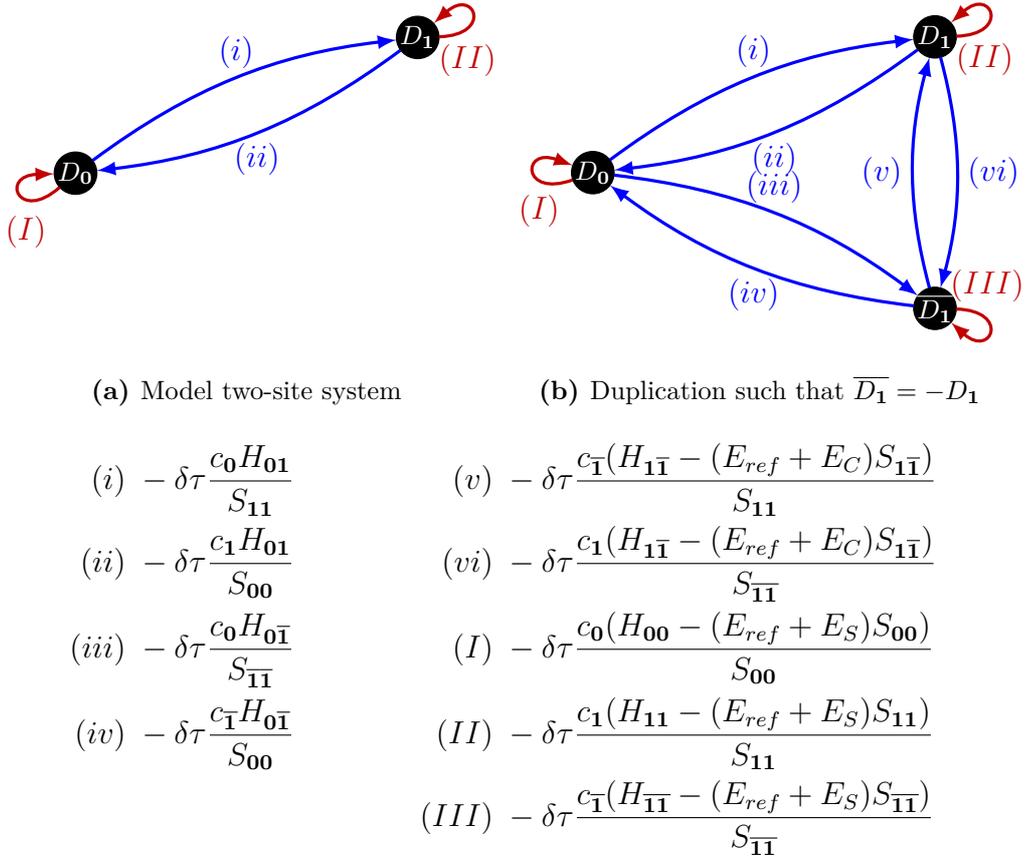


Figure 3.9: The possible spawning patterns in a model two-site system being expanded to have three sites. Inter-site spawning is notated with blue arrows, and diagonal death in red. In (b) the site $|D_1\rangle$ is duplicated such that $|\overline{D_1}\rangle = -|D_1\rangle$.

observed with $\Delta c_0 = (\Delta c_1 - \Delta c_{\overline{1}}) = 0$ and the same value of the projected energy is recovered.

The terms associated with spawning between the duplicated sites, containing E_C directly, are clearly responsible for counteracting the doubling of the spawning contribution from site $\mathbf{0}$. As a consequence, a linear combination of terms correctly represent the same region of the Hilbert space, and spawning between them is correctly considered to be equivalent to *non-local death*.

3.4.4 Algorithmic changes to suppress population explosion

For non-orthogonal FCIQMC, the derivation presented in section 2.2.2 demonstrates only that if a simulation converges, then it has converged on an eigenfunction of the Hamiltonian. It does not in any way imply that the trajectory taken will converge — there is no direct connection with the integration of the imaginary time Schrödinger equation as there is for orthogonal FCIQMC. This suggests that algorithmic changes may be required to assist the calculation through the growth phase until it has sufficiently converged.

It is clear that same spatial structure spawning is acting as non-local particle death. Consequently, an additional rule is added at the annihilation stage:

All particles spawned onto a given target site, originating from sites with the same spatial structure, are combined. If the target site is empty, the spawns are rejected. If the sign of the combined spawn is the same as the sign of particles already present, the spawns are rejected. If the sign of the combined spawn is opposite to the sign of the particles already present, and larger in magnitude to the weight of particles already present, the magnitude of the spawn is truncated to leave the site empty.

This rule may be summarised as “Ensure that the cumulative effect of same spatial structure spawns is to kill existing amplitude, not grow it”, or more directly, *enforce* that same spatial structure spawns cause *non-local death*.

This rule is extremely similar both implementationally and practically to the initiator approximation. Both place constraints on the particles which survive based on *a*) where the particle comes from, and *b*) the occupation of the target site. It is similar, also, in that in the many particle limit the approximation disappears — if the wavefunction is correctly converged, this death-like behaviour is observed, and it is notable that once a simulation has grown above a certain (generally unknown) size this restriction may be switched off with no impact. Qualitatively this has little detrimental impact on the simulation. Presumably it has a quantitative impact on the initiator error.

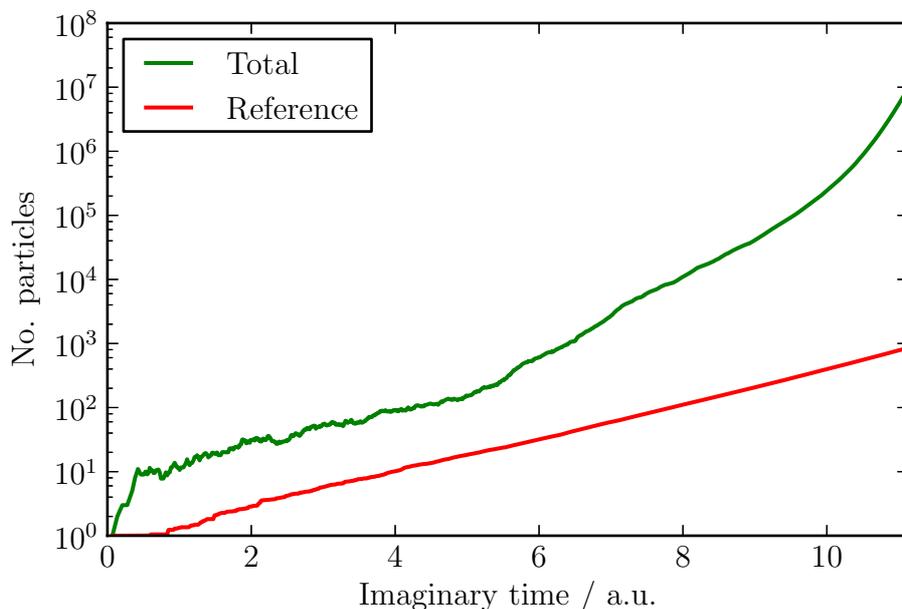


Figure 3.10: A plot of the total number of particles and the number of particles on the reference site for a simulation of N_2 in a cc-pVDZ basis set at $4.2 a_0$ using normalised spin-projected determinants. This demonstrates an unexplained convergence failure.

3.4.5 Further issues with spin-projected determinants

The derivation in section 2.2.2 demonstrates only the properties of a simulation that has converged. The openness of this scheme to simulations that never converge haunts this as a generally applicable method.

Even with the algorithmic adaptations made, there are some systems which lead to pathological behaviour. Figure 3.10 demonstrates the failure of FCIQMC using spin-projected determinants to converge for the nitrogen dimer in a cc-pVDZ basis with a bond length of $4.2 a_0$. This failure is resilient, and is not avoided by

1. suppression of the same spatial structure spawns,
2. reduction of the timestep, $\delta\tau$,
3. the use of continuous time FCIQMC (see chapter 7), or
4. any amount of the other fiddling and tweaking that have been tried.

This appears to be system specific behaviour — it is not observed at all bond lengths, or all basis sets, for N_2 — but it is certainly related to system size in some respect.

While spin-projected determinants promise a great deal of benefit in comparison to other spin eigenfunctions (largely as a result of avoiding storage of uncontrollable amounts of data) until the causes of this instability are understood they are not safely and generally applicable.

4 Hamiltonian matrix element evaluation

The FCIQMC algorithm depends strongly on two major components; *excitation generation* — which is the process of taking random steps in the Hilbert space between sites connected by non-zero Hamiltonian matrix elements, and *matrix element evaluation* — which is the process of calculating Hamiltonian matrix elements between any two arbitrary sites. Both of these processes must be efficiently implemented and integrated with each other for the overall simulation to be efficient. The implementation of Hamiltonian matrix element evaluation is explored in this chapter, along with a discussion of what information may be passed from the excitation generators to enhance efficiency.

4.1 Slater determinants and HPHF functions

The Hamiltonian matrix elements between a pair of arbitrary Slater determinants can be calculated using the list of the occupied orbitals of one of the two determinants, along with the *excitation matrix* (two lists of orbitals that are respectively only occupied in one of the two determinants) and a count of the *excitation level* between the two determinants (how many spin orbitals they differ by).

If the Slater determinants differ by more than two spin orbitals, the Hamiltonian matrix element between them is zero. Otherwise, the matrix elements are given by the expressions in table 4.1. It is worth noting that the one- and two-electron terms referenced in the table are indexed by spin-orbital. These are in a practical sense stored indexed by spatial orbitals, as this reduces the storage overhead by a factor of four. Consequently the spin restriction that $\langle ab|\hat{h}|cd\rangle = 0$ if $m_s(a) \neq m_s(c)$ or $m_s(b) \neq m_s(d)$ must be implemented explicitly.

Integral	Expression
$\langle D \hat{H} D\rangle$	$\frac{1}{2}\sum_i\sum_j(\langle ij \hat{g} ij\rangle - \langle ij \hat{g} ji\rangle) + \sum_i h_{ii}$
$\langle D \hat{H} D_m^p\rangle$	$\sum_i(\langle mi \hat{g} pi\rangle - \langle mi \hat{g} ip\rangle) + h_{mp}$
$\langle D \hat{H} D_{mn}^{pq}\rangle$	$\langle mn \hat{g} pq\rangle - \langle mn \hat{g} qp\rangle$
otherwise	0

Table 4.1: The Slater–Condon Rules^{94–96} for evaluating the Hamiltonian matrix elements of Slater determinants. The sums, $\{i, j\}$ are performed over all occupied spin-orbitals in the determinant D , and the operator $\hat{g} = r_{12}^{-1}$. The Hamiltonian operator is made up of one- and a two- electron terms. Any determinants which differ by more than two orbitals have a Hamiltonian matrix element of zero. If non-orthogonal orbitals are used, the more general formulae given by Löwdin are required.⁹⁷

HPHF functions (see section 2.7.2) are of the form

$$|X_i\rangle = \begin{cases} |D_i\rangle & \text{if closed shell} \\ \frac{1}{\sqrt{2}}(|D_i\rangle \pm |\overline{D}_i\rangle) & \text{otherwise,} \end{cases}$$

where $|\overline{D}_i\rangle$ is obtained from $|D_i\rangle$ by swapping all of the spins α and β of the unpaired electrons. This results in several simple forms for calculating the matrix elements between two arbitrary sites in an HPHF space;

$$\langle X_i|\hat{H}|X_j\rangle = \begin{cases} \langle D_i|\hat{H}|D_j\rangle & \text{if } X_i \text{ and } X_j \text{ are closed shell} \\ 0 & \text{if one of } X_i \text{ or } X_j \text{ is closed shell and the total} \\ & \text{spin is odd} \\ \sqrt{2}\langle D_i|\hat{H}|D_j\rangle & \text{if one of } X_i \text{ or } X_j \text{ is closed shell and the total} \\ & \text{spin is even} \\ \langle D_i|\hat{H}|D_j\rangle + \langle \overline{D}_i|\hat{H}|D_j\rangle & \text{if both } X_i \text{ and } X_j \text{ contain a } + \\ \langle D_i|\hat{H}|D_j\rangle - \langle \overline{D}_i|\hat{H}|D_j\rangle & \text{if both } X_i \text{ and } X_j \text{ contain a } - \\ 0 & \text{otherwise} \end{cases}$$

4.1.1 Excitation generation integration

The excitation matrix (which is formed of up to two orbitals, $\{m, n\}$, chosen to excite electrons from, and the orbitals $\{p, q\}$, chosen to excite electrons into) and the excitation level can be directly returned by the excitation generator, along with

the parity of the two excitations calculated from the newly generated (encoded) bit-representation of the target determinant. Given these and the source determinant, it is possible to calculate the diagonal matrix element without ever generating the expanded form of the target Slater determinant.

This can result in a significant computational saving, as the majority of the spawns are stochastically rejected, so the computational cost of generating a large number of determinant representations is avoided.

4.2 Configurational State Functions

When using CSFs, there are two fundamental approaches to evaluating the Hamiltonian matrix elements (which may then be developed into more specific cases for the different types of CSF); the CSFs may be expanded in terms of Slater determinants, or the matrix elements may be expanded as a sum over permutations.

Expansion in Slater determinants

It is possible to express each CSF, F_i , as a linear sum of appropriate Slater determinants, such that

$$|F_i\rangle = \sum_{\mathbf{j}} c_{\mathbf{j}}^{(i)} |D_{\mathbf{j}}\rangle.$$

It is then clear that it is possible to expand the calculation of Hamiltonian matrix elements as a double sum over the matrix elements between determinants with the same spatial structures, such that

$$\langle F_{\mu} | \hat{H} | F_{\nu} \rangle = \sum_{\mathbf{i}, \mathbf{j}} c_{\mathbf{i}}^{(\mu)*} c_{\mathbf{j}}^{(\nu)} \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{j}} \rangle. \quad (4.1)$$

Any regime that involves this double sum can at best scale as $\mathcal{O}(n_{det}^2)$, where n_{det} scales exponentially with N (see figure 4.1). Consequently, as the number of electrons is increased, the computation is going to be increasingly dominated by the matrix elements associated with highly excited states, which have the most unpaired electrons. These are also where the least meaningful wavefunction magnitude resides. This makes direct expansion an undesirable scheme (see Kotani-Yamanouchi CSFs for an example, section 4.4).

Summation over permutations

Starting from the most general formulation of the basis set obtained from a given spatial function, Φ ,

$$\langle F_\mu | \hat{H} | F_\nu \rangle = \langle \mathcal{A} \Phi_\mu \Theta_\mu | \hat{H} | \mathcal{A} \Phi_\nu \Theta_\nu \rangle,$$

and considering that the Hamiltonian operator is both spin-free (operates only on the spatial coordinates), and symmetric in the coordinates of the electrons (it commutes with every permutation and every element of the symmetric group algebra), the matrix elements simplify⁹⁸,

$$\langle F_\mu | \hat{H} | F_\nu \rangle = \frac{1}{N!} \sum_{\mathbf{P}} (-1)^P \langle \Phi_\mu | \hat{H} \mathbf{P} | \Phi_\nu \rangle \langle \Theta_\mu | \mathbf{P} | \Theta_\nu \rangle, \quad (4.2)$$

where the antisymmetriser, \mathcal{A} , has been explicitly written in terms of permutations and parity elements. The overlap matrix elements can be obtained in much the same way,

$$\langle F_\mu | F_\nu \rangle = \frac{1}{N!} \sum_{\mathbf{P}} (-1)^P \langle \Phi_\mu | \mathbf{P} | \Phi_\nu \rangle \langle \Theta_\mu | \mathbf{P} | \Theta_\nu \rangle.$$

As is outlined in figure 4.1, and the following section, algorithms that scale with the number of available permutations necessarily perform poorly. The number of permutations available scales factorially, i.e. super-exponentially, and as such a scheme that explicitly expands the CSFs in Slater determinants will actually scale better!

However, there are two specific instances in which this permutation scheme is of implementational utility;

1. if explicit, and small, limits are able to be placed on the permutations that need to be included, either through demonstration of which terms will be zero, or which ones will cancel, or
2. if group theory permits transformation of this expression into a more tractable form.

The use of Serber functions (sections 3.3.4 and 4.5), or spin-projected determinants (sections 3.3.5 and 4.6) are desirable, since their Hamiltonian matrix elements are made tractable through group theory when expressed in the form here.

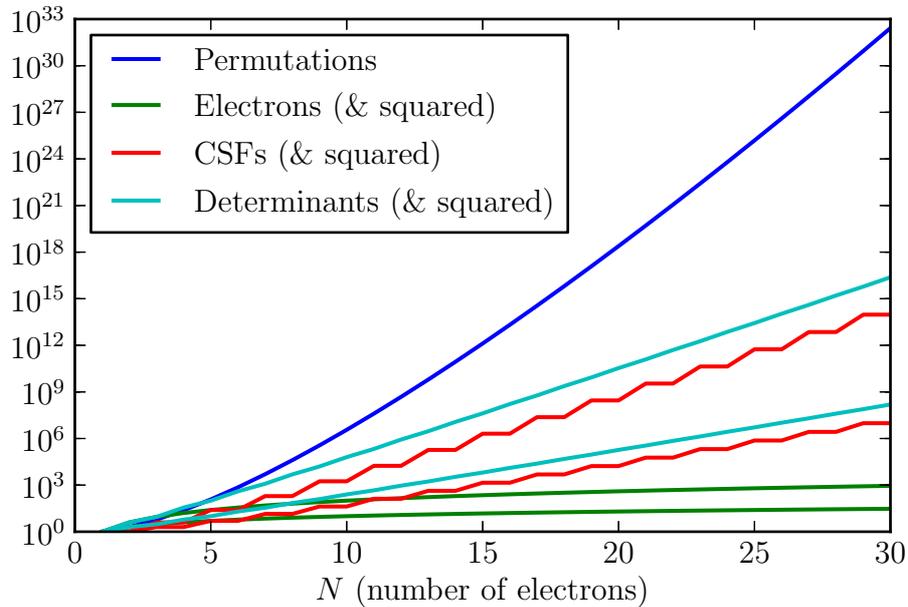


Figure 4.1: Scaling of representation components. For a given number of electrons, N , the plot shows the number of available electronic permutations, along with the number of available unique CSFs with the lowest available total spin, the number of corresponding Slater determinants and the number of electrons. The squares of these values are also plotted. Algorithms are preferred to scale depending on the number of electrons, and then CSFs present, and then the number of determinants required. Even scaling dependent on the squares of these values is preferable to being dependent on all permutations.

4.2.1 Scaling of matrix element calculations

There are numerous different parameters that could influence the scaling of matrix element evaluation, and thus be a factor in the overall scaling of an implementation of FCIQMC. It is worth noting that there is a degree of choice available, both in the choice of CSFs used, and to a lesser extent in the choice of implementation. The terms that the scaling could depend on are discussed here in order of preference here, and illustrated in figure 4.1.

Number of electrons, N

Calculating the Hamiltonian matrix elements between arbitrary Slater determinants scales in the worst-case as $\mathcal{O}(N^2)$, with the most commonly used cases being $\mathcal{O}(1)$ and $\mathcal{O}(N)$. Once the permutation representation matrices are known, the matrix elements corresponding to Serber functions (sec-

tion 3.3.4) can be reduced to a similar formulation. The matrix elements for spin-projected determinants scale similarly, but without the dependence on representation matrices.

Number of CSFs, n_{csf}

Although there are a great many elements of an FCIQMC simulation involving CSFs that depend on the number of CSFs, Hamiltonian matrix element generation is not directly one of these. The size of the permutation representation matrices is n_{csf}^2 , and the cost of calculating each of them scales in the same way. However, the potential permutational number of them required is a stricter limiting factor.

Number of Slater determinants, n_{det}

The trivial expansion of CSFs into linear combinations of determinants leads to $\mathcal{O}(n_{det}^2)$ computation, although by careful enumeration of determinants that differ by at most two spin-orbitals this can be reduced to $\mathcal{O}(n_{det})$ in a lot of cases.

Number of permutations, $N!$

The pathological scaling of the number of permutations means that no realistic implementation of matrix element calculation can rely on explicit expansion in terms of permutations. It is notable that the number of permutation representation matrices required for the use of Serber functions (in principle) scales as $N!$. In practice, due to the limits on a maximum of two spatial orbitals being excited, only a tiny subset of these are ever required, and if necessary even these could be shrunk by judicious use of symmetry (see section 4.5.3).

It should be noted from the logarithmic scale of figure 4.1 that it is *strongly* preferred to opt for scalings towards the top of this list.

4.3 Rumer-type CSF

Rumer-type CSFs are non-orthogonal, and the matrix elements are relatively computationally complicated — with the most straightforward means being an expansion into Slater determinants. Although there are means to do this more efficiently⁹⁹, it was not perceived as immediately worthwhile to implement Hamil-

tonian matrix elements for these CSFs, and will not be considered further here.

The primary purpose of considering Rumer-type spin eigenfunctions is as a means to calculate permutation representation matrices efficiently, see section 4.5.2. This is relatively efficient as the expansion into Slater determinants is more compact than for other types of spin eigenfunction.

4.4 Kotani-Yamanouchi CSFs

For two CSFs, $\{X_\mu, X_\nu\}$ (see section 3.3.3), the matrix elements, $\langle X_\mu | \hat{H} | X_\nu \rangle$, may be calculated according to the double sum over determinants expressed in equation 4.1. Various special cases for evaluating the CSF matrix elements can be used. Note that in these subsections there are many sums over one and two electron integrals of the form \sum_i^N , $\sum_i^{N_o}$ and \sum_i^g . These sums are written from index 1 for notational convenience, meaning “the i^{th} orbital”, “the i^{th} singly occupied orbital” and “the i^{th} doubly occupied orbital” respectively.

4.4.1 X_μ and X_ν share the same spatial structure

As the spatial structure of the two CSFs is identical, all contributing determinants have g doubly occupied orbitals ($2g$ paired electrons) and N_o singly occupied spatial orbitals. The Hamiltonian matrix elements can be constructed out of a number of components.

4.4.1.1 Diagonal terms in sum

Consider first the diagonal terms in equation 4.1 ($i = j$) and break the resultant expression into several pieces. Each of the component terms,

$$c_i^{(\mu)} c_i^{(\nu)} \langle D_i | \hat{H} | D_i \rangle,$$

may be considered in terms of the Slater–Condon rules given in table 4.1,

$$\frac{1}{2} \sum_{i,j} (\langle ij | \hat{g} | ij \rangle - \langle ij | \hat{g} | ji \rangle) + \sum_i h_{ii},$$

broken down into sums over spatial rather than spin orbitals.

One electron terms

The one electron terms depend only on the spatial orbitals occupied;

$$\langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{i}} \rangle = \sum_k^N h_{kk}.$$

This sum is over all occupied spin-orbitals in $D_{\mathbf{i}}$, such that any doubly occupied spatial orbitals are counted twice. As the spatial orbitals are the same for all considered determinants, this term is conserved in the sum across determinants. The overall contribution therefore becomes

$$\begin{aligned} H_{\text{contrib}} &= \sum_{\mathbf{i}} c_{\mathbf{i}}^{(\mu)} c_{\mathbf{j}}^{(\nu)} \langle D_{\mathbf{i}} | \hat{H}_{\text{one el}} | D_{\mathbf{i}} \rangle \\ &= \left(\sum_k^N h_{kk} \right) \left(\sum_{\mathbf{i}} c_{\mathbf{i}}^{(\mu)} c_{\mathbf{i}}^{(\nu)} \right) \\ &= \left(\sum_k^N h_{kk} \right) \delta_{\mu\nu}. \end{aligned} \quad (4.3)$$

Two electron terms, from the same orbital

For the diagonal terms in equation 4.1, consider the two electrons integrals in the same doubly occupied spatial orbital, k ;

$$\begin{aligned} &\frac{1}{2} \sum_{i,j} \left(\langle ij | \hat{g} | ij \rangle - \langle ij | \hat{g} | ji \rangle \right) \Big|_{i=k_{\alpha}, k_{\beta}; j=k_{\alpha}, k_{\beta}} \\ &= \frac{1}{2} \left[\langle k_{\alpha} k_{\alpha} | \hat{g} | k_{\alpha} k_{\alpha} \rangle - \langle k_{\alpha} k_{\alpha} | \hat{g} | k_{\alpha} k_{\alpha} \rangle + \right. \\ &\quad \langle k_{\alpha} k_{\beta} | \hat{g} | k_{\alpha} k_{\beta} \rangle - \langle k_{\alpha} k_{\beta} | \hat{g} | k_{\beta} k_{\alpha} \rangle + \\ &\quad \langle k_{\beta} k_{\alpha} | \hat{g} | k_{\beta} k_{\alpha} \rangle - \langle k_{\beta} k_{\alpha} | \hat{g} | k_{\alpha} k_{\beta} \rangle + \\ &\quad \left. \langle k_{\beta} k_{\beta} | \hat{g} | k_{\beta} k_{\beta} \rangle - \langle k_{\beta} k_{\beta} | \hat{g} | k_{\beta} k_{\beta} \rangle \right] \\ &= \langle kk | \hat{g} | kk \rangle. \end{aligned}$$

where k includes all g doubly occupied orbitals. Note that for singly occupied spatial orbitals the coulomb and exchange contributions cancel. As $\langle kk | \hat{g} | kk \rangle$ depends only on the spatial orbitals, this term is invariant between different determinants being considered. Thus the overall contribution is given in

spatial orbitals by

$$\begin{aligned} H_{contrib} &= \left(\sum_k^g \langle kk | \hat{g} | kk \rangle \right) \left(\sum_i c_i^{(\mu)} c_i^{(\nu)} \right) \\ &= \left(\sum_k^g \langle kk | \hat{g} | kk \rangle \right) \delta_{\mu\nu}. \end{aligned} \quad (4.4)$$

Two electron terms, doubly occupied — doubly occupied interactions

For the diagonal terms in equation 4.1, consider the two electron integrals between doubly occupied spatial orbitals k, l ;

$$\begin{aligned} & \frac{1}{2} \sum_{i,j} \left(\langle ij | \hat{g} | ij \rangle - \langle ij | \hat{g} | ji \rangle \right) \Big|_{i=k_\alpha, k_\beta; j=l_\alpha, l_\beta} \\ &= \frac{1}{2} \left[\langle k_\alpha l_\alpha | \hat{g} | k_\alpha l_\alpha \rangle - \langle k_\alpha l_\alpha | \hat{g} | k_\beta l_\alpha \rangle + \right. \\ & \quad \langle k_\alpha l_\beta | \hat{g} | k_\alpha l_\beta \rangle - \langle k_\alpha l_\beta | \hat{g} | k_\beta l_\alpha \rangle + \\ & \quad \langle k_\beta l_\alpha | \hat{g} | k_\beta l_\alpha \rangle - \langle k_\beta l_\alpha | \hat{g} | k_\alpha l_\beta \rangle + \\ & \quad \left. \langle k_\beta l_\beta | \hat{g} | k_\beta l_\beta \rangle - \langle k_\beta l_\beta | \hat{g} | k_\alpha l_\beta \rangle \right] + \text{terms } k \leftrightarrow l \\ &= 4 \langle kl | \hat{g} | kl \rangle - 2 \langle kl | \hat{g} | lk \rangle, \end{aligned}$$

where the final expression is in terms of spatial orbitals, and thus invariant between different determinants. The overall energy contribution is thus

$$H_{contrib} = \left(\sum_{k < l}^g 4 \langle kl | \hat{g} | kl \rangle - 2 \langle kl | \hat{g} | lk \rangle \right) \delta_{\mu\nu}. \quad (4.5)$$

Two electron terms, doubly occupied - singly occupied interactions

For the diagonal terms in equation 4.1, consider the two electron integrals between doubly occupied spatial orbital k and singly occupied orbital l . Each determinant has a well defined spin component for each orbital, labelled σ , taking the value α or β ($m_s = \pm \frac{1}{2}$). Thus the matrix element component for

orbitals k and l is given by

$$\begin{aligned}
& \frac{1}{2} \sum_{i,j} \left(\langle ij|\hat{g}i|ij\rangle \langle ij|\hat{g}|ji\rangle \right) \Big|_{i=k_\alpha, k_\beta; j=l_\sigma} \\
&= \frac{1}{2} \left[\langle k_\alpha l_\sigma|\hat{g}|k_\alpha l_\sigma\rangle - \langle k_\alpha l_\sigma|\hat{g}|l_\sigma k_\alpha\rangle + \right. \\
&\quad \left. \langle k_\beta l_\sigma|\hat{g}|k_\beta l_\sigma\rangle - \langle k_\beta l_\sigma|\hat{g}|l_\sigma k_\beta\rangle \right] + \text{terms } k \leftrightarrow l \\
&= 2 \langle kl|\hat{g}|kl\rangle - \langle kl|\hat{g}|lk\rangle.
\end{aligned}$$

As this expression is in terms of spatial orbitals, it is invariant between different determinants. Thus the final expression is given by

$$H_{contrib} = \left(\sum_k^g \sum_l^{N_o} 2 \langle kl|\hat{g}|kl\rangle - \langle kl|\hat{g}|lk\rangle \right) \delta_{\mu\nu}. \quad (4.6)$$

Two electron terms, singly occupied - singly occupied interactions

For the diagonal terms of equation 4.1, consider the two electron integrals between singly occupied orbitals k and l . The coulomb term remains constant, as it depends only on the spatial orbitals occupied. The exchange term, however, is only observed if the two orbitals have the same spin (i.e. $\sigma_k^i = \sigma_l^i$). To calculate this requires enumerating the determinants which will be done using the methodology described in section 4.4.1.2

$$H_{contrib} = \sum_{k<l}^{N_o} \sum_{\mathbf{i}} c_{\mathbf{i}}^{(\mu)} c_{\mathbf{j}}^{(\nu)} \left(\langle kl|\hat{g}|kl\rangle - \langle kl|\hat{g}|lk\rangle \delta_{\sigma_k^i \sigma_l^i} \right). \quad (4.7)$$

4.4.1.2 Off-diagonal terms in sum

Consider the off-diagonal terms in equation 4.1 (those with $i \neq j$). All of the relevant determinants need to be enumerated, selecting those with non-zero Hamiltonian matrix elements and summing the relevant elements.

By the Slater–Condon rules, all matrix elements are equal to zero if the determinants differ by more than two spin orbitals. The only differences between the determinants being considered are the spin structures of the unpaired electrons. The only way that a determinant can be changed by two or fewer spin orbitals, whilst maintaining the total value of M_s , is to swap the spins of one orbital with $M_s = \frac{1}{2}$ and one with $M_s = -\frac{1}{2}$. This will cause the determinants to differ by

precisely two spin orbitals.

An algorithm thus presents itself. Iterate through the list of enumerated determinants. For each determinant, iterate through the unpaired electrons, selecting all possible pairs with $M_s = \pm\frac{1}{2}$. The excitation matrix $m, n \rightarrow p, q$ is then directly obtainable, along with the corresponding matrix element.

The first step is to enumerate all possible determinants. This is equivalent to generating all the combinations of n_α electrons with spin $m_s = \frac{1}{2}$ selected from n_{open} unpaired electrons. The ordered set of combinations is generated as the *lexicographic combinations* using an algorithm presented by Knuth¹⁰⁰. These have the convenient property that, given a particular combination, its position in the list can be obtained directly. Due to the heavy use of the `choose` function, calculation of the index in the list is quite slow.

If the selected positions (*pos*) of the β electrons ($m_s = -\frac{1}{2}$) within the set of unpaired electrons are stored, then the sum $\sum_{pos} 2^{pos-1}$ is monotonically increasing through the set of combinations. This value is calculated for each of the permutations generated, and stored. Given any arbitrary permutation this value may be calculated easily giving its index in the list by performing a binary search. This is substantially faster than calculating the index directly.

