

# Quantum algorithms for simulating quantum systems

**Sam McArdle**

St. Edmund Hall

University of Oxford



A thesis submitted for the degree of:

*Doctor of Philosophy*

Michaelmas 2020

## Abstract

Conventional computers are invaluable tools for analysing and predicting the behaviour of the world around us; from the formation of the Universe at large, to the motion of subatomic particles. Unfortunately, the latter problem has proved exceedingly difficult to simulate accurately, due to the computational complexity allowed by the laws of quantum mechanics. There have been a number of proposals to ‘fight fire with fire’ – to use quantum computers to efficiently simulate quantum systems of interest. State-of-the-art resource estimates suggest that it may be possible to simulate classically intractable systems using hundreds of thousands, or millions, of physical qubits. Unfortunately, such resources are beyond our current capabilities – and may remain so for the foreseeable future.

This thesis presents a number of approaches to lessen the burden of quantum simulation. I introduce a hybrid quantum-classical algorithm for ansatz-based imaginary time evolution, and show how it can be used to find the ground states of Hermitian and non-Hermitian Hamiltonians of chemical relevance. While such hybrid algorithms have become popular in recent years, it is still unclear if they will be able to demonstrate quantum advantage, without the protection of quantum error correction. To partially address this question, I present a technique for error mitigation in chemical simulations, that can be used to reduce the effects of noise, at a modest cost (compared to quantum error correction).

The majority of studies to date have focused on finding the electronic ground states of molecular systems. It may be prudent to investigate alternative simulation targets, which may turn out to require fewer resources. In this thesis I show how two alternative physical phenomena can be simulated on digital quantum computers: the vibrations of molecules, and muon spectroscopy experiments. For the case of molecular vibrations, I show how to map the problem onto a digital quantum computer, and extract properties of physical interest. For the case of muon spectroscopy, I present a quantum algorithm to analyse the data arising from muon spectroscopy experiments, and use numerical simulations to infer the error corrected resources required to simulate classically challenging system sizes.

## Featured publications and preprints

The work presented in this thesis is that of the author, unless otherwise specified. Each chapter is largely based on the following publications and preprints:

- Chapters 2 & 3:

*Quantum computational chemistry*

S. McArdle, S. Endo, A. Aspuru-Guzik, S.C. Benjamin, & X. Yuan  
[Reviews of Modern Physics 92 \(1\), 015003 \(2020\)](#)

- Chapter 4:

*Variational ansatz-based quantum simulation of imaginary time evolution*

S. McArdle, T. Jones, S. Endo, Y. Li, S.C. Benjamin, & X. Yuan  
[Nature Partner Journal Quantum Information 5 \(1\), 1-6 \(2019\)](#)

&

*Improving the accuracy of quantum computational chemistry using the transcorrelated method*

S. McArdle & D.P. Tew  
[arXiv:2006.11181 \(2020\)](#)

- Chapter 5:

*Error-Mitigated Digital Quantum Simulation*

S. McArdle, X. Yuan, & S.C. Benjamin  
[Physical Review Letters 122 \(18\), 180501 \(2019\)](#)

- Chapter 6:

*Digital quantum simulation of molecular vibrations*

S. McArdle, A. Mayorov, X. Shan, S.C. Benjamin, & X. Yuan  
[Chemical Science, 10, 5725-5735 \(2019\)](#)

- Chapter 7:

*Learning from physics experiments, with quantum computers:*

*Applications in muon spectroscopy*

S. McArdle  
[arXiv:2012.06602 \(2020\)](#)

# Acknowledgements

This thesis would not have been possible without the help and support of a great many people, who I would like to take the time to thank here:

Simon; thank you for always being a patient, enthusiastic, and supportive supervisor. Over the last three years you've given me the freedom to research the topics that I wanted, to pursue internships and collaborations, and to attend conferences across the world. Thank you for always being willing and excited to discuss new research ideas, and to help whenever I have been stuck. I am still in awe of your ability to incisively spot the issue that I have been missing for days, or to immediately understand how a method works from my vague explanation (especially given I probably haven't understood the method fully yet, despite reading the paper). I hope that one day I can develop even a fraction of your insight and intuition!

Xiao; thank you for effectively being my co-supervisor during my first year in Oxford. You always had some maths at hand to cut through my confusion, and most of the papers in this thesis wouldn't have happened without your quantum information prowess. Thank you as well for being a great friend; your regular teasing of Suguru, Tyson and myself was a constant source of amusement, and you were a fantastic companion in our trip to New York and Toronto!

To my DPhil co-starters Tyson and Cai; I don't think I've ever been more intimidated by the abilities of my peers than when we first met – but I'm glad that we have become close friends!

Tyson; how you're able to remain so modest despite your inexhaustible list of talents (programming, physics, music, poetry – even hurdling!?) is an impressive and admirable feat. Though perhaps it is your misadventures in dining and dating that help to keep you grounded... Thank you for always being willing to walk me through what must seem like the most basic of coding tasks. Your enthusiasm for new ideas and hilarious tales have made the past three years richer and more exciting!

Cai; I've come to take for granted your Feynman-esque physical intuition – but am still blown away whenever I read one of your papers, and realise that you've invented a new maths concept to prove something that I would have described as 'left for future work'. Thank you for entertaining my misguided questions about quantum error correction, and for always having a funny story to tell. Thank you as well for switching to more basic packed lunches in year 3, as I was starting to get jealous of your daily feast!

To the group veterans Suguru and Xiaosi. Sugusu-tan; thank you for being the silliest and craziest friend that I know. Your stories of life in Japan could fill a book, which I would be able to re-read many times. Behind your playful personality lies a sharp intellect. You are one of the most creative researchers I've met; always coming up with new ideas, or spotting the hole in my logic. Thank you for always being happy to explain error mitigation to me – even over Zoom, in your pyjamas.

Xiaosi; thank you for always patiently explaining what the next DPhil step was, or why my understanding of quantum error correction wasn't quite right. Your excitement and loud exclamations brought joy to the office, and I'm looking forward to being colleagues again in the future!

To the more recent starters; Balint, Carlos, Sam, Armands, & Richard. Thank you for stimulating lunchtime conversations, for asking the difficult questions that always made me think, and for kindly sharing your time and knowledge with me. To the most recent starters; Cica, Hans, Daniel, Arthur, & Hamza. Thank you for all of your kind support during the thesis write-up period. I hope that one day soon we can all meet without the help of Zoom!

To my supervisors and collaborators during my internships with the Google Quantum team; Sergio, Jarrod, Ryan, Zhang, Bill, Matt, Tom, Nick, Kevin, Murphy, Eddie, & Vadim. Thank you for bringing me into your team, and providing so much support over my two summers with you. I learned so much from my time with Google, and am proud of the projects that we worked on. Thank you as well for being such great friends; our dinners out, chats over lunch, and surf sessions made the summer truly memorable, and I look forward to catching up again in the future!

To those in the Aspuru-Guzik group in Toronto, the Zoller group in Innsbruck, & the IBM group in Yorktown; thank you for kindly hosting me during my visits. It was fantastic to learn about the great work that you were doing, to explore your home towns and cities, and to make new like-minded friends across the globe.

To my family and friends; thank you for always taking an interest in my research (I know it hasn't always been easy). It has always meant a lot that you would send me links to news articles about quantum computing, ask me what my papers were about and how they were progressing, or try to better understand what exactly it is that I do each day. Thank you as well

for your regular reminders that I'm *still* at University!

Thank you to my parents for your constant support over the last three years – and of course, the 23 years prior. For always being excited about my new developments, and for listening patiently while I explained things in far more detail than was probably necessary. Thank you for questioning my logic and assertions, and always challenging me to explain things in the most simple way. Thank you for always coming to visit me, whether for house rental hunting in Oxford or for exciting days out in Los Angeles, and for giving me so much to look forward to. I'm so grateful that you've always given me the space and freedom to pursue my own interests, and to explore my own path. I wouldn't be where I am without your advice and encouragement, which always steers me in the right direction.

To my fiancée Katyana; you have probably been more involved in this thesis than anyone else. Thank you for being perpetually patient with me; over the last three years you have had to put up with long rants about the state of quantum computing, over-detailed explanations of my problems and solutions, and my recurring indecisiveness. That you are always happy and excited by my news is a testament to how wonderful you are – even when you have to deal with me getting up in the middle of the night to write down some new idea. By now you are probably more knowledgeable about quantum computing than many PhD students, and you have helped me sound through so many challenges, both technical and professional. I am more grateful than I can express for your endless support; you have travelled under and over oceans to visit me at conferences and internships, and have allayed so many of my stresses and worries. Thank you for always supplying the fun and excitement, and for giving me something to smile and laugh about every day.

For my parents, *Joy & Steve*, and my grandparents *Doreen, Nana, Elsie & Ron*. Without your tireless support of my education and unwavering love and encouragement, this thesis would not have been possible.

# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Computational chemistry</b>	<b>8</b>
2.1	Classical approaches . . . . .	8
2.1.1	Basis sets . . . . .	10
2.1.2	Second quantisation . . . . .	12
2.1.3	Classical simulation techniques . . . . .	15
2.1.3.1	Hartree-Fock . . . . .	15
2.1.3.2	Coupled cluster . . . . .	16
2.1.3.3	Other techniques . . . . .	17
2.1.4	Systems of interest . . . . .	19
2.2	Mapping onto quantum computers . . . . .	20
2.2.1	Mapping to qubits . . . . .	21
2.2.2	Application to H <sub>2</sub> . . . . .	22
2.2.3	Application to LiH . . . . .	24
<b>3</b>	<b>Quantum simulation</b>	<b>26</b>
3.1	Quantum simulation algorithms . . . . .	26
3.1.1	Time evolution . . . . .	27
3.1.2	Quantum phase estimation . . . . .	31
3.1.3	Variational quantum eigensolver . . . . .	33
3.1.3.1	Ansätze . . . . .	37
3.2	Quantum error mitigation and correction . . . . .	40
3.2.1	Error mitigation . . . . .	41
3.2.2	Error correction . . . . .	45
<b>4</b>	<b>Variational quantum algorithms for imaginary time evolution</b>	<b>50</b>

---

4.1	Chapter summary . . . . .	50
4.2	Introduction . . . . .	51
4.3	Theoretical background . . . . .	52
4.3.1	Ansatz-based quantum imaginary time evolution . . . . .	52
4.3.2	Explicit correlation & the transcorrelated method . . . . .	60
4.4	Results . . . . .	66
4.4.1	Initial benchmarks . . . . .	66
4.4.2	Application to explicitly correlated Hamiltonians . . . . .	71
4.5	Discussion . . . . .	78
<b>5</b>	<b>Error mitigation for quantum simulation</b>	<b>82</b>
5.1	Chapter summary . . . . .	82
5.2	Introduction . . . . .	82
5.3	Symmetry verification . . . . .	84
5.3.1	Particle number parity . . . . .	85
5.3.2	Spin parity . . . . .	88
5.3.3	Particle number . . . . .	89
5.3.4	Practical considerations . . . . .	93
5.4	Results . . . . .	94
5.5	Discussion . . . . .	98
<b>6</b>	<b>Quantum algorithms for simulating molecular vibrations</b>	<b>102</b>
6.1	Chapter summary . . . . .	102
6.2	Introduction . . . . .	102
6.3	Mapping vibrations to qubits . . . . .	106
6.4	Simulating vibrational problems . . . . .	113
6.5	Discussion . . . . .	119
<b>7</b>	<b>Quantum simulation of muon spectroscopy experiments</b>	<b>123</b>
7.1	Chapter summary . . . . .	123

---

7.2	Introduction . . . . .	124
7.2.1	Muon spectroscopy . . . . .	125
7.2.2	Muon polarisation functions . . . . .	128
7.3	Quantum algorithm for muon spectroscopy analysis . . . . .	134
7.3.1	Mapping the spin system to qubits . . . . .	136
7.3.2	Preparing the initial state of the system . . . . .	138
7.3.3	Evolving the state in time . . . . .	141
7.3.4	Measuring the polarisation . . . . .	142
7.4	Results . . . . .	144
7.4.1	Noiseless simulations . . . . .	147
7.4.1.1	Quantifying errors . . . . .	147
7.4.1.2	Analysing $\text{CaF}_2 + \mu^+$ spectra . . . . .	151
7.4.2	Noisy simulations . . . . .	155
7.4.3	Resource estimates . . . . .	161
7.4.3.1	NISQ resource estimates . . . . .	162
7.4.3.2	Error corrected resource estimates . . . . .	163
7.5	Discussion . . . . .	167
<b>8</b>	<b>Conclusion</b>	<b>171</b>
	<b>Appendix A Appendices</b>	<b>218</b>
A.1	Appendix: Error mitigation for quantum simulation . . . . .	218
A.1.1	Number of measurements and error analysis . . . . .	218
A.2	Appendix: Quantum simulation of muon spectroscopy ex- periments . . . . .	220
A.2.1	Preparing Dicke states . . . . .	220
A.2.2	Additional numerical simulations of Trotter and sam- pling error . . . . .	225

# 1 | Introduction

*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 that the exact application of these laws leads to equations much too complicated to be soluble.*

Paul Dirac, 1929 [1]

Dirac's statement still rings true today, even in the face of the digital revolution that has endowed us with computational powers unimaginable in the 1920's. Despite the exponential growth in computing power resulting from Moore's Law, and similarly impressive algorithmic developments, we have yet to overcome the challenges of simulating ever larger quantum systems. Storing a general quantum wavefunction appears to require an amount of classical memory scaling exponentially with the number of particles in the system. While generations of scientists have devised ever-improving, polynomially scaling approximate solutions to this problem, these are of limited applicability to many phenomena of interest – including high temperature superconductivity, transition metal catalysis, and photochemistry.

While superposition and entanglement underpin this problem, they may also provide its solution. In 1982, Feynman proposed using controllable quantum systems as a platform to simulate other quantum systems of interest [2]. A simulator is often defined as a device that, when run and then measured, provides insights into a mathematical model that we believe accurately describes a system of interest [3, 4]. Such a simulator could be either analog (whereby the parameters describing the system and its evolution are continuous) or digital (described by a discrete, digitized state space, with discrete transitions between states) [5]. For example, if we wanted to use a device to learn about how air flows around an aeroplane wing, we could consider both a model wing in a wind tunnel, and a com-

putational fluid dynamics calculation performed on a classical computer. We can also divide simulators into those that are ‘universal’, by which we mean that they can be reprogrammed to simulate any system of interest, and those that are not. It is typically straightforward to reprogram digital computers (quantum or classical) to simulate other phenomena, while the case is more nuanced for analog systems.