The required sum is

$$\begin{aligned} H_{contrib} &= \sum_{i \neq j} c_i^{(\mu)} c_j^{(\nu)} \langle D_i | \hat{H} | D_j \rangle \\ &= \sum_{i < j} (c_i^{(\mu)} c_j^{(\nu)} + c_i^{(\nu)} c_j^{(\mu)}) \langle D_i | \hat{H} | D_j \rangle. \end{aligned}$$

The contribution of each of the pairs of determinants only needs to be calculated once, with the $c_i^{(\nu)} c_j^{(\mu)}$ term summing in the reverse interaction. Therefore each pair of determinants need only be generated once. A condition may be added that two spins will only be swapped if the first has $m_s = \frac{1}{2}$, and the second has $m_s = -\frac{1}{2}$. This ensures that if the connection $A \rightarrow B$ is made, $B \rightarrow A$ is rejected, and double counting is avoided.

The algorithm chosen is as follows:

1. Enumerate all determinants which may be components of CSFs X_μ and X_ν .
2. Select the first determinant, D_i , in the list.

3. Generate all pairs, (s, t) , of unpaired electrons in D_i such that the first electron in the pair has $m_s = \frac{1}{2}$, and the second has $m_s = -\frac{1}{2}$.
4. For each of these pairs the associated double excitation is $s_\alpha t_\beta \rightarrow s_\beta t_\alpha$.
5. The matrix element, $\langle D_i | \hat{H} | D_j \rangle$, can be calculated using the Slater–Condon rules. It is multiplied by $(c_i^{(\mu)} c_j^{(\nu)} + c_i^{(\nu)} c_j^{(\mu)})$ and added to the overall sum.
6. If there are more determinants, select the next determinant in the list and return to item 3.

4.4.2 X_μ and X_ν differ by one spatial orbital

This is the least frequently used case, but also the hardest to simplify in any meaningful way. It is calculated according to the general formula expressed in equation 4.1.

4.4.3 X_μ and X_ν differ by two spatial orbitals

All of the component determinants of the two CSFs must differ by at least two spin orbitals. By the Slater–Condon rules (table 4.1), the only pairs of determinants that contribute to the sum in equation 4.1 are those which differ by precisely 2 spin orbitals.

Therefore the only spin orbitals which may differ between pairs of determinants being considered are the differing spatial orbitals. The remainder of the two determinants must share the same spin structure. This allows all of the contributing pairs to be directly generated, rather than performing the double sum and rejecting the components which do not contribute (a much slower process).

If one determinant has a doubly occupied orbital, which is singly occupied in the other determinant, then both of the orbitals in the second determinant can have their spin structure permuted, whereas neither of the doubly occupied orbitals may be (as this would not change the determinant, see figure 4.2). This leaves a maximum of four orbitals which may vary per determinant. The number of these is labelled γ . The orbitals which may vary can be obtained straightforwardly from the spatial structures of the CSFs provided.

The spin structure of a determinant may be represented as a binary string (for $m_s = \pm\frac{1}{2}$). Using the lexicographic combinations, a sequence of spin structures

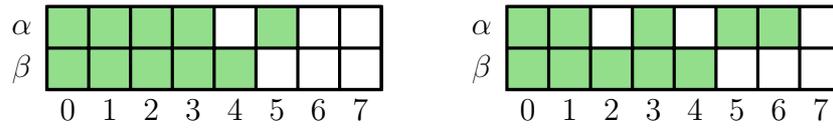


Figure 4.2: Determinants differing by one spin orbital; occupied spin orbitals are represented in green with spatial orbitals enumerated along the base. Note that due to the difference being caused by a doubly occupied orbital becoming two singly occupied orbitals, there are *two* spatial orbitals whose spins may be varied whilst the two determinants only differ by one spin orbital. In the case of CSFs differing by 2 spatial orbitals, as discussed, the maximum number of orbitals whose spins are allowed to vary is 4.

where the last bits change the most rapidly may be generated — in particular for every combination of the first $N - \gamma$ bits, all of the allowed combinations of the last γ bits will be explored before the any of first $N - \gamma$ bits change.

Assigning these rapidly changing bits to the spatial orbitals which are allowed to change within a determinant, whilst remaining connected to another determinant with the same majority spin structure, it then becomes straightforward to access all of the connected determinant pairs. Two lists of (carefully ordered) determinants for the two CSFs being considered are obtained.

All of the contributing terms in the sum (equation 4.1) are now roughly diagonal (see figure 4.3). By the properties of the lexicographic combinations, all of the majority spin structures appear in the same order in both of the two lists. Not all of the majority spin structures may be present in each list as there may be different numbers of unpaired electrons in each CSF giving subtly different constraints on the spin structures allowed. These non-allowed terms may be avoided by considering the summed M_s values of the $N - \gamma$ orbitals. The CSF with the ‘additional’ determinants will always be the one with the larger number of unpaired electrons, as this gives more combinatorial freedom for selection of determinants.

The process is as follows for CSFs X_μ, X_ν , containing sorted lists of determinants $D_{\mathbf{i}}^{(\mu)}, D_{\mathbf{j}}^{(\nu)}$.

1. Set $\mathbf{i} = \mathbf{j} = \mathbf{1}$
2. Consider determinants $D_{\mathbf{i}}^{(\mu)}, D_{\mathbf{j}}^{(\nu)}$. If $\sum_k^{N-\gamma(\mu,k)} m_s(k, \mu) = \sum_l^{N-\gamma(\nu,l)} m_s(l, \nu)$ go to step 4.
3. If X_μ has more unpaired electrons than X_ν , increment \mathbf{i} unless \mathbf{i} is the last determinant of list $D_{\mathbf{i}}^{(\mu)}$ when the sum is complete. If X_ν has more unpaired

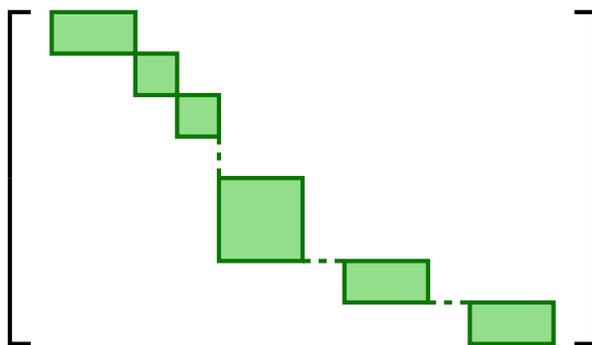


Figure 4.3: Summation of matrix element components for X_μ and X_ν which differ by two spatial orbitals. Terms are summed roughly diagonally between determinants. Only those orbitals which differ spatially are allowed to vary in terms of spin across the summed blocks (shaded green). When any change, other than within a block, is made to one CSF, another block is required for the other CSF, making a move roughly diagonally. This avoids calculation of a (potentially) large number of terms which are known to be equal to zero.

electrons, then similarly for \mathbf{j} . Return to step 2.

4. Perform the sum $\sum_{\mathbf{i}'} \sum_{\mathbf{j}'} c_{\mathbf{i}'}^{(\mu)} c_{\mathbf{j}'}^{(\nu)} \langle D_{\mathbf{i}'}^{(\mu)} | \hat{H} | D_{\mathbf{j}'}^{(\nu)} \rangle$, where $\mathbf{i}' \geq \mathbf{i}, \mathbf{j}' \geq \mathbf{j}$ and \mathbf{i}', \mathbf{j}' are constrained such that the majority components of the spin structure are not allowed to change. Add this to the overall matrix element.
5. Set \mathbf{i}, \mathbf{j} to the values $\mathbf{i}' + \mathbf{1}$ and $\mathbf{j}' + \mathbf{1}$ from the last term in the above sum (i.e. change the spin structure of both determinants at once). If this moves past the end of either list, then exit the sum. Otherwise go to step 2.

The matrix elements $\langle D_{\mathbf{i}}^{(\mu)} | \hat{H} | D_{\mathbf{j}}^{(\nu)} \rangle$ depend only on two 4-index, 2-electron integrals by the Slater–Condon rules, characterised by the excitation matrix $m, n \rightarrow p, q$. The values of $\langle mn | \hat{g} | pq \rangle$ and $\langle mn | \hat{g} | qp \rangle$ are numerically invariant to changes within the same spatial orbitals. Some of these values are, however, negative and some disappear due to constraints imposed by the m_s values.

The value required may be expressed as a standard coulomb interaction $\langle mn | \hat{g} | pq \rangle$ and an exchange interaction $\langle mn | \hat{g} | qp \rangle$. When the exchange interaction is forbidden by the component m_s values, and when the overall sign must be inverted to represent the spin orbital excitation matrix rather than the ‘standard’ excitation matrix (which assumes all unpaired electrons have $m_s = -\frac{1}{2}$ as a result of the standard representation ordering, see section 3.4.1) needs to be determined. This can be considered in two parts, the exchange integral components and the overall parity.

Exchange Integral

The m_s values for the target orbitals p, q are given by the determinantal excitation matrix. The m_s values for m, n can be obtained by examining the two determinants. There are three cases.

1. The excitation is from a singly occupied orbital. The m_s value depends on the m_s value of the orbital in the second determinant being considered as part of the pair.
2. The excitation is from a doubly occupied orbital. The m_s value is opposite to the m_s value of the remaining half of the pair in the same determinant.
3. Both excitations have occurred from the same doubly occupied orbital. The m_s value choices for m, n are then arbitrary, as long as one is assigned $\frac{1}{2}$ and the other $-\frac{1}{2}$.

When considering a pair of determinants $D_i^{(\mu)}, D_j^{(\nu)}$, the exchange interaction can now be included if $m_s(m) = m_s(q)$ and $m_s(n) = m_s(p)$.

Overall Parity

The overall sign of the interaction is known as the parity. If the orbitals are sorted in increasing numerical order (this is arbitrary but must be the same order used when storing integrals), then the number of permutations of adjacent orbitals required to line up the determinants will either be even (in which case the parity is positive) or odd (in which case the parity is negative).

Initially the parity is calculated for the ‘standard’ representation with all electrons in the beta position, this is done by considering the lineup operation in full. For every singly occupied orbital in X_μ which corresponds to a doubly occupied orbital in X_ν (see figure 4.4b), if that orbital has $m_s = \frac{1}{2}$ then the parity is inverted. The effect on the parity of each of the determinants in both lists can be calculated by considering this, and the parity of the overall element in the sum is then given by the product of these values.

This process is substantially quicker than calculating the parity from scratch for every pair of determinants.

Once the parity has been obtained and the exchange interaction determined, the

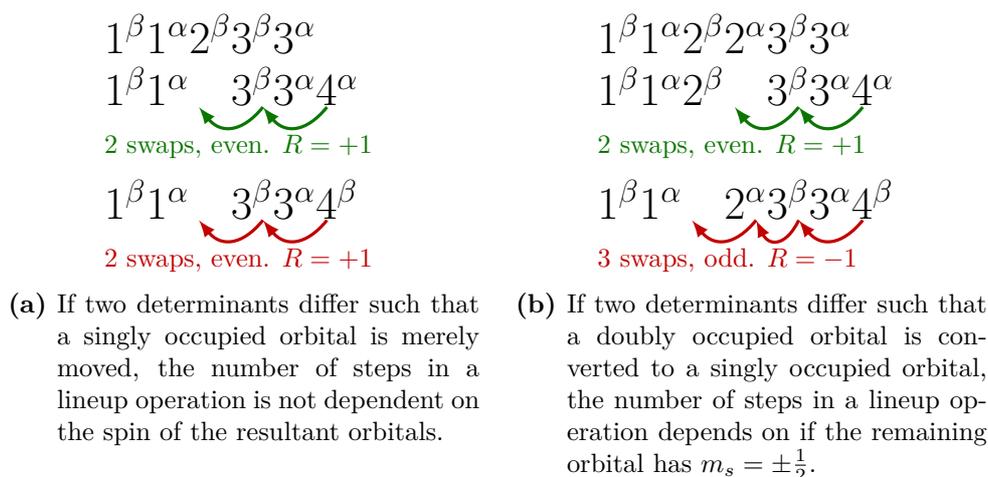


Figure 4.4: Dependence of the parity of matrix elements on the relative spin configuration of the associated determinants.

contribution of each orbital is given by

$$\langle D_i^{(\mu)} | \hat{H} | D_j^{(\nu)} \rangle = R \cdot \left(\langle m' n' | \hat{g} | p' q' \rangle - \delta_{m_s(m) m_s(q)} \delta_{m_s(n) m_s(p)} \langle m' n' | \hat{g} | q' p' \rangle \right)$$

where $R = \pm 1$ indicates the parity, and m', n', p' and q' index the spatial orbitals associated with the excitation (i.e. these are constant for all $D_i^{(\mu)}$ and $D_j^{(\nu)}$).

4.5 Serber-type CSFs

Salmon et al.⁸³ have demonstrated that the Hamiltonian matrix elements between Serber CSFs can be expressed by formulae which are extremely similar in structure to those for Slater determinants as displayed in table 4.1. In particular, the matrix elements are expressed as a sum over electronic terms of the Hamiltonian matrix elements between spatial orbitals — the only difference is that these terms are multiplied by factors dependent on the line-up permutation between the two CSFs and the spin eigenfunction labels applied.

There are several pieces of information that are required to generate the matrix elements;

- the ordered list of spatial orbitals occupied in one of the two CSFs,
- the Serber function label for each of the CSFs,
- the list of orbitals differing between the two sets of spatial orbitals, and

- the line-up permutation between the (natural orders of the) two sets of spatial orbitals.

All of these pieces of information may be obtained directly from the excitation generators, or as explained in the next sections. These can then be passed to the matrix element calculation routine proper (section 4.5.4).

4.5.1 Line-up permutation

The spatial component of the CSFs is represented as a list of occupied spatial orbitals, in a defined order (g doubly occupied orbitals followed by N_o singly occupied orbitals, see section 3.4.1). For the two CSFs X_μ and X_ν , the line-up permutation, \mathcal{L} is represented as a vector, \mathbf{L} , of indices such that L_i is

- the index of the i^{th} orbital of X_ν in X_μ , if the orbital is found in both functions, or
- the index of the i^{th} orbital in X_μ which is not found in X_ν , if this is the i^{th} orbital in X_ν that which is not found in X_μ .

This permutation can be found in $\mathcal{O}(N)$ time for any two arbitrary, ordered CSFs. It may also be found directly when generating an excitation.

4.5.2 Generation of permutation representation matrices

For the calculation of these matrix elements, the permutation representation matrix elements,

$$[\mathbf{P}]_{\theta_i \theta_j}^{NS} = \langle \theta_i | \mathbf{P} \theta_j \rangle, \quad (4.8)$$

are required. Calculation of these matrix elements is time-consuming. The most obvious manner to do this is to expand the Serber functions in terms of determinants, and then to permute the electrons and to sum the components of the resultant overlap terms.

Given that the transformation matrix between Serber and Kotani-Yamanouchi spin eigenfunctions is readily available, it is possible to generate the permutation representation matrices for the Serber functions from those for the Kotani-Yamanouchi functions, $U_{\theta_i \theta_j}(\mathbf{P})$. There are several efficient means to calculate the representation matrices for the Kotani-Yamanouchi functions, and the $\mathcal{O}(n_{csf}^2)$ operation

to transform the matrices is more efficient than the $\mathcal{O}(n_{det}^2)$ operation involved in expanding into determinants.

Wu and Zhang¹⁰¹ presented an algorithm for determining the representation matrices of the Kotani-Yamanouchi functions as a product of three matrices, two of which are triangular, and all of which are trivially obtained. The representation matrices of arbitrary permutations may also be obtained by combining the transpositions (i, j) . These can in turn be reduced to a dependence on the primitive transpositions

$$(i, j+1) = (j, j+1)(i, j)(j, j+1) \quad j \neq i$$

$$U_{\theta_l \theta_k}(i, j+1) = \sum_{m,n} U_{\theta_l \theta_m}(j, j+1) U_{\theta_m \theta_n}(i, j) U_{\theta_n \theta_k}(j, j+1) \quad j \neq i.$$

Rettrup¹⁰² provides a direct means to obtain the representation matrix elements for the primitive transpositions, and permit the above double sum to be reduced to (a maximum of four) contributing terms;

$$U_{\theta_l \theta_k}(i, j+1) = \rho_{j,j+1}^l \rho_{j,j+1}^k U_{\theta_l \theta_k}(i, j) - \rho_{j,j+1}^l \gamma_{j,j+1}^k U_{\theta_l \theta_n}(i, j) \\ - \rho_{j,j+1}^k \gamma_{j,j+1}^l U_{\theta_m \theta_k}(i, j) + \gamma_{j,j+1}^l \gamma_{j,j+1}^k U_{\theta_m \theta_n}(i, j)$$

where n, m are only included for terms such that

$$\theta_n = (j, j+1)\theta_k, \\ \theta_m = (j, j+1)\theta_l$$

are valid Kotani-Yamanouchi functions, and

$$U_{\theta_l \theta_l}(j, j+1) = \frac{1}{d_{j,j+1}^l} \equiv -\rho_{j,j+1}^l \\ U_{\theta_l \theta_m}(j, j+1) = \begin{cases} \sqrt{1 - (\rho_{j,j+1}^l)^2} \equiv \gamma_{j,j+1}^l & \text{if } S_m = (j, j+1)S_l \\ 0 & \text{otherwise} \end{cases}$$

where $d_{j,j+1}^l$ is the axial distance between the numbers j and $j+1$ in the l^{th} standard Young tableau^{103,104} (corresponding to θ_l). Other authors have also provided efficient means for calculating these elements^{105,106}, or for representations of cyclic permutations^{107,108}. The use of all of these methods for arbitrary permutations is difficult.

Alternatively, the permutation matrices for the Kotani-Yamanouchi functions may be obtained by transformation from those for the Rumer functions. As the expansion of Rumer functions into combinations of Slater determinants is especially compact, this route to generating the matrices is particularly efficient.

A highly efficient code, called SPINS, has been written by Karadakov et al.¹⁰⁹ to manipulate different forms of spin function and associated permutation representation matrices. Sections of this code have been integrated into the Alavi group code, called NECI, to facilitate working with arbitrary permutation matrices of Serber functions. This code generates entire permutation representation matrices, first in a Rumer function basis, before transforming them into a Kotani-Yamanouchi function basis and finally into a Serber function basis. There is no straightforward means to generate only specific elements of these matrices efficiently.

4.5.3 Storage of permutation representation matrices

Although the computational cost for producing each element of the permutation representation matrices is relatively low, it is necessary to calculate entire matrices at a time. The size of these matrices is n_{csf}^2 , and the cost of calculating multiple entire permutation representation matrices for each evaluated Hamiltonian matrix element quickly becomes prohibitive as the number of electrons increases.

Consequently it is important to only calculate each of the representation matrices once, and to store them for later access. In principle there are $N!$ available permutation matrices, but as the non-zero Hamiltonian matrix elements correspond to CSFs differing by a maximum of two spatial orbitals only a tiny fraction of the permutations are ever required (see figures 4.5 and 4.6).

Permutation representation matrices are generated on a calculate-on-demand basis — a process otherwise known as *memoisation*. A unique identifier, η , for each permutation may be calculated as

$$\eta(\mathbf{P}) = \sum_i^N (P_i - 1) \times N^{i-1}$$

and is used as the index to a `map*` which allows efficient lookup of previously

*The `map` function used is either `std::map`¹¹⁰ or, if a sufficiently recent compiler is available, `std::unordered_map`¹¹¹ as defined by the C++ standard. Wrapper code is available to

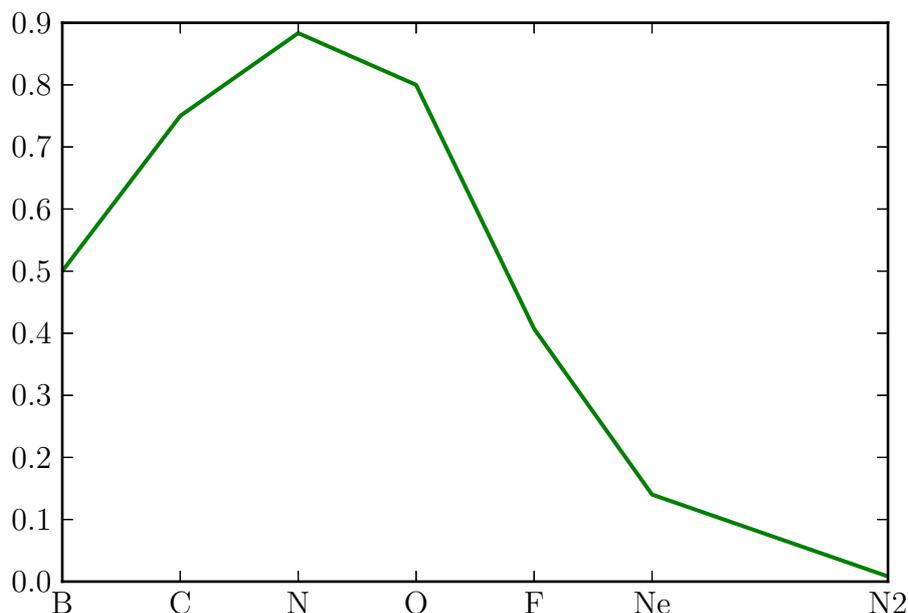
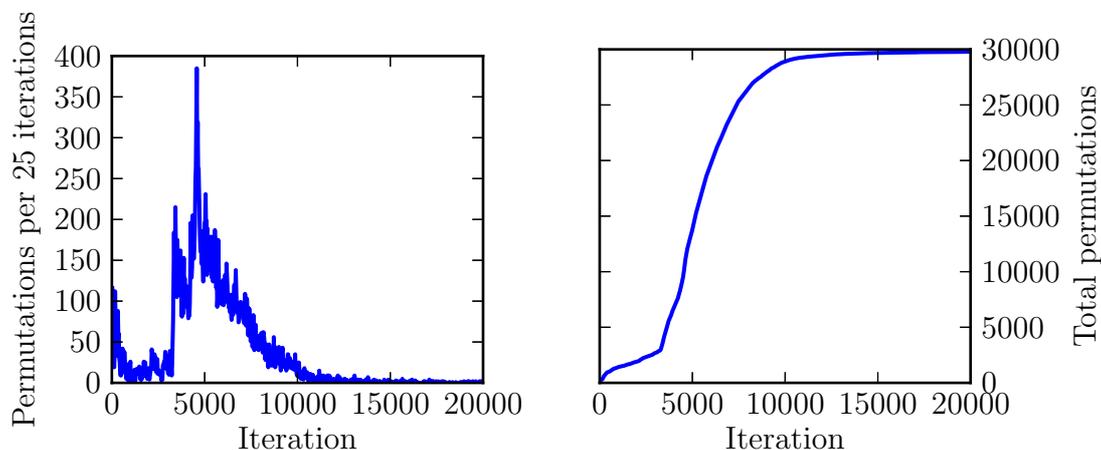


Figure 4.5: The proportion of the available N -electron permutations required in FCIQMC simulations as the number of electrons is increased. All of these calculations used cc-pVDZ basis sets with core electrons frozen, such that N varies from three to ten.

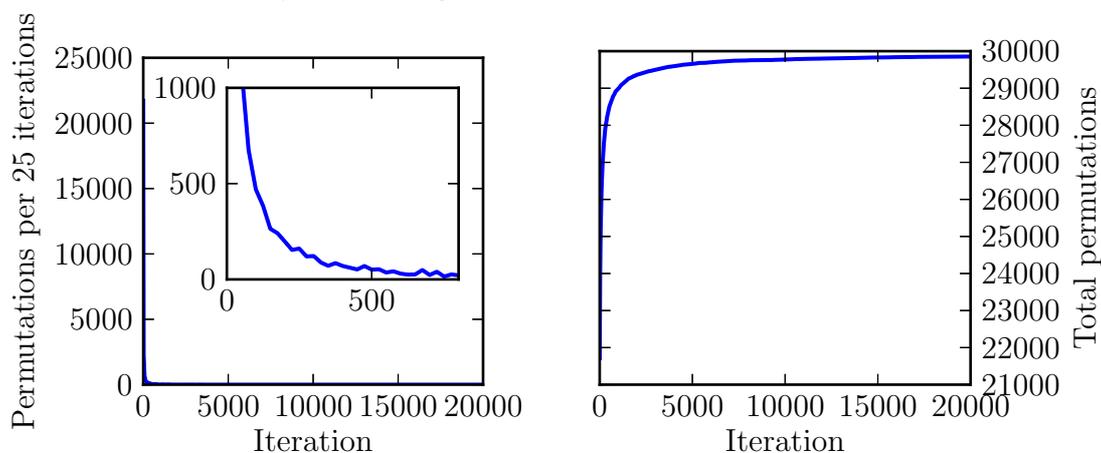
calculated matrices. If 64-bit integers are used, all possible permutations can be given a unique value with no possibility of collisions caused by integer overflows for up to 16 electrons. This permits each of required permutations to be only calculated (and stored) once*, saving substantially on resources. See figure 4.6.

integrate this efficiently with Fortran dynamic memory management.

*‘Once’ in this context means ‘once per MPI process’. In reality the vast majority of permutations required on each processor are shared across all of them. It is perfectly possible to share this memory across all processors situated on each physical node. This requires careful use of resource locking and synchronisation, as only one process can be permitted to write to the shared data structure at a time. Fortunately writing is a rare occurrence.



(a) Generation of permutations in an *i*-FCIQMC run using Serber functions. It is worth noting the shoulder in the total permutation count graph, showing how the number of permutations required rises rapidly as a new region of the Hilbert space becomes accessible.



(b) Generation of permutations after resuming the calculation in 4.6a from a dump file after 43225 iterations. Note the huge spike as most required permutations are calculated in one iteration. Resuming these calculations takes longer than for other types of basis function.

Figure 4.6: The on-the-fly generation of permutations for a stretched N_2 system in a cc-pVDZ basis set, in a calculation using 10 electrons. The number of permutations available is $10! \approx 3.6 \cdot 10^6$, whereas only just under $3 \cdot 10^4$ permutations are made use of (i.e. fewer than 1%).

4.5.4 Evaluation of matrix elements

The normalisation constant for the matrix elements is given by

$$N(\phi_a^i, \phi_b^i, \phi_r^j, \phi_s^j) = \left(\frac{n(\phi_a, \Phi_i)n(\phi_b, \Phi_i)n(\phi_r, \Phi_j)n(\phi_s, \Phi_j)}{[1 + \delta(\phi_a, \phi_b)]^3[1 + \delta(\phi_r, \phi_s)]^3} \right)^{\frac{1}{2}},$$

where $n(\phi_a, \Phi_i)$ is defined to be the occupancy number of orbital ϕ_a in the spatial function Φ_i . For most of the cases considered this simplifies into a more tractable form as given in each of the expressions below. The formulae used to calculate matrix elements, $\langle X_i | \hat{H} | X_j \rangle$, depend on how many spatial orbitals differ between the two CSFs⁸³. If they differ by more than two spatial orbitals, the matrix elements must equal zero. The remainder of this section presents formulae for the matrix elements between Serber functions differing by different numbers of spatial orbitals.

Same spatial structure

If the two spatial structures are identical, $\Phi_i = \Phi_j$, the two basis functions can differ only by their Serber functions,

$$\begin{aligned} \langle X_i | \hat{H} | X_j \rangle &= \langle \Phi_i \theta_i | \hat{H} | \Phi_j \theta_j \rangle \\ &= \delta(\theta_i, \theta_j) \sum_{\phi_a} \{ n(\phi_a) h_{ii} + [n(\phi_a) - 1] \langle \phi_a \phi_a | \hat{g} | \phi_a \phi_a \rangle \} \\ &\quad + \sum_{\phi_a < \phi_b} n(\phi_a) n(\phi_b) \{ \delta(\theta_i, \theta_j) \langle \phi_a \phi_b | \hat{g} | \phi_a \phi_b \rangle \\ &\quad \quad \quad - [(a, b)_{\theta_i, \theta_j}^{N, S} \langle \phi_a \phi_b | \hat{g} | \phi_b \phi_a \rangle] \}. \end{aligned}$$

Differ by one spatial orbital

If the two spatial structures differ by one orbital, such that ϕ_μ in Φ_i becomes ϕ_σ in Φ_j ,

$$\begin{aligned} \langle X_i | \hat{H} | X_j \rangle &= \langle \Phi_i \theta_i | \hat{H} | \Phi_j \theta_j \rangle \\ &= \sqrt{n(\phi_\mu, \Phi_i) n(\phi_\sigma, \Phi_j)} \\ &\quad \times \left[(-1)^l [\mathcal{L}]_{\theta_i, \theta_j}^{N, S} \{ h_{\phi_\mu \phi_\sigma} + [n(\phi_\mu, \Phi_i) - 1] \langle \phi_\mu \phi_\mu | \hat{g} | \phi_\sigma \phi_\mu \rangle \right. \\ &\quad \quad \quad \left. + [n(\phi_\sigma, \Phi_j) - 1] \langle \phi_\sigma \phi_\sigma | \hat{g} | \phi_\mu \phi_\sigma \rangle \right] \\ &\quad + \sum_{\phi_j \neq \phi_\mu, \phi_\sigma} n(\phi_j, \Phi_i) \left\{ (-1)^l [\mathcal{L}]_{\theta_i, \theta_j}^{N, S} \langle \phi_\mu \phi_j | \hat{g} | \phi_\sigma \phi_j \rangle \right. \\ &\quad \quad \quad \left. - (-1)^l [(\mu j) \mathcal{L}]_{\theta_i, \theta_j}^{N, S} \langle \phi_\mu \phi_j | \hat{g} | \phi_j \phi_\sigma \rangle \right\}. \end{aligned}$$

Differ by two spatial orbitals

If the two spatial structures differ by two spatial orbitals, such that ϕ_μ, ϕ_ν in Φ_i becomes ϕ_σ, ϕ_τ in Φ_j ,

$$\begin{aligned} \langle X_i | \hat{H} | X_j \rangle &= \langle \Phi_i; \theta_i | \hat{H} | \Phi_j; \theta_j \rangle \\ &= N(\phi_\mu, \phi_\nu; \phi_\sigma, \phi_\tau) \left\{ (-1)^l [\mathcal{L}]_{\theta_i, \theta_j}^{N, S} \langle \phi_\mu \phi_\nu | \hat{g} | \phi_\sigma \phi_\tau \rangle \right. \\ &\quad \left. - (-1)^l [(\mu, \nu) \mathcal{L}]_{\theta_i, \theta_j}^{N, S} \langle \phi_\mu \phi_\nu | \hat{g} | \phi_\tau \phi_\sigma \rangle \right\}. \end{aligned}$$

These formulae are extremely useful, as they reduce the calculation of the Hamiltonian matrix elements to a form which is essentially the same as the standard Slater–Condon rules modified only by the addition of components from the representation matrices.

4.6 Spin-projected determinants

Using spin-projected determinants as a basis set introduces a number of complexities, most notably those associated with non-orthogonality. The reason for considering this additional complexity is due to the efficiency of the matrix element calculation, as presented by Harris¹¹².

It is necessary to both be able to calculate matrix elements between arbitrary spin-projected determinants — mainly for debugging purposes — and more normally to evaluate diagonal matrix elements and off-diagonal matrix elements in conjunction with excitation generation. In order to evaluate these matrix elements a significant amount of information is required. In this chapter, the means of generating this information are explored first, followed by the matrix element calculation itself.

The first step of a calculation is to select which of the determinants, F_i, F_j , has the most pairs of electrons. The determinant with the most paired electrons is labelled as F_μ , and the other determinant is labelled as F_ν . If both of the determinants have the same number of unpaired electrons, it doesn't matter which way around this selection is made.

As described in section 3.3.5.1, if the spin-projected determinants are being explicitly normalised, such that $S_{\nu\nu} = \langle F_\nu | F_\nu \rangle = 1$, the matrix elements (along with

the overlap elements) must be adjusted to take this into account, such that

$$H_{\mu\nu} \longrightarrow \frac{1}{\sqrt{S_{\mu\mu}S_{\nu\nu}}} H_{\mu\nu}$$

$$S_{\mu\nu} \longrightarrow \frac{1}{\sqrt{S_{\mu\mu}S_{\nu\nu}}} S_{\mu\nu}.$$

The efficiency of this scheme for evaluating matrix elements depends very strongly on tight integration with the excitation generator. In particular the most computationally expensive step in the general case is aligning F_μ to meet the requirements of the following section. Within the excitation generator, the number of orbitals which change are known along with all details of the excitation — and as a consequence the information outlined below may be obtained directly. See appendix 4.6.4.