When considering analog quantum simulators, a number of different platforms have been developed, that are capable of simulating a multitude of scientific phenomena of interest [6, 7]. For example, we can use photons propagating in waveguides to simulate excitation transfer in biological systems [8], cold atoms in optical lattices to simulate phase transitions in condensed matter systems [9] or non-equilibrium dynamics [10], trapped ions to simulate interacting spin systems [11], and networks of coupled superconducting resonators to simulate open quantum systems [12]. In some cases these simulators are able to provide information that is beyond the capability of classical digital computers [13, 14]. For example, recent simulations of non-equilibrium dynamics of spin- $\frac{1}{2}$  particles simulated by trapped ion [15] and cold atom [16] simulators, appear beyond the classical simulation threshold. Very recently, simulations have been performed with up to 256 cold atoms [17], considerably beyond what would be possible with brute force classical computation. It is often possible to engineer a number of different interactions of interest in these systems, and our ability to address individual particles is gradually improving. Nevertheless, there are two main challenges that may restrict the applicability of analog quantum simulators in the long term. The first, is that not all of these simulators are straightforwardly re-programmable (while recent work has shown that even simple spin models are universal for quantum simulation [18, 19], it may be extremely challenging to engineer the control required). Secondly, it is not currently known how to protect analog quantum systems from

the effects of environmental noise and decoherence. It is often claimed that analog quantum simulators should be considered more robust to noise than their digital counterparts. There are typically two reasons given for this. The first, is that the properties we wish to extract are often expected to be resilient to noise; for example, the magnetization of a spin system, or the phase diagram of a condensed matter system. It has been noted that such ‘macroscopic’ observables do often appear robust to noise [20], as they may correspond to thermodynamic macrostates consistent with many microstates, allowing them to withstand small perturbations [5, 21]. However, it has also been shown that ‘microscopic’ observables do appear to be sensitive to noise [20]. The second argument often given as to why noise is less problematic for analog simulators, is that noise is likely also present in the real system that we wish to investigate. Unfortunately, counter-examples exist that show that noise present in the simulator does not always correspond to physical noise in the target system [22]. Given the difficulty in isolating analog simulators from errors, it is challenging to predict the sizes to which these systems will be able to scale, as well as whether we are able to trust the results of a given calculation.

Alongside the development of analog quantum simulation, Lloyd [23], Aspuru-Guzik [24], and many others [7] have formalised how digital quantum computers may be used to efficiently simulate problems in physics, chemistry, and materials science. These algorithms typically exploit the ability of qubits to exhibit the same superposition and entanglement that appear to make classical simulations so challenging. Exploiting this correspondence requires suitable mappings, which I discuss for the case of quantum chemistry in Chapter 2. Algorithms protected by quantum error correction have been designed to extract properties of interest, as discussed in Chapter 3. Error correction in digital quantum computers is made possible by a discretization of the continuum of possible errors. This is facilitated by

restricting ourselves to a discrete set of quantum gates. Achieving both error correction, and computational universality with restricted gatesets is a delicate balancing act, as discussed in Sec. 3.2.2. Provable accuracy bounds on both algorithms, and methods of error correction, enable us to perform estimates of the resources required to simulate classically intractable systems – and, perhaps more importantly, to trust the outputs of the computation. This independent verifiability has previously been suggested as a distinction between a ‘simulation method’ and a ‘computation method’ [3]. Unfortunately, even the best approaches developed to date appear to suggest that hundreds of thousands, if not millions, of physical qubits may be required [25–27]. These dauntingly high resources have led some to question whether it might be prudent to develop algorithms amenable to noisy quantum computers, foregoing error correction. These NISQ (noisy intermediate-scale quantum) [21] algorithms are typically heuristic in nature, and often trade circuit depth for an increased number of circuit repetitions. While there has been significant interest in such approaches over the last 5 years, it is currently unclear whether classically intractable simulations can be performed without quantum error correction. In addition, it is important to note that while NISQ devices may appear to be faulty digital quantum computers, they are in fact analog machines, because of the use of a continuous gateset. As I shall discuss throughout this thesis, NISQ devices simulate physical systems using gate sequences that typically break up the complicated system dynamics into more tractable parts (for example, the Trotterization technique discussed in Sec. 3.1.1). One might question why we would expect this approach to have benefits over analog quantum simulators, which do not require such a decomposition of the problem. One potential advantage of NISQ devices, is that they are straightforwardly re-programmable. For example, I show in this thesis how the same device could be used to investigate the electronic structure of Fermi-Hubbard

models, vibrations in molecules, and muon spectroscopy experiments. But does this suggest that NISQ devices will be a ‘Jack of all trades, Master of none’? As the progenitor of error corrected devices, NISQ machines will need to have precise control over individual qubits, and low gate and measurement errors. This may benefit the accurate determination of electronic energy levels in molecules, which looks difficult to achieve with recently proposed analog approaches [28]. Moreover, the control afforded by the quantum circuit model allows the construction of error mitigation techniques (see Refs. [29, 30] and Chapter 5) that can be used to further improve the simulation results. Ultimately, the best resource for a given simulation will depend on what exactly we wish to simulate, and what we want to learn from the simulation. It will be more constructive to view analog and NISQ quantum simulations as complementary techniques in a scientist’s toolbox, rather than as competing platforms in a winner-takes-all scenario.

This thesis investigates a number of ways to reduce the resources required for quantum simulation on both NISQ, and error corrected, quantum computers. In Chapter 4 I present a hybrid quantum-classical algorithm for ansatz-based imaginary time evolution, which can be used to find the ground states of quantum systems, and appears competitive with leading NISQ approaches. In Chapter 5 I introduce an error mitigation technique that can be used to reduce the impact of noise in quantum simulations, which uses modest additional resources. In Chapters 6 & 7 I discuss approaches to simulate novel phenomena on digital quantum computers; specifically, molecular vibrations and muon spectroscopy. By considering less widely studied areas, we may identify new targets for quantum simulation that require fewer resources than existing proposals (which are largely based on finding the electronic ground states of chemical systems). I discuss this topic at length in Chapter 7.

The investigations presented in this thesis rely on a combination of analytic work, and numerical simulations. These simulations used conventional computers to emulate running quantum circuits. The cost of these emulations scales exponentially with the size of the simulated system (as the emulations store the wavefunction or density matrix of the quantum computer). However, a number of software packages have been developed that dramatically reduce the cost of these calculations, compared to naively performing the underlying linear algebraic operations. The availability and capabilities of such packages have evolved rapidly over the past three years, which has made possible work that would have been prohibitive in the years prior. In particular, I have made use of:

- QuEST [31], a high performance quantum circuit emulator written in C. QuEST can be efficiently parallelised using OpenMP, MPI, or GPUs. Its rapid execution speed made the 29 qubit simulations presented in Chapter 7 possible.
- Cirq [32] and ProjectQ [33], quantum circuit emulators written in Python. These packages provide flexible frameworks for manipulating quantum circuits, which proved invaluable for emulating the hybrid quantum-classical algorithms presented in Chapters 4 & 5.
- OpenFermion [34], a Python package for generating fermionic electronic structure Hamiltonians, and mapping them to qubits. This package was used to generate the fermionic Hamiltonians used in Chapters 4 & 5, to develop the qubit-mapped vibrational Hamiltonians discussed in Chapter 6, and to construct the spin Hamiltonians used in Chapter 7.

The rapid emergence of such tools, matched by the equally swift progress in quantum computing research, has made the last three years an exciting

time to begin research in this field. The next few years will likely be just as exciting. In particular, as quantum processors in the cloud become easier to access, and of higher quality, it may soon become possible to run experiments on real hardware that would be too challenging to numerically emulate.

This thesis assumes little knowledge of chemistry, and I introduce the relevant background material on computational chemistry in Chapter 2. I assume a basic knowledge of quantum computing, as provided by the initial chapters of Ref. [35]. I do not presume knowledge of quantum algorithms, error mitigation techniques, or error correction, which are introduced in Chapter 3.

## 2 | Computational chemistry

### Contents

---

<b>2.1 Classical approaches</b> . . . . .	<b>8</b>
2.1.1 Basis sets . . . . .	10
2.1.2 Second quantisation . . . . .	12
2.1.3 Classical simulation techniques . . . . .	15
2.1.4 Systems of interest . . . . .	19
<b>2.2 Mapping onto quantum computers</b> . . . . .	<b>20</b>
2.2.1 Mapping to qubits . . . . .	21
2.2.2 Application to H <sub>2</sub> . . . . .	22
2.2.3 Application to LiH . . . . .	24

---

In this chapter, I first introduce the classical approaches to computational chemistry that are pertinent to this thesis. In particular, I focus on techniques for finding the electronic ground states of chemical systems. I then discuss how this problem can be mapped onto a quantum computer. The information presented in this chapter is largely adapted from Ref. [36]:

*Quantum computational chemistry*  
S. McArdle, S. Endo, A. Aspuru-Guzik, S.C. Benjamin, & X. Yuan  
[Reviews of Modern Physics 92 \(1\), 015003 \(2020\)](#)

### 2.1 Classical approaches

The Coulomb Hamiltonian (in atomic units) of a molecule consisting of  $G$  nuclei and  $N$  electrons is given by

$$H = - \sum_i \frac{\nabla_i^2}{2} - \sum_I \frac{\nabla_I^2}{2M_I} - \sum_{i,I} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{I \neq J} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|}. \quad (2.1)$$

where  $M_I$ ,  $\mathbf{R}_I$ , and  $Z_I$  denote the mass (in units of the electron mass), position (in units of Bohr radii), and atomic number of the  $I^{\text{th}}$  nucleus, and

$\mathbf{r}_i$  is the position (in units of Bohr radii) of the  $i^{\text{th}}$  electron. This Hamiltonian measures the energy in units of ‘Hartree’ (1 Hartree = 27.211 eV). Eq. (2.1) neglects relativistic effects (e.g. spin-orbit coupling, and the increased mass of the electron at high speeds), and is a valid approximation for molecules containing light atoms. Analytic solutions for the molecular Hamiltonian are only known for the most simple systems, such as the Hydrogen atom. Consequently, we must rapidly turn to approximate numerical methods when dealing with more complex systems. It is common to incorporate further approximations to lessen the computational burden. A widely applied simplification is the Born-Oppenheimer approximation, which separates the nuclear motion from the electronic motion. Protons and neutrons are three orders of magnitude more massive than an electron, and so typically move on a slower timescale. As a result, in many systems it is reasonable to consider the nuclei as being classical point charges, that are fixed in place. The electrons then interact with each other in the potential caused by these fixed nuclei, which determines the equilibrium electronic structure of the system. The electronic structure Hamiltonian for a given nuclear configuration is

$$H_e(\{\mathbf{R}_I\}) = -\sum_i \frac{\nabla_i^2}{2} - \sum_{i,I} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}. \quad (2.2)$$

Finding the lowest lying energy levels for a range of nuclear configurations enables us to map out the potential energy surfaces of the molecule. These potential energy surfaces can then be taken as the basis for nuclear motion (vibrational and rotational motion), which provide corrections to the electronic structure. The vibrational structure of molecules will be discussed in more detail in Chapter 6. The Born-Oppenheimer approximation breaks down in a number of scenarios, including at degeneracies between the electronic potential energy surfaces. I do not consider relativistic or non-Born-

Oppenheimer effects in this thesis. We typically seek energy eigenvalues to an accuracy of at least  $1.6 \times 10^{-3}$  Hartree, known as ‘chemical accuracy’. If the energy is known to chemical accuracy, then the chemical reaction rate at room temperature can be approximated to within an order of magnitude using the Eyring equation [37, 38]

$$\text{Rate} \propto e^{-\Delta E/k_B T}, \quad (2.3)$$

where  $T$  is the temperature of the system, and  $\Delta E$  is the energy difference between the initial and final states.

Eq. (2.2) describes electrons interacting in continuous, real space – and so corresponds to an infinite dimensional Hilbert space. In order to make calculations more tractable, it is necessary to discretize the system. One approach to this discretization would be to use a finite sized grid. While this approach offers the benefit of being able to treat the nuclei and electrons on an equal footing (i.e. without the Born-Oppenheimer approximation), the large number of grid points required restricts this approach to small system sizes. A more efficient approach is to exploit our knowledge of the system by projecting onto a set of basis functions.

### 2.1.1 Basis sets

The most common approach to computational chemistry is to consider a basis set consisting of  $M$  single-particle atomic or molecular spin-orbitals (although alternative basis sets that utilise multi-particle orbitals are discussed in Sec. 4.3.2). The spin-orbitals are functions that seek to approximate the ‘true’ orbitals of the system, such as: exponential functions (Slater-type orbitals), Gaussian functions (Gaussian-type orbitals), or plane waves. Different types of functions have their own strengths and weaknesses. For

example, Slater-type orbitals better approximate the far-field behaviour of atomic systems than Gaussian-type orbitals. However, the latter functions provide more efficient evaluation of the integrals required to construct the Hamiltonian. Similarly, while plane waves are appropriate for describing periodic systems, they provide an inferior resolution for describing molecular systems than Gaussian-type orbitals, and so require a larger number of basis functions. A common approach to molecular chemistry (and the approach taken in this thesis) is to construct a set of approximate Slater-type orbitals from linear combinations of Gaussian-type orbitals. The larger the basis set, the smaller the discretization error resulting from projecting onto a finite basis set.

In order to satisfy the Pauli exclusion principle the electronic wavefunction must be antisymmetric under the exchange of any two electrons in the system. This can be achieved by constructing the wavefunction from a linear combination of Slater determinants

$$\Psi(\mathbf{x}_0 \dots \mathbf{x}_{N-1}) = \sum_i c_i \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_0^i(\mathbf{x}_0) & \phi_1^i(\mathbf{x}_0) & \dots & \phi_{N-1}^i(\mathbf{x}_0) \\ \phi_0^i(\mathbf{x}_1) & \phi_1^i(\mathbf{x}_1) & \dots & \phi_{N-1}^i(\mathbf{x}_1) \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \phi_0^i(\mathbf{x}_{N-1}) & \phi_1^i(\mathbf{x}_{N-1}) & \dots & \phi_{N-1}^i(\mathbf{x}_{N-1}) \end{vmatrix} \quad (2.4)$$

where  $c_i$  are complex amplitudes, and  $\phi_j^i$  specifies the  $j^{\text{th}}$  basis function occupied in the  $i^{\text{th}}$  Slater determinant. Swapping the positions of any two electrons is equivalent to interchanging two rows of each determinant, which changes the overall sign of the wavefunction. Each Slater determinant contains only the  $N$  occupied orbitals, and not the  $M - N$  virtual orbitals.

Rather than expanding the wavefunction as a sum of antisymmetrised products of single-particle spin orbitals, it is more efficient to just keep track of which spin-orbitals are occupied in a given Slater determinant, while ‘remembering’ that the Slater determinants represent antisymmetrised products.

### 2.1.2 Second quantisation

It is possible to reformulate our problem into one that involves the manipulation of Fock space occupation number vectors  $|f\rangle$  that serve as a shorthand for Slater determinants. For a given Slater determinant,  $\psi(\mathbf{x}_0 \dots \mathbf{x}_{N-1})$ , we write that

$$\psi(\mathbf{x}_0 \dots \mathbf{x}_{N-1}) := |f_{M-1}, \dots, f_p, \dots, f_0\rangle = |f\rangle, \quad (2.5)$$

where  $f_p = 1$  when  $\phi_p$  is occupied (and therefore present in the Slater determinant), and  $f_p = 0$  when  $\phi_p$  is empty (and therefore not present in the determinant). Electrons are excited into or out of the single electron spin-orbitals by fermionic creation/annihilation operators,  $a_p^\dagger/a_p$ , respectively. These operators obey fermionic anti-commutation relations which are required to maintain the correct exchange statistics of the electronic wavefunction

$$\begin{aligned} \{a_p, a_q^\dagger\} &= a_p a_q^\dagger + a_q^\dagger a_p = \delta_{pq}, \\ \{a_p, a_q\} &= \{a_p^\dagger, a_q^\dagger\} = 0. \end{aligned} \quad (2.6)$$

The actions of these operators on the determinants are given by

$$\begin{aligned} & a_p |f_{M-1}, f_{M-2}, \dots, f_0\rangle \\ &= \delta_{f_p, 1} (-1)^{\sum_{i=0}^{p-1} f_i} |f_{M-1}, f_{M-2}, \dots, f_p \oplus 1, \dots, f_0\rangle, \end{aligned} \quad (2.7)$$

$$\begin{aligned} & a_p^\dagger |f_{M-1}, f_{M-2}, \dots, f_0\rangle \\ &= \delta_{f_p, 0} (-1)^{\sum_{i=0}^{p-1} f_i} |f_{M-1}, f_{M-2}, \dots, f_p \oplus 1, \dots, f_0\rangle, \end{aligned}$$

where  $\oplus$  denotes addition modulo 2. The phase term  $(-1)^{\sum_{i=0}^{p-1} f_i}$  enforces the exchange anti-symmetry of fermions. The spin-orbital occupation operator is given by

$$\begin{aligned} \hat{n}_i &= a_i^\dagger a_i, \\ \hat{n}_i |f_{M-1}, \dots, f_i, \dots, f_0\rangle &= f_i |f_{M-1}, \dots, f_i, \dots, f_0\rangle, \end{aligned} \quad (2.8)$$

and counts the number of electrons in a given spin-orbital.

The second quantised form of the electronic structure Hamiltonian is obtained by projecting Eq. (2.2) onto the single-particle basis functions [39, 40]

$$H = \sum_{p,q} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{p,q,r,s} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s, \quad (2.9)$$

with

$$\begin{aligned} h_{pq} &= \int d\mathbf{x} \phi_p^*(\mathbf{x}) \left( -\frac{\nabla^2}{2} - \sum_I \frac{Z_I}{|\mathbf{r} - \mathbf{R}_I|} \right) \phi_q(\mathbf{x}), \\ h_{pqrs} &= \int d\mathbf{x}_1 d\mathbf{x}_2 \frac{\phi_p^*(\mathbf{x}_1) \phi_q^*(\mathbf{x}_2) \phi_r(\mathbf{x}_2) \phi_s(\mathbf{x}_1)}{|\mathbf{r}_1 - \mathbf{r}_2|}. \end{aligned} \quad (2.10)$$

The first integral represents the kinetic energy terms of the electrons and their Coulomb interaction with the nuclei. The second integral is due to the electron-electron Coulomb repulsion. In a Gaussian basis set, this Hamilto-

nian typically contains  $\mathcal{O}(M^4)$  terms.

The Slater determinants are the eigenstates of a Hermitian operator (typically the non-interacting or mean-field Hamiltonian for the system), and so form a basis set for the problem Hilbert space. Correspondingly, the occupation vectors form a basis for the Fock space that describes the system. We can therefore expand the wavefunction as linear combination of occupation number vectors (representing Slater determinants)

$$|\Psi\rangle = \sum_f \alpha_f |f\rangle, \quad (2.11)$$

where  $\alpha_f$  are complex coefficients. If all  $\binom{M}{N}$  determinants are included, the wavefunction is known as the full configuration interaction (FCI) wavefunction. However, this wavefunction contains a number of determinants which scales exponentially with the number of electrons, making large calculations classically intractable. Consequently, many classical computational approaches to the electronic structure problem seek to approximate the FCI wavefunction using polynomially scaling methods. There are two main sources of complexity that make it challenging to construct good approximations to the FCI wavefunction. The wavefunction must account for correlation between the electrons, which can loosely be partitioned into ‘static’ and ‘dynamic’ correlation. Static correlation arises when multiple determinants are equally important in the wavefunction – such as during bond breaking, or when describing low-spin open-shell excited states [41]. These effects are important in many systems of interest, such as transition metal catalysts [42]. Methods designed to treat static correlation typically store as many of the important determinants as possible, which leads to a large memory cost. In turn, this makes these methods too expensive to apply in large basis sets.

Dynamic correlation arises from the Coulomb repulsion between electrons, which tends to anti-correlate their positions. The most severe manifestation of dynamic correlation occurs at short inter-electronic distances. As two electrons approach each other, the Coulomb potential diverges, leading to sharp features in the wavefunction, known as the ‘electron-electron cusps’. In order to accurately resolve these cusps, a large set of single-particle basis functions are required. Methods to treat dynamic correlation often circumvent the need to store a wavefunction, but are typically designed to work from single determinant reference states. Consequently, we see that accurately treating both static and dynamic correlation becomes a computationally difficult task. In Sec. 4.3.2 I discuss an alternative formulation of the electronic structure problem designed to address this challenge. However, for now I will summarise the conventional classical approaches to approximate the FCI wavefunction.

### 2.1.3 Classical simulation techniques

#### 2.1.3.1 Hartree-Fock

The Hartree–Fock (HF) method is a mean-field technique which aims to find the dominant Slater determinant in the wavefunction. The HF technique considers each of the  $N$  electrons moving in an effective potential generated by the other  $N - 1$  electrons in the system. The method iteratively constructs the mean-field Hamiltonian for the system, and then minimises the energy of the single determinant wavefunction by optimising the spatial form of the single-particle orbitals. This produces a set of orthonormal ‘canonical molecular orbitals’ from the initially supplied set of atomic orbitals. As a mean-field method, the HF technique does not account for either static or dynamic correlation. The procedure has a modest cost, and can provide approximate solutions to systems with thousands of spin-orbitals [39].

### 2.1.3.2 Coupled cluster

The coupled cluster (CC) method considers excitations above a single determinant reference, typically taken as the Hartree-Fock state. The CC wavefunction is given by

$$|\Psi_{\text{CC}}\rangle = e^T |\Psi_{\text{HF}}\rangle, \quad (2.12)$$

where  $T = \sum_i T_i$ ,

$$\begin{aligned} T_1 &= \sum_{i \in \text{virt}, \alpha \in \text{occ}} t_{i\alpha} a_i^\dagger a_\alpha, \\ T_2 &= \sum_{i, j \in \text{virt}, \alpha, \beta \in \text{occ}} t_{ij\alpha\beta} a_i^\dagger a_j^\dagger a_\alpha a_\beta, \\ &\dots, \end{aligned} \quad (2.13)$$

where *occ/virt* denotes orbitals that are occupied/unoccupied in the Hartree-Fock state, and  $t$  are excitation amplitudes. The excitation operator is truncated, typically at the level of single and double excitations (CCSD), to make the method computationally tractable. The CCSD method does not store the wavefunction generated by Eq. (2.12), and instead uses this form as the basis for deriving coupled non-linear equations. Solving these equations yields the CCSD approximation to the FCI ground state [39, 43]. There has been significant work to reduce the computational costs of CCSD, which in its canonical form scales as  $\mathcal{O}((M - N)^4 N^2)$  [43]. Approximations which exploit the locality of dynamical electron correlation in certain systems [44, 45] can reduce the scaling to be, in some cases, linear [46].

The CC method is effective at recovering dynamic correlation in the wavefunction, and has been successfully applied to systems containing hundreds of spin-orbitals [47–49]). However, the canonical CC method described above utilises a single determinant reference state, and does not

perform well when applied to systems with a multireference character. As a result, many CC methods struggle to treat systems with strong static correlation. While there have been efforts to develop multireference CC approaches, these have their own limitations, and are not in widespread use, as discussed by Ref. [50]. While CC methods have been successfully applied to some systems displaying static correlation [51, 52], higher excitation degrees are typically required, increasing the computational cost of the method.

In Sec. 3.1.3.1 I discuss a modified form of the CC method, known as unitary coupled cluster (UCC). This method is both variational and suitable for use with multireference states. While no efficient classical algorithms for the UCC method are known, it can be efficiently implemented on a quantum computer.

### 2.1.3.3 Other techniques

There are numerous other methods for approximating the FCI wavefunction. In this thesis I will make reference to multiconfigurational self-consistent field (MCSCF) methods, tensor network techniques, and quantum Monte Carlo methods. I briefly elaborate on these approaches here.

The MCSCF approach considers a multideterminant wavefunction, and variationally optimises both the amplitudes of each determinant, and the underlying form of the single-particle orbitals [53]. Because of its multireference nature, MCSCF is often effective at describing static correlation. To reduce the cost of the method, we typically consider an ‘active space’ of only the most important orbitals. We then carry out a MCSCF calculation on all of the determinants that could be generated from distributing the active space electrons in these orbitals. This is known as the complete active space self-consistent field (CASSCF) method [53], and can be applied

to active spaces with up to 40 spin-orbitals [54].

Tensor network methods, such as density matrix renormalisation group (DMRG), consider low rank approximations of the FCI wavefunction generated by contracting networks of tensors. These methods are particularly effective at dealing with systems with geometrically restricted entanglement, such as the use of matrix product states for one dimensional systems [55]. In recent years, tensor network techniques have become competitive with CASSCF methods for describing chemical systems with strong static correlation. They have been applied to metalloenzyme complexes with active spaces of over 70 spin-orbitals [56, 57], or Fermi-Hubbard models with around 100 sites [52]. Nevertheless, challenges still remain for applying tensor network techniques to 3D systems with long-range interactions, which can lead to a larger spread of entanglement in the system. For more information on the use of tensor network methods in quantum chemistry, I refer the reader to Refs. [58, 59].

There are a number of variants of Quantum Monte Carlo, and I will only consider the variant known as full configuration interaction Quantum Monte Carlo (FCIQMC) [60] here. This approach considers a set of ‘signed random walkers’ in the Fock space of occupation vectors, which spawn and annihilate stochastically according to an update rule that mimics propagation in imaginary time. The weight of walkers on a given determinant gives its amplitude in the wavefunction. As discussed in Chapter 4, propagating a system in imaginary time drives it towards its ground state. FCIQMC (and related variants) have been used to obtain FCI-level results for both molecular systems with tens of spin-orbitals [61, 62], and Fermi-Hubbard systems with around 100 sites [52]. Nevertheless, these methods have their own shortcomings, including the Monte Carlo ‘sign problem’ [63] (in the case of FCIQMC, this arises from needing to accurately resolve the sign of weakly present determinants, which may require a large number of ran-

dom walkers[64–66]). This leads to difficulties associated with the memory constraints of storing the coordinates of a large number of random walkers in some systems.

#### 2.1.4 Systems of interest

As discussed above, it has proven challenging to classically simulate systems with more than around 100 spin-orbitals that display both static and dynamic correlation. These features appear to arise in many systems of physical interest.

For example, the open-shell and spatially degenerate nature of many transition metals make them difficult to simulate classically [50]. These compounds form crucial parts of many catalysts, particularly transition-metal-based biological catalysts [42, 67]. Prominent examples that have caught the attention of the quantum computing community include: the FeMoco molecule present in the enzyme nitrogenase [27, 68, 69], which is responsible for biological nitrogen fixation at ambient temperatures and pressures [70, 71], and ruthenium catalysts for converting carbon dioxide to methanol [72]. Simulations of these systems could be performed with 100–200 logical qubits, although it has been noted that these calculations would be in small basis sets, that do not fully capture the dynamic correlation in the system [73].

Other archetypal examples of strongly correlated systems include high temperature superconductors. These systems are not well described by the Bardeen-Cooper-Schrieffer theory of superconductivity, and in many cases, the driver of superconductivity is still not fully understood. In the case of cuprate superconductors, it is widely believed that superconductivity is driven by the physics of the copper-oxygen planes in the system. Simplified models have been developed that attempt to capture the physics

of high temperature superconductors [74], including the widely studied Fermi-Hubbard model [75–77]. The Fermi-Hubbard model considers fermions hopping between nearest-neighbour lattice sites with strength  $t$ . These fermions feel a repulsive (or attractive) force  $U$  when they occupy the same lattice site,  $i$ . The Hamiltonian is given by

$$H = -t \sum_{\langle i,j \rangle, \sigma} \left( a_{i,\sigma}^\dagger a_{j,\sigma} + a_{j,\sigma}^\dagger a_{i,\sigma} \right) + U \sum_i n_{i,\uparrow} n_{i,\downarrow} \quad (2.14)$$

where  $\langle i, j \rangle$  denotes a sum over nearest-neighbour lattice sites, and  $\sigma$  is a spin-coordinate. This Hamiltonian has only  $\mathcal{O}(M)$  terms, where  $M$  is the number of spin-sites. Classical methods have not yet been able to completely solve this Hamiltonian. Close to half filling, at intermediate interaction strengths, it becomes difficult to accurately resolve the phase diagram of the system, due to multiple competing orders [78]. This is the same regime that is believed to be most relevant to understanding cuprate superconductors [52]. The most effective classical methods are able to accurately approximate the ground states of 2D Fermi-Hubbard models with up to around 100 lattice sites, far from the thermodynamic limit [52]. As a result, there have been a number of investigations into applying quantum algorithms to simulate the Fermi-Hubbard model [25, 79–84].

## 2.2 Mapping onto quantum computers

In this section I will discuss methods for mapping second quantised chemistry calculations onto quantum computers. The quantum algorithms used to solve the resulting problems will be discussed in Chapter 3. For a discussion of other ways of mapping the electronic structure problem onto quantum computers (e.g. real space or first quantised mappings), I refer the reader to Ref. [36].

### 2.2.1 Mapping to qubits

In this thesis I consider two ways of mapping the second quantised electronic structure Hamiltonian to qubits. The first is the Jordan-Wigner (JW) encoding [85], which maps Fock space occupation vectors directly onto computational basis states. More formally,

$$\begin{aligned} |f_{M-1}, f_{M-2}, \dots, f_0\rangle &\rightarrow |q_{M-1}, q_{M-2}, \dots, q_0\rangle, \\ q_p &= f_p \in \{0, 1\}. \end{aligned} \quad (2.15)$$

The fermionic creation and annihilation operators are mapped to

$$\begin{aligned} a_p &= \frac{1}{2}(X_p + iY_p) \otimes Z_{p-1} \otimes \dots \otimes Z_0, \\ a_p^\dagger &= \frac{1}{2}(X_p - iY_p) \otimes Z_{p-1} \otimes \dots \otimes Z_0. \end{aligned} \quad (2.16)$$

The string of  $Z$  operators recovers the exchange phase factor  $(-1)^{\sum_{i=0}^{p-1} f_i}$ . Consequently, the second quantised Hamiltonian in Eq. (2.9) is mapped to a linear combination of products of single-qubit Pauli operators

$$H = \sum_j h_j P_j = \sum_j h_j \bigotimes_i \sigma_i^j, \quad (2.17)$$

where  $h_j$  is a real scalar coefficient,  $\sigma_i^j$  represents one of  $I, X, Y$  or  $Z$ ,  $i$  denotes which qubit the operator acts on, and  $j$  denotes the term in the Hamiltonian. Each  $P_j$  term is referred to as a ‘Pauli string’, and the number of non-identity single-qubit Pauli operators in a given string is called its ‘Pauli weight’. The JW mapping produces Pauli strings with weight  $\mathcal{O}(M)$ .