4.6.1 Alignment

All spin-projected determinants, F_i , are stored in a standardised form, such that the first $N - N_o$ electrons are pairs of doubly-occupied orbitals in increasing numerical order, followed by the remaining N_o unpaired electrons also in increasing numerical order.

Given a pair of spin-projected determinants, F_μ, F_ν , ordered in this way, the permutation \mathbf{P}_0^* is defined such that

1. $\mathbf{P}_0\theta_\mu$ contains the sequence $\alpha\beta\alpha\beta\cdots$ for all spins corresponding to doubly occupied orbitals in Φ_ν .
2. subject to (1), all spatial orbitals in $\mathbf{P}_0\Phi_\mu$ correspond as closely as possible to those in Φ_ν .

The parity is given by $\epsilon_P = (-1)^P$ where P is defined as the number of primitive transpositions required to construct the permutation \mathbf{P} . P is not unique, but is well defined as odd or even, and can be obtained by an $\mathcal{O}(N^2)$ operation (best case $\mathcal{O}(N)$). The following algorithm performs a line-up operation, shuffling the remaining terms up as each successive orbital is found and removed.[†]

*In the literature the reverse of this permutation is defined, such that the given conditions apply to \mathbf{P}_0^{-1} . As we are not attempting derive the expressions for evaluating the matrix elements, and are only interested in the parity and value of k (see later), the simpler expression is used.

[†]If the extra line $\Phi_{tmp}(i) = orb$ were added after line 10, then Φ_{tmp} would become ordered.

```

1:  $\Phi_{tmp} \leftarrow \mathbf{P}_0 \Phi_\mu$ 
2:  $par \leftarrow 1$ 
3: for  $i = 1$  to  $N$  do
4:    $orb \leftarrow \Phi_\mu(i)$ 
5:   for  $j = i$  to  $N$  do
6:     if  $\Phi_{tmp}(j) = orb$  then
7:       break
8:     end if
9:     if  $j \neq i$  then
10:       $\Phi_{tmp}(i + 1 : j) = \Phi_{tmp}(i : j - 1)$ 
11:      if  $\text{mod}(j - i, 2) = 1$  then
12:         $par \leftarrow -par$ 
13:      end if
14:    end if
15:  end for
16: end for

```

Implementationally, performing an actual line-up operation, and calculating the parity long-hand is extremely inefficient. Given that the system is restricted to double spatial excitations, once the excitation generator has selected a target spin structure, determining the optimum aligned determinant (and hence the parity) is straightforward — although it requires considering all possible ‘shapes’ of excitation.

4.6.2 Sanibel coefficient

It can be shown¹¹³ that matrix elements between F_μ and F_ν depend only on the difference in the spatial structures and on the value k , which is defined as the number of α spins in $\mathbf{P}_0\theta_\mu$ that match with β spins in θ_ν (and vice versa).

This value enters the energy expressions through the Sanibel coefficient.* These have been extensively studied, and various formulae have been given for calculating them. The Sanibel coefficient may be defined as the overlap of the k^{th} spin-

*The different formulae for these coefficients were presented and discussed at a conference at Sanibel, on Sanibel Island, in 1962.

projected determinant with a reference one,

$$C_k(S, M_s, N_o) = \langle \theta_k | \hat{O}_S \theta_0 \rangle,$$

where $\theta_0(n_\alpha, n_\beta) = [\alpha^{n_\alpha}][\beta^{n_\beta}]$, and to construct θ_k , the final k β terms are swapped with the final k α terms; i.e. $\theta_k(n_\alpha, n_\beta) = [\alpha^{n_\alpha - k}][\beta^k][\beta^{n_\beta - k}][\alpha^k]$. Löwdin¹¹⁴ obtained recurrence formulae for these coefficients, but more direct approaches have been obtained since. Sasaki and Ohno¹¹⁵ found

$$C_k(S, M_s, N_o) = \frac{(2S+1)(N_o - M_s - k)!k!}{(N_o + S + 1)!} \\ \times \sum_l (-1)^{k+l} \frac{(N_o + M_s - k + l)!(S - M_s + k - l)!}{l!(k-l)!(S - M_s - l)!(N_o - S - k + l)!}.$$

A similar derivation was given by Pauncz¹¹⁶, although his expression was determined only up to a constant (with the exception of $M_s = 0$). Via a (somewhat exciting) series of integrals, Smith^{117,118} obtains

$$C_k(S, M_s, N_o) = \frac{4S+2}{2S+N_o+2} \\ \times \sum_{j=0}^{S-M_s} (-1)^{S-M_s+k-j} \binom{S-M_s}{j} \binom{S+M_s}{S-M_s-j} \left(\frac{N_o}{2} + S \right)^{-1}$$

for $n_\alpha \geq n_\beta$, to which Harris¹¹² agrees, although giving a subtly different form,

$$C_k(S, M_s, N_o) = \frac{4S+2}{2S+N_o+2} \sum_{j=0}^{S-|M_s|} (-1)^{k+j} \binom{S-M_s}{j} \binom{S+M_s}{j} \left(\frac{N_o}{2} + S \right)^{-1}.$$

Smith and Harris¹¹⁹ demonstrate that all these different forms for the Sanibel coefficient are equivalent.

As the spin-pure subspace is always smaller than the deterministic Hilbert space with a given M_s , it is always beneficial to work in the regime where $S = M_s$. This gives the smallest possible deterministic space to project. As such the relevant Sanibel coefficients simplify to an expression without a sum ($j = 0$),

$$C_k(S, N_o) = \frac{4S+2}{2S+N_o+2} (-1)^k \left(\frac{N_o}{2} + S \right)^{-1} \binom{N_o}{k}.$$

Particular permutations that are required to be considered in deriving the energy

expressions may modify the value of k . $k(tu|vw)$ is defined as the k value resulting when the spins associated with orbitals t and u in θ_μ are exchanged, where v and w are the orbitals corresponding to t and u in $\Phi_\nu\theta_\nu$. The changes in the value of k are dependent on the spins associated with t, u, v, w as follows;

$$\begin{aligned} k(\alpha\alpha|\alpha\alpha) &= k(\alpha\alpha|\beta\beta) = k(\beta\beta|\alpha\alpha) = k(\beta\beta|\beta\beta) = k \\ k(\alpha\beta|\alpha\beta) &= k(\beta\alpha|\beta\alpha) = k + 1 \\ k(\alpha\beta|\beta\alpha) &= k(\beta\alpha|\alpha\beta) = k - 1 \end{aligned}$$

4.6.3 Element classification

Once the spin-projected determinants, F_μ, F_ν , have been chosen, aligned and the spin-mismatch parameter, k , measured then the *category* of the matrix element must be determined. The matrix elements for spin-projected determinants were presented by Harris (see table 4.2) as a large number of special cases (and all symmetry related cases) that are non-zero, with specific expressions provided for the non-zero cases.

Classifying these matrix elements requires certain information about the relationship between the determinants F_μ and F_ν ;

i, j correspond to (up to two) spatial orbitals that are doubly occupied in F_ν and not in F_μ . If these spatial orbitals are singly occupied in F_μ then it is required to keep track of the spin of this singly occupied orbital, which is stored as i_{pair} or j_{pair} .

r, s correspond to (up to two) spin orbitals belonging to singly occupied spatial orbitals in F_μ that are not occupied in F_ν .

l, m correspond to (up to two) spin orbitals belonging to singly occupied spatial orbitals in F_ν , that are required for categorising the pair. This can be for several reasons;

- they could be unoccupied in F_μ , or
- they could be doubly occupied in F_μ , or
- they could have a specified spin in F_μ by having a well defined spin relative to $i_{\text{pair}}, j_{\text{pair}}, r$ or s .

An example of the last of these cases is given by

$$\begin{aligned}
 1^\beta 1^\alpha 2^\beta &\longrightarrow 1^\beta 2^\beta 3^\alpha && \text{matrix element is zero} \\
 &\longrightarrow 1^\beta 2^\alpha 3^\beta && \text{category 8, } l = 2^\alpha \text{ (see table 4.2),} \\
 &\longrightarrow 1^\alpha 2^\beta 3^\beta && \text{category 3}
 \end{aligned}$$

where the spin of the orbital matched with l after alignment enters into the $k(\alpha\beta|lm)$ expression for the energy, and so an assignment of orbital l is required.

Once these lists of orbitals are generated, the category may be found by reference to the first column of table 4.2, and the expressions for the Hamiltonian matrix elements by the third column. If a suitable category does not exist for the lists generated (nor for any case where the spins are all inverted), then the matrix element is zero.

4.6.4 Integration with excitation generation

If matrix elements are being calculated in the process of spawning, the required information can be directly extracted from the excitation generator.* The process of categorisation involves identifying which of a large number of special cases applies — see table 4.2, and returning the data specified in section 4.6.5.

At the end of the excitation generation process, there are a number of pieces of information that are directly available:

- The excitation matrix. That is the (up to two) electrons, m, n , that are the source of the spatial excitation and their corresponding orbitals, ϕ_m, ϕ_n , along with the target orbitals, p, q .
- The source spin projected determinant, $F_{\mathbf{I}}$, in both its expanded form in the standard order presented in section 3.4.1 and its bit representation.
- The target spin projected determinant, $F_{\mathbf{J}}$, in an expanded form that is *not* sorted in the standard order, but in the order corresponding to $F_{\mathbf{I}}$ with the

*If multiple-structure spawning (see section 5.5) is in use, this process is complicated by having to pass parallel information about all the spawns independently. As the choice of matrix element category can depend on the spin structure chosen, all of the terms within a multiple-structure spawn may have different data. ‘Blob-to-blob’ (section 5.5.1) spawning hasn’t been attempted for spin-projected determinants.

Category	Element ¹	Expression ^{2,3}
<i>Overlap</i> ⁴		
1	()	C_k
<i>One-electron operator</i> ⁵		
1	()	$C_k [2 \sum_i h_{ii} + \sum_l h_{ll}]$
2	$(r_\alpha l)$	$C_k h_{rl}$
3	$(i_\alpha r \beta i_\alpha i_\beta)$	$C_k h_{ri}$
3	$(i_\alpha l_\beta l_\alpha \beta i_\alpha i_\beta l)$	$C_k h_{li}$
<i>Two-electron operator</i> ⁶		
1	()	$C_k \{ \sum_{ij} [2 \langle ij \hat{g} ij \rangle - \langle ji \hat{g} ij \rangle] + \sum_{il} [2 \langle il \hat{g} il \rangle - \langle li \hat{g} il \rangle] + \sum_{l < m} \langle lm \hat{g} lm \rangle - \sum_{l < m} C_{k(ml ml)} \langle ml \hat{g} lm \rangle \}$
2	$(r_\alpha l)$	$C_k \{ \sum_i [2 \langle ir \hat{g} il \rangle - \langle ri \hat{g} il \rangle] + \sum_{m \neq l} \langle mr \hat{g} ml \rangle - \sum_{m \neq l} C_{k(rm lm)} \langle rm \hat{g} ml \rangle \}$
3 ⁷	$(i_\alpha r_\beta i_\alpha i_\beta)$	$C_k \{ \sum_j [2 \langle jr \hat{g} ji \rangle - \langle rj \hat{g} ji \rangle] + \sum_l \langle lr \hat{g} li \rangle - \sum_l \langle rl_\beta \hat{g} li \rangle \}$
4 ⁷	$(i_\alpha l_\beta l_\alpha i_\alpha i_\beta l)$	$C_k \{ \sum_j [2 \langle jl \hat{g} ji \rangle - \langle lj \hat{g} ji \rangle] + \sum_m \langle lm \hat{g} im \rangle + \sum_m [C_{k(\alpha\beta lm)} - C_k] \langle m_\beta l \hat{g} im \rangle \}$
5	$(rs lm)$	$C_k \langle rs \hat{g} lm \rangle - C_{k(rs lm)} \langle sr \hat{g} lm \rangle$
6	$(r_\alpha s_\beta i_\alpha i_\beta)$	$C_k \langle rs \hat{g} ii \rangle$
7	$(r_\alpha r_\beta i_\alpha i_\beta)$	$C_k \langle rr \hat{g} ii \rangle$
8 ⁷	$(i_\alpha l_\beta r_\alpha i_\alpha i_\beta l)$	$\sum_m C_{k(\alpha\beta lm)} \langle m_\beta r \hat{g} im \rangle$
9	$(i_\alpha r_\beta s_\beta i_\alpha i_\beta l)$	$C_k [\langle rs \hat{g} il \rangle - \langle sr \hat{g} il \rangle]$
10	$(i_\alpha r_\beta s_\alpha i_\alpha i_\beta l)$	$C_k \langle rs \hat{g} il \rangle$
11	$(i_\alpha r_\beta r_\alpha i_\alpha i_\beta l)$	$C_k \langle rr \hat{g} il \rangle$
12	$(i_\alpha l_\beta l_\alpha r_\alpha i_\alpha i_\beta lm)$	$C_k \langle lr \hat{g} im \rangle$
13	$(i_\alpha l_\beta l_\alpha r_\beta i_\alpha i_\beta lm)$	$C_k \langle lr \hat{g} im \rangle + [C_{k(\alpha\beta lm)} - C_k] \langle rl \hat{g} im \rangle$
14	$(i_\alpha r_\beta j_\alpha s_\beta i_\alpha i_\beta j_\alpha j_\beta)$	$C_k [\langle rs \hat{g} ij \rangle - \langle sr \hat{g} ij \rangle]$
15	$(i_\alpha r_\beta s_\alpha j_\beta i_\alpha i_\beta j_\alpha j_\beta)$	$C_k \langle rs \hat{g} ij \rangle$
16	$(i_\alpha r_\beta r_\alpha j_\beta i_\alpha i_\beta j_\alpha j_\beta)$	$C_k \langle rr \hat{g} ij \rangle$
17	$(i_\alpha l_\beta j_\alpha r_\beta l_\alpha i_\alpha i_\beta j_\alpha j_\beta l)$	$C_k [\langle lr \hat{g} ij \rangle - \langle rl \hat{g} ij \rangle]$
18	$(l_\alpha i_\beta j_\alpha r_\beta l_\beta i_\alpha i_\beta j_\alpha j_\beta l)$	$C_k \langle lr \hat{g} ij \rangle$
19	$(i_\alpha l_\beta j_\alpha m_\beta l_\alpha m_\alpha i_\alpha i_\beta j_\alpha j_\beta lm)$	$C_k [\langle lm \hat{g} ij \rangle - \langle ml \hat{g} ij \rangle]$
20	$(i_\alpha l_\beta m_\alpha j_\beta l_\alpha m_\alpha i_\alpha i_\beta j_\alpha j_\beta lm)$	$C_k \langle lm \hat{g} ij \rangle - C_{k(\alpha\beta lm)} \langle ml \hat{g} ij \rangle$

¹ These terms describe the relationship between the two spin-projected determinants being considered, those corresponding to F_μ are on the left, and F_ν on the right. i, j, l, m, r, s are as described in the text in section 4.6.3.

² C_k is the Sanibel coefficient with parameters $C_k(S, M_s, N_o)$, as defined in section 4.6.2. All entries are to be multiplied by the parity of \mathbf{P}_0 , $(-1)^{P_0}$.

³ All terms which are symmetry (spin-inverted, or μ and ν reversed) related to those given are also non-zero.

⁴ To calculate the overlap term, the operator $\mathcal{B} = \mathcal{I}$, a zero-electron operator.

⁵ The terms h_{ia} are the one electron integrals.

⁶ The terms $\langle ij | \hat{g} | ab \rangle$ are the two electron integrals, with $\hat{g} = r_{12}^{-1}$.

⁷ The term l_β and m_β indicate that m, l should only be included in the sum if they have the same spin as that labelled β in the 'element' term.

Table 4.2: Non-zero spin-projected Hamiltonian matrix elements

source and target orbitals swapped. It is also available in the bit representation form.

From this information, there is a direct but extremely complicated mapping to the information required to generate the Hamiltonian matrix elements. The length and complexity of this mapping prevent it from being effectively presented in full here. Further (more explicit and long winded) details are available on request.

There are a number of specific component pieces of information that are required, and the general approach to obtaining these is consistent, whatever the ‘shape’ of the particular excitation being considered:

Number of unpaired electrons

The number of paired and unpaired electrons in the target determinant is directly available by knowledge of the source site, and which orbitals are being excited.

Choosing a Hamiltonian matrix element category

Table 4.2 presents an extremely specific list of categories of Hamiltonian matrix element that are non-zero. The excitations need to be mapped onto this list carefully, taking note of all of the symmetry related cases. This is done through an (extremely large) branching structure, that takes account of:

- The number of spatial orbitals that are excited.
- If these excitations are coming from, or creating, singly occupied or doubly occupied orbitals.
- If two excited electrons are coming from, or going to, the same spatial orbital as each other.
- The m_s relationship between any orbitals that are excited from or to, and the residual orbitals that they interact with.

Assigning μ and ν

One of the symmetries that needs careful attention is that either of the spin projected determinants, $F_{\mathbf{I}}, F_{\mathbf{J}}$, can be selected as F_{μ} or F_{ν} (and vice versa). In the branching structure required to choose the category, there is normally only one way of making this choice at each entry. In some cases, this choice can be made in multiple ways (sometimes corresponding to two different categories that could be used to calculate the same Hamiltonian

matrix element). In these cases the implementationally simpler choice is made.

Orbital assignment

The orbitals ϕ_m, ϕ_n, p, q , and those directly spin-related to them, can be directly assigned to the relevant terms $i, j, i_{\text{pair}}, j_{\text{pair}}, l, m, r, s$ as required. There are a number of cases where other orbitals are required to correctly provide the description needed for the entry in table 4.2. In a few cases the position of this orbital in $F_{\mathbf{I}}$ is determined by the position of one of the target orbitals in the *aligned* determinant $F_{\mathbf{J}}$. The $\mathcal{O}(N^2)$ alignment operation can be avoided by considering all of the available ‘shapes’ of excitation explicitly, which reduces the possible searching for the correct index to a maximum of an $\mathcal{O}(N)$ operation.

Mismatch parameter, k

The generated determinant, $F_{\mathbf{J}}$, is not sorted into the standard order presented in section 3.4.1. A baseline mismatch parameter is calculated by finding the mismatch parameter for these mis-aligned determinants, which is then corrected for according to the specific ‘shape’ of the excitation corresponding to its location in the branching structure described above.

Parity, ϵ_P

The alignment permutation is never explicitly needed, and as such never needs to be generated (in contrast to the use of Serber functions). The parity of this permutation can be obtained directly from a few pieces of information. The offsets between the source indices of excitations and the target locations need to be found, and carefully considered. In particular, in the same way (and at times in the same step) as the searching required for orbital assignment, the $\mathcal{O}N^2$ alignment operation can be avoided by specific consideration of all the possible ‘shapes’ of the excitation including:

- Do the excited electrons move into or out of the doubly and singly occupied regions of the standard representation?
- If there are multiple electrons being excited, does the path of one electron cross the other?
- The movements of other orbitals between sections of the standard representation, in particular between the doubly occupied and singly occupied regions, need to be accounted for.

- The possibility of out-by-one errors needs to be carefully managed when either the source, or target, orbitals (or both) are at the ends of the sorted list of orbitals in either representation.

Similarly to finding particular orbitals as described above, the worst case examples for determining the parity scale as $\mathcal{O}(N)$. This is the slowest element of the classification process, but it is a substantial improvement on the explicit alignment operations required otherwise.

4.6.5 Matrix element generation

Harris shows that the matrix element of an arbitrary operator, \mathcal{B} , can be reduced to

$$\begin{aligned} \langle F_\mu | \mathcal{B} | F_\nu \rangle &= \langle \mathcal{A} \hat{O}_S \Phi_\mu \theta_\mu | \mathcal{B} | \mathcal{A} \hat{O}_S \Phi_\nu \theta_\nu \rangle \\ &= \langle \Phi_\mu \theta_\mu | \hat{O}_S \mathcal{B} \mathcal{A} | \Phi_\nu \theta_\nu \rangle \\ &= \frac{1}{N!} \sum_{\mathbf{P}} (-1)^P \langle \Phi_\mu | \mathcal{B} \mathbf{P} | \Phi_\nu \rangle \langle \theta_\mu | \hat{O}_S \mathbf{P} | \theta_\nu \rangle. \end{aligned}$$

Note that this is the same as equation 4.2, but expressed for an arbitrary operator, \mathcal{B} . The general spin function, Θ_μ , is replaced with $\hat{O}_S \theta_\mu$, and the expression is simplified by noting that \hat{O}_S is idempotent and commutes with an arbitrary permutation.

Harris continues to show that the spin and spatial parts of this expression may be sensibly reduced into forms that are $\mathcal{O}(1)$ and (at worst) $\mathcal{O}(N^2)$ respectively. Once the relationship between two determinants has been determined, and categorised according to the ‘Element’ column of table 4.2, then all of the terms required to calculate the matrix elements (also given in the same table) are known;

1. The spin-projected determinants, F_μ and, if necessary for the given matrix element category, F_ν . Note that the newly spawned determinant does not need to be aligned.
2. The excitation category.
3. The spin mismatch parameter, k .
4. The parity of the line-up operation, $\epsilon_P = (-1)^P$. Neither the actual line-up permutation nor the aligned determinant are required.

5. A list of pairs of doubly occupied orbitals, i, j , which are present in F_ν , and if relevant, a list of the corresponding singly occupied orbitals, $i_{\text{pair}}, j_{\text{pair}}$, in F_μ .
6. A list of orbitals r, s that are singly occupied in F_μ and do not appear in F_ν .
7. A list of orbitals l, m that are singly occupied in F_ν , and are required for categorising the pair, either by being doubly occupied, not present or of specified (possibly different) spin in F_μ .
8. A list of the orbitals in F_μ corresponding to orbitals l, m in F_ν , if required for calculating $k(tu|vw)$.

These may be provided as elucidated above, or through the excitation generation routines.

Once this information is known, the matrix element can be trivially calculated given the entries in table 4.2.

4.6.6 The overlap matrix

As spin-projected determinants are non-orthogonal, the overlap matrix elements enter the spawning terms, as described in equations 2.10a and 2.17. As spin-projected determinants can be written as a linear sum of Slater determinants with the same spatial structure, it is trivial to see that the overlap matrix elements are only non-zero between spin-projected determinants with the same spatial structure (this can also be seen from the section for zero-electron operators, i.e. $\mathcal{B} = \mathcal{I}$ only containing *category 1* in table 4.2).

In the case of same-spatial-structure spawning, the overlap matrix element is given by the Sanibel coefficient,

$$S_{\mu\nu} = C_k(S, M_s, N_o),$$

for the spin-mismatch value k associated with this pair of determinants. For single and double excitations the overlap matrix element is zero.

4.7 Mixed Hilbert spaces

If a mixed representation of the Hilbert space is being used, such that the space is split into regions represented by different types of basis function, the calculation of matrix elements between sites in the different regions requires careful consideration.

The mixed spaces considered in this thesis are those with CSFs representing regions of the space with fewer than $N_{o,\max}$ unpaired electrons, and either Slater determinants or HPHF functions representing the remainder. These matrix elements are most straightforwardly broken down in terms of the expansion of the CSFs in Slater determinants,

$$\langle D_\mu | \hat{H} | F_\nu \rangle = \sum_{\mathbf{j}} c_{\mathbf{j}}^{(\nu)} \langle D_\mu | \hat{H} | D_\nu \rangle,$$

where the matrix elements between Slater determinants can be found as usual. This operation can be performed in $\mathcal{O}(n_{csf})$ time.

If HPHF functions are used instead of Slater determinants, the same coefficient symmetry is necessarily found between the determinants within a spin eigenfunction as within the HPHF function, and as such the matrix elements may be obtained by straightforward multiplication by $\sqrt{2}$ for all cases where the HPHF function being considered is not closed shell.

4.8 Summary

Due to poor computational scaling, the use of Kotani-Yamanouchi CSFs is not realistically viable. This leaves us with two different schemes, both of which are computationally tractable, and scale roughly as $\mathcal{O}(N^2)$ in the worst case. They have different strengths and weaknesses;

Serber CSFs

Once the pre-computation cost of generating the desired permutation representation matrices has been paid (this can be done once, and the resultant output stored) Serber function matrix elements are extremely efficient to calculate. The use of Serber functions makes a deliberate trade-off to use

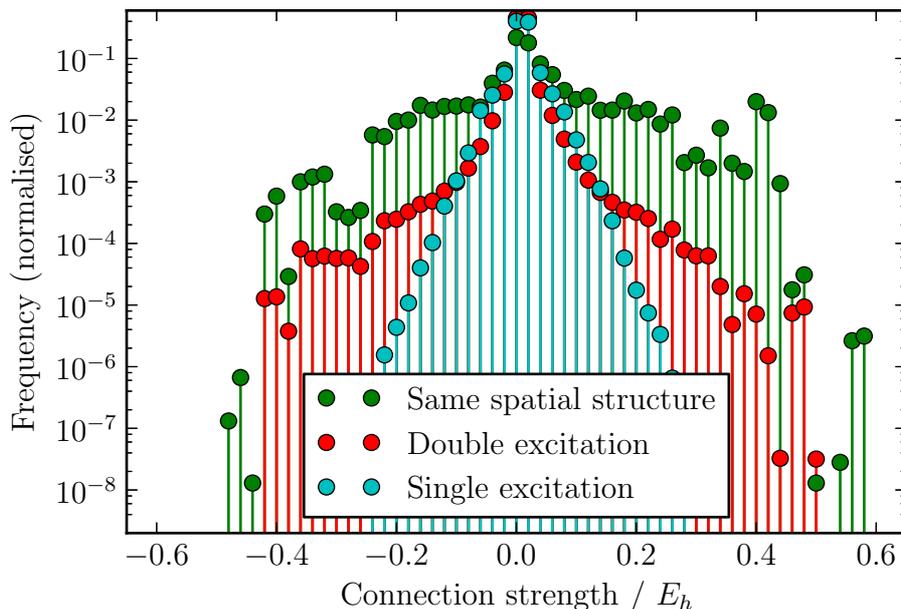


Figure 4.7: A histogram of the spawning connections strength, $K_{ij} - E_S \delta_{ij}$, using Serber functions for N_2 with a bond length of $6.00 a_0$ in a cc-pVDZ basis set. This histogram should be compared to figure 3.7.

substantially more memory than other schemes, in return for a simplification of the matrix element calculation. There are two problems with their more general use; *a)* the permutation representation matrices get rapidly larger and more numerous, such that the maximum size of system that can be considered is probably 14 or 16 electrons, and *b)* Integration with the excitation generator is less efficient than for spin-projected determinants, as the actual line-up permutation is required, rather than just its parity.

Figure 4.7 shows the main advantage of Serber functions over spin-projected determinants. The distribution of Hamiltonian matrix elements corresponding to same spatial structure spawns is compact and similar in magnitude to those for the single and double excitations. This is in contrast to spin-projected determinants (see figure 3.7, it is worth noting that the density of these two histograms is the same) for which the same spatial structure spawns have wildly vary magnitudes. As a consequence, particle dynamics for Serber functions are substantially more manageable.

Spin-projected determinants

Hamiltonian matrix element evaluation for spin-projected determinants is extremely efficient from a computation point of view, and it integrates well with excitation generation. In particular, they pose no additional storage overhead. Their most obvious weakness is their non-orthogonality — this leads to problems controlling the simulation as the system size increases (see sections 3.4.2 and 3.4.5).

5 Excitation generation using CSFs

Within FCIQMC, ‘excitation generation’ is one of the most interesting and important considerations. An implementation of the propagation equation 2.10a using stochastic spawning depends on efficient and random exploration of the Hilbert space. There are two properties that are required for this to be efficient;

1. the excitation generator must be able to generate all sites, $\{F_{\mathbf{j}}\}$, that are connected to the starting site, $F_{\mathbf{i}}$, by a non-zero Hamiltonian matrix element, $\langle F_{\mathbf{j}}|\hat{H}|F_{\mathbf{i}}\rangle$, and
2. the generation probability, $p_{\text{gen}}(\mathbf{j}|\mathbf{i})$ of a connected site must be calculable. This includes accounting for all possible ways that a connection could be generated.

Within this scheme it is permissible for the excitation generator to abort an attempted generation — normally as a result of internal choices that lead to there being no available target sites. In this context, the result is treated as though it were a site with a zero connecting Hamiltonian matrix element, and so no particles are spawned. It is not possible to attempt spawning again, as this would require renormalising all of the generation probabilities to account for all possible ways of having to retry.

There are several additional details which are desirable, but not in any sense required;

- the distribution of connections generated should be *sensible*. The permitted value of the time-step, $\delta\tau$, is dependent on the combined distribution of matrix elements and generation probabilities. It is beneficial if the excitation generator assists in smoothing this combined distribution. See section 5.3. Also,
- the excitation generator should make as much use of symmetry as possible.

Given the symmetry of the problem, a great deal is known about which matrix elements, $\langle F_j | \hat{H} | F_i \rangle$ must be zero. Avoiding generating these increases the generation probabilities, maximising the possible value of $\delta\tau$, and minimising wasted computational time,

- the excitation generators should be fast — in particular they should use as few random numbers as possible, and should scale as favourably as possible with increasing system size, and
- as no excitation is generated for which the associated Hamiltonian matrix element is not also generated, any information which can be easily determined during the process of excitation generation, and communicated directly to the matrix element generation routines is beneficial.

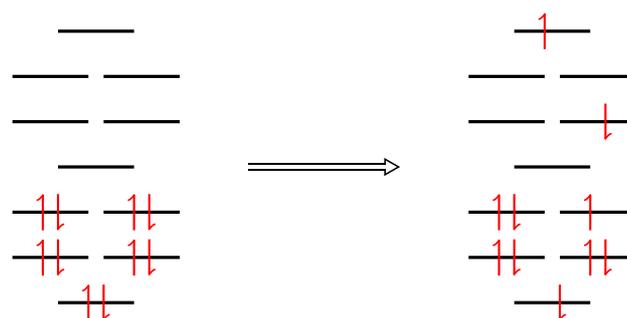
A scheme for efficiently generating excitations from arbitrary spin eigenfunctions is presented in this chapter. An overview is given in flowchart form in figure 5.2. The boxes in the flowchart are numbered — these boxes are referenced as numbers in rounded brackets in the text.

5.1 Comparison to determinantal excitation generators

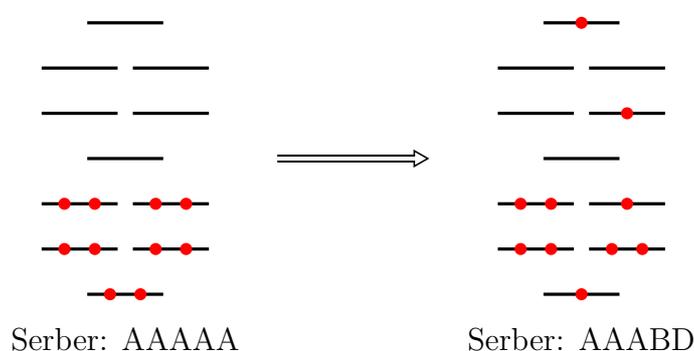
Within all of the CSF schemes used, each basis function is an antisymmetrised product of a spatial and a spin eigenfunction.* As such the process for generating excitations is the same in all of these cases. An overview of the whole process is shown in figure 5.2.

The excitation generation process is similar to that used in determinantal and HPHF bases, which have been published and discussed previously⁷³. An example determinantal excitation and CSF excitation are shown in figure 5.1, demonstrating the similarities and differences. In the determinantal case, two spin orbitals are chosen to excite from and to, and in the spin eigenfunction case the same process

*The spin-projected determinant scheme differs slightly, in that a representation of the occupied spin-orbitals in the determinant to be projected is maintained, whereas in the other cases the spatial orbital structure and the spin eigenfunction label associated with the are stored separately. The excitation process is, however, logically the same, with a spatial excitation being made and a spin structure ‘chosen’ as the final step.