The second mapping considered is the Bravyi-Kitaev (BK) mapping [86–88] and related BK-tree mapping [89]. Under these mappings, the qubits store partial sums of orbital occupation numbers. I show how this can be applied to the LiH molecule in Ref. [36]. These mappings again yield electronic

structure Hamiltonians that can be written as a linear combination of Pauli strings, each with weight  $\mathcal{O}(\log_2 M)$ . A thorough comparison of the BK and JW mappings was performed in Ref. [90] for 86 molecular systems.

### 2.2.2 Application to $\text{H}_2$

In Chapters 4 & 5 I perform electronic structure calculations on the  $\text{H}_2$  molecule, in a minimal STO-3G basis set. This basis set includes a single approximate Slater-type orbital, constructed from a linear combination of three Gaussian-type orbitals, for each Hydrogen atom. The four spin-orbitals are denoted

$$|1s_{A\uparrow}\rangle, |1s_{A\downarrow}\rangle, |1s_{B\uparrow}\rangle, |1s_{B\downarrow}\rangle, \quad (2.18)$$

where the subscript  $A$  or  $B$  denotes which of the two Hydrogen atoms the spin-orbital is centred on, and the  $\uparrow / \downarrow$  denotes the  $s_z$  value of the electron in the spin-orbital. Following convention, I work in the molecular orbital basis for  $\text{H}_2$

$$\begin{aligned} |\sigma_{g\uparrow}\rangle &= \frac{1}{\sqrt{N_g}}(|1s_{A\uparrow}\rangle + |1s_{B\uparrow}\rangle), \\ |\sigma_{g\downarrow}\rangle &= \frac{1}{\sqrt{N_g}}(|1s_{A\downarrow}\rangle + |1s_{B\downarrow}\rangle), \\ |\sigma_{u\uparrow}\rangle &= \frac{1}{\sqrt{N_u}}(|1s_{A\uparrow}\rangle - |1s_{B\uparrow}\rangle), \\ |\sigma_{u\downarrow}\rangle &= \frac{1}{\sqrt{N_u}}(|1s_{A\downarrow}\rangle - |1s_{B\downarrow}\rangle), \end{aligned} \quad (2.19)$$

where  $N_{g/u}$  are normalisation factors that depend on the overlap between the atomic orbitals,  $N_g = 2(1 + \langle 1s_A | 1s_B \rangle)$ ,  $N_u = 2(1 - \langle 1s_A | 1s_B \rangle)$ . In the second quantised formalism, the occupation vector describing the Slater determinants that can be created with these spin-orbitals is given by

$$|\psi\rangle = |f_{\sigma_{u\downarrow}}, f_{\sigma_{u\uparrow}}, f_{\sigma_{g\downarrow}}, f_{\sigma_{g\uparrow}}\rangle. \quad (2.20)$$

I calculated the integrals that determine the second quantised Hamiltonian using OpenFermion [34]. The JW mapped Hamiltonian for 4 qubit  $H_2$  is given by

$$\begin{aligned}
H = & h_0 I + h_1 Z_0 + h_2 Z_1 + h_3 Z_2 + h_4 Z_3 \\
& + h_5 Z_0 Z_1 + h_6 Z_0 Z_2 + h_7 Z_1 Z_2 + h_8 Z_0 Z_3 + h_9 Z_1 Z_3 \\
& + h_{10} Z_2 Z_3 + h_{11} Y_0 Y_1 X_2 X_3 + h_{12} X_0 Y_1 Y_2 X_3 \\
& + h_{13} Y_0 X_1 X_2 Y_3 + h_{14} X_0 X_1 Y_2 Y_3,
\end{aligned} \tag{2.21}$$

where  $h_i$  are the coefficients determined by the Hamiltonian integrals. In the JW encoding, the Hartree-Fock state for  $H_2$  is given by

$$|\Psi_{\text{HF}}^{H_2}\rangle = |0011\rangle. \tag{2.22}$$

This occupation vector represents the Slater determinant

$$\Psi_{\text{HF}}^{H_2}(\mathbf{r}_1, \mathbf{r}_2) = \frac{1}{\sqrt{2}}(\sigma_{g\uparrow}(\mathbf{r}_1)\sigma_{g\downarrow}(\mathbf{r}_2) - \sigma_{g\uparrow}(\mathbf{r}_2)\sigma_{g\downarrow}(\mathbf{r}_1)), \tag{2.23}$$

where  $\mathbf{r}_i$  is the position of electron  $i$ . This is the model for  $H_2$  that I investigate in Chapter 5.

In Chapter 4, I instead consider a BK mapping of the  $H_2$  Hamiltonian, which gives

$$\begin{aligned}
H = & h'_0 I + h'_1 Z_0 + h'_2 Z_1 + h'_3 Z_2 + h'_4 Z_0 Z_1 + h'_5 Z_0 Z_2 + h'_6 Z_1 Z_3 + h'_7 X_0 Z_1 X_2 \\
& + h'_8 Y_0 Z_1 Y_2 + h'_9 Z_0 Z_1 Z_2 + h'_{10} Z_0 Z_2 Z_3 + h'_{11} Z_1 Z_2 Z_3 + h'_{12} X_0 Z_1 X_2 Z_3 \\
& + h'_{13} Y_0 Z_1 Y_2 Z_3 + h'_{14} Z_0 Z_1 Z_2 Z_3.
\end{aligned} \tag{2.24}$$

We can see that this Hamiltonian only acts off-diagonally (i.e. with an operator that is not  $I$  or  $Z$ ) on qubits 0 and 2 [91–93]. As a result, by following

the procedure detailed in Ref. [93] it can be reduced to a two-qubit Hamiltonian

$$H = g_0 I + g_1 Z_0 + g_2 Z_1 + g_3 Z_0 Z_1 + g_4 Y_0 Y_1 + g_5 X_0 X_1. \quad (2.25)$$

### 2.2.3 Application to LiH

In Chapter 4 I consider simulations of the Lithium Hydride (LiH) molecule, in a minimal STO-3G basis. This basis contains  $\{1s, 2s, 2p_x, 2p_y, 2p_z\}$  functions for Lithium, and a single  $\{1s\}$  orbital for Hydrogen – a total of 12 spin-orbitals. I reduced the resources required by considering only the most relevant orbitals for LiH – known as an ‘active space’. This was achieved by first calculating the one electron reduced density matrix (1-RDM) for LiH, using a classically tractable configuration interaction with single and double excitations reference state. The 1-RDM with respect to a state  $|\phi\rangle$  is given by  $\rho_{ij}^1 = \langle \phi | a_i^\dagger a_j | \phi \rangle$ . For LiH at an internuclear distance of 1.45 Å, the 1-RDM (combining the occupancies of spin-up and down orbitals) is given by

$$\begin{pmatrix} 1.99991 & -0.00047 & 0.00047 & 0 & 0 & -0.00120 \\ -0.00047 & 1.95969 & 0.06691 & 0 & 0 & 0.00842 \\ 0.00047 & 0.06691 & 0.00968 & 0 & 0 & -0.01385 \\ 0 & 0 & 0 & 0.00172 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.00172 & 0 \\ -0.00120 & 0.00842 & -0.01385 & 0 & 0 & 0.02728 \end{pmatrix}. \quad (2.26)$$

The diagonal elements of the 1-RDM are the occupation numbers of the corresponding canonical Hartree-Fock orbitals. To determine the active space,

I diagonalised the Hermitian 1-RDM

$$\begin{pmatrix} 1.99992 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.96201 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.03459 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.00005 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.00172 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.00172 \end{pmatrix}, \quad (2.27)$$

to give the 1-RDM in terms of natural molecular orbitals (NMOs). The diagonal entries are called the natural orbital occupation numbers (NOONs). I used the same unitary matrix that diagonalised the 1-RDM to rotate the LiH orbitals. This was equivalent to performing a change of basis from the canonical orbitals to the NMOs. In the NMO basis, the first orbital had a NOON close to two, and so could be considered doubly occupied. I removed any terms containing  $a_0^\dagger, a_0, a_1^\dagger, a_1$  from the LiH fermionic Hamiltonian, where 0 and 1 denote the spin-orbitals that correspond to the core spatial orbital. Similarly, the fourth spatial orbital had a NOON close to zero, and so could be considered unoccupied. I removed the two corresponding fermionic operators from the Hamiltonian. This yielded a fermionic Hamiltonian acting on 8 spin-orbitals. I mapped this fermionic Hamiltonian to a qubit Hamiltonian using the JW transformation. This gave the 8 qubit LiH Hamiltonian investigated in Chapter 4. All of these steps were performed using OpenFermion.

## 3 | Quantum simulation

### Contents

---

<b>3.1 Quantum simulation algorithms</b> . . . . .	<b>26</b>
3.1.1 Time evolution . . . . .	27
3.1.2 Quantum phase estimation . . . . .	31
3.1.3 Variational quantum eigensolver . . . . .	33
<b>3.2 Quantum error mitigation and correction</b> . . . . .	<b>40</b>
3.2.1 Error mitigation . . . . .	41
3.2.2 Error correction . . . . .	45

---

In this chapter, I discuss some of the main quantum simulation techniques that are referenced and utilised in this thesis. In Sec. 3.1, I introduce several quantum simulation algorithms: time evolution methods (Sec. 3.1.1), quantum phase estimation (Sec. 3.1.2), and the variational quantum eigensolver (Sec. 3.1.3). I then discuss techniques that can be used to protect quantum simulations from noise in Sec. 3.2. The information presented in this chapter is in part adapted from Ref. [36]:

*Quantum computational chemistry*  
S. McArdle, S. Endo, A. Aspuru-Guzik, S.C. Benjamin, & X. Yuan  
[Reviews of Modern Physics 92 \(1\), 015003 \(2020\)](#)

### 3.1 Quantum simulation algorithms

In this section I discuss quantum simulation algorithms that feature throughout this thesis. As discussed in Chapter 2, finding the low-lying energy levels of chemical systems is a widely anticipated application of future quantum computers. The first quantum algorithms proposed to solve this problem used subroutines that evolve the chemical system in time. I discuss a number of time evolution algorithms in Sec. 3.1.1. I then discuss in Sec. 3.1.2 how these subroutines are incorporated into the quantum phase

estimation algorithm, which can be used to find the eigenvalues and eigenvectors of chemical Hamiltonians. Unfortunately, initial resource estimates showed that quantum phase estimation would likely necessitate the use of deep circuits. The algorithm must then be protected with quantum error correction, to prevent noise from corrupting the results of the simulation. An alternative solution to this problem was put forwards in Ref. [94], which introduced a new quantum algorithm for finding ground states; the variational quantum eigensolver (VQE). As I discuss in Sec. 3.1.3, the VQE involves approximating the ground state using a short depth circuit, and then directly measuring the expectation value of the Hamiltonian, term by term. By reducing the circuit depth required (compared to quantum phase estimation) the VQE seeks to reduce the level of noise, such that quantum error correction may not be required. Whether this goal can be achieved is still an open question, which I discuss in Chapter 5 of this thesis.

### 3.1.1 Time evolution

Algorithms to propagate a system in time are typically referred to as performing ‘Hamiltonian simulation’. The first such algorithm, by Lloyd [23], used a product formula approximation (also referred to as Trotterization) to divide the propagator into a sequence of operators with a known decomposition into quantum hardware primitives. This algorithm was specialised to the simulation of fermionic Hamiltonians in Ref. [95], and has received considerable attention in the years since. Many works have derived tighter bounds on the simulation error arising from Trotterization, as summarised in Ref. [96]. Below, I report Trotter error bounds from Ref. [96] that are pertinent to the work presented in this thesis. I consider the case where the Hamiltonian can be decomposed into a sum of Pauli strings, as  $H = \sum_{\alpha} h_{\alpha} H_{\alpha}$ , where  $h_{\alpha}$  are real coefficients. A first-order Trotter decom-

position of the time evolution operator is given by

$$U_1(t) := \left( \prod_j e^{-\frac{it}{n} h_j H_j} \right)^n, \quad (3.1)$$

where  $n$  is the number of Trotter steps used. The error in the first-order product formula is upper bounded by

$$\|U(t) - U_1(t)\| := \epsilon \sim \mathcal{O}\left(\frac{(L\Lambda t)^2}{n}\right) \quad (3.2)$$

where  $U(t)$  is the true time evolution operator,  $\| \cdot \|$  denotes the spectral norm,  $L$  is the number of terms in  $H$ , and  $\Lambda = \max_\alpha \|h_\alpha\|$ . A tighter bound on the error is given by

$$\epsilon \sim \frac{t^2}{2n} \sum_i \left\| \sum_{j>i} [h_j H_j, h_i H_i] \right\|, \quad (3.3)$$

which is known to be tight, up to an application of the triangle inequality. Higher-order product formulae can be used to obtain improved error-scaling. The second-order product formula is given by

$$U_2(t) := \left( \prod_{j=1}^L e^{-\frac{it}{2n} h_j H_j} \prod_{k=L}^1 e^{-\frac{it}{2n} h_k H_k} \right)^n. \quad (3.4)$$

An upper-bound on the second-order Trotter error is given by

$$\|U(t) - U_2(t)\| := \epsilon \sim \mathcal{O}\left(\frac{(L\Lambda t)^3}{n^2}\right), \quad (3.5)$$

and a tighter bound (tight up to an application of the triangle inequality) is given by

$$\begin{aligned} \epsilon \sim \frac{t^3}{n^2} & \left[ \frac{1}{12} \left( \sum_i \left\| \sum_{j>i} \sum_{k>i} [h_k H_k, [h_j H_j, h_i H_i]] \right\| \right) \right. \\ & \left. + \frac{1}{24} \left( \sum_i \left\| \sum_{j>i} [h_i H_i, [h_i H_i, h_j H_j]] \right\| \right) \right]. \end{aligned} \quad (3.6)$$

It has also been shown that introducing aspects of randomized compilation into these algorithms can lower the gate counts required to obtain results of a given accuracy. For example, randomly permuting the ordering of the product formula terms in each step can reduce the error scaling obtained [97].