(a) An example of a determinant excitation. Two electrons are chosen, and placed into two vacant spin-orbitals, whilst maintaining spatial symmetry and projected spin value M_s .



(b) An example of a CSF excitation. A spatial excitation is generated by choosing two electrons and placing them into two vacant orbitals, whilst maintaining symmetry. A new spin eigenfunction label (in this case for a Serber function) is selected and applied. Note that the projected spin of the electrons is not considered for the spatial excitation.

Figure 5.1: The difference between determinantal and CSF excitation generators.

is performed with spatial orbitals and in addition a spin structure corresponding to the correct number of unpaired electrons is chosen.

5.2 Generating an excitation

When generating an excitation a process of several steps is followed;

1. If a mixed scheme is in use, the correct determinantal, HPHF or CSF excitation generator is selected depending on the portion of the Hilbert space currently in use.
2. Symmetry information associated with the spatial orbitals for the current

site is calculated, in $\mathcal{O}(N)$ time.

3. A choice is made between a same-spatial-structure, single or double spatial excitations (1—2).
4. The spatial excitation is performed (3—14 and 8—14).
5. A spin structure corresponding with the generated spatial structure is chosen at random (16—17).

The following sections explore these actions in detail. If they are all considered sequentially, the overall generation probability is given by

$$p_{\text{gen}}(\mathbf{j}|\mathbf{i}) = p_{\text{type}} \times p_{\text{spatial}} \times p_{\text{spin}}.$$

For the work contained in this thesis, the systems considered are molecular, with spatial point-group symmetry. In principle, CSFs can be implemented for other types of systems with other symmetry requirements (such as the Uniform Electron Gas, Hubbard model, solids or anything requiring non-abelian symmetries),* but this is not explored further here.

5.2.1 Spatial symmetry information

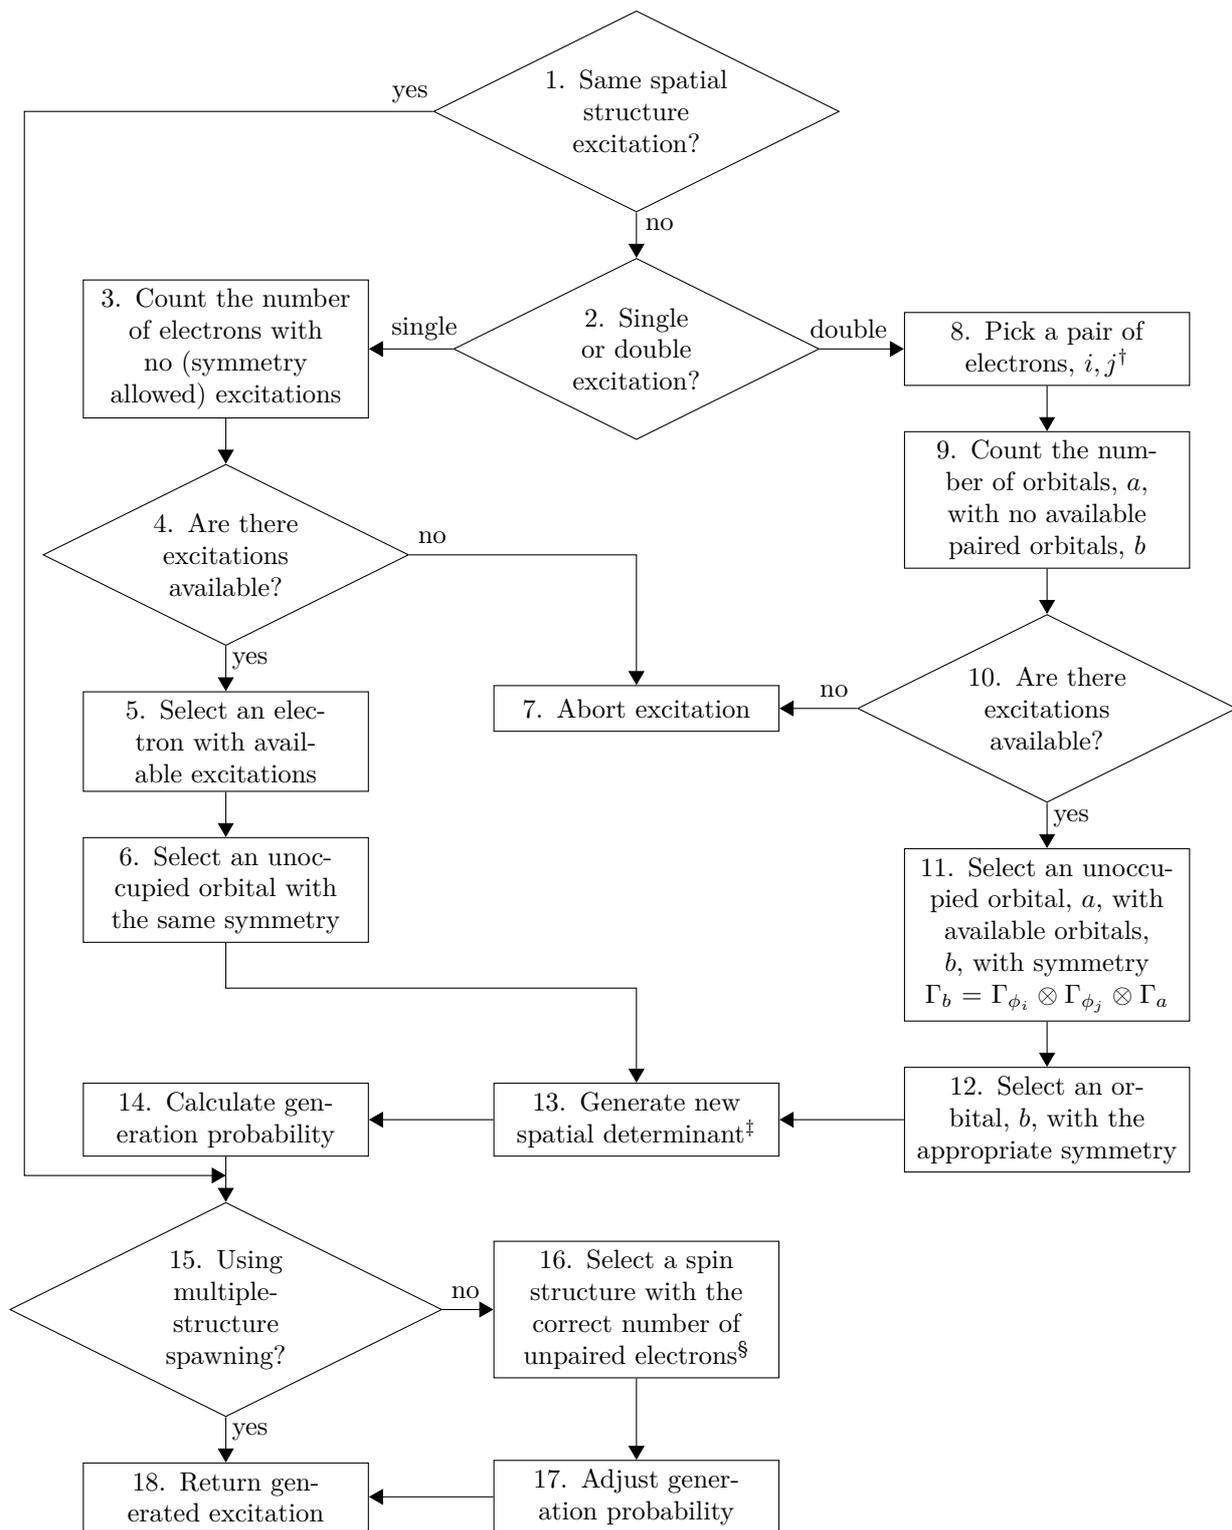
Working with molecular systems, for a CSF with spatial component $\Phi_{\mathbf{s}} = \mathcal{A} \prod_{\alpha}^N \phi_{\alpha}^{(\mathbf{s})}$ the symmetry of the spin eigenfunction is given by $\Gamma_{\mathbf{s}} = \otimes_{\alpha}^N \Gamma_{\alpha s}$. Symmetry indicates that the Hamiltonian matrix element between two CSFs $F_{\mathbf{s}}$ and $F_{\mathbf{t}}$ can only be non-zero if $\Gamma_{\mathbf{s}} \otimes \Gamma_{\mathbf{t}} = \Gamma_1$, or alternatively

$$\langle F_{\mathbf{s}} | \hat{H} | F_{\mathbf{t}} \rangle = 0 \quad \forall \Gamma_{\mathbf{s}} \neq \Gamma_{\mathbf{t}}, \quad (5.1)$$

where only spatial symmetries are considered.

The spatial symmetries of each orbital in the basis set are known, and three lists containing the number of singly occupied, doubly occupied and vacant spatial orbitals associated with each symmetry class may be constructed for any given CSF in $\mathcal{O}(N)$ time (the list of occupied orbitals is constructed, and their symmetries

*For exploring the properties of the nitrogen atom, conserving the magnetic quantum number, M_l , is required to distinguish between states. This has been implemented purely by rejecting generated excitations with incorrect total values of M_l . It can be implemented more fully as discussed for determinantal systems by Booth *et al.* [73].



[†] Unlike for determinantal excitations there is no unique triangular mapping to select a pair of electrons with one random number, as some electrons are spatially equivalent. This must be considered in the generation probabilities.

[‡] It is at this stage that alignment permutations and parity for Serber function excitation generation can be extracted. This is done by generating the correctly aligned structure explicitly from the original structure.

[§] It is at this stage that characterising the excitation type, and obtaining the parity, is performed for spin-projected determinants.

Figure 5.2: Random excitation generation overview

used to increment the required symmetry lists). For symmetry Γ , these are labelled as $N_s(\Gamma)$, $N_d(\Gamma)$ and $N_v(\Gamma)$.

Where there are multiple particles on one site, these lists only need to be calculated once, and may be stored between multiple calls to the excitation generators.

5.2.2 Choice of excitation type

When generating an excitation, the first step involved is to choose the type of excitation to be generated. All known excitations to unconnected sites, i.e. $\langle F_i | \hat{H} | F_j \rangle = 0$, are excluded, leaving only three categories of excitation as a consequence of the Slater–Condon rules⁹⁶. Given an initial basis function, a new one may be selected with either *a*) the same structure, *b*) a structure which differs by one spatial orbital, or *c*) a structure which differs by two spatial orbitals. A random selection is made according to a pre-selected distribution which is constant across the entire Hilbert space such that the probability is given by

$$p_{\text{type}} = \begin{cases} p_{\text{same}} & \text{if same spatial structure} \\ (1 - p_{\text{same}}) \times p_{\text{single}} & \text{if single (spatial) excitation} \\ (1 - p_{\text{same}}) \times p_{\text{double}} & \text{if double (spatial) excitation.} \end{cases}$$

The probability p_{same} is dependent on the properties of the source site being considered. If it has connected sites with the same spatial structure, then it is non-zero and specified as a global constant, otherwise it is zero (1). The probabilities $p_{\text{single}} + p_{\text{double}} = 1$ and determine the type of excitation to be selected given that the same spatial structure is not being generated (2). This scheme is used because p_{same} is required to be significant to reduce the impact of this form of excitation on $\delta\tau$, potentially resulting in substantial waste of resources when considering sites that have no excitations to other sites with the same spatial structure.

If a single or double excitation is selected, then either one or two occupied orbitals and one or two unoccupied orbitals needs to be selected, as discussed in sections 5.2.3 and 5.2.4. The bit-representation of the site is then easily updated by setting and clearing the relevant bits. The decoded representation never needs to be directly generated, as all of the information required for generating the off-diagonal matrix elements can be obtained without it.

Once the target spatial structure has been selected, then a spin structure must be selected (see section 5.2.5).

5.2.3 Generation of (spatial) single excitations

To generate a single excitation an electron, i , located in orbital ϕ_i , must be selected along with a vacant or singly occupied (spatial) orbital, a , for it to be excited to. This is performed in several stages.

Selecting an electron (3)

In order to satisfy the condition in equation 5.1, that $\Gamma(F_i) = \Gamma(F_j)$, then it is necessary that $\Gamma_{\phi_i} = \Gamma_a$. As a consequence of the basis set being finite, for any given choice of electron, i , there may not be any available target orbitals to choose from. The number of these electrons is counted (3). Iterating over the symmetry classes, Γ , there are three cases to consider;

1. if $N_v(\Gamma) = 0$ and $N_s(\Gamma) = 1$, there are no excitations available from the singly occupied orbital as it cannot excite to itself,
2. if $N_v(\Gamma) = N_s(\Gamma) = 0$, there are no excitations from the doubly occupied orbitals, and
3. only one electron out of each of the doubly occupied orbitals may be selected as both would lead to the same spatial excitations.

Combining these gives a total number of electrons which cannot be excited of

$$\begin{aligned}
 n_{\text{no excit}} &= \sum_{\Gamma} N_d(\Gamma) \\
 &+ N_s(\Gamma) \quad \text{if } N_v(\Gamma) = 0 \text{ and } N_s(\Gamma) = 1 \\
 &+ N_d(\Gamma) \quad \text{if } N_v(\Gamma) = N_s(\Gamma) = 0.
 \end{aligned}$$

If none of the electrons have a valid excitation (4), i.e. $n_{\text{no excit}} = N$, then the excitation is aborted. Otherwise an electron, i , may be selected uniformly from those permitted by the above constraints with probability

$$p_{elec} = \frac{1}{N - n_{\text{no excit}}}.$$

Selecting an Orbital

Once an electron i has been chosen from an orbital ϕ_i with symmetry Γ_{ϕ_i} , a spatial orbital with the same symmetry needs to be picked at random (6).

The only permitted choices are either

1. vacant orbitals, or
2. singly occupied orbitals (other than the source orbital).

This gives the available number of choices as

$$n_{avail} = \begin{cases} N_s(\Gamma) + N_v(\Gamma) - 1 & \text{if exciting from a singly occupied orbital} \\ N_s(\Gamma) + N_v(\Gamma) & \text{otherwise.} \end{cases}$$

and an associated probability of

$$p_{orb} = \frac{1}{n_{avail}}.$$

Constructing the final function

The final spatial structure is now constructed by clearing the bit associated with orbital ϕ_i , and setting the bit associated with the target orbital a (13).

The generation probability (14) is given by

$$\begin{aligned} p_{spatial}(\phi_i \rightarrow a) &= p_{elec} \times p_{orb} \\ &= \frac{1}{(N - n_{no\ excit})n_{avail}}. \end{aligned}$$

5.2.4 Generation of (spatial) double excitations

To generate a double excitation a pair of electrons, i and j , located in orbitals ϕ_i and ϕ_j , must be selected along with a pair of vacant or singly occupied (spatial) orbitals, a and b , for them to be placed in to. This is performed in several stages.

Symmetry constraints

To retain the overall symmetry of the molecule, the constraint on the orbital symmetries for the excitation $\phi_i, \phi_j \rightarrow a, b$ becomes

$$\Gamma_{\phi_i} \otimes \Gamma_{\phi_j} = \Gamma_a \otimes \Gamma_b. \quad (5.2)$$

Picking a pair of electrons (8)

Excluding the cases of extremely small basis sets, there are very few pairs of electrons with no allowed excitations. The symmetry constraint expressed in equation 5.2 is more easily met than in the single excitation case, and it is now not worthwhile to count the cases with no electrons and renormalise the probabilities to take this into account.

Defining N_{sing} , N_{doub} and N_{vac} as the number of singly-, doubly- and unoccupied spatial orbitals respectively, which may be calculated as $\sum N_s(\Gamma)$, $\sum N_d(\Gamma)$, and $\sum N_v(\Gamma)$ (section 5.2.1), the number of possible electrons to pick for the first, i , and second, j , choices respectively are given by

$$N_{\text{avail}}(i) = N_{\text{sing}} + N_{\text{doub}}$$

$$N_{\text{avail}}(j|i) = \begin{cases} N_{\text{sing}} + N_{\text{doub}} - 1 & \text{if } i \text{ singly occupied} \\ N_{\text{sing}} + N_{\text{doub}} & \text{otherwise.} \end{cases}$$

Unless the two electrons picked are from the same spatial orbital, there are two ways to generate both pairs of electrons — i, j is equivalent to j, i — and this must be accounted for in the generation probability. Noting that $N_{\text{avail}}(i) = N_{\text{avail}}(j)$,

$$p_{\text{pair}}(i, j) = \begin{cases} \frac{1}{N_{\text{avail}}(i)} \cdot \frac{1}{N_{\text{avail}}(j|i)} & \text{if } i, j \text{ from same spatial orbital} \\ \frac{1}{N_{\text{avail}}(i)} \left[\frac{1}{N_{\text{avail}}(j|i)} + \frac{1}{N_{\text{avail}}(i|j)} \right] & \text{otherwise.} \end{cases}$$

The orbitals ϕ_i, ϕ_j from which the excitation occurs are obtained from these electrons, and the required symmetry product, $\Gamma_p = \Gamma_{\phi_i} \otimes \Gamma_{\phi_j}$, is calculated. As singly occupied orbitals which are being excited from are not available to excite to, an array, $\Delta N_s(\Gamma)$, indicating the change in the number of singly occupied spatial orbitals available to excite to for each symmetry class, Γ , is created.

Picking the first orbital

An orbital a is required which satisfies the following conditions:

- A spatial orbital cannot be excited to if it is being excited from.
- A doubly occupied orbital cannot be excited into.
- An orbital, a , cannot be selected if there are no orbitals, b , of the correct

symmetry such that $\Gamma_b = \Gamma_p \otimes \Gamma_a$.

Iterating over all symmetries, Γ , the number of orbitals which cannot be chosen, $N_{\text{no pair}}$, is counted (9). If $N_{\text{no pair}}$ is equal to the number of available orbitals (10), the excitation is rejected (7).

Orbitals a are then selected at random (11) until one is found which meets the above conditions. The probability of having picked orbital a is then given by

$$p(a) = \frac{1}{N_{\text{sing}} + N_{\text{vac}} - N_{\text{no pair}}}$$

Picking the second orbital

Due to the selection above, there exists at least one suitable b orbital where $\Gamma_b = \Gamma_p \otimes \Gamma_a$. Orbitals with symmetry Γ_b are now picked at random until one is found which matches the same selection criteria as for orbital a (12).

In addition, if $\Gamma_a = \Gamma_b$ then the location of orbital a must be considered. If orbital a has been excited into a singly occupied orbital, it becomes a doubly occupied orbital, and that orbital is no longer available to choose for orbital b . This leaves a probability of

$$p(b|a) = \begin{cases} \frac{1}{N_v(\Gamma_b) + N_s(\Gamma_b) + \Delta N_s(\Gamma_b)} & \text{if } \Gamma_a \neq \Gamma_b, \\ & \text{or } a \rightarrow \text{vacant orbital} \\ \frac{1}{N_v(\Gamma_b) + N_s(\Gamma_b) + \Delta N_s(\Gamma_b) - 1} & \text{if } \Gamma_a = \Gamma_b, \\ & \text{and } a \rightarrow \text{singly occupied.} \end{cases}$$

Constructing the final function (13)

If a, b are in a closed pair, i.e. a new doubly occupied orbital has been created, a, b can only have been picked in one way. Therefore the generation probability is given by (14)

$$p_{\text{spatial}}(\phi_i, \phi_j \rightarrow a, b) = p_{\text{pair}}(i, j) \cdot p(a) \cdot p(b|a).$$

If a, b are not in a closed pair, then the choices a, b and b, a are equivalent and both must be considered in the generation probability;

$$p_{\text{spatial}}(\phi_i, \phi_j \rightarrow a, b) = p_{\text{pair}}(i, j) [p(a)p(b|a) + p(b)p(a|b)].$$

5.2.5 Selection of spin structure or label

In all cases of excitations involving spin eigenfunctions, the choice of spin structure is independent of the spatial excitation that has been made. The number of spin structures available, n_{spin} , depends on the number of unpaired electrons, N_o , and on the type of spin functions in use,

$$n_{spin} = \begin{cases} n_{csf}(N_o) & \text{if using Kotani-Yamanouchi or Serber CSFs} \\ \frac{n_{det}(N_o)}{2} & \text{if using truncated spin-projected determinants} \\ n_{det}(N_o) & \text{if using spin-projected determinants.} \end{cases}$$

In all cases except for the spin-projected determinants, these values are equal. If the target CSF has the same spatial structure as the source function, then the spin structure must change in order to ensure that an off-diagonal (spawning) term is generated (16). The probability of selecting a given spin structure is given by (17)

$$p_{spin} = \begin{cases} \frac{1}{n_{spin}-1} & \text{if same spatial structure excitation} \\ \frac{1}{n_{spin}} & \text{otherwise.} \end{cases}$$

The generated excitation is now returned to the main spawning loop (18).

If multiple-structure spawning is used (see sections 5.5 and 5.5.1), then the selection of a specific spin structure is omitted (15).

5.3 Determining type probabilities

As described in section 2.6, the maximum permissible value of the imaginary timestep, $\delta\tau$, is inversely proportional to the maximum spawning strength. Because the generation probability, p_{gen} appears in the denominator of the spawning expressions, it is important to select the values of p_{single} , p_{double} and p_{same} such that the smallest generation probability values are maximised, and more directly, such that the range of the largest values of n_s are minimised.

The choice of these values impacts the overall efficiency, but so long as they are within the correct ballpark region the simulation is relatively insensitive to the absolute values. Empirically it is found that enumerating the number of connected

single and double determinantal excitations from the reference site, and setting the single and double excitation generation probabilities such that

$$p_{\text{single}} + p_{\text{double}} = 1.0$$

works well, with

$$p_{\text{single}} = \frac{n_{\text{single}}}{n_{\text{single}} + n_{\text{double}}}.$$

For determinants with existent connections to same spatial structure sites, the setting of p_{same} is complicated by the massive range of values that the Hamiltonian matrix elements may take (in some basis sets, see particularly figure 3.7 for spin-projected determinants). As such, p_{same} should be chosen such that the maximum extent of the distribution of n_s values for same spatial structure spawns (when histogrammed) remains within the bounds of the values for single and double excitations.

As the role of same spatial structure spawns appears to be somewhat significant in correctly converging local structure within the evolving wavefunction, it may be beneficial to increase the proportion of spawns which are to sites with the same spatial structure. Certainly a value of $p_{\text{same}} = 0.05$ is more than adequate for most cases, while having little proportional impact on the dynamics of the single and double excitations, and is a safe default value in lieu of system-specific examination.

5.4 Integration of excitation generators and Hamiltonian matrix element calculation

The process of spawning a new particle requires three pieces of information; *a*) the source basis function, *b*) the target basis function, and *c*) the Hamiltonian matrix element between these sites. Once the excitation generation process has been completed, the first two of these are directly available. Calculating the Hamiltonian matrix element can, however, be a relatively involved process.

The very nature of the excitation generation process means that a great deal of information about the two basis functions involved, and their relationship, is already known. This knowledge can make a great deal of difference to the computational

complexity of calculating the Hamiltonian matrix elements.

The precise information that is required is discussed in the subsections of chapter 4, as this is dependent on the particular choice of spin functions in use.

5.5 Multiple-structure spawning

When CSFs are used instead of determinants or HPHF functions, this has substantial impacts on the spawning dynamics. Determinants are permitted to differ by a maximum of two spin orbitals, but CSFs are permitted to have spatial structures which differ by two *spatial* orbitals. Generally the number of spatial excitations from a CSF is similar to the number of spin excitations from a determinant. As a consequence the potential number of CSFs connected to a starting spin eigenfunction is larger than the number of connected determinants by a factor of approximately $n_{\text{csf}}(N_o, S)$, where N_o is the number of unpaired electrons of the target site.

As explained in section 2.6, the maximum value for the time step, $\delta\tau$, is inversely proportional to the maximum strength of each of the connections in equation 2.17, and as a consequence is proportional to the minimum generation probability produced.

The restrictions on $\delta\tau$ tend to be due to sites with large numbers of unpaired electrons, which are extremely highly excited in comparison to the reference site (and also to the majority of occupied sites which contribute substantially to the calculated wavefunction). This is a consequence of spatial structures with large numbers of unpaired electrons having many available spin structures, and thus correspondingly lower generation probabilities. A perverse consequence of this is that a lot of unnecessary computational effort is spent treating the majority of occupied sites with a $\delta\tau$ value orders of magnitude below that which they require. For an alternative approach to dealing with restrictions on $\delta\tau$, see chapter 7.

If the restriction on one-to-one spawning is lifted, this problem may be mitigated. This corresponds to letting the set $\{\mathbf{k}\}$ considered in equation 2.15 contain all sites which share the same spatial structure. From an occupied site, a spatial excitation is made, and then spawning is attempted to *all* sites associated with the target

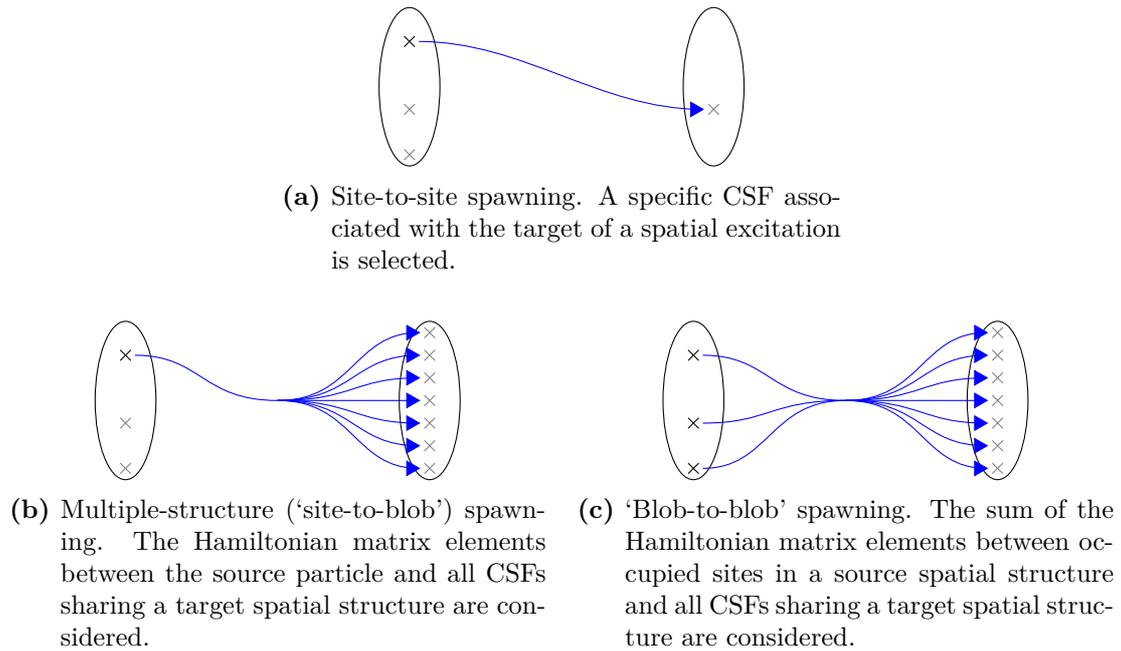


Figure 5.3: Spawning between different combinations of the CSFs within the source and target spatial structures associated with a spatial excitation. This generates spawning between multiple different sites during the same spawning step. The source spatial structure is shown to be sparsely occupied (only a few crosses). If a cross is black it is being considered as the source of a spawn, if it is grey it is not. All of the target sites being considered are indicated in grey.

spin structure* as demonstrated in figure 5.3b. The consequence of this is that the probability of generating site \mathbf{j} given site \mathbf{i} ,

$$p_{\text{gen}}(\mathbf{j}|\mathbf{i}) = p_{\text{gen}}(\mathbf{J}|\mathbf{i}) = p_{\text{gen}}(\mathbf{J}|\mathbf{I}),$$

is now given by the probability of making the spatial excitation from spatial structure \mathbf{I} to spatial structure \mathbf{J} , increasing the generation probabilities by a factor of $n_{\text{csf}}(N_o)$.

This has the advantage that for the highly excited sites, with many unpaired electrons, which dominate the time step dependence, the generation probabilities are increased by the most — preventing these highly excited states from causing an increase in the cost of the low-lying states that do not require such a small time step. It is worth noting that this also improves the overall cost scaling of the

*If the spawn being attempted is a same spatial structure spawn, spawning is attempted at all sites that would give spawning-like behaviour, i.e. the source site is excluded from attempting to spawn to itself. This behaviour is already covered under the death step.

system — the total computational cost is inversely proportional to the time step, and as a consequence multiple-structure spawning removes an $\mathcal{O}(n_{csf}) \approx \mathcal{O}(e^N)$ term from the computational scaling. Figure 5.4 demonstrates clearly the benefit of using multiple-structure spawning on the time step values that may be used. The permitted time steps are still smaller than those available for determinantal calculations, but they are substantially larger than if multiple-structure spawning were not available. As a consequence, Serber functions become comparable to HPHF functions in terms of computational cost.

When using Serber CSFs it is especially advantageous to make use of a multiple-structure spawning scheme. As described in section 4.5, once a Hamiltonian matrix element between two CSFs has been generated, generating the matrix element between different CSFs with the same spatial structures is the same, with only the indices into the permutation matrices used changing. The majority of the computational cost involved is spent generating the spatial excitation, and manipulating the line-up permutations. It is trivial to return the matrix elements corresponding to one column of each of the permutation matrices used — giving substantially more convergence ‘bang’ for your computational ‘buck’ than otherwise. This modification of the algorithm is essentially computationally free, while permitting comparable values of $\delta\tau$ to determinantal and HPHF calculations.

As an aside, it is worth noting that the total weight of particles spawned per unit imaginary time remains roughly the same if multiple-structure spawning is used. A substantial number of additional spawning attempts are made, but the generation probabilities are increased in proportion. As a consequence of this, the acceptance ratio (the ratio of spawned particles to attempted spawns) appears to plummet when multiple-structure spawning is in use.

5.5.1 ‘Blob-to-blob’ spawning

A logical extension of multiple-structure spawning is to consider all possible connections between source and target spatial structures in each spawning step. The sum over coefficients, c_j , in equation 2.10a is considered in re-approximating equa-

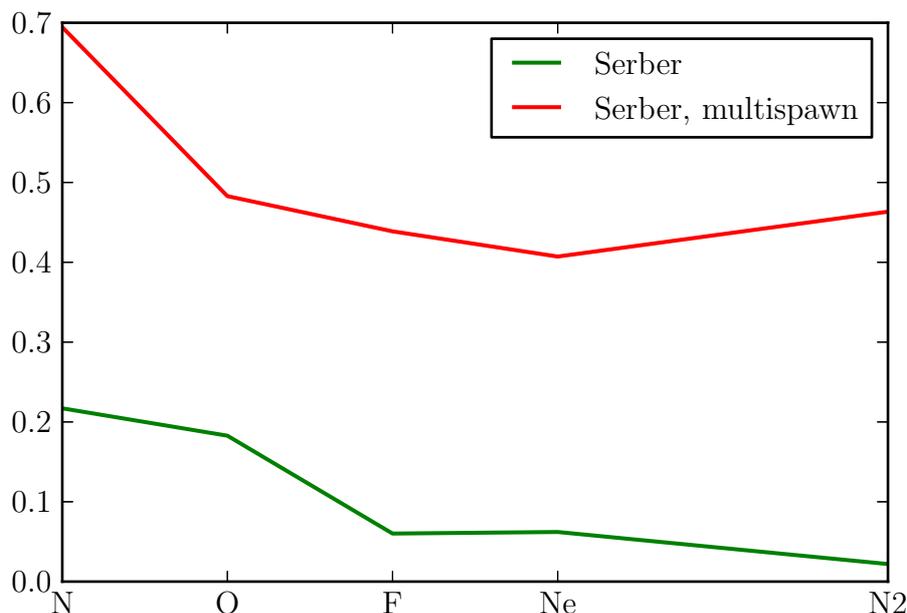


Figure 5.4: Variation in the ratio between the maximum $\delta\tau$ permitted without blooms of three or more particles occurring between determinantal calculations and those involving Serber functions, with or without multiple-structure spawning, as the number of electrons is increased. All of these systems are modelled using a cc-pVDZ basis set. $\delta\tau$ is found by searching, as described in section 2.6. The lighter first row atoms are not included, as determining a bound on $\delta\tau$ for such a small system becomes meaningless.

tion 2.16, to give

$$\frac{\sum_{\mathbf{j} \in \mathbf{J}} |c_{\mathbf{j}}|}{\gamma} \times \left[-\delta\tau\gamma \frac{\sum_{\substack{\mathbf{j} \in \mathbf{J} \\ \mathbf{j} \neq \mathbf{k}}} (K_{\mathbf{k}\mathbf{j}} - ES_{\mathbf{k}\mathbf{j}}) c_{\mathbf{j}}}{p_{\text{gen}}(\mathbf{K}|\mathbf{J}) \sum_{\mathbf{j} \in \mathbf{J}} |c_{\mathbf{j}}|} \longrightarrow \Delta c_{\mathbf{k}} \quad \forall \mathbf{k} \in \mathbf{K} \in \{\mathbf{I} \leftarrow \mathbf{J}\} \right].$$

As shown in figure 5.3c, all of the occupied sites, \mathbf{j} , in a source spatial structure, \mathbf{J} , are considered in generating the spawns to each of the available sites, \mathbf{k} , associated with a target spatial structure, \mathbf{K} . The total number of spawns attempted is determined from the cumulative weight on all of the occupied sites being spawned from, with the magnitude of each spawning attempt, γ , being determined as before.