Alternative time evolution algorithms have been developed that can, in some cases, perform time evolution asymptotically more efficiently than Trotterization. In particular, the number of gates required for these ‘post-Trotter’ methods scales poly-logarithmically with the inverse of the desired simulation error. One such approach is to expand the propagator as a truncated Taylor series [98–100]. There are a number of ways that this approach can be realised. For a Hamiltonian that can be expanded as a linear combination of unitary operators (e.g. as a sum of Pauli strings), the terms in its truncated Taylor expansion will also be unitary operators. Oracle circuits (which must be explicitly constructed for the problem at hand) can be used to implement these unitary operators in superposition, thus realising time evolution under the Hamiltonian. Alternatively, one can develop oracle circuits that give access to the non-zero elements of sparse Hamiltonians, and use these to realise the Taylor expansion. Because a sum of unitary operators is not necessarily unitary, this algorithm will have a non-zero failure probability. This can be controlled using techniques such as amplitude amplification [101]. Another post-Trotter method is the combination of qubitization [102] with quantum signal processing [103, 104]. Qubitization provides access to the Hamiltonian via a ‘block encoding’, and implements a quantum walk operator with eigenvalues  $e^{-i\arcsin(E_k/\lambda)}$  where  $E_k$  is the  $k^{\text{th}}$  eigenvalue of the Hamiltonian  $H$  and  $\lambda$  is the 1-norm of the Hamiltonian. Quantum signal processing can be used to invert the arcsin function, recovering the desired evolution  $e^{iHt}$ .

A recent approach to Hamiltonian simulation, known as qDRIFT [105], can be viewed through either a Trotter or post-Trotter lens. The technique probabilistically selects a number of terms from the Hamiltonian according to their strength, and then evolves under these terms (each for an equal duration in time). In this way, qDRIFT resembles a randomized product formula, where all terms are applied with the same strength, and not every Hamiltonian term is necessarily included. However, qDRIFT can also be seen as a kind of ‘incoherent post-Trotter’ approach<sup>1</sup>. The error in qDRIFT depends on the 1-norm of the Hamiltonian, much like post-Trotter methods. A single step of qDRIFT is equivalent to implementing the channel

$$\mathcal{E}(\rho) = \sum_j \frac{h_j}{\lambda} e^{-i\frac{\lambda t}{N} H_j} \rho e^{i\frac{\lambda t}{N} H_j}, \quad (3.7)$$

where  $\lambda = \sum_j h_j$  is the 1-norm of the Hamiltonian (when applying qDRIFT, we shift the signs from the  $h_j$  coefficients to the  $H_j$  operators, such that  $h_j$  are all real and positive), and  $N$  is the number of qDRIFT steps used in this simulation. The error scaling of qDRIFT is given by

$$\epsilon \sim \frac{\lambda^2 t^2}{N}. \quad (3.8)$$

Most notably, the error scaling of qDRIFT is independent of the number of terms in the Hamiltonian, making it an interesting candidate for systems with a large number of weakly interacting terms. Techniques have also been developed that interpolate between Trotterization and qDRIFT [106].

<sup>1</sup>As discussed by N. Wiebe: <https://www.youtube.com/watch?v=mP7HHWk0H5M>

### 3.1.2 Quantum phase estimation

Quantum phase estimation can be used to obtain the eigenvectors and eigenvalues of a given unitary operator [107]. In its canonical form, phase estimation evolves a given state under the specified unitary operator, controlled upon an ancillary register. This process is repeated for evolutions of varying duration, to learn the phase induced by the evolution, which can be related to the eigenvalues of the operator. If the state used is an eigenstate of the unitary operator, then this approach calculates the corresponding eigenvalue. If the state is a superposition of several eigenstates, then this approach will output a superposition of eigenvalue-eigenstate pairs. Measuring the ancillary register causes this superposition to collapse. The ancillary register then yields the eigenvalue, while the main register will have collapsed to the corresponding eigenstate.

In Ref. [108], it was shown that using the time evolution operator as the black-box unitary enables us to learn the eigenvalues of the Hamiltonian that generates the time evolution. This idea was extended to finding the ground state energies of electronic structure Hamiltonians in Ref. [24]. Using the canonical approach to phase estimation presented in Ref. [35], the total evolution time required to achieve an error of  $\epsilon$  in the eigenvalue estimate is  $T = \frac{8\pi}{\epsilon}$ . This assumes a success probability of 50%, and neglects circuit synthesis and unitary approximation errors. In electronic structure calculations, we typically seek  $\epsilon \sim 10^{-3} - 10^{-4}$ . As a result, even if each unit of time evolution could be implemented with a single gate, we would need thousands of gates to realise phase estimation. In reality, we will need many gates to implement each timestep. Phase estimation has been experimentally demonstrated for a number of toy chemical systems, including on: trapped ions [109], superconducting qubits [91], nitrogen-vacancy centre qubits [110], silicon photonic processors [111, 112], and nuclear mag-

netic resonance devices [113, 114].

A number of studies have investigated the resources required to implement phase estimation for molecular electronic structure Hamiltonians. Initial investigations used Trotterization to implement the time evolution operator, and found that in a Gaussian orbital basis set,  $\mathcal{O}(M^{11})$  gates would be required [115]. While this was later reduced to  $\mathcal{O}(M^{5.5})$  gates [116], the high gatecount necessitates the use of quantum error correction. In Ref. [68], it was shown that when considering the overhead of quantum error correction, applying phase estimation to a 108 spin-orbital active space of the FeMoco molecule would require around 200 million physical qubits (assuming physical error rates of  $10^{-3}$ ), and take months to implement (using an optimistic assumption of 100 ns to perform each T gate).

Subsequent investigations have further reduced the complexity of phase estimation, using a number of techniques. Replacing Trotterization with the Taylor series approach to time evolution resulted in asymptotically more efficient algorithms (although potentially with a larger constant factor overhead) [117–119]. It was later observed that using qubitization directly (i.e. without using quantum signal processing to fix the phases), would learn the eigenvalues of  $\arcsin(H)$  [120, 121]. As this operator is isospectral to  $H$ , it is possible to apply the sine function to the measured eigenvalue, to recover the desired energy value. This technique has since been used to reduce the number of gates required to perform phase estimation on electronic structure Hamiltonians, both for plane wave basis sets [26], and for Gaussian basis sets [27, 69, 72]. These latter works make use of efficient tensor decompositions of the Coulomb operator to reduce the gate complexity of implementing the qubitization oracle. Using the most efficient of these approaches, applying phase estimation to a 152 spin-orbital active space of the FeMoco molecule would require around 4 million physical qubits (assuming physical error rates of  $10^{-3}$ ), and take four days to implement

(assuming a more realistic surface code cycle time of  $10\ \mu\text{s}$ ) [27]. Despite these significant improvements, it is clear that the deep circuits required for quantum phase estimation will require the availability of fault-tolerant quantum hardware. As a result, it is interesting to consider if there are algorithms for finding the ground states of chemical systems that do not require error correction.

### 3.1.3 Variational quantum eigensolver

The VQE in its canonical form is outlined in Fig. 3.1. The algorithm was proposed and experimentally demonstrated on a photonic processor in Ref. [94], and later elaborated upon in Ref. [122]. The VQE is a hybrid quantum-classical algorithm, and typically exchanges an increased number of circuit repetitions for a reduced circuit depth, compared to quantum phase estimation. At the core of the algorithm is the Rayleigh-Ritz variational principle, which states that for a parameterized trial wavefunction

$$\langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle \geq E_0, \quad (3.9)$$

where  $E_0$  is the lowest energy eigenvalue of the Hamiltonian. As a result, if we are able to efficiently prepare good trial states, measure their energy, and search over the parameter space, we can attempt to find an optimal approximation to the ground state. We use a quantum computer to prepare trial states and measure their energy – with the justification that a quantum processor may be able to prepare more powerful trial states than a classical algorithm. A classical optimisation routine is then used to optimise the system parameters, based on their previous values, and measurements of the energy of the trial state (or related quantities, such as energy gradients).

The trial state is generated using a parameterized quantum circuit known as the ‘ansatz’ circuit. I denote the ansatz circuit as  $U(\vec{\theta}) = U_N(\theta_N) \dots U_k(\theta_k) \dots U_1(\theta_1)$ .

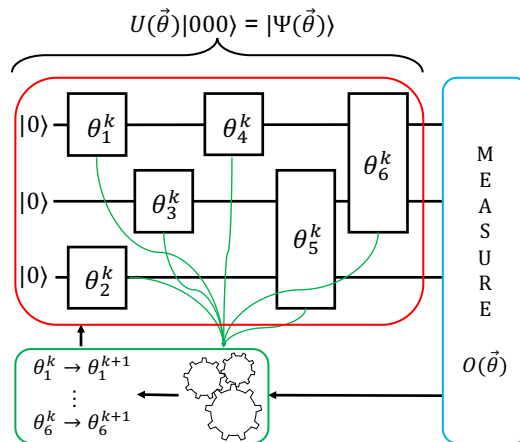


Figure 3.1: A schematic of the VQE. The VQE attempts to find the ground state of a given problem Hamiltonian, by classically searching for the optimal parameters  $\vec{\theta}$  which minimise  $\langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle$ . The state preparation and measurement subroutines (red, upper left, and blue, right) are performed on the quantum computer. The measured observable  $O(\vec{\theta})$  and parameter values are fed into a classical optimisation routine (green, lower), which outputs new values of the parameters. The new parameters are then fed back into the ansatz circuit. The ansatz circuit can include both parameterized and non-parameterized gates. This loop is repeated until the energy converges. Reproduced with permission from Ref. [36].

Here,  $U_k(\theta_k)$  is the  $k^{\text{th}}$  single- or two-qubit unitary gate, controlled by parameter  $\theta_k$ . Not every gate in the circuit needs to be parameterized, and parameters can be set to a fixed value as desired. The trial state is then given by  $|\Psi(\vec{\theta})\rangle = U(\vec{\theta})|\bar{0}\rangle$ . I discuss some of the ansatz circuits considered in this thesis in Sec. 3.1.3.1.

In order to measure the energy, or related quantities such as the energy gradient, a procedure known as ‘Hamiltonian averaging’ is used. As discussed in Sec. 2.2.1, the Jordan-Wigner and Bravyi-Kitaev mappings convert second quantised fermionic Hamiltonians into a weighted sum of Pauli strings. We can calculate the energy of a trial state by summing the expectation values of each Pauli string in the Hamiltonian (multiplied by its corresponding

coefficient). To obtain each expectation value, we repeatedly reinitialise the processor, and prepare and measure the trial state. More formally, for a Hamiltonian decomposed into a sum of Pauli strings,  $H = \sum_{\alpha} h_{\alpha} H_{\alpha}$

$$E(\vec{\theta}_k) = \sum_j^N h_j \langle \Psi(\vec{\theta}_k) | H_j | \Psi(\vec{\theta}_k) \rangle. \quad (3.10)$$

Refs. [122, 123] showed that the number of measurements  $N_m$ , required to estimate the energy to a precision  $\epsilon$ , is bounded by

$$N_m = \frac{(\sum_i |h_i|)^2}{\epsilon^2}, \quad (3.11)$$

This leads to a scaling of  $\mathcal{O}(M^6/\epsilon^2)$  measurements in a Gaussian basis set [124, 125]. A number of techniques have recently been developed to reduce the large number of measurements that may be required to simulate realistic systems of interest. These techniques typically group together terms that commute, and are discussed in more detail in Ref. [36].

As discussed above, a classical optimisation technique is used to update the parameters in the ansatz circuit. A large number of different approaches have been introduced, seeking to find the global optimum of the system with either fewer measurements, and/or greater noise resilience. A selection of these results are reviewed in Refs. [36, 122, 123, 126]. These include direct search methods such as: the Nelder-Mead simplex algorithm, simulated annealing, or particle-swarm optimisation. A number of gradient-based methods have also been applied, including: vanilla gradient descent, L-BFGS-B (an approximate Hessian method), and ADAM [127]. In Chapter 4 I introduce an alternative approach to optimising parameterized quantum circuits: ansatz-based imaginary time evolution. A key consideration for any optimisation routine is what the starting values of the parameters

should be. It was shown in Ref. [128] that when a randomly initialised ansatz circuit forms a 2-design (i.e. the ansatz matches the Haar distribution to the second moment), a given observable will concentrate to a parameter-independent value, which causes energy gradients that are essentially zero in all directions in parameter space. This is known as the ‘barren plateau’ problem, and was originally shown numerically for ‘hardware efficient’-type ansätze (see Sec. 3.1.3.1) composed of single qubit rotations on each qubit, followed by a ladder of entangling gates, repeated in a number of layers [128]. It was originally claimed that more structured ansätze, such as the unitary coupled cluster ansatz (see Sec. 3.1.3.1) may also be able to form a 2-design for systems with more than 10 spin-orbitals, if the parameters were completely randomised [128]. However, subsequent work [129] has suggested that the structured unitary coupled cluster and Hamiltonian variational ansätze discussed in Sec. 3.1.3.1 are ‘usually trainable even when randomly initialized’. Nevertheless, it is important to bear in mind that this could be an artefact of the modest system sizes classically simulated to date, and that it may be prudent to consider analytically if such ansätze do indeed fulfil the conditions required to be a 2-design. A number of methods have been suggested for mitigating the barren plateau problem, including: using structured ansätze with good initial guesses [128, 130], training circuits in a layerwise fashion [131, 132], or correlating related parameters in the ansatz [133].

The VQE has been experimentally demonstrated on many platforms, including: photonic systems [94], trapped ion processors [134–137], and superconducting qubits [91, 92, 138, 139]. While the VQE has enabled accurate calculations of ground state energies of systems with up to twelve qubits, it is still unclear if it will be able to surpass classical algorithms. In particular, the ansatz circuits used need to be deep enough to create a

powerful trial state, but not so deep that noise corrupts the calculation. Moreover, the number of measurements required may be very large, leading to an unfeasibly long runtime, if the algorithm is not parallelised across multiple processors [80].

### 3.1.3.1 Ansätze

In this thesis I utilise three different ansatz circuits for approximating ground states on quantum computers. In Chapter 4, Sec. 4.4.1 I utilise ‘hardware efficient’ ansätze [92] for the simulation of the  $H_2$  and LiH molecules. These ansätze are composed of repeated, dense circuit blocks, and seek to build a flexible trial state using as few gates as possible. A major limitation of hardware efficient ansätze is that it is unclear how to best initialise their parameters in such a way that barren plateaus are avoided. In addition, as the circuit is not tailored to the chemical problem at hand, it may be less effective than a more specialised ansatz circuit.

In Chapter 4, Sec. 4.4.2 I use an alternative ansatz circuit, known as the Hamiltonian variational ansatz [140], to investigate the Fermi-Hubbard model. This ansatz was inspired by adiabatic state preparation and the quantum approximate optimisation algorithm (a quantum-classical hybrid algorithm for combinatorial optimisation problems that is similar to the VQE [141]). The idea is to Trotterize an adiabatic evolution to the ground state, using a number of Trotter steps that may be insufficient for accurate results. One can then variationally optimise the Trotter evolution times to create an ansatz for the ground state. The ansatz is given by

$$U(\vec{\theta}) = \prod_s \prod_j e^{i\theta_j^s P_j} \quad (3.12)$$

where  $P_j$  are the Pauli strings in the Hamiltonian, and  $S$  is the total number

of steps used. Previous work has shown that this ansatz can prepare good approximations to the ground state of the Fermi-Hubbard model using a modest number of layers [81, 84, 140], and that each Trotter step can be implemented with a low circuit depth [80–83].

In Chapter 5 I use a unitary coupled cluster (UCC) ansatz to simulate the  $H_2$  molecule. This ansatz circuit is inspired by the classical coupled-cluster (CC) method described in Sec. 2.1.3.2. While the CC method is non-variational, and struggles to converge when applied to multireference states, the UCC ansatz is able to overcome these limitations. As with the CC method, the UCC ansatz considers excitations above an initial reference state

$$U(\vec{\theta}) = e^{T-T^\dagger}, \quad (3.13)$$

where  $T = \sum_i T_i$ , and

$$\begin{aligned} T_1 &= \sum_{i \in \text{virt}, \alpha \in \text{occ}} t_{i\alpha} a_i^\dagger a_\alpha, \\ T_2 &= \sum_{i, j \in \text{virt}, \alpha, \beta \in \text{occ}} t_{ij\alpha\beta} a_i^\dagger a_j^\dagger a_\alpha a_\beta, \\ &\dots \end{aligned} \quad (3.14)$$

and *occ* are occupied orbitals in the reference state, and *virt* are orbitals that are initially unoccupied in the reference state. The UCC ansatz was originally proposed as a classical technique [142, 143], but is not efficient to simulate on classical computers, as  $e^{T-T^\dagger}$  has a non-terminating Baker-Campbell-Hausdorff expansion. However, the ansatz has an efficient implementation on quantum hardware – as first noted by Ref. [94]. A comprehensive review of the UCC method is given in Ref. [123]. One approach to implementing the UCC operator is through Trotterization. It is currently unclear as to whether the optimisation step of the VQE is able to compen-

sate for the Trotter error introduced, and different conclusions have been reported in the literature for strongly correlated systems [123, 144, 145].

For the simulations of the H<sub>2</sub> molecule presented in Sec. 5.4, I ordered the spin-orbitals as

$$|f_{\sigma_{u\downarrow}}, f_{\sigma_{g\downarrow}}, f_{\sigma_{u\uparrow}}, f_{\sigma_{g\uparrow}}\rangle. \quad (3.15)$$

The most general state (with the same spin and electron number as the Hartree-Fock state) for the 4 qubit Jordan-Wigner mapped H<sub>2</sub> molecule is given by

$$|\psi^{\text{H}_2}\rangle = \alpha |0101\rangle + \beta |1010\rangle + \gamma |1001\rangle + \delta |0110\rangle. \quad (3.16)$$

This state can be generated using a singlet UCC ansatz, considering single and double excitations above the Hartree-Fock state (UCCSD). For H<sub>2</sub>, the only operators which do not change the  $s_z$  value of the molecule when acting upon the Hartree-Fock state are:  $a_1^\dagger a_0$ ,  $a_3^\dagger a_2$ ,  $a_3^\dagger a_1^\dagger a_2 a_0$ . Other valid operators are equivalent to these operators, and can be combined with them. The singlet UCCSD operator then takes the form

$$U = e^{t_{10}(a_1^\dagger a_0 - a_0^\dagger a_1) + t_{32}(a_3^\dagger a_2 - a_2^\dagger a_3) + t_{3120}(a_3^\dagger a_1^\dagger a_2 a_0 - a_0^\dagger a_2^\dagger a_1 a_3)}. \quad (3.17)$$

Using the Jordan-Wigner encoding gives

$$\begin{aligned} (a_1^\dagger a_0 - a_0^\dagger a_1) &= \frac{i}{2}(X_1 Y_0 - Y_1 X_0) \\ (a_3^\dagger a_2 - a_2^\dagger a_3) &= \frac{i}{2}(X_3 Y_2 - Y_3 X_2) \\ (a_3^\dagger a_1^\dagger a_2 a_0 - a_0^\dagger a_2^\dagger a_1 a_3) &= \frac{i}{8}(X_3 X_2 Y_1 X_0 + Y_3 X_2 X_1 X_0 + Y_3 Y_2 Y_1 X_0 + Y_3 X_2 Y_1 Y_0 \\ &\quad - X_3 Y_2 Y_1 Y_0 - Y_3 Y_2 X_1 Y_0 - X_3 Y_2 X_1 X_0 - X_3 X_2 X_1 Y_0). \end{aligned} \quad (3.18)$$

I split the UCCSD operator using a single Trotter step, resulting in a product of exponentiated Pauli strings. Each of these terms was implemented using

the decomposition given in Ref. [35]. The resulting singlet UCCSD ansatz circuit for  $H_2$  contains 92 single-qubit gates and 56 two-qubit gates.

### 3.2 Quantum error mitigation and correction

So far, I have considered noiseless quantum circuits, which can be described with unitary operators applied to pure state wavefunctions. In order to treat noise in its fully generality, it becomes necessary to consider quantum circuits as being represented by completely positive trace preserving maps, known as quantum channels [35]. These channels are applied to density operators, which can represent either pure or mixed states. Deviation from the intended unitary dynamics of the circuit applied can occur for a number of reasons, which I loosely divide into two categories: decoherence caused by interaction with the environment, and control errors. In reality, there will be significant overlap between these two categories. Accounting for these errors from first principles modelling is a challenging and complex task [146]. This motivates the use of simple models, designed to capture the effects of common noise channels. For example, one could approximate some systematic control errors as coherent over-rotations – as might result from a miscalibrated microwave pulse. Such an error model could be described by operations like  $R_z(\theta) \rightarrow R_z(\theta + \delta)$ . Alternatively, we can consider stochastic error models that account for Markovian, incoherent errors in a system. A simple and commonly used error model, that features throughout this thesis, is the depolarising channel. The single-qubit depolarising channel is given by

$$\mathcal{E}(\rho) = (1 - p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z), \quad (3.19)$$

where  $X, Y, Z$  are the Pauli gates applied to the noisy qubit, and  $p$  is the probability of an error occurring. The two-qubit depolarising channel is

given by

$$\mathcal{E}(\rho) = (1 - p)\rho + \frac{p}{15} \sum_{i,j} (O_{i,1} O_{j,2} \rho O_{i,1}^\dagger O_{j,2}^\dagger), \quad (3.20)$$

where  $O_{i,q}$  is the operator  $O_i$  acting on qubit  $q$ , and  $O_i$  runs over  $I, X, Y, Z$ , and the possibility  $O_i = O_j = I$  is not included (this possibility, of no error occurring, is described by the term outside of the sum). Note that the two-qubit depolarising channel is not the tensor product of two single-qubit depolarising channels, and thus represents a correlated noise model. In some scenarios, depolarising noise has been found to provide a reasonable approximation for the noise present in real devices – for example, in the case of random circuits running on superconducting processors [147]. As discussed above, we can think of the value  $p$  as being the error rate of the hardware in such a noise model. Typical values of  $p$  for existing hardware range from  $\mathcal{O}(10^{-2})$  [148] to  $\mathcal{O}(10^{-3})$  [147, 149, 150], but have not been improved below this threshold since it was first achieved.

### 3.2.1 Error mitigation

If we consider a stochastic and Markovian error model, such as the depolarising model discussed above, then one interpretation of the model is that after each gate there is an independent chance of an error occurring. If we consider a circuit with 1000 gates, each with an error rate of  $p = 10^{-3}$ , then on average, we would expect one error to occur each time the circuit is run. Nevertheless, on some executions of the circuit, we would expect no errors to occur. If our algorithm is one where the results are averaged across multiple circuit executions (as with the VQE), then it is desirable to extract the noiseless signal from the noisy data. This is the goal of error mitigation techniques. These techniques can be distinguished from quantum error correction, in that they do not correct errors, and so are not scalable techniques. If too many noisy gates are applied, then the noiseless signal will

be negligibly small. A typical rule of thumb is that on the order of one error per circuit is an appropriate target for error mitigation techniques [30, 151]. Similarly, unlike dynamical decoupling or spin echos, the approaches discussed below do not improve the error rate of the gates applied. Nevertheless, these techniques *effectively* reduce the error rate of the calculation, by combining the results of an increased number of calculations in such a way that the noiseless component of the signal is amplified.

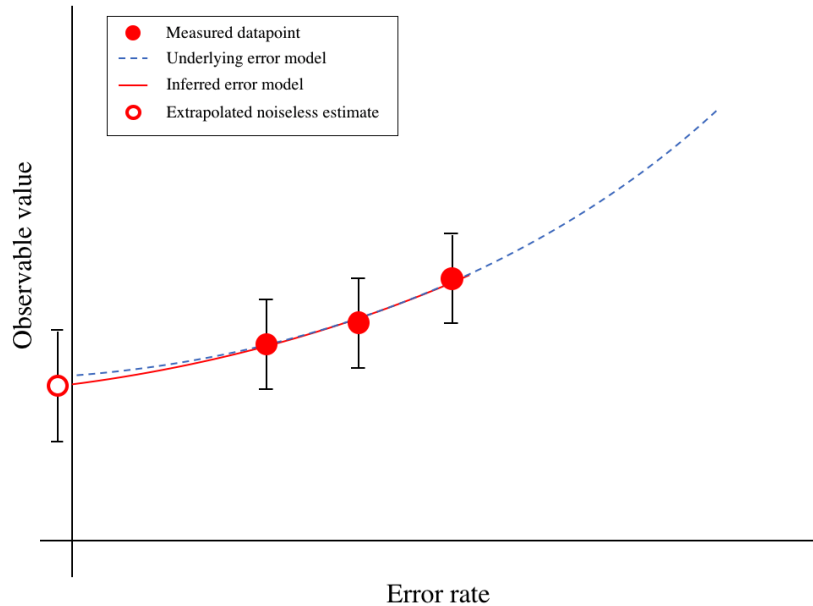


Figure 3.2: The extrapolation method of error mitigation. The noise in the system is deliberately increased to measure the value of the observable at a number of error rates. We infer an underlying error model based on these results, and use this to predict the noiseless observable value. This extrapolation increases the variance of the observable expectation value.

One of the first error mitigation techniques developed was the extrapolation technique [29, 152]. The motivation for error extrapolation is that the expectation value of the measured observable should change smoothly with the magnitude of noise in the system. If one source of error can be

considered as ‘dominant’, then we can imagine a curve plotted out by the observable value as a function of the error rate, as shown in Fig. 3.2. By deliberately ‘boosting’ the level of noise in the system, we can measure a number of data points along this underlying curve. We can then try to devise a model for the data, and use this to model to estimate the noiseless value of the observable. Initial proposals considered a linear extrapolation, motivated by a Taylor expansion of the observable value [29, 152]. Subsequent works have shown that better fits can be obtained by using exponential extrapolations [30], or multi-exponential functions [151, 153]. There are a number of possible approaches to boosting the dominant noise rate. One initial proposal was to convert the noise into a stochastic Pauli channel through Pauli twirling, and then to ‘drop in’ additional Pauli errors with the appropriate probabilities [152]. Subsequent works have considered physically stretching the duration of gates implemented via microwave pulses [138], or by inserting gate-inverse identity resolutions after each gate (e.g. replacing a CNOT gate by three CNOTs) [154]. I employ the linear extrapolation technique in Chapter 5, and exponential extrapolation in Chapter 7.

An alternative error mitigation method known as the ‘quasiprobability technique’, was introduced at the same time as error extrapolation [29]. The quasiprobability technique requires the user to learn an accurate error model for all of the noisy processes in the circuit. It is then possible to express an inverse channel for each of these noisy channels. In general, these inverse channels may be unphysical operations. However, the expectation value with respect to the noiseless circuit (obtained by following each noisy physical gate with its potentially unphysical inverse channel) can be expanded as a sum of expectation values, each taken with respect to a different noisy circuit. By sampling from these circuits, we can estimate the noiseless ex-

pectation value. The classical process of weighting and summing the different expectation values resembles simulating a probabilistic operation. However, in this case, some of the probabilities can be negative – thus the name ‘quasiprobability technique’. A number of follow-up investigations have developed the error model learning process [30, 155, 156]. I do not make use of the quasiprobability technique in this thesis, but will briefly discuss how it could be useful for the work presented in Chapter 7.

A third approach to error mitigation is the quantum subspace expansion [157, 158]. While this technique was primarily introduced to find the excited states of quantum systems, it was noted that the method could also suppress the effects of errors. The subspace expansion works by considering a set of excitation operators applied to the eigenstate of interest, and using the states generated to form an incomplete basis. For example, we could consider the set of all single-qubit Pauli operators. By measuring an overlap matrix  $S_{ij} = \langle \psi | P_i P_j | \psi \rangle$ , and the Hamiltonian matrix in this basis  $H_{ij}^{\text{QSE}} = \langle \psi | P_i H P_j | \psi \rangle$ , we can construct and solve a generalised eigenvalue problem to better approximate the desired state. If the effect of noise is small, then the noisy states may live in the subspace spanned by the states  $\{P_i | \psi \rangle\}$ , allowing us to mitigate the noise with this procedure.

A number of recent works have considered how error mitigation could be tailored to the problem at hand. In particular, these works have considered how we might be able to use our knowledge of the symmetries present in simulations of physical systems to reduce the impact of noise. I discuss my proposal for such an approach (Ref. [159]) in Chapter 5. Such approaches typically seek to filter out circuit executions which lead to a violation of known symmetries. This can be achieved in a number of ways: by using stabiliser checks with additional ancilla qubits to detect changes in particle

number [159, 160], by effectively projecting the density matrix to a ‘physical state’ manifold through taking additional measurements [139, 160–162], or by changing into a basis where all Hamiltonian terms commute with the particle number operator, allowing one to directly post-select on particle number [163].

### 3.2.2 Error correction

As discussed above, error mitigation techniques do not provide a scalable solution for tackling noise. In order to run arbitrarily long circuits it is necessary to be able to correct errors during the calculation. This can be achieved through quantum error correction (QEC). In this section I provide a brief overview of QEC, and direct the reader to Refs. [164–167] for further details. Initial approaches to QEC were heavily influenced by classical error correction techniques. Many of these techniques redundantly store information by encoding an unprotected bit into an error correcting code, constructed from several bits. The code is constructed such that measuring combinations of ‘check’ operators on groups of bits flags when errors have occurred. These ‘syndrome measurements’ are used in conjunction with a decoding algorithm to infer the most likely errors that could have occurred, and thus what corrective action should be taken.

Although protecting quantum information utilises many of the same ideas as protecting classical information, QEC is more complicated than its classical counterpart. While classical information only needs to be protected against errors which flip or lose bits, qubits can undergo bit flips, loss, leakage from the computational subspace, and phase flip errors. Furthermore, while classical information can be copied to provide redundancy, the no-cloning theorem of quantum mechanics means that redundancy must be created in quantum codes more carefully. Finally, the syndrome measurements must be carefully constructed such that no information about the

logical state of the system is learnt, as this would collapse the superposition that we are trying to protect. The number of single-qubit errors that a code can detect and/or correct is determined by the ‘distance’ of the code.