By including as many terms, containing information from as many occupied sites and as many matrix elements as possible, blob-to-blob spawning acts to smooth out the stochastic changes in the wavefunction and thus to minimise the statistical noise.

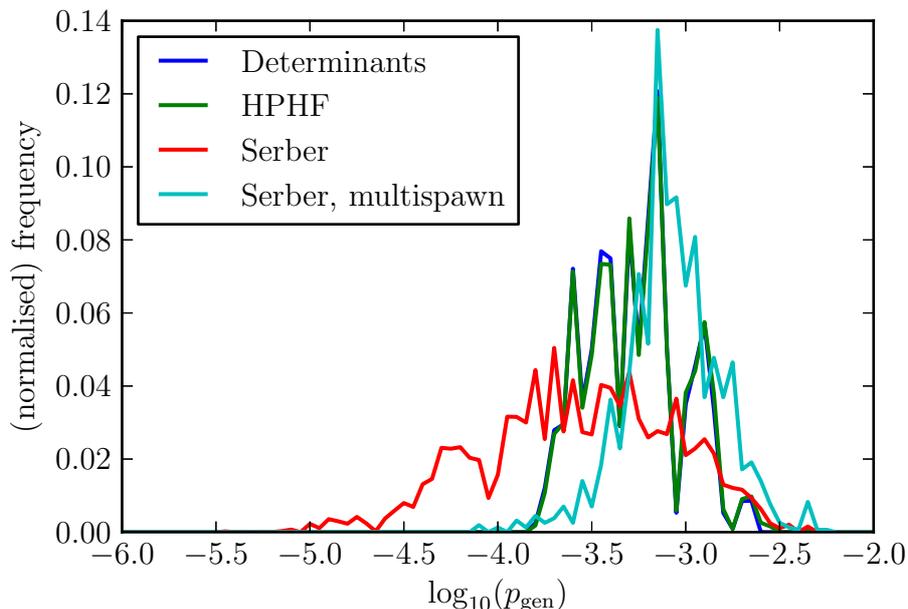


Figure 5.5: Histogram of p_{gen} values for stretched N_2 (bond length $4.2a_0$) in a cc-pVDZ basis, using Slater determinants, HPHF functions and Serber CSFs (as a prototypical spin eigenfunction). Note the substantial improvement in generation probabilities when making use of multiple-structure spawning with Serber functions.

It is worth noting that such ‘blob-to-blob’ spawning has no substantial impact on the maximum time step that may be used compared to the normal multiple-structure spawning, as the generation probability depends only on the spatial excitation such that $p_{\text{gen}}(\mathbf{j}|\mathbf{i}) = p_{\text{gen}}(\mathbf{J}|\mathbf{I})$, and is therefore unchanged.

5.6 Spawning in mixed representation schemes

When using a mixed representation scheme, any spawning which crosses the boundary needs careful consideration. In the mixed schemes in use with CSFs, the partitioning of the space is performed according to the number of unpaired electrons, N_o . Below a threshold the basis functions are CSFs, and above this threshold either determinants or HPHF functions are used.

Special considerations must be made for exciting out of the CSF portion of the space (increasing N_o) or back into it (decreasing N_o).

Increasing N_o

If spawning is occurring from a site within the CSF portion of the space, the normal CSF excitation generator is used. Once the spatial structure has been selected, and before the eigenfunction label is applied (see flowchart in figure 5.2), the change in the number of unpaired electrons, ΔN_o , is considered. If this does not take the excitation over the threshold, the normal excitation process is continued. If the new value of $N_o > N_{o,\max}$, a random spin structure is selected and applied to the spatial structure, and it is now considered to be a Slater determinant.

As the CSF excitation generators specifically generate spatial structures with the correct probability, there is only one way of generating each Slater determinant and so the generation probability is given by $p_{\text{gen}} = p_{\text{spatial,csf}} \times \frac{1}{n_{\text{det}}(N_{o,\text{new}})}$.

Decreasing N_o

If spawning is occurring from a site within the Slater determinant portion of the space, the normal Slater determinant excitation generator is used. Once a target Slater determinant is selected the change in the number of unpaired electrons, ΔN_o , is considered. If this reduces N_o such that it falls below $N_{o,\max}$, the occupied spatial orbitals are extracted from the spin-orbital representation of the generated determinant, and then a random spin function (with the correct number of unpaired electrons) is applied.

As the Slater determinant excitation generators must work in *spin* orbitals, there are multiple ways in which the same *spatial* structure may be generated from a given starting determinant. This degeneracy must be considered in the generation probability — the probabilities of generating each of the possible determinants with the same spatial structure must be summed. Fortunately, as demonstrated in figure 5.6 where all of the classes of excitation are enumerated, all of the cases where the number of unpaired electrons is reduced have a degeneracy of one. As such the generation probability for these excitations can be given as $p_{\text{gen}} = p_{\text{gen,det}} \times \frac{1}{n_{\text{csf}}(N_{o,\text{new}})}$.

This same analysis may be performed for HPHF functions (section 2.7.2) instead of Slater determinants. As HPHFs are represented by a subset of Slater determinants (matched with their spin-paired determinants with the same spatial structure) the only change to the excitation generation process is that the number of permit-

ted spin functions to choose from when increasing N_o is halved, and as such the generation probabilities are doubled. As the parity within a HPHF matches that within any CSFs within the same space, the matrix element is simply multiplied by a factor of $\sqrt{2}$.

Single excitations	ΔN_o	Degen.	Double Excitations (from opposing singles)	ΔN_o	Degen.
	0	1		-2	1
	-2	1		0	2
	+2	2		-2	1
	0	1		-4	1
Double excitations (from closed pair)			Double Excitations (from matched singles)		
	0	1		0	1
	+2	2		-2	1
	0	1		-4	1
	-2	1	Double Excitations (1 closed, 1 open)		
Double Excitations (from two closed pairs)				0	1
	+2	2		+2	3
	+4	6		0	2
	+2	3		0	1
	0	2		-2	1
	0	1		-2	1

Figure 5.6: Degeneracy of spatial excitations within the determinantal excitation generator considering all of the classes of excitation which may be generated, categorised by their effect on *spatial*-excitations. Each pair of spin orbitals is represented by two stacked boxes. On the left, electrons which may be selected are in green, with alternatives in blue. Electrons which are present but may not be selected are in grey. On the right, electrons which are already present are in grey, and newly placed electrons are in green. The change in unpaired electrons and the degeneracy of the choice are indicated. All selections which reduce the number of unpaired electrons are listed in **bold**.

6 Nitrogen

6.1 The nitrogen atom

The nitrogen atom is particularly interesting from the perspective of using spin eigenfunctions with FCIQMC. It has two unusual properties;

High spin ground state

The ground state of the molecule has configuration $1s^2 2s^2 2p^3$, with all of the unpaired electrons spin-aligned. As a consequence, this state has a spin of $\frac{3}{2}$, with excited states existing with lower total spin. This means that the excited states cannot be isolated solely by restricting the M_s value of the Slater determinants used.

Odd number of electrons

Because Nitrogen has 7 electrons, it does not have any $M_s = 0$ states, and as a consequence HPHF functions may not be used. This means that any non-ground state spin states obtained are not otherwise obtainable by FCIQMC.

The nitrogen molecule has been extensively characterised experimentally, and some appropriate experimental results are summarised in table 6.1. The ground state is the 4S state, with $S = \frac{3}{2}$ and orbital angular momentum $L = 0$. Restricting the total spin to $S = \frac{1}{2}$, should extract the first excited state of the system labelled 2D with $L = 2$. This state could also be obtained by restricting M_l in a determinantal calculation. Simultaneous restriction of M_s or $S = \frac{3}{2}$ and $L = 1$ or 2 should extract the 4P and 4D states respectively. It is expected that as finite basis sets are being used these states are likely to be less quantitatively accurately obtained than the lower states.

In the cc-pVDZ¹²² basis set, the behaviour of spin-projected determinants and Serber functions are compared to Slater determinants, and it is demonstrated that the different energy levels may be obtained with each. The values obtained are

Configuration	Term		Energy / E_h
	symbol	Energy / cm^{-1}	
$2s^2 2p^3$	4S	0.000	0.000 000
$2s^2 2p^3$	$^2D^\dagger$	19 224.464	0.087 593
		19 233.177	0.087 633
$2s^2 2p^3$	$^2P^\dagger$	28 838.920	0.131 400
		28 839.306	0.131 402
$2s^2 2p^2(^3P)3s$	$^4P^\dagger$	83 284.070	0.379 470
		83 317.830	0.379 624
		83 364.620	0.379 837
$2s^2 2p^2(^3P)3s$	$^2P^\dagger$	86 137.350	0.392 471
		86 220.510	0.392 850
$2s 2p^4$	$^4P^\dagger$	88 107.260	0.401 446
		88 151.170	0.401 646
		88 170.570	0.401 735
$2s^2 2p^2(^3P)3s$	2S	93 581.550	0.426 389
$2s^2 2p^2(^3P)3s$	$^4D^\dagger$	94 770.880	0.431 808
		94 793.490	0.431 911
		94 830.890	0.432 081
		94 881.820	0.432 313

[†] Multiple energetic values associated with a single configuration and term symbol are caused by spin-orbit coupling. In particular the total angular momentum, J , has permitted values in integer steps in the range $|S - L| \leq J \leq |S + L|$. This effect is not observed with a non-relativistic Hamiltonian, and so the coupling is not reproduced in further calculation here.

Table 6.1: Experimentally obtained values^{120,121} for selected energy levels of N. The term symbols are included to indicate which states are obtainable by making use of the various symmetries available within NECL.

compared to those obtained by diagonalising the Hamiltonian matrix explicitly. The underlying contracted gaussian basis set is then increased from the double-zeta cc-pVDZ basis set through to the quintuple-zeta cc-pV5Z basis set¹²², and the convergence relative energies of the atomic states obtained are compared to the experimental values available.

All Hartree–Fock calculations were performed using Q-CHEM¹²³ with modifications to dump integrals with \hat{L}_z symmetry preserved by Alex Thom. At present, the version of Q-CHEM available does not support angular momentum functions high enough to go beyond cc-pV5Z basis sets. In all of the FCIQMC simulations the core (1s) orbitals are frozen, to reduce the calculation to one involving only

five electrons.

6.1.1 Different types of CSF

Before any further considerations can be made about the different available CSFs, it is important to demonstrate that FCIQMC works in the different regimes available. In the case of the nitrogen atom, it is important to demonstrate that the simulation is able to converge onto all of the states that can be selected using the spin and angular momentum symmetry available using all of the available types of CSF.

The trajectories of the projected energy converging onto the 4S and 2D states are presented in figure 6.1 for each of the basis sets. In imaginary time, the behaviour of each of these CSFs is essentially indistinguishable – all of them converge at roughly the same rate with the same degree of accuracy to the same energy levels.

One of the presumed benefits of using CSFs was the potentially more compact representation of the converged wavefunction, as a consequence of the radically smaller region of the Hilbert space being considered after implicitly block diagonalising the Hamiltonian matrix. If the coefficients are plotted in reverse order of occupation size, this would be expected to result in a dramatic reduction in the ‘tail’ of the wavefunction, implying that the representation has become more compact. Figure 6.2 presents the structure of the coefficients associated with the two lowest energy states that are available.

It is clear that the wavefunction representations are to all practical purposes exactly the same for the ground state. As the ground state is high spin, and the determinantal space is restricted to only use determinants where $M_s = \frac{3}{2}$, the size of the CSF space is similar to the size of the determinantal space (there are still determinants with five unpaired electrons, and the spin eigenfunction space with $S = \frac{3}{2}$ and five unpaired electrons is smaller than the determinantal one). This suggests that the restriction of the space to high spin states by use of M_s has a very similar impact on the simulation to restricting the total spin by using spin eigenfunctions, resulting in canonical determinantal FCIQMC already evolving its coefficients to produce an efficient representation of these spin states.

The representation of the excited 2D state using Serber functions is also essentially identical to that for Slater determinants. In this case, the means used to restrict the space are different. Using Serber functions, the total spin is restricted such

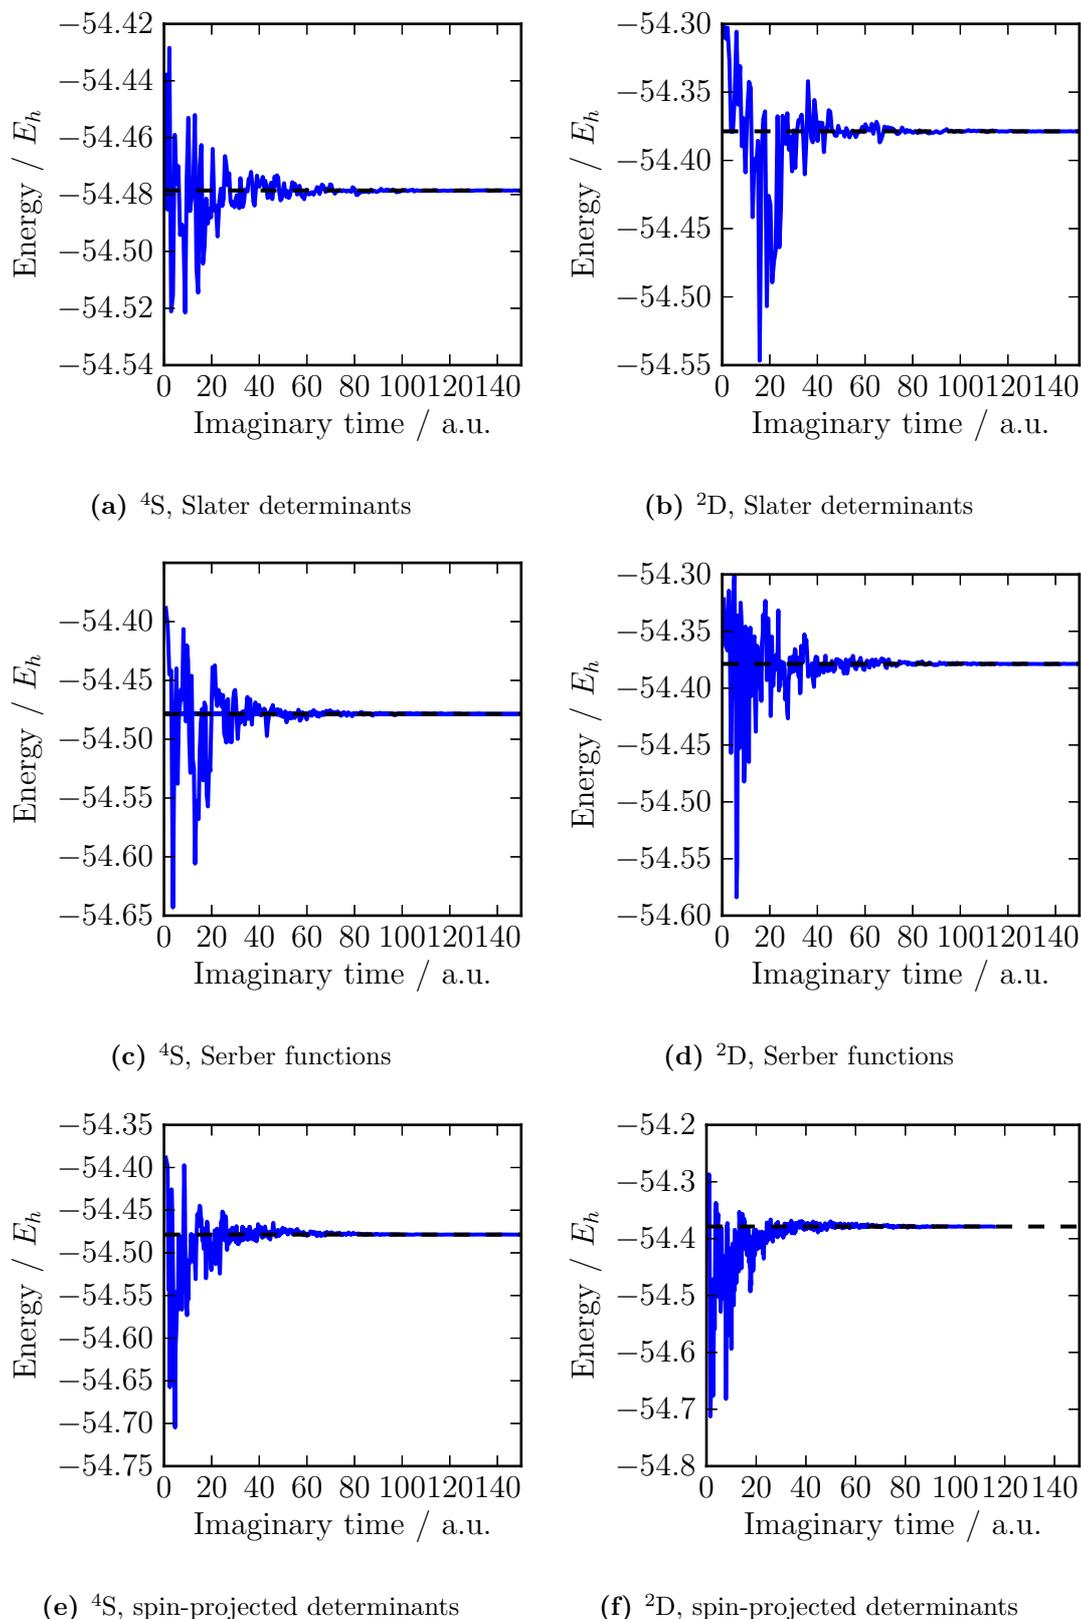


Figure 6.1: The convergence of the projected energy onto the available eigenstates for N in a cc-pVDZ basis set. This demonstrates convergence onto multiple available states using the different available spin eigenfunctions with S restricted, and Slater determinants with M_l restricted. The known FCI energy is a dotted black line.

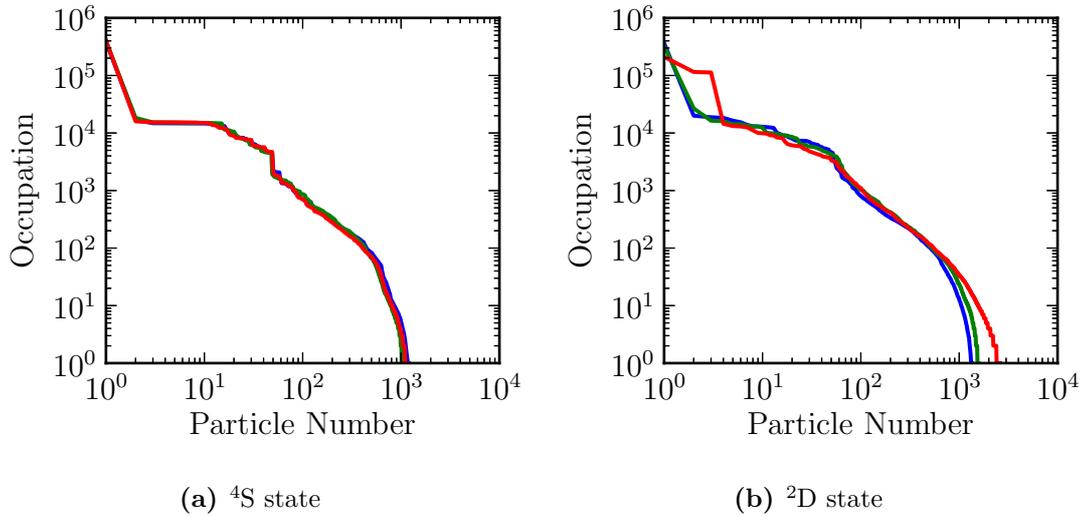


Figure 6.2: Plots of the coefficient on the n^{th} most populated site for the different CSFs compared to the determinantal solutions. The first excited state is obtained by restricting S for the spin eigenfunctions, and by restricting M_l for the determinantal solution. The occupations for Slater determinants are plotted in blue, Serber CSFs in green and spin-projected determinants in red.

that $S = \frac{1}{2}$. When using Slater determinants, the projected magnetic quantum number is restricted such that $M_l = 2$. It is possible to conclude, therefore, that the effect of the two symmetries on the Hilbert space is extremely similar.

The representation of the excited $2D$ state using spin projected determinants is worse than the representations using Slater determinants or Serber functions, requiring substantially more particles to be located on a larger number of sites. This is probably a consequence of the over-complete nature of the basis set, and suggests this as a disadvantage for further use of spin-projected determinants.

The remainder of calculations performed have made use of Serber spin eigenfunctions, as they have the best computational performance, and the most controlled behaviour, of any of the implemented CSFs.

6.1.2 Convergence of errors

There are three primary sources of error involved in an i -FCIQMC calculation that need to be controlled to obtain meaningful results; errors associated with the statistical analysis of the results, the initiator error which depends on the number

of particles in the simulation, and the errors associated with the basis set.

Statistical accuracy

The first element of improving the statistical quality of the results is to reduce the magnitude of the fluctuations in the results. Primarily this has been approached by using a mixed coefficient representation scheme, with the coefficients that contribute to the projected energy, and those immediately connected to them, represented by real numbers ($\omega = 0$), and all other coefficients fully discretised as integers (see section 2.4.1).

Beyond this, the approach is depressingly empirical and twofold. Firstly, it is important to ensure, by visual inspection, that data is only considered after the wavefunction has converged. Secondly, if the error bars in the collected results are too large, the calculation should be resumed and run for longer to reduce the errors.

Initiator error

As discussed previously, in section 2.4.4, the only systematic approach to eliminating the initiator error is to run increasingly large calculations until it is clear that the error has been eliminated. Fortunately the nitrogen atom is a relatively small system, having only five electrons, and even in the larger basis sets does not require vast numbers of particles to converge effectively.

Figure 6.3 includes some sample plots demonstrating the convergence of the resultant energy onto the FCI energy obtained by diagonalisation.

Basis set error

No FCI method can obtain a larger fraction of the physical, experimental, correlation energy than may be represented by the basis set.* It can be observed as the basis set is increased that a larger proportion of the total energy is resolved.

This analysis of the nitrogen atom makes use of the sequence of basis sets developed by Dunning¹²². These present a systematic increase in complexity of the representation of the wavefunction, and can in principle provide the basis for extrapolation to the infinite basis set limit¹²⁴ by writing the energy

$$E(X) = E_{\text{exact}} + AX^{-3}$$

*A larger proportion may be obtained, along with some relativistic corrections, by application of perturbative corrections in addition to the FCI methodology.^{21,36}

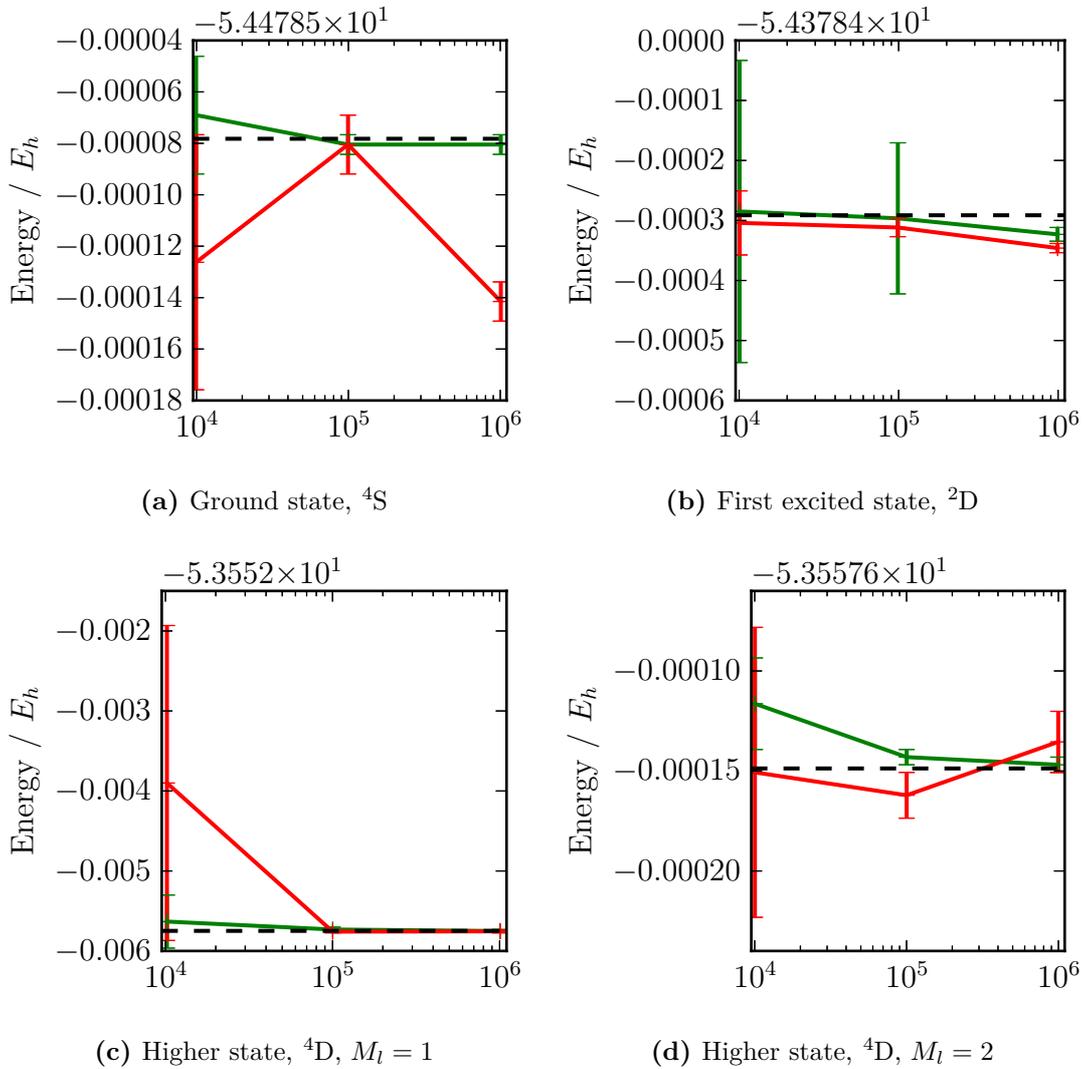


Figure 6.3: Convergence of energies for the nitrogen atom with particle count in a cc-pVDZ basis set. This demonstrates the reduction in both random and initiator error as the number of particles is increased. Projected energy values are plotted in green, shift values in red and the known FCI energy is a dotted black line. Note that the shift energy estimator is much less well behaved than the projected energy.

Basis set	State / E_h			
	4S	2D	$^4D, M_l = 1$	$^4D, M_l = 2$
cc-pVDZ	-54.478 581(5)	-54.378 72(1)	-53.557 75(2)	-53.557 746(5)
cc-pVTZ	-54.514 891(6)	-54.422 688(9)	-53.8013(2)	-53.801 61(1)
cc-pVQZ	-54.525 037(9)	-54.435 58(2)	-53.904 06(2)	-53.908 02(4)
cc-pV5Z	-54.528 13(2)	-54.439 84(3)	-53.9986(8)	-54.001 87(8)
Extrapolated V(Q5)Z	-54.531 38	-54.444 29	-54.097 89	-54.100 34

Basis set	Excitation energy / E_h		
	$^4S \rightarrow ^2D$	$^4S \rightarrow ^4D, M_l = 1$	$^4S \rightarrow ^4D, M_l = 2$
cc-pVDZ	0.099 85(2)	0.920 82(2)	0.920 83(1)
cc-pVTZ	0.092 20(2)	0.7135(1)	0.713 27(1)
cc-pVQZ	0.089 44(3)	0.620 97(3)	0.617 01(5)
cc-pV5Z	0.088 29(6)	0.5294(8)	0.5262(1)
Extrapolated V(Q5)Z	0.087 08	0.4334	0.4310
Experimental	0.087 61	0.432 03	0.432 03

Table 6.2: A progression of converged state and excitation energies for the Nitrogen atom in a series of Dunning basis sets. The extrapolated limit is included for comparison. It is worth noting that the 4P state is not resolved in any of these simulations, implying that within the basis sets considered it is not straightforwardly accessible, or has crossed to a higher energy than the 4S state. As such, both the simulations with $M_l = 1$ and $M_l = 2$ converged to the same state. The absence of this state agrees with the FCI result from explicit diagonalisation in the cc-pVDZ basis set. The experimental results provided for comparison are the averaged values of those provided for the excitation (ignoring spin-orbit coupling) in table 6.1.

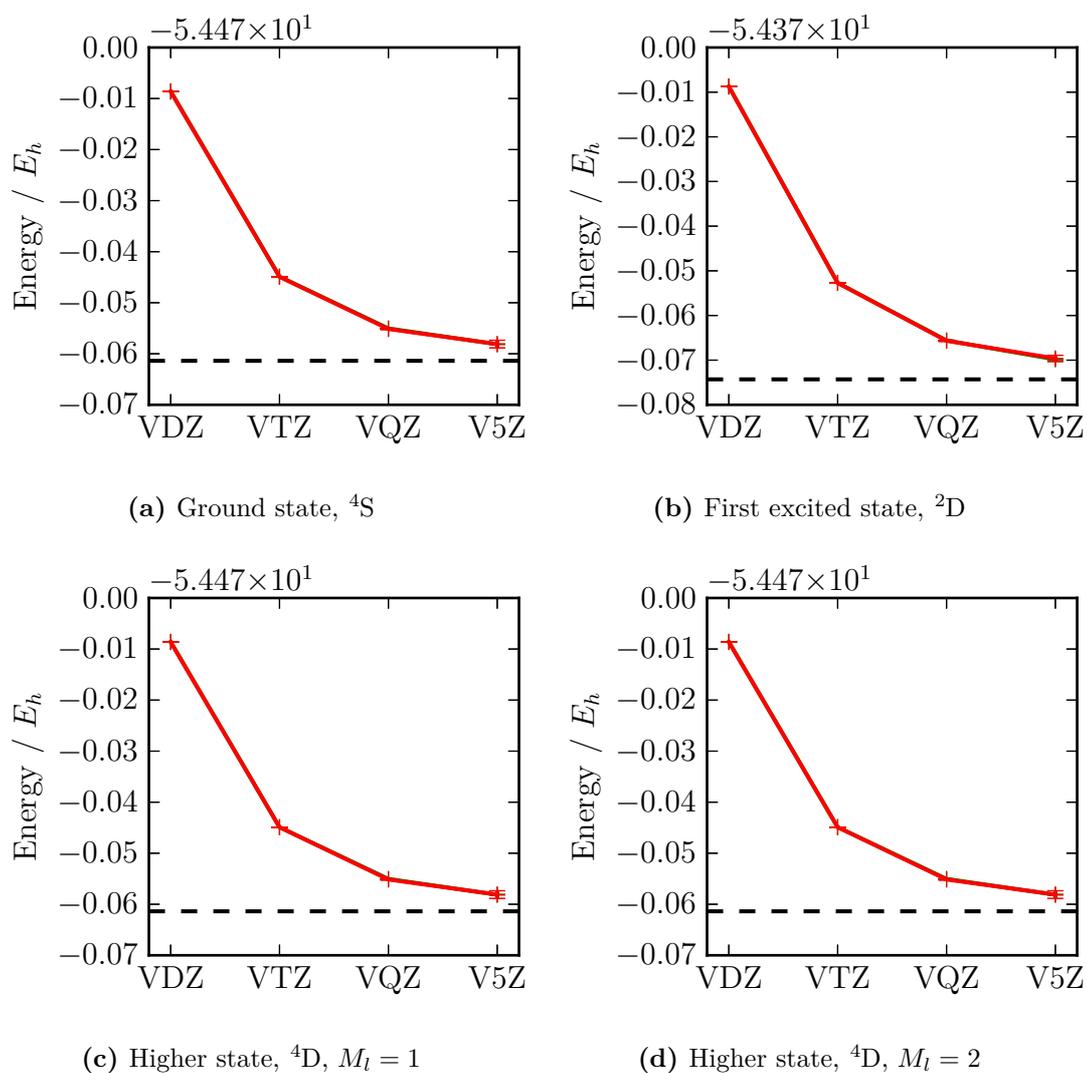


Figure 6.4: Convergence of energies for several states of the nitrogen atom with basis set size. The projected energy is plotted in green, and the shift in red. Error bars are included. The extrapolated infinite basis set limit is included as a dashed black line.

where X is the cardinal number of the basis set. This can then be fitted to the data being considered.*

Table 6.2, with results plotted in figure 6.4, presents the convergence of the total energy associated with the four states being considered with the increase in basis set size. The asymptotic behaviour can be clearly seen, although the V5Z basis has not yet reached the converged value.