In order to ascertain how effective a proposed quantum error correcting code may be, we need a model for the noise in the system. It has been shown that if the noise is sufficiently well behaved (e.g. noise is local, and correlations are limited) then error correction may be performed in a ‘fault tolerant’ manner [168–170]. This means that even if we assume that all quantum operations (i.e. syndrome measurement, corrective operations) are noisy [171], we are still able to suppress noise to an arbitrarily low value by scaling up the code – provided that the noise is below a code-dependent value known as the code threshold. ‘Scaling up’ the code can be achieved by concatenating multiple codes, or by growing some codes (e.g. the surface code).

Fault tolerant gates are typically achieved by using constructions that limit the spread of physical errors. An example is the use of transversal logical gates, which can be implemented by applying a physical gate independently on the individual physical qubits that encode the logical qubit. Assuming only one physical gate error occurs, then the logical qubit contains only one erred physical qubit. This can be identified and corrected using QEC. For example, the Steane code can implement all single-qubit Pauli and Clifford operations, and the CNOT gate, transversally [172]. However, the Eastin-Knill theorem [173] shows that quantum error correcting codes cannot possess a universal set of gates that can be implemented transversally. As a result, more complex fault tolerant constructions are required to achieve universal fault tolerant quantum computing (see Refs. [174, 175] for recent summaries of these approaches).

One of the most widely studied error correcting codes is the surface code [176]. Its popularity stems from its comparatively high fault tolerant threshold of around  $p = 10^{-2}$  [177–179], as well as its compatibility with a 2D grid of physical qubits with nearest-neighbour connectivity. This is a realistic architecture for many solid state qubits. The surface code possesses separate distances against  $X$  and  $Z$  errors, determined by the length and width of the 2D array of physical qubits. Universality is typically achieved in the surface code using a Clifford + T or Clifford + Toffoli gateset. The Clifford operations can be implemented through approaches including defect braiding [180] or lattice surgery [181], while T/Toffoli gates can be implemented using a process known as magic state distillation and injection [182]. Magic states are logical qubit states that can be consumed through gate teleportation to execute non-Clifford operations on other qubits in the circuit. The gate teleportation process can be implemented using only Clifford operations. An example magic state is  $|A\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\pi/4}|1\rangle)$ , which can be used to implement a T gate. High quality magic states can be generated through a process known as distillation. This procedure uses a circuit encoded in a higher-level code with a transversal T/Toffoli gate. The circuit takes as input a number of noisy magic states (encoded in surface code logical qubits), injects them to perform a transversal non-Clifford gate, and – subject to check operations of the higher-level code passing – outputs a less noisy magic state. This higher quality magic state can be input to further rounds of distillation, or injected into the main computation. If the checks fail, the qubits are discarded, and the process is repeated. Arbitrary angle single-qubit rotations can be synthesised from these T/Toffoli gates. The number of T/Toffoli gates required per rotation depends logarithmically on the inverse of the synthesis error [183]. A reasonable assumption for calculations of the size considered in this thesis is that around 100 T/Toffoli gates per single-qubit rotation would suf-

face [184]. Magic state distillation was previously believed to account for the majority of resources in most fault tolerant computations, even if the T gates were performed in a serial fashion. Recent work has shown that this is not necessarily the case [185, 186]. However, we can consider adding additional physical qubits for distilling magic states in parallel to speed up computations (so-called ‘space-time’ tradeoffs).

Such tradeoffs may be necessary because the clock-speed of error corrected quantum computers may be significantly slower than that of the underlying physical hardware. For example, in the surface code the time taken to implement logical operations scales with the code distance (because syndrome measurements may be noisy, and so must be repeated a number of times to determine the most likely error). This is not a property of all error correcting codes [187] – but clearly implementing logical operations must take at least as long as implementing the same operations on physical qubits.

Given the potentially high cost of implementing non-Clifford operations (like the T or Toffoli gates) in the surface code, the number of such operations in an algorithm is typically taken as an appropriate cost metric. Together with details such as how quickly magic states can be distilled and consumed, and the error rate of the physical hardware, this value can be used to estimate how many physical qubits would be required to protect the logical qubits from errors for the duration of the computation. Based on current physical error rates of  $p = 10^{-3}$ , many estimates suggest that around  $10^3 - 10^4$  physical qubits per logical qubit may be required to perform interesting tasks in a fault-tolerant manner [174, 188, 189]. This results in estimates of around  $3 \times 10^9$  T/Toffoli gates and  $2 \times 10^7$  physical qubits to factor 2048-bit RSA integers using Shor’s algorithm [190], or around  $7 \times 10^9$  T/Toffoli gates and  $4 \times 10^6$  physical qubits to simulate classically intractable models of the FeMoco molecule [27]. Clearly it will be exceptionally chal-

lenging to build quantum computers of this size, and so it is valuable to ask if there are useful applications that could be tackled with fewer resources. I discuss one such proposal – analysing muon spectroscopy data – in Chapter 7.

## 4 | Variational quantum algorithms for imaginary time evolution

### Contents

---

<a href="#">4.1 Chapter summary</a>	50
<a href="#">4.2 Introduction</a>	51
<a href="#">4.3 Theoretical background</a>	52
<a href="#">4.3.1 Ansatz-based quantum imaginary time evolution</a>	52
<a href="#">4.3.2 Explicit correlation &amp; the transcorrelated method</a>	60
<a href="#">4.4 Results</a>	66
<a href="#">4.4.1 Initial benchmarks</a>	66
<a href="#">4.4.2 Application to explicitly correlated Hamiltonians</a>	71
<a href="#">4.5 Discussion</a>	78

---

### 4.1 Chapter summary

This chapter investigates a hybrid quantum-classical algorithm for ansatz-based imaginary time evolution. This algorithm was originally reported in Ref. [191]:

*Variational ansatz-based quantum simulation of imaginary time evolution*  
S. McArdle, T. Jones, S. Endo, Y. Li, S.C. Benjamin, & X. Yuan  
[Nature Partner Journal Quantum Information 5 \(1\), 1-6 \(2019\)](#)

My contributions in that publication were: conceiving of the idea of implementing imaginary time evolution on a quantum computer, working with Xiao Yuan to develop the algorithm, performing the initial numerical simulations on the  $H_2$  and LiH molecules, assisting with the follow-up numerical simulations on LiH (performed by Tyson Jones), and analysing the asymptotic costs of the algorithm. This algorithm was later generalised to consider non-Hermitian Hamiltonians, for the purpose of finding the ground states of ‘transcorrelated’ Hamiltonians, in Ref [192]:

*Improving the accuracy of quantum computational chemistry using the transcorrelated method*  
S. McArdle & D.P. Tew  
[arXiv:2006.11181 \(2020\)](#)

My contributions to this work were: proposing and developing the idea to use the ansatz-based imaginary time algorithm to simulate the transcorrelated Hamiltonian, and carrying out all numerical simulations.

## 4.2 Introduction

This thesis is concerned with the quantum simulation of quantum systems, with a particular focus on finding ground states. A common classical method for calculating the ground states of quantum systems is to evolve them in imaginary time [193]. If we evolve a state  $|\psi\rangle = \sum_j c_j |E_j\rangle$  (where  $|E_0\rangle$  is the non-degenerate ground state) in imaginary time

$$\begin{aligned} & \lim_{\tau \rightarrow \infty} \frac{e^{-H\tau} |\psi\rangle}{\sqrt{\langle \psi | e^{-2H\tau} | \psi \rangle}} \\ &= \lim_{\tau \rightarrow \infty} \frac{\sum_j c_j e^{-\tau E_j} |E_j\rangle}{\sqrt{\sum_k |c_k|^2 e^{-2\tau E_k}}} \\ &= |E_0\rangle. \end{aligned} \tag{4.1}$$

The denominator is required to enforce normalisation of the state. The amplitudes of higher energy states decay exponentially with the magnitude of their energy above the ground state. A number of classical methods seek to efficiently approximate this evolution for certain systems, such as the time evolving block-decimation algorithm [194] applied to tensor network approaches, or the full configuration interaction quantum Monte Carlo (FCIQMC) technique [60]. An interesting feature of these imaginary time-based methods, is that they do not rely on the Rayleigh-Ritz variational principle. This means that these methods can be compatible with non-Hermitian Hamiltonians, which arise in a number of settings in many-body physics – including the explicit treatment of electron correlation in quantum chemistry. I will discuss this application in detail in Sec. 4.3.2.

Unfortunately, it appears inefficient to exactly simulate imaginary time evo-

lution of a general quantum system, using a classical computer. This stems from the widely held belief that the memory required to store a general quantum state grows exponentially with the size of the simulated system. This is normally the point at which quantum computing is invoked as the path forwards, as it enables us to store general wavefunctions with polynomially scaling resources. However, implementing imaginary time evolution on a quantum computer is not straightforward. Imaginary time evolution is a non-unitary, but deterministic, process. In contrast, the gates applied to quantum computers are unitary, while measurements realise non-unitary transforms with probabilistic outcomes. In this chapter, I will discuss a hybrid quantum-classical algorithm for implementing ansatz-based imaginary time evolution. This approach exploits the ability of quantum hardware to store many-body quantum states, and the ability of classical computers to simulate arbitrary (including unphysical) processes. I will then apply this algorithm to the problem of finding the ground states of electronic structure Hamiltonians, including non-Hermitian explicitly correlated Hamiltonians.

## 4.3 Theoretical background

### 4.3.1 Ansatz-based quantum imaginary time evolution

The imaginary time algorithm was inspired by a similar approach for simulating real time evolution [152]. These ‘variational quantum simulation’ algorithms were later formalised in a unified framework in Ref. [195]. They proceed by using an ansatz circuit,  $U(\vec{\theta}_{\tau/t})$ , to represent the state of the quantum system at a given point  $\tau/t$  on its imaginary/real time trajectory. To push the state forwards in time, the circuit parameters are evolved according to an update rule. This rule is derived from an appropriate variational principle applied to the corresponding time-dependent Schrödinger

equation. The algorithms seek the state that is closest in distance to the state obtained from ‘true’ time evolution, but that can still be prepared by the ansatz circuit. For the case of imaginary time evolution, the minimisation of an energy cost function thus happens as a corollary of a sufficiently long propagation, rather than due to an application of the Rayleigh-Ritz variational principle. As such, the method should be able to converge to the ground states of both Hermitian and non-Hermitian Hamiltonians.

Here, I re-derive the evolution of parameters formula corresponding to imaginary time evolution under a Hamiltonian that may be non-Hermitian,  $H'$ . We seek to propagate an initial state in imaginary time, according to the equation

$$|\phi(\tau)\rangle = \frac{e^{-H'\tau} |\phi(0)\rangle}{\sqrt{\langle\phi(0)| e^{-H'^\dagger\tau} e^{-H'\tau} |\phi(0)\rangle}}. \quad (4.2)$$

We can verify that this state satisfies a modified version of the imaginary time Schrödinger equation

$$\frac{\partial |\phi(\tau)\rangle}{\partial\tau} = -[H' - \Re(E_\tau)] |\phi(\tau)\rangle, \quad (4.3)$$

where  $\Re(E_\tau)$  is the real part of  $E_\tau = \langle\phi(\tau)|H'|\phi(\tau)\rangle$ , and is necessary to maintain normalisation. McLachlan’s variational principle [196] applied to Eq. (4.3), gives

$$\delta|(\partial/\partial\tau + H' - \Re(E_\tau)) |\phi(\tau)\rangle| = 0 \quad (4.4)$$

where

$$|(\partial/\partial\tau + H' - \Re(E_\tau)) |\phi(\tau)\rangle| = \left( (\partial/\partial\tau + H' - E_R) |\phi(\tau)\rangle \right)^\dagger \left( (\partial/\partial\tau + H' - E_R) |\phi(\tau)\rangle \right), \quad (4.5)$$

and  $E_R = \Re(\langle\phi(\tau)|H'|\phi(\tau)\rangle)$ . We restrict the algorithm to states that can be created by an ansatz circuit  $|\Phi(\tau)\rangle = |\Phi(\vec{\theta}_\tau)\rangle$ . McLachlan’s principle then

aims to minimise the Euclidean distance between the updated parameterized state, and the state after true evolution in imaginary time. Replacing  $|\phi(\tau)\rangle$  with  $|\Phi(\tau)\rangle$ , yields

$$\begin{aligned}
& |(\partial/\partial\tau + H' - E_R) |\Phi(\tau)\rangle | \\
&= \left( (\partial/\partial\tau + H' - E_R) |\Phi(\tau)\rangle \right)^\dagger \left( (\partial/\partial\tau + H' - E_R) |\Phi(\tau)\rangle \right) \\
&= \left( \frac{\partial \langle \Phi(\tau) |}{\partial \tau} + \langle \Phi(\tau) | H'^\dagger - \langle \Phi(\tau) | E_R \right) \left( \frac{\partial |\Phi(\tau)\rangle}{\partial \tau} + H' |\Phi(\tau)\rangle - E_R |\Phi(\tau)\rangle \right) \\
&= \sum_{i,j} \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_j} \dot{\theta}_i \dot{\theta}_j + \sum_i \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} (H' - E_R) |\Phi(\tau)\rangle \dot{\theta}_i \\
&\quad + \sum_i \langle \Phi(\tau) | (H'^\dagger - E_R) \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i} \dot{\theta}_i + \langle \Phi(\tau) | (H'^\dagger - E_R) (H' - E_R) |\Phi(\tau)\rangle.
\end{aligned} \tag{4.6}$$

Choosing to minimise the function with respect to  $\dot{\theta}_i$ , we obtain

$$\begin{aligned}
& \frac{\partial |(\partial/\partial\tau + H' - E_R) |\Phi(\tau)\rangle |}{\partial \dot{\theta}_i} \\
&= \sum_j \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_j} + \frac{\partial \langle \Phi(\tau) |}{\partial \theta_j} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i} \right) \dot{\theta}_j \\
&\quad + \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} (H' - E_R) |\Phi(\tau)\rangle + \langle \Phi(\tau) | (H'^\dagger - E_R) \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i}.
\end{aligned} \tag{4.7}$$

McLachlan's variational principle requires

$$\frac{\partial |(\partial/\partial\tau + H' - E_R) |\Phi(\tau)\rangle |}{\partial \dot{\theta}_i} = 0, \tag{4.8}$$

We then let

$$\begin{aligned}
A_{ij} &= \Re \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_j} \right) \\
&= \frac{1}{2} \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_j} + \frac{\partial \langle \Phi(\tau) |}{\partial \theta_j} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i} \right), \\
C_i &= \Re \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} H' |\Phi(\tau)\rangle \right) \\
&= \frac{1}{2} \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} H' |\Phi(\tau)\rangle + \langle \Phi(\tau) | H'^{\dagger} \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i} \right),
\end{aligned} \tag{4.9}$$

and use

$$\begin{aligned}
-E_R \left( \frac{\partial \langle \Phi(\tau) |}{\partial \theta_i} |\Phi(\tau)\rangle + \langle \Phi(\tau) | \frac{\partial |\Phi(\tau)\rangle}{\partial \theta_i} \right) &= -E_R \left( \frac{\partial}{\partial \theta_i} \langle \Phi(\tau) | \Phi(\tau) \rangle \right) \\
&= -E_R \left( \frac{\partial}{\partial \theta_i} 1 \right) \\
&= 0
\end{aligned} \tag{4.10}$$

to obtain an equation for the evolution of the parameters  $\vec{\theta}$

$$\sum_j A_{ij} \dot{\theta}_j = -C_i. \tag{4.11}$$

For the case of a Hermitian Hamiltonian  $H' = H'^{\dagger}$ , these equations are identical to those derived in Ref. [191].