It is worth noting that the ^4P state has not been resolved. This state was also missing from the full diagonalisation performed in the cc-pVDZ basis set, which implies that it is either not well represented in these bases or that it has crossed the ^4D state to be higher in energy, and therefore not appearing in the correct place. As such, all of the $M_l = 1$ and $M_l = 2$ simulations with $S = \frac{3}{2}$ have converged on the ^4D state.

For comparison with experiment, the differences between the energy levels need to be considered. The convergence of the difference between two energy levels does not occur at the same rate as the convergence of the energy levels themselves. Figure 6.5 presents the convergence on these energy gaps with basis sets size relative to the known experimental values. The extrapolated limits for all of the states considered are within ‘chemical accuracy’⁶⁰ of the experimental values (1 kcal mol^{-1}).

6.1.3 Summary

It is clear that FCIQMC using CSFs is able to obtain results to an accuracy useful for comparison to experiment, including some excited states. This is quite exciting for further work.

In the context of the nitrogen atom, it turns out that all of the available states may be obtained by applying restrictions on the total values of M_s and M_l associated with the determinants in use. While this is extremely convenient for the benchmarking of spin-pure FCIQMC, it fundamentally reduces its effectiveness. To be noticeably useful for restricting the state that is converged to, it would be necessary to implement \hat{L}^2 eigenfunctions to work in conjunction with the \hat{S}^2

*In this case only the data corresponding to the highest two cardinal numbers calculated are used. More sophisticated extrapolation schemes may be used, in principle, and may improve accuracy.

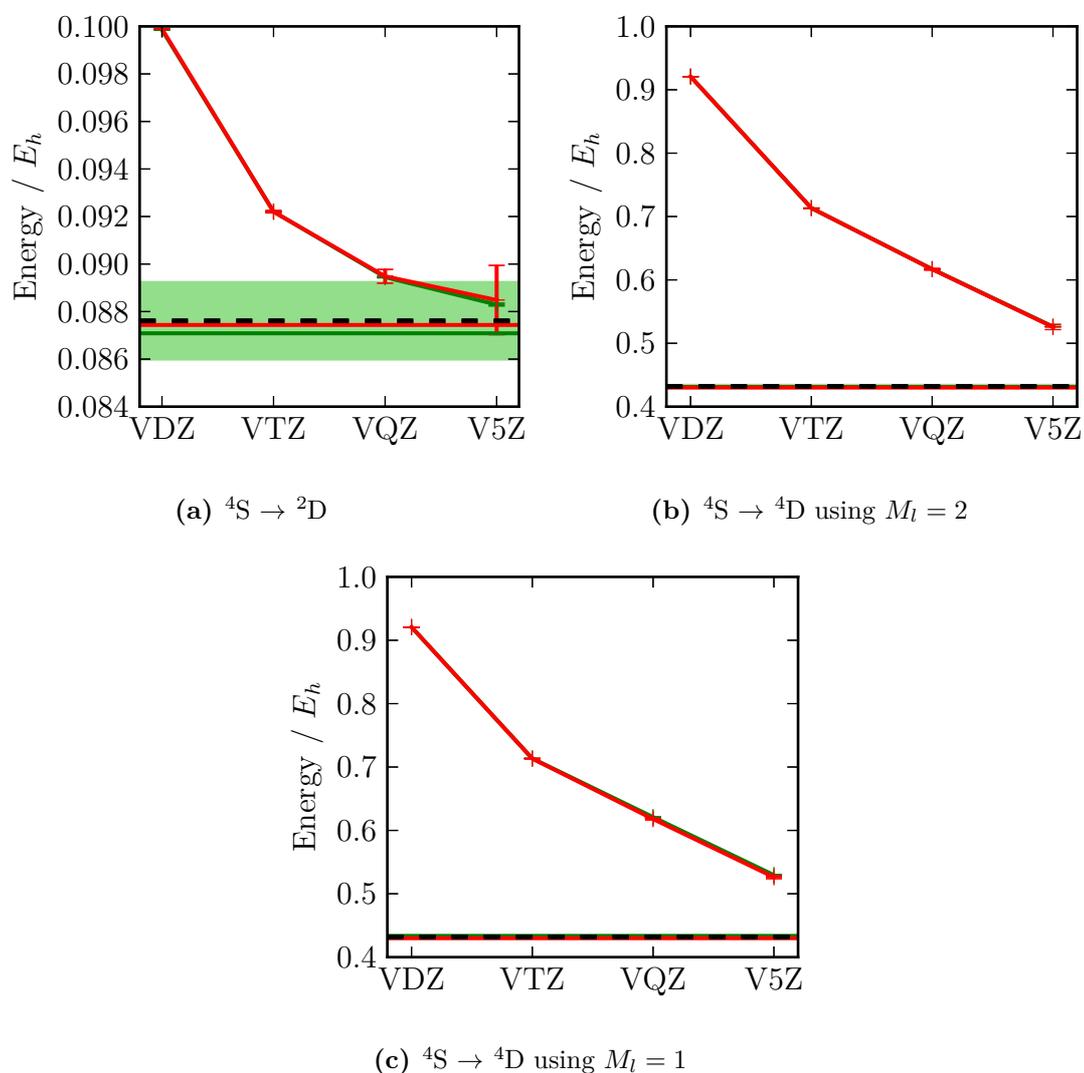


Figure 6.5: The convergence of the excitation energies for the states being considered with increases in basis set size. The projected energy and the value obtaining by extrapolating the *cc*-pVQZ and *cc*-pV5Z (V(Q5)Z) results are plotted in green, the shift likewise in red. The experimental values are plotted using a black dashed line, surrounded by a band which indicates ‘chemical accuracy’⁶⁰ of 1 kcal mol^{-1} . The energies associated with the different spin-orbit couplings are indistinguishable on the scale of stochastic error and convergence in these simulations.

Note that the ground state of the atom is better represented in the basis set than the excited state, and hence its energy is more accurate. As the energies converge from above, so does the difference between the two states.

eigenfunctions currently available.

It is also disappointing that the changes caused to the representation of the wavefunction by restricting the space to CSFs do not appear to have any greater impact on the representation of the wavefunction than using other forms of symmetry to attain the same states.

6.2 The nitrogen dimer

The nitrogen dimer, N_2 , is a very strongly bonded homonuclear diatomic (it has an effective triple bond). Its ground state electron configuration is $(1s)\sigma_g^2(1s)\sigma_u^2(2s)\sigma_g^2(2s)\sigma_u^2(2p)\pi_u^4(2p)\sigma_g^2$, and it has several excited states while still maintaining the overall configuration $1s^42s^42p^6$, the relevant configurations are, with distinct S values, are shown in figure 6.6.

By inspection of the orbitals structure, it would be expected that as the required value of the total spin is increased, the energy of the excited states will increase in this order. This would mean that the energy levels could all be obtained by making use of normal determinantal FCIQMC and restricting the projected spin value, M_s — the opposite of what made the nitrogen atom interesting. There are, however, several reasons why the binding curves of nitrogen are potentially interesting.

Four different spin states

As a consequence of the relatively high number of electrons in p-orbitals, and consequently the potential number of electrons that can singly occupy p-orbitals, there are four different spin states that are physically and computationally interesting to obtain the binding curves for (and two different spatial symmetries associated with some of these spin states). This provides an excellent test for the ability of FCIQMC to operate correctly in different CSF regimes.

Dissociates to two nitrogen atoms

The dissociated limit under all of the total molecular spin states is two non-interacting nitrogen atoms in their ground state. This is of interest for several reasons:

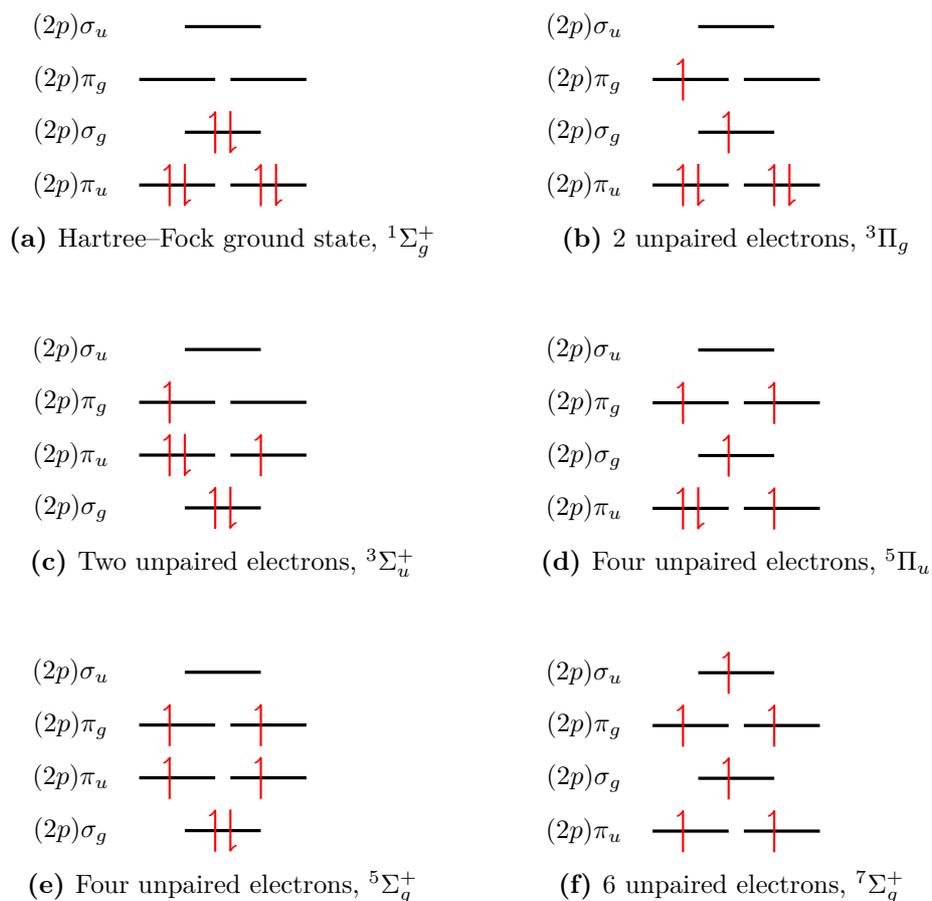


Figure 6.6: Reference Slater determinants for N_2 with different M_s values and symmetries. These correspond to the different excited spin states, with the associated value of S . Only the electrons in $2p$ orbitals have been considered. Note that at short bond lengths, the σ_g orbital is higher in energy than the π_u orbitals due to sp mixing, resulting in the ${}^3\Pi_g$ and ${}^5\Pi_u$ being lower in energy. At longer bond lengths, this mixing reduces, and the ${}^3\Sigma_u^+$ and ${}^5\Sigma_g^+$ states become lower respectively (as drawn).

1. These are the atoms studied above, and as a consequence the ground state is well characterised.
2. This ground state requires all of the electrons to be unpaired. As a consequence, only the $S = 3$ state should dissociate correctly in the Hartree–Fock calculation, as its electrons are dissociated throughout the binding curve as shown in figure 6.6f.
3. The energy levels converge as the bond length becomes long, making some of the calculations extremely difficult using Slater determinants. Using CSFs removes the low-lying excited states and accelerates convergence.

4. In the fully stretched limit, the two Nitrogen atoms exist in a four fold degenerate state, in which $S = 0, 1, 2, 3$ are all permitted.

Do the states cross?

In the Hartree–Fock solutions the excited states corresponding to the same value of S , but different spatial symmetries, cross. It appears to be a reasonable assumption that within spatial symmetries the states corresponding with different values of S do not cross — even though this is not true for the Hartree–Fock states. Considering the spin states separately with CSFs gives a great deal more confidence in the results in this respect.

All of the calculations in this basis set make use of the cc-pVDZ basis set, with the underlying Hartree–Fock calculations performed by MOLPRO^{125,126}. In all of the FCIQMC simulations the core (1s) orbitals are frozen, to reduce the calculation to one involving only ten electrons.

6.2.1 Hartree–Fock solutions

Before FCIQMC simulations can be performed, the antisymmetrised basis set is generated through Hartree–Fock simulations for each of the bond lengths. In principle, as the underlying basis set is the same in each case, there is no benefit to running Hartree–Fock simulations for the different M_s values. FCIQMC will generate the same overall wavefunction. These simulations are, however, useful for three primary reasons.

1. They obtain an optimised Hartree–Fock ground state, rather than just what is left over in the virtual orbitals. This can be used to demonstrate the failure of Hartree–Fock methods at obtaining useful binding curves.
2. The optimised Hartree–Fock states for each M_s value and spatial symmetry can be used as the reference sites in FCIQMC simulations. This ensures maximum weight on the reference site and improves the statistical quality of the results.
3. If the ground state Hartree–Fock solution is used to generate all FCIQMC results, it is not trivially obvious what the reference site should be. Obtaining this from the Hartree–Fock solver simplifies calculations, especially as this site changes across the space as the molecule dissociates.

Bond Length / a_0	Energy / E_h					
	$^1\Sigma_g^+$	$^3\Pi_g$	$^3\Sigma_u^+$	$^5\Pi_u$	$^5\Sigma_g^+$	$^7\Sigma_g^+$
0.90	-102.004 558	-101.474 672	-100.815 801	-100.953 134	-99.505 009	-99.606 125
1.00	-104.125 525	-103.507 401	-103.009 579	-103.036 524	-101.834 707	-101.817 955
1.10	-105.641 086	-105.002 042	-104.650 290	-104.570 638	-103.620 182	-103.450 042
1.20	-106.716 470	-106.092 594	-105.836 150	-105.692 887	-104.934 189	-104.611 210
1.30	-107.475 068	-106.882 645	-106.694 785	-106.510 569	-105.907 383	-105.509 203
1.40	-108.005 991	-107.451 838	-107.317 145	-107.104 883	-106.632 382	-106.170 565
1.50	-108.372 715	-107.859 195	-107.767 561	-107.535 218	-107.174 705	-106.657 809
1.60	-108.620 490	-108.147 774	-108.091 800	-107.844 551	-107.581 177	-107.016 435
1.70	-108.781 847	-108.348 989	-108.322 855	-108.064 105	-107.885 809	-107.279 951
1.80	-108.880 433	-108.485 906	-108.484 859	-108.216 786	-108.113 715	-107.473 581
1.90	-108.933 613	-108.575 599	-108.595 687	-108.295 491	-108.283 681	-107.534 462
2.00	-108.954 210	-108.630 781	-108.668 680	-108.402 047	-108.409 873	-107.692 558
2.10	-108.951 709	-108.660 940	-108.713 820	-108.478 503	-108.502 998	-107.828 491
2.20	-108.933 108	-108.673 151	-108.738 567	-108.532 514	-108.571 140	-107.951 445
2.30	-108.903 553	-108.672 672	-108.748 483	-108.569 838	-108.620 398	-108.064 465
2.40	-108.866 811	-108.663 396	-108.747 709	-108.594 797	-108.655 373	-108.166 648
2.50	-108.825 632	-108.648 184	-108.739 330	-108.610 627	-108.679 538	-108.256 921
2.60	-108.782 016	-108.629 126	-108.725 640	-108.619 744	-108.695 522	-108.335 335
2.70	-108.737 400	-108.607 731	-108.708 340	-108.623 948	-108.705 318	-108.402 742
2.80	-108.692 801	-108.585 075	-108.688 684	-108.624 569	-108.710 430	-108.460 331
2.90	-108.648 921	-108.561 911	-108.667 588	-108.622 591	-108.711 995	-108.509 349
3.00	-108.606 226	-108.538 754	-108.645 709	-108.618 736	-108.710 868	-108.550 975
4.00	-108.270 952	-108.345 254	-108.442 590	-108.547 577	-108.641 523	-108.737 146
6.00	-107.962 749	-108.607 731	-108.219 370	-108.440 476	-108.506 778	-108.776 132
8.00	-107.852 633	-108.086 814	-108.133 456		-108.445 499	-108.776 842
10.00	-107.807 876	-108.058 348	-108.100 918		-108.425 155	-108.776 829

Table 6.3: Hartree–Fock energies for N_2 in a cc-pVDZ basis set at a variety of bond-lengths with specified values of the projected spin eigenvalue, $M_s = 0, 1, 2, 3$. The data in this table are plotted in figure 6.7.

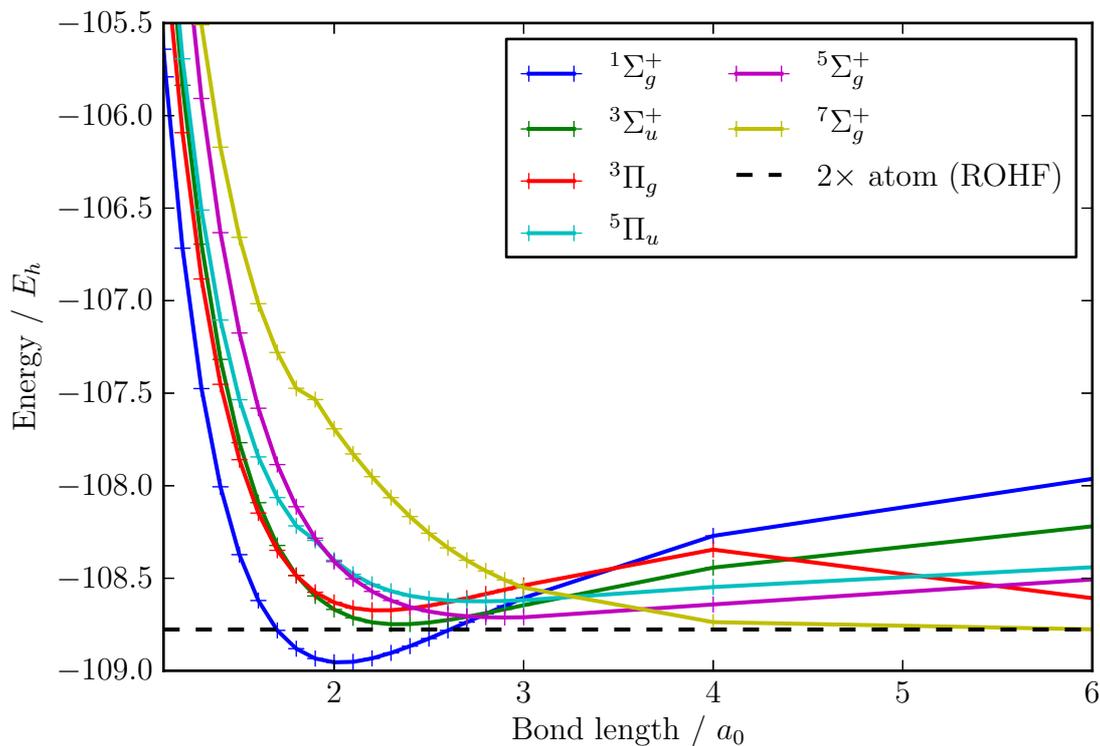


Figure 6.7: Binding curves for N_2 in a cc-pVDZ basis set using Hartree–Fock. Note that the ground state bond length is reasonable. However, only the (unbound) ${}^7\Sigma_g^+$ dissociates to the correct energy. Note also the energy crossings, where the different states with the same M_s value cross.

The Hartree–Fock energies obtained are listed in table 6.3, and plotted as binding curves in figure 6.7. Entirely apart from quantitative errors, it is clear that there are fundamental issues with the Hartree–Fock solution. In particular, the dissociation behaviour of restricted Hartree–Fock calculations is known to be extremely perverse. RHF calculations will not split electrons that are paired at the equilibrium geometry, resulting in energies for the dissociated species that are substantially higher than two times the atomic ROHF energy, and increasing with separation.

The severity of this problem increases with the number of paired electrons in the problem, resulting in a rough inversion of the energetic order of the orbitals. Only a state with all of the valence electrons unpaired is capable of dissociating into two high-spin nitrogen atoms, resulting in the only state with plausible dissociation behaviour being the $S = 3$ state.

There are also couple of slight discontinuities in the binding curves, most notably in the ${}^7\Sigma_g^+$ state. These correspond to the points in the simulation where the distortion of the geometry causes the orbital energies to cross further — in particular, at very short bond lengths, the 2s orbitals cross with the lowest bonding 2p orbitals, resulting in a change of symmetry. In practice, this involves leaving the configurations that we are interested in solving for.

Once the simulation is no longer near the Hartree–Fock ground state, the Fock space becomes a soup of different local minima. It is entirely possible that using Hartree–Fock metadynamics¹²⁷ would produce better and smoother curves. Fortunately, FCIQMC should be able to obtain the ground state in a given symmetry irrespective of the quality of the Hartree–Fock solution.

6.2.2 Binding curves for excited states

Using the Hartree–Fock solutions obtained above as the input data for the calculations, binding curves for N_2 can be obtained using FCIQMC. In this case Serber CSFs have been used to restrict the total spin value to be equal to the M_s value used in obtaining the Hartree–Fock solutions. This provides a degree of protection from any potential energy crossings, and simplifies convergence in the long bond length regime where the energy levels are extremely close together.

The best estimates of the energies are included in table 6.4 and plotted in figure 6.8. The estimated stochastic error bars are included in the binding curve plot, although they are not noticeably visible on the scale of the overall binding curve. There are several things to notice about these results:

‘Missing’ data points

The aim of this project was to determine the binding curves of the lowest energy state for each value of S . It turns out that some of the binding curves corresponding to the same value of S but different spatial symmetries cross. In the regions where a binding curve is the higher energy of the states with a given S , and the geometry is at the more extreme distortions of the molecule, the calculations become extremely difficult in a number of ways;

- The Hartree–Fock state is no longer the correct reference state for the FCI solution. It is difficult to find a reasonable reference site to use.*

*A UHF solution would work well, but cannot be used with the CSF paradigm used here.

Bond Length / a_0	Energy / E_h					
	$^1\Sigma_g^+$	$^3\Pi_g$	$^3\Sigma_u^+$	$^5\Pi_u$	$^5\Sigma_g^+$	$^7\Sigma_g^+$
0.90	-102.198 69(2)	-101.675 36(3)		-101.105 80(7)		
1.00	-104.325 92(1)	-103.723 46(1)		-103.193 62(5)		
1.10	-105.850 155(8)	-105.229 90(3)		-104.734 83(5)		
1.20	-106.935 569(7)	-106.331 24(3)		-105.866 07(1)		
1.30	-107.704 83(1)	-107.132 09(3)	-106.929 47(9)	-106.692 398(9)		
1.40	-108.246 776(8)	-107.711 94(3)	-107.553 22(8)	-107.295 70(8)		
1.50	-108.624 707(9)	-108.129 68(3)	-108.0102(1)	-107.734 26(1)	-107.371 32(2)	
1.60	-108.883 89(1)	-108.428 37(4)	-108.341 45(3)	-108.051 71(1)	-107.779 97(2)	-107.177 701(9)
1.70	-109.056 93(1)	-108.639 21(7)	-108.579 47(8)	-108.279 15(1)	-108.087 12(3)	-107.446 87(2)
1.80	-109.167 56(2)	-108.785 53(6)	-108.748 71(7)	-108.439 60(2)	-108.317 72(3)	-107.646 36(1)
1.90	-109.233 26(2)	-108.884 04(6)	-108.866 87(4)	-108.540 88(2)	-108.490 48(2)	-107.757 53(1)
2.00	-109.266 93(4)	-108.9480(1)	-108.9476(1)	-108.651 15(2)	-108.619 64(2)	-107.911 51(2)
2.10	-109.278 13(2)	-108.9865(1)	-109.0006(1)	-108.731 09(2)	-108.715 73(3)	-108.043 85(2)
2.20	-109.273 87(4)	-109.007 01(9)	-109.0336(2)	-108.788 25(5)	-108.787 02(3)	-108.163 43(2)
2.30	-109.259 36(4)	-109.0145(1)	-109.052 70(7)	-108.828 50(4)	-108.839 47(3)	-108.273 04(2)
2.40	-109.238 40(5)	-109.0130(1)	-109.0614(1)	-108.856 02(3)	-108.877 74(3)	-108.371 96(2)
2.50	-109.213 67(5)	-109.0059(1)	-109.062 95(9)	-108.874 12(3)	-108.905 19(4)	-108.459 31(2)
2.60	-109.187 07(4)	-108.9947(1)	-109.0594(1)	-108.885 35(3)	-108.924 67(3)	-108.535 11(1)
2.70	-109.160 29(5)	-108.981 16(8)	-109.0528(2)	-108.891 52(4)	-108.938 00(3)	-108.600 21(1)
2.80	-109.134 11(6)	-108.966 77(8)	-109.0441(3)	-108.894 13(4)	-108.946 81(4)	-108.655 73(1)
2.90	-109.109 19(1)	-108.951 98(9)	-109.0356(2)	-108.894 07(4)	-108.952 29(5)	-108.702 91(1)
3.00	-109.086 13(5)	-108.9376(1)	-109.0260(4)	-108.892 34(6)	-108.955 31(5)	-108.742 91(2)
4.00	-108.9720(4)		-108.9653(8)		-108.9507(1)	-108.920 62(2)
6.00	-108.959(1)		-108.957(1)		-108.9570(6)	-108.956 79(1)
8.00	-108.958(1)		-108.953(2)		-108.9559(5)	-108.957 27(2)

Table 6.4: FCIQMC energies for N_2 in a cc-pVDZ basis at a variety of bond-lengths for a number of different states. Spin is restricted using Serber CSFs, and spatial symmetries in the normal way. The data in this table are plotted in figure 6.8.

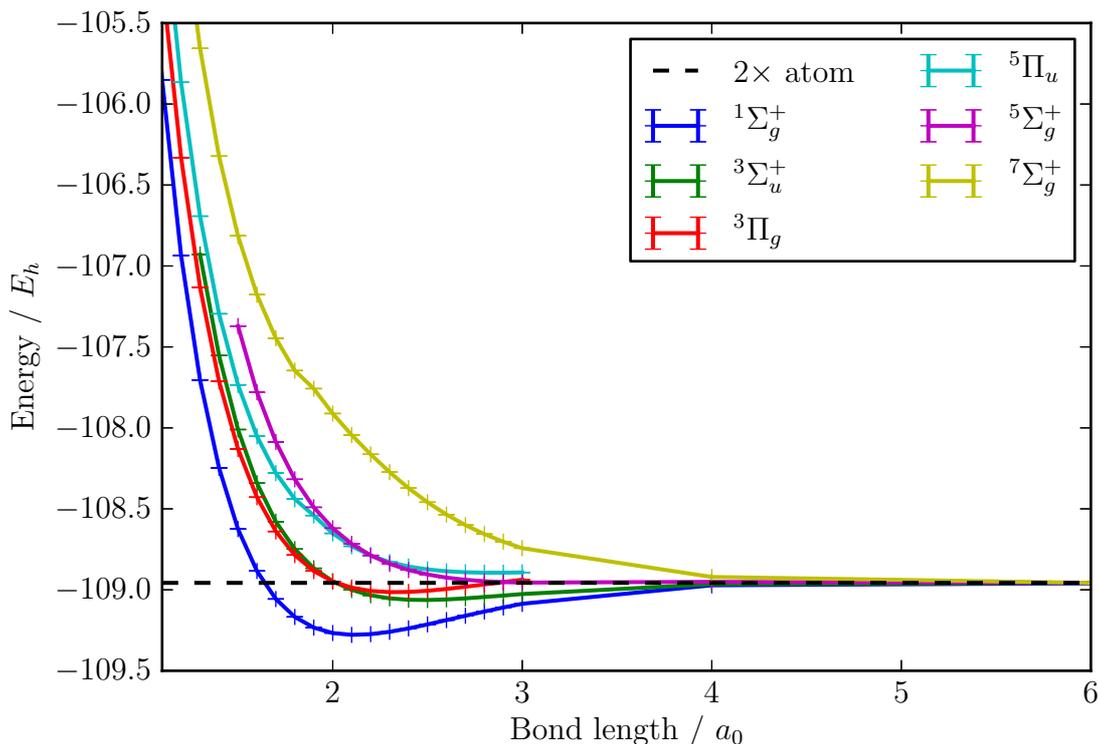


Figure 6.8: Binding curves for several different states of N_2 in a cc-pVDZ basis set. The statistical errors are included on the plot. Note that near the extreme geometries, some of the states where two curves with the same total spin S cross do not have all points considered. As the geometries become more extreme, orbital crossings in the higher energy states cause the configuration to change, and convergence of an FCIQMC simulation becomes extremely difficult.

- The solutions become highly multi-reference, with many sites carrying similar, large coefficients.
- The permitted time steps begin to fall dramatically, and the number of particles required to obtain a statistically accurate result becomes extremely high.

As such, some of these curves have been truncated in these limits, such that only the important region is considered.

Dissociation behaviour

The simulations become noticeably larger and more difficult as large bond lengths are approached, with the converged wavefunction taking on substantially more multi-reference character. Through all of this, however, the

lowest energy state corresponding to each value of S correctly tends to twice the atomic energy of Nitrogen in a cc-pVDZ basis as obtained in the previous section.

Bond lengths

It is worth noting that the equilibrium bond length in the ground state matches the experimental value well.

Comparison to experiment

As these results are performed in a cc-pVDZ basis set, even though qualitatively they are good, they do not provide a particularly good quantitative description of the energy levels in the Nitrogen molecule. It has already been demonstrated that FCIQMC can obtain good results for this system, although to get physically meaningful results requires going beyond cc-pVQZ basis sets¹²⁸. As the purpose of the study was to demonstrate the functionality of FCIQMC using CSFs, it is difficult to justify spending the large amount of computational time this would require.

6.2.3 Summary

Although doing so has consumed a substantial amount of computational time, it is clear that Serber functions are able to be used to effectively generate binding curves for multiple spin states in non-trivial sized systems.

This opens up the developmental capacity to use Serber functions for larger systems and to apply them to (arbitrarily sized) model systems such as the Uniform Electron Gas and the Hubbard model. These systems will be interesting to study, as their Hilbert spaces contain a very large number of sites of similar energy to the reference state, and as such their converged wavefunctions are very multi-configurational. However, to do this will require a substantial amount more development work to overcome the strict size limitations of the systems that can be considered using Serber functions before running out of memory to store the permutation representation matrices.

7 A brief aside — Continuous Time FCIQMC

7.1 Motivation

The performance of FCIQMC is strongly dependent on the magnitude of the (imaginary) time step, $\delta\tau$. If this is too small, significant computational resources are wasted to progress the simulation sufficiently far. Conversely, if it is too large the imaginary time integration in eqn. 2.2 is not well approximated by 2.4, and the simulation never accurately settles on the target wavefunction. This results in an incorrect energy being reported. The restrictions on the value of $\delta\tau$ are described in section 2.6.

In this section we develop a variant of FCIQMC, known as Continuous Time FCIQMC, or *ct*-FCIQMC, which removes the direct dependence of the dynamics of the simulation on the time step. This new methodology asks the question “when is the next particle to be spawned” rather than the more usual question of “how many particles should I spawn this iteration”. The consequence of this is to concentrate the computational effort on those sites which spawn the most particles — the equivalent of making the time step a property which varies depending on the site in the canonical scheme.

7.2 Derivation

As suggested by Spencer¹²⁹, instead of calculating how many particles are spawned from a given particle within a time period, the length of time until a given particle will spawn a child may be calculated.

Starting from the imaginary–time Schrödinger equation (eqn. 2.1) adjusted for the energy of the reference determinant and the shift,

$$\frac{\partial \Psi}{\partial \tau} = (-\hat{H} + E_{\text{ref}} + S)\Psi,$$

the coefficient of each determinant D_i is projected out to give the diffusion equation

$$\begin{aligned} \frac{dc_i}{d\tau} &= -\sum_j [H_{ij} - (E_{\text{ref}} + S)S_{ij}] c_j \\ &= -\sum_j \Xi_{ij} R_{ij} c_j, \end{aligned}$$

where Ξ_{ij} and R_{ij} are the sign and magnitude of $H_{ij} - (E_{\text{ref}} + S)S_{ij}$ respectively. R_{ij} gives the rate at which a particle on determinant \mathbf{j} spawns onto site \mathbf{i} . Considering a time interval $\Delta\tau$, the probability, p_{ij} , of the first child of site \mathbf{j} being spawned onto site \mathbf{i} within the n th period, $(n-1)\Delta\tau < \tau = n\Delta\tau$, is

$$p_{ij} = \underbrace{\left(\prod_{\mathbf{k}} (1 - R_{\mathbf{kj}} \Delta\tau) \right)^{n-1}}_{\text{prob. no children spawned in } n-1 \text{ time periods}} \times \underbrace{\prod_{\mathbf{k} \neq \mathbf{i}} (1 - R_{\mathbf{kj}} \Delta\tau)}_{\text{prob. no children spawned on determinants } \mathbf{k} \neq \mathbf{i} \text{ in } n\text{-th period}} \times \underbrace{R_{ij} \Delta\tau}_{\text{prob. child spawned on determinant } \mathbf{i} \text{ in } n\text{-th period}} \quad (7.1)$$

$$= \prod_{\mathbf{k}} (1 - R_{\mathbf{kj}} \Delta\tau)^n \left(\frac{R_{ij} \Delta\tau}{1 - R_{ij} \Delta\tau} \right) \quad (7.2)$$

$$\approx \prod_{\mathbf{k}} \left(1 - \frac{R_{\mathbf{kj}} \tau}{n} \right)^n R_{ij} \Delta\tau \quad \text{For small } \Delta\tau. \quad (7.3)$$

Allowing $\Delta\tau \rightarrow 0$, and correspondingly $n \rightarrow \infty$, this becomes

$$p_{ij} = \prod_{\mathbf{k}} \exp(-R_{\mathbf{kj}}) R_{ij} \Delta\tau \quad (7.4)$$

$$= \exp\left(-\sum_{\mathbf{k}} R_{\mathbf{kj}} \tau\right) R_{ij} \Delta\tau. \quad (7.5)$$

Thus

$$p_{ij}(\text{1st spawn to determinant } \mathbf{i} \text{ is between } \tau \text{ and } \tau + d\tau) = e^{-R_j \tau} R_{ij} d\tau \quad (7.6)$$

where $R_j = \sum_{\mathbf{k}} R_{\mathbf{kj}}$. Summing this over all target determinants \mathbf{i} gives

$$p_j(\text{1st spawn is between } \tau \text{ and } \tau + d\tau) = e^{-R_j \tau} R_j d\tau, \quad (7.7)$$

which may be used as a probability distribution function for when the next spawn occurs. The cumulative distribution function is

$$\begin{aligned} P_{\mathbf{j}}(\tau) &= \int_0^{\tau} p_{\mathbf{j}}(\tau') d\tau' \\ &= 1 - e^{-R_{\mathbf{j}}\tau}, \end{aligned} \quad (7.8)$$

which can be seen to be normalised. This is easily inverted so that the time until the next spawning from determinant \mathbf{j} can be selected as

$$\tau = -\frac{1}{R_{\mathbf{j}}} \ln(u) \quad (7.9)$$

where u is a random number selected from the uniform distribution on $[0, 1)$.

7.3 Implementing continuous time FCIQMC

7.3.1 Performing annihilation

As has been demonstrated by Spencer et al.⁶⁴, annihilation is necessary for convergence of the wavefunction. If each particle is tracked through imaginary time individually, along with its progeny, annihilation between particles that are simultaneously present in imaginary time must be performed explicitly.

To resolve this, periodic “annihilation barriers” are positioned throughout imaginary time, separated by a period $\delta\tau$, and the simulation paused at each of these. Each time the next spawn from a particle is calculated to occur after the next annihilation barrier, no spawn is performed, and the particle’s time is set to be at the barrier. This does not affect the statistics of spawning as the probability of a particle at time τ spawning its next particle at time τ' , after the barrier located at time τ_a , is

$$\begin{aligned} p(\tau') &= R_{\mathbf{j}} e^{-R_{\mathbf{j}}(\tau' - \tau)} d\tau \\ &= R_{\mathbf{j}} e^{-R_{\mathbf{j}}(\tau_a - \tau)} e^{-R_{\mathbf{j}}(\tau' - \tau_a)} d\tau \\ &= p(\tau' > \tau_a) \times p(\tau' | \tau_a), \end{aligned}$$

(see equation 7.8) such that the likelihood of spawning at any specific time is

unaffected.

Synchronising the annihilation of all particles at given times, facilitates the implementation of this algorithm across multiple MPI processes. All newly generated particles may be communicated to the correct MPI process and annihilated with particles found there. The calculation is then resumed as normal. Whilst the simulation is extremely insensitive to the period of imaginary time between each of the annihilation barriers, the number of new particles spawned is exponential with passing imaginary time, so a substantial amount of memory is required to store these particles if the time interval is too large.

7.3.2 Generating new particles

The derivation above requires knowledge of an overall spawning rate from any given site (including particle death), and the ability to generate spawns which are distributed according to the distribution of connection strengths between particles. Both of these imply enumeration of the connections from each of the sites in the system when they are considered:

Overall spawning rate

To obtain the overall spawning rate, *all* connections must be enumerated, and the Hamiltonian matrix elements between the source and all connected sites must be generated and their absolute values summed. The sum of the absolute values of the off-diagonal matrix elements, and separately the diagonal matrix element, can be stored with the particles such that the total spawning rate for a particle can be given as

$$R_i = |K_{ii} - E_S| + R_{i,\text{off-diag}}.$$

If non-orthogonal and non-orthonormal basis functions are used, this expression is further complicated as the weighted energy estimator enters the off-diagonal terms, and

$$R_i = \left| \frac{K_{ii} - E_S S_{ii}}{S_{ii}} \right| + \sum_{\substack{j \leftarrow i \\ j \neq i}} \left| \frac{K_{ij} - E_{\text{weighted}} S_{ij}}{S_{ii}} \right|$$

must be frequently regenerated.

Spawning site selection

To generate the correct distribution of new particles, given a correct attempted spawning rate, a random number on $[0, 1)$ should be chosen, and then the connected sites enumerated one by one, summing in $\frac{R_{ij}}{R_i}$ until this sum exceeds the random number chosen. The site then being considered is spawned to. This enumeration may be optimised slightly by considering the diagonal (death-like) spawn first, as the diagonal matrix element is likely to be the largest and this reduces the likelihood of needing to enumerate lots of connections.

As a consequence of this, as noted by Spencer, the algorithm is unsuitable for use in any systems other than those model systems such as the Hubbard model where the summations may be performed analytically. It is certainly highly impractical for molecular systems where some of the matrix element calculations are relatively costly, and efficient enumeration is complicated by considerations of symmetry and so forth.

Fortunately a subtly different approach may be taken. One advantage of the continuous time scheme is that all spawns are accepted, whereas in the normal FCIQMC scheme the vast majority of spawns are rejected. The distribution of particles may be generated computationally more efficiently by loosening this advantage; if a parameterised (and adjustable) guess of the correct total spawning rate is made, such that we oversample the required spawning rate, then random excitation generators may be used to generate connections and the correct distribution generated by discarding excess spawns (see figure 7.1).

This scheme rewards the use of maximally uniform excitation generators, as the rate of oversampling required is determined by the combination of lowest non-uniform generation probability combined with the highest related Hamiltonian matrix element, so improving the worst case (very low generation probabilities) has disproportionate benefit. The more parameterised the guess for the oversampling rate is, the more efficiently the scheme works as the discard rate is very strongly linked to it. For the testing in this thesis the following scheme was used:

- As the diagonal matrix elements are stored anyway, these and the shift are included directly.
- If same spatial structure spawns are permitted, i.e. non-orthogonal basis functions are being used, the sampling rate for same spatial structure exci-

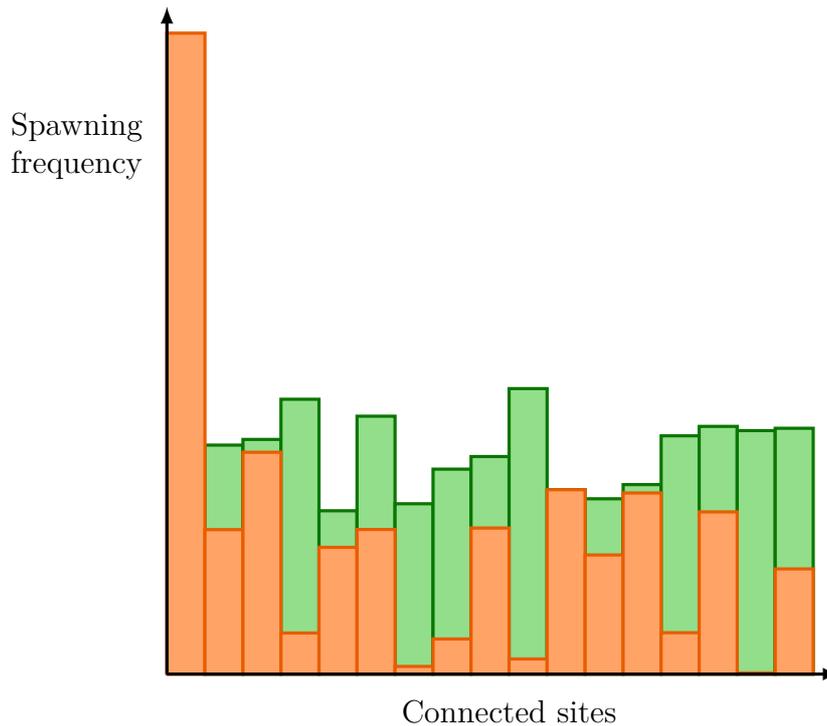


Figure 7.1: Oversampling of spawning rates in *ct*-FCIQMC. The required distribution of spawned particles is coloured orange, with the stochastically generated one in green. The largest spawning rate is from a site to itself — this is equivalent to particle death in normal FCIQMC, and can be sampled exactly. For the remainder of the connections, the (non-uniform) random excitation generator is called sufficiently frequently to generate each of the connections at least as frequently as required, and the remaining excess spawns in green are discarded. So long as the excitation generators are sufficiently uniform, the discard rate is less wasteful than the acceptance rates in normal FCIQMC.

tations is calculated as a multiple of the diagonal Hamiltonian matrix element with a factor parameterised by the number of unpaired electrons present and the excitation level relative to the reference site.

- Single and double excitations are considered separately, as the excitation generators are parameterised by the relative rate of single and double excitations. These are categorised by the number of unpaired electrons present, and the excitation level relative to the reference determinant.

For each excitation generated the generation probability, $p_{\text{gen}}(\mathbf{j}|\mathbf{i})$, returned by the excitation generator can be multiplied by the sampling rate to generate the frequency $f(\mathbf{j}|\mathbf{i})$, which can be compared to the required frequency $R_{\mathbf{i}\mathbf{j}}$. The particle

is then spawned with probability $\frac{R_{ij}}{f(\mathbf{j}|\mathbf{i})}$ or discarded otherwise. If $f(\mathbf{j}|\mathbf{i}) < R_{ij}$ then the relevant parameter (indexed by the number of unpaired electrons and the excitation level of the source of the excitation relative to the reference determinant) must be updated to ensure that on future iterations $f(\mathbf{j}|\mathbf{i}) = R_{ij}$.

7.3.3 Algorithm overview

This algorithm is implemented in a recursive way, such that each particle generates particles which are processed similarly. This tree of generated particles is interrupted when an annihilation barrier is reached. Most elegantly, if a particle dies before it reaches the annihilation barrier, the routines back out of the recursive structure — if this is a newly spawned particle then no further reference to it is required, and if it was in the main lists, its coefficient can be decremented.

For each of the particles stored in the occupied list, on site \mathbf{i} , with associated time τ and attempted spawning rate $R_{\mathbf{i}}$, the following process is performed;

1. The time until the next spawning attempt from this particle is calculated according to equation 7.9 such that

$$\tau \rightarrow \tau - \frac{1}{R_{\mathbf{i}}} \ln(u),$$

where u is a random number on $[0, 1)$.

2. If the calculated time, τ_{spawn} , is after the next annihilation barrier, the time, τ , associated with the source particle is advanced to the annihilation barrier. If this particle is a newly spawned one, it is stored in the spawned list. The processing loop now ends.
3. A connected site, \mathbf{j} , is selected with the probability $p_{\text{gen}}(\mathbf{j}|\mathbf{i}) = \frac{R_{ij}}{R_{\mathbf{i}}}$. See section 7.3.2.
 - If a null site is returned, nothing is done.
 - If the same site is returned, by analogy with diagonal death, the particle is destroyed, and the processing loop ends.
 - If a new particle is obtained, its attempted spawning rate $R_{\mathbf{j}}$ is calculated if not using oversampling, and the new particle is processed in the same way as its parent.

4. If the particle has not been destroyed, its associated time, τ , is advanced to the newly calculated value of τ_{spawn} , and these steps are repeated by returning to step 1.

Once this process is complete, annihilation between the newly spawned list and the remaining particles in the main particle list is performed as normal. It is worth noting that due to the tree of spawned particles and their further progeny that is generated prior to annihilation taking place, the amount of memory required to store particles is exponentially dependent on the time period between annihilation barriers. In real systems, this memory requirement is likely to provide the upper limit on the choice of inter-annihilation barrier time period, rather than anything fundamental about the algorithm dynamics.

7.3.4 The initiator approximation

Implementing the initiator approximation is trickier for *ct*-FCIQMC than for normal FCIQMC. The two primary criteria, as described in section 2.4.4, are difficult to implement, as particle survival depends on:

Status of source site

The outcome depends on if the source site is an initiator, which in turn depends on the number of particles present on the site. During normal FCIQMC these statuses are determined after annihilation, where all particles on a given site are located on the same MPI process and the main particle list is compressed to one entry per site. As *ct*-FCIQMC recursively explores the tree of spawned particles and their further progeny, it is not possible to know the status of the sites associated with second or further generation spawning.

Status of target site

If spawns occur from non-initiator sites, then the survival of their progeny is dependent on the occupation status of the target site. As *ct*-FCIQMC propagates the spawning recursively prior to annihilation taking place, it is not possible to perform these tests before second or further generation spawning.

This conundrum can be resolved by a couple of changes to the implementation of the initiator approximation:

- All particles are assumed to survive until the next annihilation barrier. The spawn aborting behaviour is considered to be a function of annihilation, rather than of spawning.
- All second and further generation spawns are considered to come from non-initiator source sites.

These modifications have the advantage that the only time information about the overall occupation of sites is required is during annihilation, when this information is available. Similarly to normal FCIQMC, in the large particle limit all target sites with non-zero coefficients are occupied, and the approximation tends to the exact value.

In principle, due to the differing consideration of second and further generation spawns, this approximation becomes less good as the gap between the annihilation barriers increases — this is similar to the general behaviour of *ct*-FCIQMC in decreasing the efficiency of annihilation. In practice, the simulation still appears to be remarkably insensitive to gaps between annihilation barriers substantially larger than the values of $\delta\tau$ required for normal FCIQMC.

7.4 Acceptance ratios

The predominance of rejected particles in FCIQMC provided the primary impetus behind the development of *ct*-FCIQMC. In both FCIQMC and *ct*-FCIQMC, particle spawns are accepted with a certain probability, and rejected otherwise. However, in the two different cases, the meaning of this value is different.

FCIQMC

Particles are accepted or rejected to round the value of n_s , the magnitude of each spawn, up to the minimum acceptable value or down to zero as appropriate. As such it is roughly a stochastic estimator of the average value of $\frac{n_s}{c_{\min}}$.

As $\delta\tau$ is an adjustable parameter, this metric becomes a measure of how large the maximal spawn size is relative to the average spawn size, and is a measure of how extreme the uppermost outlier is.

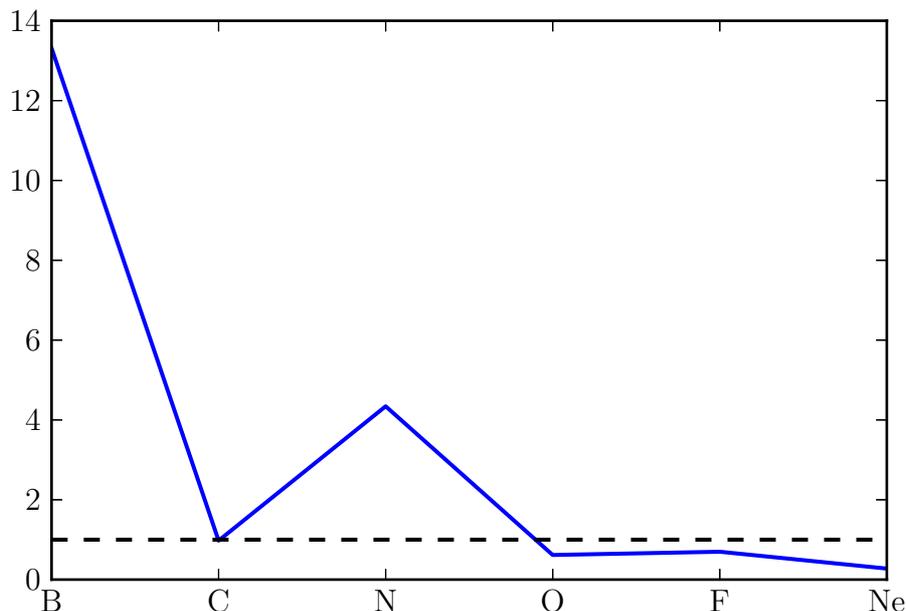


Figure 7.2: A plot of the ratio of the spawn acceptance ratio for *ct*-FCIQMC divided by that for FCIQMC for calculations of the ground state of the first row atoms in a cc-pVDZ basis set. The dashed line is positioned at 1.0, where the acceptance ratio is the same for both methods.

ct-FCIQMC

All particles spawned according to the distribution $p_{\text{gen}}(\mathbf{j}|\mathbf{i}) = \frac{R_{\mathbf{j}\mathbf{i}}}{R_{\mathbf{i}}}$ are accepted, barring abortion via the initiator approximation. The acceptance ratio in this context is the proportion of spawns generated by the random excitation generator that are kept to generate this required distribution. This is a metric of the mismatch between two distributions, both of which are fairly non-uniform, and as such is substantially more dependent on the number of electrons and the basis set size than the acceptance ratio in canonical FCIQMC.

It is notable from figure 7.2 that, as the system size is increased, the acceptance ratios in *ct*-FCIQMC deteriorate relative to those in canonical FCIQMC. This implies that the rate at which the mismatch between the generated distribution of spawns and the required distribution of spawns deteriorates is worse than the behaviour of the number of rejected spawns caused by non-uniformity in the matrix elements in FCIQMC. A consequence of this is that *ct*-FCIQMC radically improves the efficiency of small calculations, that can already be easily performed, but

rapidly makes larger (and more interesting) calculations more difficult than they already were.

In principle the acceptance ratio, and hence the computational cost of *ct*-FCIQMC simulations, is systematically improvable. The more accurately the generated distribution of particles matches the required distribution, the more efficient the simulation. The estimation of the oversampling factors and the distribution of generation probabilities given those factors are in principle arbitrarily parameterisable. In practice this is tricky, but it could be interesting to investigate what parameters could be used.

7.5 Comparison of FCIQMC and *ct*-FCIQMC

The primary consideration of any methodological adaptation is whether it works. In this case, *ct*-FCIQMC works well, producing accurate FCI energies in much the same way as canonical FCIQMC. This shifts the primary interest to whether the adaptation works better than the original method. There are a number of different components to be considered in this case. The dynamics of the simulations should be compared, as well as the behaviour of the output variables. Ultimately, the computational cost of the new method is of interest.

General calculation profile

The overall calculation profile of a *ct*-FCIQMC calculation is extremely similar to that of a canonical FCIQMC calculation. Figure 7.3 plots the total number of particles and the number of particles on the reference site for an FCIQMC simulation, with the largest timestep permissible to avoid particle blooms, and a *ct*-FCIQMC simulation with $\delta\tau = 0.01$, nearly ten times larger.

What is evident is that, apart from the initial phase of the calculation when the oversampling factors in *ct*-FCIQMC are rapidly changing as they are discovered empirically, the dynamics are essentially the same. This demonstrates that these dynamics, with rapid initial growth and wavefunction convergence during an annihilation plateau, are a property of integrating the imaginary time Schrödinger equation in a discrete antisymmetrised space, rather than being an artefact of the canonical FCIQMC algorithm.

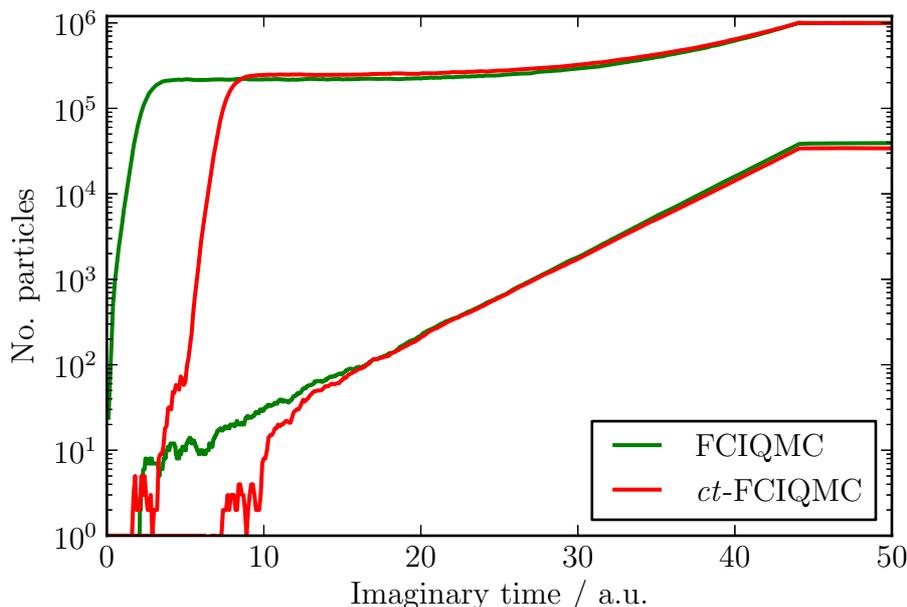


Figure 7.3: The total number of particles and the number of particles on the reference site in both a canonical FCIQMC and a ct -FCIQMC simulation of Ne in an aug-cc-pVDZ basis set. For canonical FCIQMC, $\delta\tau = 0.00217$, which is the largest value possible to avoid particle blooms. For ct -FCIQMC, $\delta\tau = 0.01$.

When the initiator method is used, from the perspective of overall calculation profile, i - ct -FCIQMC and i -FCIQMC are also identical.

Height and length of the annihilation plateau

The ct -FCIQMC algorithm works by spawning a ‘tree’ of particles, sign incoherently, for a period of imaginary time, and periodically generating sign coherence through annihilation. As the time periods are increased, an increasing distribution of incoherent particles presents itself for annihilation on each occasion, and the effectiveness of annihilation is reduced. This has a tendency to increase the height of the plateau, as demonstrated in figure 7.4.

Interestingly, the length of the plateau in imaginary time does not display monotonic behaviour, instead displaying a distinct optimum which is shorter than in canonical FCIQMC. Consequently, at this optimum, ct -FCIQMC both requires less imaginary time to escape the plateau, and fewer iterations per unit imaginary time, than FCIQMC.

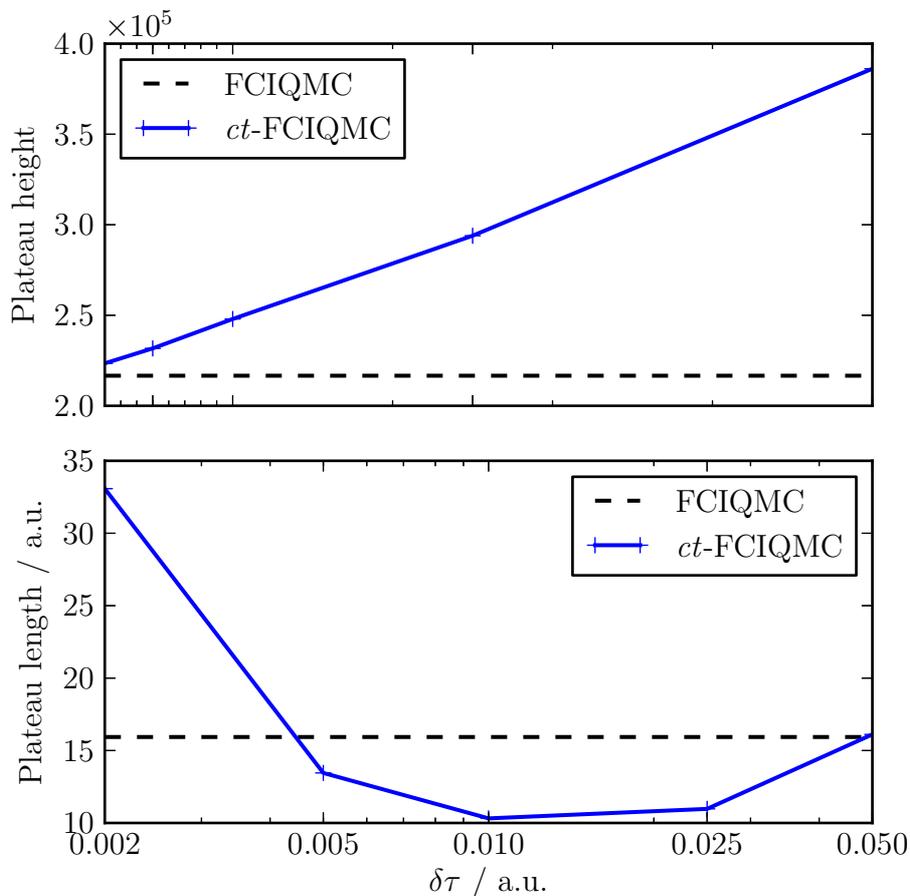


Figure 7.4: Dependence of the plateau height and length on $\delta\tau$ in *ct*-FCIQMC compared to FCIQMC for Ne in an aug-cc-pVDZ basis set. For the FCIQMC simulation, $\delta\tau = 0.00217$, which is the largest value possible to avoid particle blooms.

Does the initiator method still make sense?

The initiator method appears to work as effectively in *i-ct*-FCIQMC as in *i*-FCIQMC in terms of suppression of the annihilation plateau.

Computational cost

The computational cost of the simulation may be measured in three strongly interlinked ways; *a*) the cost per iteration, *b*) the cost per unit of imaginary time, and *c*) the cost for a given statistical accuracy. The cost per iteration is certainly going to increase as the gap between the annihilation barriers is increased. In principle, an exponential growth of progeny from each particle could be observed, leading to an exponential increase in computational cost

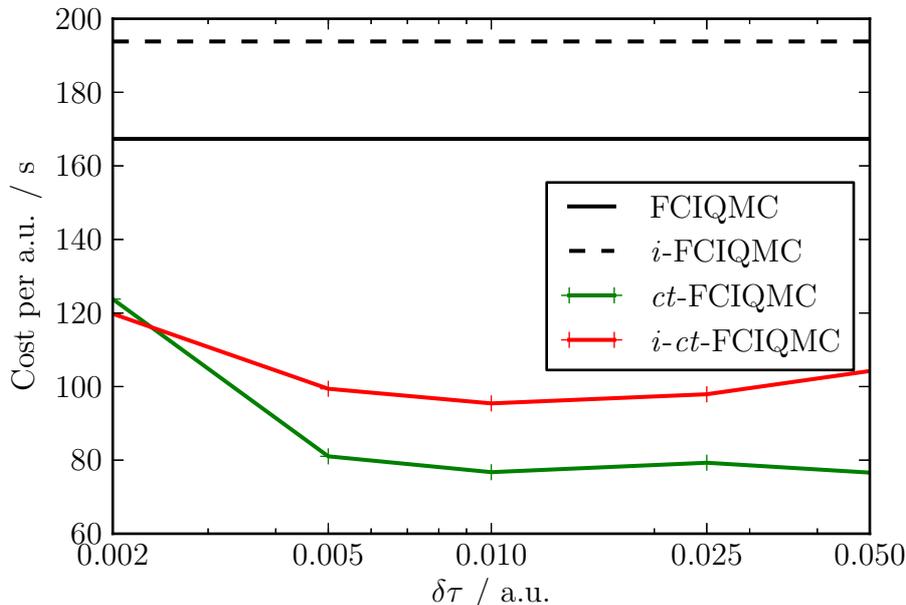


Figure 7.5: The computation cost per unit imaginary time of *ct*-FCIQMC in both its full and initiator approximation implementations compared to canonical FCIQMC for differing values of the annihilation barrier separation, $\delta\tau$. The system is Ne in an aug-cc-pVDZ basis set with the particle count constrained to one million particles by the shift. For canonical FCIQMC, the largest $\delta\tau$ to avoid particles blooms is 0.00217. The initiator approximation is sensitive to the initial increase in $\delta\tau$ as the aborted particles are only removed at the annihilation barriers giving an efficiency boost to simulations which do this more often. This is also responsible for the increase in cost of the initiator approximation past an optimum, which is less observable in the full scheme.

with $\delta\tau$.

In practice, at low values of $\delta\tau$, the majority of sites do not spawn progeny before the next annihilation barrier. As a consequence, increasing $\delta\tau$ does not necessarily result in any additional work, and may even reduce the number of calls to the excitation generator for this site. As such, the cost per iteration initially grows sub-linearly, with exponential growth only picking up at higher values of $\delta\tau$.

This is observed in the behaviour of the computational cost per unit imaginary time, as shown in figure 7.5, which initially falls with $\delta\tau$, before levelling out and eventually turning up again (this latter behaviour is only observed in the initiator approximation plots).

As $\delta\tau$ increases, the reduction in the number of annihilation barriers increases the correlation between different iterations, and also decreases the number of data points output. As a result of this, a longer period of imaginary time is required to get statistically useful results (an effect clearly seen in figure 7.6a, where the stochastic error associated with the last two data points increases rapidly). The consequence of this is that, although the cost per unit imaginary time is decreasing, there is an optimum value for $\delta\tau$ which obtains useful results as rapidly as possible. For Ne in an aug-cc-pVDZ basis, this appears to be approximately ten times larger than for canonical FCIQMC.

An interesting point worth noting from figure 7.5 is that the cost per unit imaginary time is substantially higher when using the initiator approximation than otherwise. This is unexpected, as the only algorithmic change is to throw away some more of the attempted spawns. This may be a consequence of the wavefunction obtained in the initiator approximation — by restricting spawning outside of the well converged region, to obtain the same number of particles in the simulation as a simulation without the initiator method means that there are more particles in the ‘core’ region of the wavefunction. This results in a different subset of the oversampling factors being dominant in the overall cost and a different spawning pattern.

Accuracy of the energy estimators

The accuracy of a simulation’s numerical results depends on a number of factors. The first of these factors is wavefunction convergence. For all of the results considered here, the shift has been permitted to vary once the simulation contains 1 million particles — this is above the annihilation plateau in all cases, and is sufficient to ensure that the wavefunction is converged when the initiator approximation has been used.

Once there are enough particles in the simulation, the next criterion is whether enough data has been collected to produce results with a sufficiently high statistical accuracy. This is influenced by how long it takes to reach a converged wavefunction, i.e. the height and length of the plateau, the cost per iteration, the cost per unit time, how large the statistical fluctuations are and the degree of correlation between different data points.

In an attempt to gain some insight into the combination of all of these factors, the results in figure 7.6 are obtained from simulations that were all run for

the same amount of real computational time; five hours across four MPI processes.

For *ct*-FCIQMC, in figure 7.6a, the stochastic error in the values is roughly the same as for canonical FCIQMC (marked by the green band) up to $\delta\tau = 0.01$, and all of these values agree with the FCIQMC energy estimator and the FCI energy to within error bars. Once $\delta\tau$ increases above this, the errors increase extremely rapidly, as discussed above.

For *i-ct*-FCIQMC, it is notable in figure 7.6b that at small values of $\delta\tau$, the stochastic error is smaller than for *i*-FCIQMC, with the values still in agreement with the FCI energy. The errors rapidly become large past the same limit as *ct*-FCIQMC.

Finally, note that as with canonical FCIQMC simulations, the shift energy estimator is much less reliable than the projected energy estimator.

In conclusion, *ct*-FCIQMC and *i-ct*-FCIQMC simulations are similar to their canonical parents, with the exception of being extremely insensitive to the value of $\delta\tau$ chosen as the separation between the annihilation barriers. There is an optimum value both in terms of efficiency of propagating the simulation through imaginary time, and in terms of the cost of convergence of the stochastic errors, but there is substantial leeway for determining this value with mild computational cost and small stochastic error penalties.

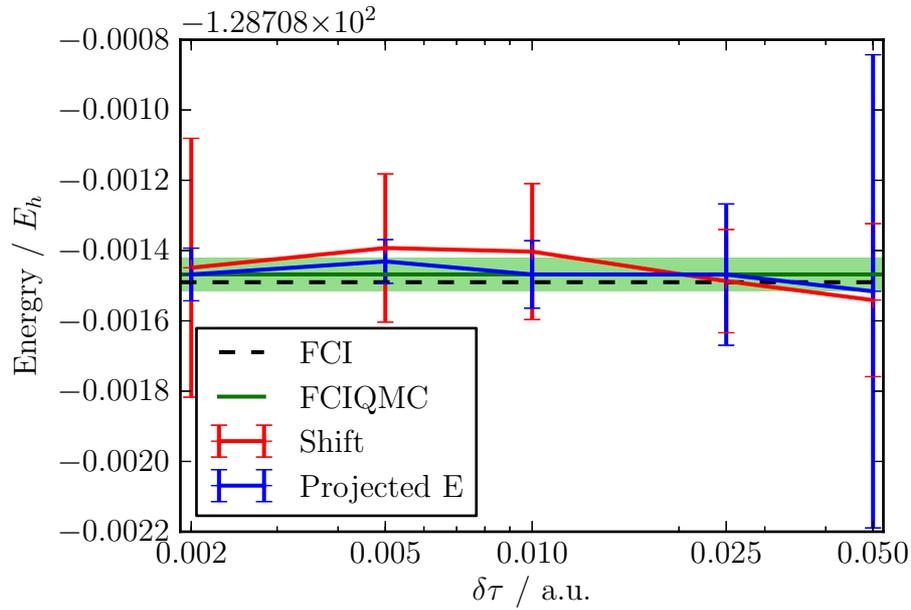
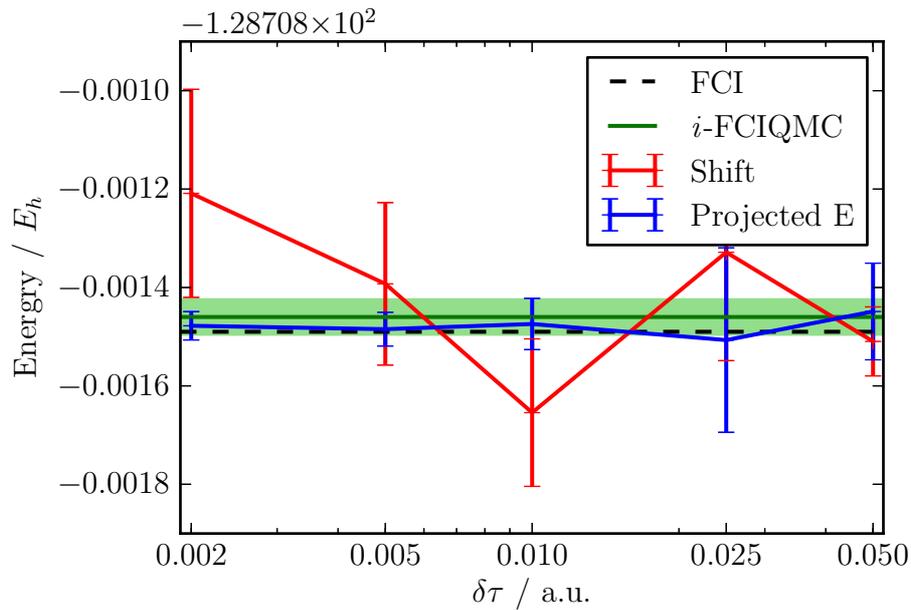
(a) Canonical FCIQMC and ct -FCIQMC(b) i -FCIQMC and i - ct -FCIQMC

Figure 7.6: The projected and shift energy estimators, with error bars, for ct -FCIQMC with differing values of the annihilation barrier separation, $\delta\tau$. The system is Ne in an aug-cc-pVDZ basis set with the particle count constrained to one million particles by the shift — this is above the annihilation plateau. For canonical FCIQMC, the largest $\delta\tau$ to avoid particle blooms is 0.00217. Except at the largest value tested, the simulation is fairly insensitive to the value of $\delta\tau$. The values obtained for the shift are better than those obtained for canonical FCIQMC.

7.6 Summary

This section has presented a novel implementation of *ct*-FCIQMC which avoids the need for enumeration of all connections to a site at any point in the simulation, and an application of the initiator approximation to this method. It appears that this method offers the opportunity to relax the dependence of FCIQMC on the time step, $\delta\tau$, and avoid some of the pitfalls associated with large variations in connection strength across the Hilbert space.

However, as a consequence of the oversampling required to generate the correct distribution of spawns, the acceptance ratio falls rapidly as the system size is increased, aggressively cutting the efficiency. As a consequence, *ct*-FCIQMC improves the efficiency of those calculations that were already straightforward to perform. Unfortunately, in its current configuration, calculations that are challenging, and therefore interesting, become more difficult.

As a consequence, this method does not yet present itself for effective general application. Hopefully it will provide insight and a source of ideas for further algorithmic development.

8 Concluding remarks

This thesis has presented methods for the application of Hilbert spaces constructed from spin eigenfunctions to FCIQMC, including the requisite algorithmic changes. This permits control over convergence to states other than the natural ground state in the basis set. Of the basis sets considered, Serber type spin eigenfunctions are clearly the most applicable to calculations due to the extremely efficient expressions for the Hamiltonian matrix elements and their suitability for spawning to multiple related spin structures simultaneously.

Unfortunately, as a method, this is currently limited to an absolute maximum of sixteen electrons. In practice, on most available computational resources, memory limits will be reached well before this*. It is clear that further implementational work is required to share this effectively read-only data across multiple processes, but this will not have a substantial qualitative impact on the sizes of systems that can be considered. At present the full N -electron permutations are required, but in principle this can be reduced to $4 + N_{o,max}$ -electron permutations for truncated calculations. It is an open question as to how much of the space need be represented in CSFs to be useful.

If the convergence issues associated with spin-projected determinants can be overcome, they may provide the route to overcoming the memory induced upper bound on the system size. The first step to understanding this probably involves implementation of the necessary algorithmic components to use Rumer type CSFs, in order to determine the extent to which the pathological behaviour is a consequence of the non-orthogonality, and to what extent it is a property of the spin-projected determinants themselves.

It is worth noting that, as would be expected, using CSFs does not give a more

*Operating on commonly available compute clusters, with 1024Mb of memory available per MPI process, required only using a subset of the processors available to increase this memory availability for the N_2 calculations, in a cc-pVDZ basis set with core electrons frozen and ten active valence electrons, used in this thesis.

general solution to finding excited states. It is clear that an approach based on subdividing the space according to symmetries can only ever reveal a very small number of additional states. These functions also do not assist in representing the short range electronic behaviour — the electronic cusps — which are as poorly represented by linear sums of diffuse spatial functions when using CSFs as otherwise. To some extent the inter-electronic behaviour will need to be considered directly. Work is ongoing to address these issues using new methodologies.

There are a number of useful model systems, in particular the Uniform Electron Gas and the Hubbard model, which are parameterisable to arbitrary size and computational difficulty. These models are particularly dependent on capturing the static correlation between a vast number of energetically similar basis functions and states. The extent to which imposed spin structure might assist in generating useful solutions is unclear, and needs investigation. This will involve a certain amount of further implementational work, as the internal connectivity of the relevant Hilbert spaces is notably different. This difference might lead to more productive behaviour in conjunction with CSFs than has been observed in molecular systems.

Returning to the question asked in the introduction; *does restricting the basis set to prevent spin-incoherent growth assist in the convergence of the wavefunction in the remaining space, or is spin primarily useful as a metric of convergence?* It is clear that using CSFs does not, for molecular systems, substantially improve the convergence of the wavefunction except in the specific, and rare, cases when it permits convergence on states that are otherwise inaccessible. At least in molecular systems, spin is far more useful as a metric of convergence than as a generator of convergence.

Bibliography

1. P. A. M. Dirac. *Quantum Mechanics of Many-Electron Systems*. Proceedings of the Royal Society of London Series A **123**, 714 (1929).
2. D. R. Hartree. *The Wave Mechanics of an Atom with a Non-Coulomb Central Field. Part I. Theory and Methods*. Mathematical Proceedings of the Cambridge Philosophical Society **24**, 89 (1928).
3. V. Fock. *Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems*. Zeitschrift für Physik **61**, 126 (1930).
4. A. Szabo & N. S. Ostlund. *Modern Quantum Chemistry. Introduction to Advanced Electronic Structure Theory* 111–230. ISBN: 0-486-69186-1 (Dover Publications, 1996).
5. P. Hohenberg & W. Kohn. *Inhomogeneous Electron Gas*. Physical Review **136**, B864 (1964).
6. S. Lundqvist & N. H. March. *Theory of the Inhomogeneous Electron Gas* (Plenum Press, 1983).
7. C. Møller & M. S. Plesset. *Note on an Approximation Treatment for Many-Electron Systems*. Physical Review **46**, 618 (1934).
8. M. Head-Gordon, J. A. Pople & M. J. Frisch. *MP2 energy evaluation by direct methods*. Chemical Physics Letters **153**, 503 (1988).
9. R. Krishnan & J. A. Pople. *Approximate fourth-order perturbation theory of the electron correlation energy*. International Journal of Quantum Chemistry **14**, 91 (1978).
10. W. Duch & G. H. F. Diercksen. *Size-extensivity corrections in configuration interaction methods*. The Journal of Chemical Physics **101**, 3018 (1994).
11. J. A. Pople, M. Head-Gordon & K. Raghavachari. *Quadratic configuration interaction. A general technique for determining electron correlation energies*. The Journal of Chemical Physics **87**, 5968 (1987).

12. K. Raghavachari, G. W. Trucks, J. A. Pople & M. Head-Gordon. *A fifth-order perturbation comparison of electron correlation theories*. Chemical Physics Letters **157**, 479 (1989).
13. P. J. Knowles & N. C. Handy. *A new determinant-based full configuration interaction method*. Chemical Physics Letters **111**, 315 (1984).
14. C. D. Sherrill & H. F. S. III. *The Configuration Interaction Method: Advances in Highly Correlated Approaches*. Advances in Quantum Chemistry **34**, 143 (1999).
15. G. Booth, A. Thom & A. Alavi. *Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space*. The Journal of Chemical Physics **131**, 054106 (2009).
16. G. H. Booth, D. Cleland, A. J. W. Thom & A. Alavi. *Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum Monte Carlo methods*. The Journal of Chemical Physics **135**, 084104 (2011).
17. T. D. Crawford & H. F. Schaefer. in *Reviews in Computational Chemistry* 33–136 (Kohn Wiley & Sons, Inc.). ISBN: 9780470125915. doi:10.1002/9780470125915.ch2.
18. E. Schrödinger. *An Undulatory Theory of the Mechanics of Atoms and Molecules*. Physical Review, 1049 (1926).
19. P. Atkins & R. Friedman. *Molecular Quantum Mechanics* 4th ed., 23. ISBN: 0-19-927498-3 (Oxford University Press, 2005).
20. P. Dirac. *The Quantum Theory of the Electron*. Proceedings of the Royal Society of London Series A **117**, 610 (1928).
21. M. Reiher. *Relativistic Douglas-Kroll-Hess theory*. Wile Interdisciplinary Reviews: Computational Molecular Science **2**, 139 (2011).
22. M. Born & R. Oppenheimer. *Zur Quantentheorie der Molekeln*. Annalen der Physik **389**, 457–484 (1927).
23. P. Atkins & R. Friedman. *Molecular Quantum Mechanics* 4th ed., 226–228. ISBN: 0-19-927498-3 (Oxford University Press, 2005).
24. W. Pauli. *Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren*. Zeitschrift für Physik **31**, 765 (1925).
25. J. C. Slater. *Note on Hartree's Method*. Physical Review **35**, 210 (1930).

26. T. H. D. Jr. & P. J. Hay. in *Methods of Electronic Structure Theory* 1–27 (Springer US, 1977). ISBN: 978-1-4757-0889-9. doi:10.1007/978-1-4757-0887-5_1.
27. C. C. J. Roothaan. *New Developments in Molecular Orbital Theory*. Reviews of Modern Physics **23**, 69 (1951).
28. J. A. Pople & R. K. Nesbet. *Self-Consistent Orbitals for Radicals*. The Journal of Chemical Physics **22**, 571 (1954).
29. W. Kutzelnigg & W. Klopper. *Wave functions with terms linear in the inter-electronic coordinates to take care of the correlation cusp. I. General theory*. The Journal of Chemical Physics **94**, 1985 (1991).
30. E. Schrödinger. *Quantisierung als Eigenwertproblem*. Annalen der Physik **385**, 437 (1926).
31. R. Krishnan, J. A. Pople, E. S. Replogle & M. Head-Gordon. *Fifth order Moeller-Plesset perturbation theory: comparison of existing correlation methods and implementation of new methods correct to fifth order*. The Journal of Physical Chemistry **94**, 5579 (1990).
32. J. A. Pople, J. S. Binkley & R. Seeger. *Theoretical models incorporating electron correlation*. International Journal of Quantum Chemistry **10**, 1 (1976).
33. J. A. Pople, R. Seeger & R. Krishnan. *Variational configuration interaction methods and comparison with perturbation theory*. International Journal of Quantum Chemistry **12**, 149 (1977).
34. J. Goldstone. *Derivation of the Brueckner Many-Body Theory*. Proceedings of the Royal Society of London Series A **239**, 267 (1957).
35. C. Schwartz. *Importance of Angular Correlations between Atomic Electrons*. Physical Review **126**, 1015 (1962).
36. W. Klopper, F. R. Manby, S. Ten-No & E. F. Valeev. *R12 methods in explicitly correlated molecular electronic structure theory*. International Reviews in Physical Chemistry **25**, 427 (2006).
37. L. Bytautas & K. Ruedenberg. *Correlation energy extrapolation by intrinsic scaling. IV. Accurate binding energies of the homonuclear diatomic molecules carbon, nitrogen, oxygen and fluorine*. The Journal of Chemical Physics **122**, 154110 (2005).
38. G. H. Booth, D. Cleland, A. Alavi & D. P. Tew. *An explicitly correlated approach to basis set incompleteness in full configuration interaction quantum Monte Carlo*. Journal of Chemical Physics **137**, 164112 (2012).

39. S. R. Langhoff & E. R. Davidson. *Configuration interaction calculations on the nitrogen molecule*. International Journal of Quantum Chemistry **8**, 61 (1974).
40. D. Prendergast, M. Nolan, C. Filippi, S. Fahy & J. C. Greer. *Impact of electron-electron cusp on configuration interaction energies*. The Journal of Chemical Physics **115**, 1626 (2001).
41. T. Kato. *On the eigenfunctions of many-particle systems in quantum mechanics*. Communications on Pure and Applied Mathematics **10**, 151 (1957).
42. J. K. Cullum & R. A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations: Vol. I: Theory* ISBN: 78-0898715231 (Society for Industrial and Applied Mathematics, 2002).
43. E. R. Davidson. *Super-matrix methods*. Computer Physics Communications **53**, 49 (1989).
44. E. Rossi, G. L. Bendazzoli, S. Evangelisti & D. Maynau. *A full-configuration benchmark for the N_2 molecule*. Chemical Physics Letters **310**, 530 (1999).
45. A. Szabo & N. S. Ostlund. *Modern Quantum Chemistry. Introduction to Advanced Electronic Structure Theory* 287–291. ISBN: 0-486-69186-1 (Dover Publications, 1996).
46. H. Nakatsuji & K. Hirao. *Cluster expansion of the wavefunction. Symmetry-adapted-cluster expansion, its variational determination, and extension of open-shell orbital theory*. The Journal of Chemical Physics **68**, 2053 (1978).
47. A. J. W. Thom. *Stochastic Coupled Cluster Theory*. Physical Review Letters **105**, 263004 (2010).
48. W. L. McMillan. *Ground State of Liquid He*. Physical Review **138**, A442 (1965).
49. R. P. Feynman. *Atomic Theory of the λ Transition in Helium*. Physical Review **91**, 1291 (1953).
50. M. H. Kalos. *Monte Carlo Calculations of the Ground State of Three- and Four-Body Nuclei*. Physical Review **128**, 1791 (1962).
51. A. J. W. Thom & A. Alavi. *Stochastic Perturbation Theory: A Low-Scaling Approach to Correlated Electronic Energies*. Physical Review Letters **99**, 1434001 (2007).
52. D. Ceperley, G. V. Chester & M. H. Kalos. *Monte Carlo simulation of a many-fermion study*. Physical Review B **16**, 3081 (1977).
53. S. A. Baeurle. *Computation within the auxiliary field approach*. Journal of Computational Physics **184**, 540 (2003).

54. R. Blankenbecler, D. J. Scalapino & R. L. Sugar. *Monte Carlo calculation of coupled boson-fermion systems. I.* Physical Review D **24**, 2278 (1981).
55. N. S. Blunt, T. W. Rogers, J. S. Spencer & W. M. C. Foulkes. *Density matrix quantum Monte Carlo.* pre-print. eprint: [arXiv:1303.5007](https://arxiv.org/abs/1303.5007) (2013).
56. W. Foulkes, L. Mitas, R. Needs & G. Rajagopal. *Quantum Monte Carlo simulations of solids.* Reviews of Modern Physics **73**, 33–83 (Jan. 2001).
57. J. B. Anderson. *A random-walk simulation of the Schrödinger equation: H^+_3 .* The Journal of Chemical Physics **63**, 1499 (1975).
58. J. B. Anderson. *Quantum chemistry by random walk.* The Journal of Chemical Physics **65**, 4121 (1976).
59. D. J. Klein & H. M. Pickett. *Nodal hypersurfaces and Anderson's random-walk simulation of the Schrödinger equation.* The Journal of Chemical Physics **64**, 4811 (1976).
60. J. A. Pople. *Nobel Lecture: Quantum chemical models.* Reviews of Modern Physics **71**, 1267 (1999).
61. G. H. Booth, A. Grüneis, G. Kresse & A. Alavi. *Towards an exact description of electronic wavefunctions in real solids.* Nature **493**, 365 (2012).
62. J. Spencer. *Hellmann-Feynman sampling in FCIQMC.* Personal communication (2012).
63. J. Spencer. *Evaluation of expectation values in full configuration interaction quantum Monte Carlo.* Presented at the 2013 March Meeting of the American Physical Society (2013).
64. J. S. Spencer, N. S. Blunt & W. M. C. Foulkes. *The sign problem and population dynamics in the full configuration interaction quantum Monte Carlo method.* The Journal of Chemical Physics **136**, 054110 (2012).
65. M. Troyer & U.-K. Wiese. *Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations.* Physical Review Letters, 170201 (2005).
66. N. Kawashima, J. E. Gubernatis & H. G. Evertz. *Loop algorithms for quantum simulations of fermion models on lattices.* Physical Review B **50**, 136 (1994).
67. H. Flyvbjerg & H. G. Petersen. *Error estimates on averages of correlated data.* The Journal of Chemical Physics **91**, 461 (1989).
68. C. Overy, D. Cleland, G. H. Booth, J. J. Shepherd & A. Alavi. *Reduced Density Matrices in Full Configuration Interaction Quantum Monte Carlo.*

- Presented at the 2013 March Meeting of the American Physical Society (2013).
69. C. J. Umrigar, M. P. Nightingale & K. J. Runge. *A diffusion Monte Carlo algorithm with very small time-step errors*. The Journal of Chemical Physics **99**, 2865 (1993).
 70. F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale & C. J. Umrigar. *Semistochastic Projector Monte Carlo Method*. Physical Review Letters **109**, 230201 (2012).
 71. G. H. Booth *et al.* *Full Configuration Interaction Quantum Monte Carlo: A Game of Life, Death and Annihilation*. (Poster) presented at Advanced Methods and Applications in Quantum Chemistry 2012: An international workshop with hands-on exercises using Molpro (2012).
 72. C. Overy. *Real Coefficients*. Personal communication (2012).
 73. G. H. Booth, S. D. Smart & A. Alavi. *Linear-scaling and parallelizable algorithms for stochastic quantum chemistry*. pre-print. eprint: arXiv:1305.6981 (2013).
 74. D. Cleland, G. Booth & A. Alavi. *Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum Monte Carlo*. The Journal of chemical physics **132**, 041103 (2010).
 75. J. J. Shepherd, G. H. Booth & A. Alavi. *Investigation of the full configuration interaction quantum Monte Carlo method using homogeneous electron gas models*. The Journal of Chemical Physics **136**, 244101 (2012).
 76. T. Helgaker, P. Jørgensen & J. Olsen. *Molecular Electronic-Structure Theory* 567–568. ISBN: 9781118531471 (John Wiley & Sons Ltd., 2012).
 77. J. Olsen, B. O. Roos, P. Jørgensen & H. J. A. Jensen. *Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces*. The Journal of Chemical Physics **89**, 2185 (1988).
 78. Y. G. Smeyers & L. Doreste–Suarez. *Half-Projected and Projected Hartree-Fock Calculations for Singlet Ground States. i. four-Electron Atomic Systems*. International Journal of Quantum Chemistry **7**, 687 (1973).
 79. W. Gerlach & O. Stern. *Das magnetische Moment des Silberatoms*. Zeitschrift für Physik **9**, 353 (1922).
 80. T. Helgaker, P. Jørgensen & J. Olsen. *Molecular Electronic-Structure Theory* 46–51. ISBN: 9781118531471 (John Wiley & Sons Ltd., 2012).
 81. F. A. Matsen. *Spin-Free Quantum Chemistry*. Advances in Quantum Chemistry **1**, 59–114 (1964).

82. T. Helgaker, P. Jørgensen & J. Olsen. *Molecular Electronic-Structure Theory* 38–41. ISBN: 9781118531471 (John Wiley & Sons Ltd., 2012).
83. W. I. Salmon & J. Ruedenberg. *Many-Electron Wavefunctions Expanded in Spin-Adapted Antisymmetrized Products, and Their Expectation Values*. The Journal of Chemical Physics **57**, 2776 (1972).
84. G. Rumer. *Zur Theorie der Spinvalenz*. Nachrichten von der Gessellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse, 337 (1932).
85. G. Rumer, E. Teller & H. Weyl. *Eine für die Valenztheorie geeignete Basis der binären Vectorinvarianten*. Nachrichten von der Gessellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse, 499 (1932).
86. M. Simonetta, E. Gianinetti & I. Vandoni. *Valence-Bond Theory for Simple Hydrocarbon Molecules, Radicals and Ions*. The Journal of Chemical Physics **48**, 1579 (1968).
87. R. Pauncz. *Spin Eigenfunctions. Construction and Use* 77–89. ISBN: 0-306-40141-X (Plenum Press, 1979).
88. M. Kotani, A. Amemiya, E. Ishiguro & T. Kimura. *Tables of Molecular Integrals* (Maruzen Co. Tokyo, 1963).
89. R. Pauncz. *Spin Eigenfunctions. Construction and Use* 14–32. ISBN: 0-306-40141-X (Plenum Press, 1979).
90. J. H. V. Vleck & A. Sherman. *The Quantum Theory of Valence*. Reviews of Modern Physics **7**, 167 (1935).
91. R. Serber. *Extension of the Dirac Vector Model to Include Several Configurations*. Physical Review **45**, 461 (1934).
92. P.-O. Löwdin. *Quantum Theory of Many-Particle Systems III. Extension of the Hartree-Fock Scheme to Include Degenerate Systems and Correlation Effects*. Physical Review **97**, 1509 (1955).
93. P.-O. Löwdin. *Angular Momentum Wavefunctions Constructed by Projector Operators*. Reviews of Modern Physics **36** (1964).
94. J. C. Slater. *The Theory of Complex Spectra*. Physical Review **34**, 1293 (1929).
95. E. U. Condon. *The Theory of Complex Spectra*. Physical Review **36**, 1121 (1930).
96. A. Szabo & N. S. Ostlund. *Modern Quantum Chemistry. Introduction to Advanced Electronic Structure Theory* 70. ISBN: 0-486-69186-1 (Dover Publications, 1996).

97. P.-O. Löwdin. *Quantum Theory of Many-Particle Systems I. Physical Interpretations by Means of Density Matrices, Natural Spin-Orbitals, and Convergence Problems in the Method of Configurational Interaction*. Physical Review **97**, 1474 (1955).
98. R. Pauncz. *Spin Eigenfunctions. Construction and Use* 172–173, 210. ISBN: 0-306-40141-X (Plenum Press, 1979).
99. C. Reeves. *An Algorithm for Generating Projective Reduction Formulas for Matrix Elements of Many-Electron Wavefunctions*. Communications of the ACM **9**, 276 (1966).
100. D. E. Knuth. *The Art of Computer Programming. Generating All Combinations and Partitions* 4–6. ISBN: 1-201-85394-9 (Pearson Education, 2005).
101. W. Wu & Q. Zhang. *An efficient algorithm for evaluating the standard Young-Yamanouchi orthogonal representation with two-column Young tableau for symmetric groups*. Journal of Physics A: Mathematical and General **25**, 3737–3747 (1992).
102. S. Rettrup. *A recursive formula for Young's orthogonal representation*. Chemical Physics Letters **47**, 59–60 (1977).
103. A. J. Coleman. *The Symmetric Group Made Easy*. Advances in Quantum Chemistry **4**, 83 (1968).
104. D. E. Rutherford. *Substitutional Analysis* ISBN: 048649120X (Edinburgh University Press (reprint Dover Publications 2013), 1948).
105. J. Paldus & P. E. S. Wormer. *Calculation of permutation matrices using graphical methods of spin algebras: Explicit expressions for the Serber-coupling case*. Physical Review A **18**, 827 (1978).
106. S. Wilson & J. Gerratt. *Direct evaluation of elements of the representation matrices of the spin permutation group for transpositions*. Journal of Physics B: Atomic and Molecular Physics **12**, 339 (1979).
107. P. J. A. Ruttink. *On the evaluation of CI matrix elements for a canonically ordered basis*. Theoretica Chimica Acta **49**, 223 (1978).
108. G. G. Sahasrabudhe, K. V. Dinesha & C. R. Sarma. *Some Simplification in Spin-Free Configuration Interaction Studies*. Theoretica Chimica Acta **54**, 333 (1980).
109. P. B. Karadakov, J. Gerratt, D. L. Cooper & M. Raimondi. *SPINS: A collection of algorithms for symbolic generation and transformation of many-electron spin eigenfunctions*. Theoretica Chimica Acta **90**, 51–73 (1995).

110. ISO/IEC. *Programming languages — C++* ISO/IEC 14882:1998 (International Organization for Standardization and International Electrotechnical Commission, 1998).
111. ISO/IEC. *Programming languages — C++* ISO/IEC 14882:2003 (International Organization for Standardization and International Electrotechnical Commission, 2003).
112. F. E. Harris. *Open-Shell Orthogonal Molecular Orbital Theory*. The Journal of Chemical Physics **46**, 2769 (1967).
113. R. Pauncz, J. de Heer & P. O. Lowdin. *Studies on the Alternant Molecular Orbital Method. I. General Energy Expression for an Alternant System with Closed-Shell Structure*. The Journal of Chemical Physics **36**, 2247 (1962).
114. P.-O. Löwdin. *Quantum theory of cohesive properties of solids*. Advances in Physics **5**, 1 (1956).
115. F. Sasaki & K. Ohno. *Spin-Component Analysis of Single-Determinant Wavefunctions*. Journal of Mathematical Physics **4**, 1141 (1963).
116. R. Pauncz. *Alternant Molecular Orbital Method* 34. ISBN: 978-0721671208 (W.B. Saunders, 1967).
117. F. E. Harris. *Molecular Orbital Theory*. Advances in Quantum Mechanics **3**, 61–127 (1967).
118. V. H. Smith. *Construction of Exact Spin Eigenfunctions*. The Journal of Chemical Physics **41**, 277 (1964).
119. J. Vedene H. Smith & F. E. Harris. *Projection of Exact Spin Eigenfunctions*. Journal of Mathematical Physics **10**, 771 (1969).
120. *NIST Atomic Spectra Database (ver. 5.0)* (eds J. R. A. Kramida Yu Ralchenko & N. A. Team) (National Institute of Standards and Technology, Gaithersburg, MD, 2012). <<http://physics.nist.gov/asd>>.
121. C. E. Moore. *Tables of Spectra of Hydrogen, Carbon, Nitrogen, and Oxygen Atoms and Ions* (ed J. W. Gallagher) 339 (CRC Press, 1993).
122. T. H. Dunning. *Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen*. The Journal of Chemical Physics **90**, 1007 (1989).
123. Y. Shao *et al.* *Advances in methods and algorithms in a modern quantum chemistry program package*. Physical Chemistry Chemical Physics **8**, 3172 (2006).
124. T. Helgaker, P. Jørgensen & J. Olsen. *Molecular Electronic-Structure Theory* 322–323. ISBN: 9781118531471 (John Wiley & Sons Ltd., 2012).

125. H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby & M. Schütz. *Molpro: a general-purpose quantum chemistry program package*. Wiley Interdisciplinary Reviews: Computational Molecular Science **2**, 242 (2012).
126. H.-J. Werner *et al.* *MOLPRO, version 2012.1, a package of ab initio programs* 2012. <<http://www.molpro.net>>.
127. A. J. W. Thom & M. Head-Gordon. *Locating Multiple Self-Consistent Field Solutions: An Approach inspired by Metadynamics*. Physical Review Letters **101**, 193001 (2008).
128. D. Cleland, G. H. Booth, C. Overy & A. Alavi. *Taming the First-Row Diatomics: A Full Configuration Interaction Quantum Monte Carlo Study*. Journal of Chemical Theory and Computation **8**, 4138 (2012).
129. J. Spencer. *Continuous-Time FCIQMC*. Personal communication (2010).