We can evolve the state forwards by a timestep  $\delta\tau$  in imaginary time by updating the parameters according to an Euler update rule

$$\vec{\theta}(\tau + \delta\tau) \simeq \vec{\theta}(\tau) + \dot{\vec{\theta}}(\tau) \delta\tau = \vec{\theta}(\tau) - A^{-1}(\tau) \cdot C(\tau) \delta\tau. \tag{4.12}$$

For this equation to be applied, the matrix  $A$  needs to be invertible. It is often the case that  $A$  is underdetermined, in which case regularisation methods (such as a truncated singular value decomposition, or Tikhonov

regularisation) are applied. By repeating this process (of calculating  $A$  and  $C$ , and updating the parameters)  $N_T = \tau_{total}/\delta\tau$  times, we can simulate imaginary time evolution over a duration  $\tau_{total}$ .

Following the release of Ref. [191], it was shown that it is necessary to include an additional parameter, corresponding to the global phase of the state [195]. This is necessary because even if the trial state is able to approximate the true evolution up to a global phase, the derivatives of the two states may differ substantially. More formally, the consideration of global phase is required to ensure that the metric matrix  $A$  is gauge invariant under the global phase gauge transform (see Ref. [197]). As a result, we must evolve the global phase of the trial state in order to match it to that of the true evolution. The global phase can be included by either adding an additional  $R_z$  gate to the beginning of the circuit (i.e. applied to a qubit in the  $|0\rangle$  state, upon which it acts trivially), or by including additional terms in the  $A$  matrix, as discussed in Ref. [195].

When these additional terms are included in the  $A$  matrix, it is equivalent to the Fubini-Study metric, which is an analogue of the Fisher information applied to the ansatz. The matrix  $A$  is a metric on the ansatz space, which encodes how sensitive the ansatz state is to changes in the individual parameters. When the Hamiltonian  $H'$  is Hermitian, the vector  $C$  is the gradient vector of the energy cost function. I have not found a similarly simple interpretation of  $C$  for the case of non-Hermitian Hamiltonians. Focusing on the case of a Hermitian Hamiltonian, we see that our method can be viewed as updating the parameters according to a gradient descent rule, with the addition of a metric matrix that re-parameterizes the ansatz space. Subsequent works have shown that the method is equivalent to a quantum analogue of ‘natural gradient descent’, a powerful optimisation technique in classical machine learning [127, 198].

In addition to being able to prepare the ansatz state on a quantum computer, we need to be able to measure the elements  $A_{ij}$  and  $C_i$  using quantum hardware. We can often write our ansatz circuit in the form  $U(\vec{\theta}) = \prod_j U_j(\theta_j)$ , with parameterized gates of the form  $e^{i\theta_j P_j}$ , where  $P_j$  is a Pauli string. In this case, if  $|\Phi(\vec{\theta})\rangle = U(\vec{\theta})|\bar{0}\rangle$ , then

$$\frac{\partial |\Phi(\vec{\theta})\rangle}{\partial \theta_j} = \frac{\partial}{\partial \theta_j} U(\vec{\theta})|\bar{0}\rangle = i \left( \prod_{k>j} U_k(\theta_k) \right) P_j U_j(\theta_j) \left( \prod_{i<j} U_i(\theta_i) \right) |\bar{0}\rangle = i \bar{V}_j |\bar{0}\rangle. \quad (4.13)$$

That is, we can generate the derivative vector using a unitary circuit  $\bar{V}_j$  that differs from the ansatz unitary only in the insertion of the gate  $P_j$  at position  $j$  in the circuit. Given quantum circuits that generate the unitaries  $W, \Lambda$ , we can measure the real part of  $\langle \bar{0} | W \Lambda | \bar{0} \rangle$  using the Hadamard test circuit shown in Fig. 4.1.




Figure 4.1: A quantum circuit to evaluate the real part of  $\langle \bar{0} | W \Lambda | \bar{0} \rangle$ , which forms a core component of the imaginary time algorithm.

Because the circuits used to generate the derivative wavefunctions only differ from the original ansatz (or circuits that generate derivatives with respect to other parameters) in a small number of locations, we can significantly reduce the number of controlled gates present in the construction of Fig. 4.1. In order to apply this method to calculate the elements of the vector  $C$ , we must first decompose the Hamiltonian into a sum of Pauli strings. We can then measure the elements of  $C$  using circuits such as that shown in Fig. 4.2. Note that it is necessary to measure the  $Y$  operator on the ancilla qubit when measuring the terms  $C_j$ , due to the prefactor of  $i$  that enters Eq. (4.13). We can use a similar construction to measure the terms  $A_{ij}$ , as

shown in Fig. 4.3. These techniques can also be applied to ansätze with more complicated structure than those discussed above – but the simple case is sufficient for the examples presented herein.




Figure 4.2: A circuit to measure an element of the gradient vector,  $C$ . In this case, we seek  $C_j$ , where the  $j^{\text{th}}$  gate in the circuit is  $e^{i\theta_j P_j}$ , and are measuring the Pauli string  $H_\lambda$  in the Hamiltonian. It is also possible to measure this term without introducing an additional probe qubit [199, 200].




Figure 4.3: A circuit to evaluate the  $A_{jk}$  element of the metric matrix. When measuring these terms, we only need to simulate gates up to position  $k$  in the original ansatz circuit, which can reduce the required circuit depth.

These circuits must be repeated many times in order to evaluate each element of the  $A$  matrix or  $C$  vector. Recent work has examined the optimal way to distribute a fixed number of measurements between the elements, observing that it is necessary to obtain the highest precision on the elements of  $C$  [201]. As the expectation values are obtained by direct sampling, the error in our estimate will scale as  $\mathcal{O}(1/\sqrt{M})$ , where  $M$  is the number of samples used for a given element. On each iteration of our algorithm, populating  $C$  requires  $\mathcal{O}(M_C N_H N_p)$  measurements, where  $M_C$  is the number of measurements required to ascertain a Hamiltonian term to the required precision,  $N_H$  is the number of terms in the Hamiltonian, and  $N_p$  is the number of parameters in the ansatz. Populating the metric  $A$  requires an additional  $\mathcal{O}(N_p^2 M_A)$  measurements, where  $M_A$  is the number of

measurements required to ascertain an element of the  $A$  matrix to the required precision. Typical ansätze for the electronic structure problem can have a comparable number of parameters to Hamiltonian terms. In these cases, the additional measurements required to populate the metric do not alter the asymptotic scaling of the algorithm. In addition, numerical results indicate that often  $M_A \ll M_C$ , further reducing the additional cost of this method over standard gradient descent [191, 202].

A limitation of the algorithm is that the ansatz may not be able to accurately describe all states on the true imaginary time trajectory. If we seek the ground state of the Hamiltonian, this may not be as problematic as if we wish to faithfully simulate the imaginary time dynamics. Provided errors due to a constrained ansatz do not cause the simulation to become trapped in local minima, we do not mind if the evolution deviates from the path of true imaginary time evolution, as ultimately, it will still be driven towards the ground state. A similar argument can be made regarding the effect of both gate and measurement errors on the algorithm; provided the deviations from true imaginary time evolution caused by noise are sufficiently small, the algorithm should still drive the state towards the ground state. Although noise is not explicitly investigated in this chapter, it is likely the case that both gate and measurement errors could be treated on a similar footing, from an algorithmic perspective. In cases where measurement error dominates, the algorithm may benefit from the fact that we only measure a single qubit to populate the  $A$  and  $C$  matrices, thus reducing the number of fault locations. Moreover, measurement errors can often be addressed by efficient error mitigation strategies (even if an accurate model for the readout noise is not known), for example, those presented in Ref. [203].

This is expected to be the case for both Hermitian and non-Hermitian Hamil-

tonians. A class of non-Hermitian Hamiltonians that may be of interest are ‘transcorrelated’ electronic structure Hamiltonians, which can be used to reduce the resources required in quantum chemistry calculations. In the following section, I review the transcorrelated method, and other related explicitly correlated techniques. I then apply ansatz-based imaginary time evolution to such Hamiltonians in Sec. 4.4.2. Benchmark simulations of Hermitian electronic structure Hamiltonians are presented in Sec. 4.4.1.

### 4.3.2 Explicit correlation & the transcorrelated method

As discussed in Chapter 2, when solving the electronic structure problem we typically utilise a basis of single-particle functions. The wavefunction is then given by an expansion in antisymmetrised products of these functions (Slater determinants). This approach makes it costly to account for dynamic correlation in the wavefunction, as this parametrization is inefficient at resolving the electron-electron cusps. This inefficiency was formalised by Kato [204], who specified the conditions that wavefunctions must fulfil to accurately describe the dynamic correlation in the system. Wavefunctions constructed from products of single particle orbitals do not fulfil these conditions [205]. Superior convergence can be obtained by including functions with a dependence on the inter-electronic distances. These methods are collectively known as ‘explicitly correlated methods’, and have been reviewed in Refs. [41, 206, 207]. These techniques were employed from the earliest days of computational chemistry [208, 209]. Unfortunately, direct implementations of such approaches are constrained by the requirement to carry out  $N!$   $N$ -electron integrals, which restricts calculations to systems with at most around 6 electrons [41]. Related methods, such as R12 and F12 methods [210], circumvent this limitation and have been successfully applied to large system sizes, but introduce approximations that can complicate calculations [41].

An alternative approach to addressing dynamic correlation in the wavefunction (and the method considered in this chapter) is the transcorrelated (TC) method. The TC method was introduced by Boys and Handy [211, 212], following the earlier work of Ref. [213]. These works observed that rather than considering the wavefunction to be transformed by an auxiliary function that describes dynamic correlation, it is equivalent to consider the Hamiltonian to be transformed by the auxiliary function. This can be likened to transforming between the Schrödinger and Heisenberg pictures of quantum mechanics. To apply the TC method, we use a trial wavefunction

$$|\psi\rangle = e^{\sum_{i<j} f(\mathbf{r}_i, \mathbf{r}_j)} |\phi\rangle = e^{\hat{g}} |\phi\rangle \quad (4.14)$$

where  $|\psi\rangle$  is the wavefunction of the system prior to projection onto a basis,  $f(\mathbf{r}_i, \mathbf{r}_j)$  is a symmetric, real function of the positions of electrons  $i$  and  $j$  (referred to as a Jastrow factor), and  $|\phi\rangle$  is a wavefunction that does not explicitly depend on inter-electronic distances. Considering the solutions of the real space electronic structure Hamiltonian,  $H$  (i.e. before projection onto a basis set), we see that

$$H |\psi_i\rangle = E_i |\psi_i\rangle \quad (4.15)$$

$$\rightarrow H e^{\hat{g}} |\phi_i\rangle = E_i e^{\hat{g}} |\phi_i\rangle \quad (4.16)$$

$$\rightarrow e^{-\hat{g}} H e^{\hat{g}} |\phi_i\rangle = E_i |\phi_i\rangle. \quad (4.17)$$

I define  $H' = e^{-\hat{g}} H e^{\hat{g}}$  as the transcorrelated Hamiltonian. While the explicitly correlated wavefunction  $|\psi_i\rangle$  is an eigenstate of the original Hamiltonian, we can obtain the same eigenvalue by finding the wavefunction  $|\phi_i\rangle$ , which is an eigenstate of  $H'$ . As  $|\phi_i\rangle$  is not explicitly correlated, it may be easier to determine than  $|\psi_i\rangle$ . As the transformation  $e^{\hat{g}}$  is not unitary, the TC Hamiltonian  $H'$  is no longer Hermitian. This leads to a number of issues, most notably a lack of variational lower bound on the ground state eigen-

value, which arises because non-eigenstates may have complex energy expectation values. The lack of a variational lower bound precludes the use of many conventional techniques in classical computational chemistry, as well as quantum algorithms like the variational quantum eigensolver.

The Jastrow factor is chosen to regularise the Hamiltonian, such that the TC Hamiltonian is free from singularities, and to ensure that the function  $e^{\hat{g}}|\phi_i\rangle$  satisfies the cusp conditions in the region of the electron-electron and electron-nucleus coincidences. Since the TC transformation is performed prior to projection onto a basis set, the projected TC Hamiltonian is not isospectral with the original projected Hamiltonian, except in the limit of an infinite basis set. If both the TC Hamiltonian and the unmodified Hamiltonian are projected onto the same single particle basis set, the TC Hamiltonian will yield energies closer to those obtained in the basis set limit. As a result, transcorrelated calculations are able to achieve a similar accuracy to an unmodified calculation, while using fewer spin-orbitals. In a quantum setting, using fewer spin-orbitals would correspond to reducing the number of logical qubits required for the simulation. In a Gaussian basis, the TC Hamiltonian contains additional two and three-body operators such that the Hamiltonian contains  $\mathcal{O}(M^6)$  terms [214].

The TC transformation can also be applied after projection onto single-particle basis functions – for example, when treating the Fermi-Hubbard model [215, 216]. We can consider a similarity transformation with a Gutzwiller factor that acts to suppress double occupancies of lattice sites [216]

$$H' = \left( e^{-J \sum_i n_{i,\uparrow} n_{i,\downarrow}} \right) H \left( e^{J \sum_j n_{j,\uparrow} n_{j,\downarrow}} \right), \quad (4.18)$$

where  $n_{i,\sigma}$  is the number operator for the spin-lattice site indexed by  $i, \sigma$ . While the resulting TC Hamiltonian is still isospectral to the unmodified

Hamiltonian in this case, Ref. [216] observed that the TC Hamiltonian had more ‘compact’ right-hand eigenvectors than the regular Hamiltonian. More precisely, lower weight fermionic excitations were required to obtain a good overlap with the ground state in the TC case than when considering the unmodified Hamiltonian. In a quantum setting, the increased compactness of the right-hand eigenvectors may mean that they can be prepared using lower depth quantum circuits than eigenstates of the unmodified Hamiltonian.

Unfortunately, while the right-hand eigenvectors are made more compact, the left-hand eigenvectors gain additional dynamic correlation, and thus may become more difficult to construct from a single-particle basis expansion [205]. The left-hand eigenvector is given by

$$\langle \psi_i | H = \langle \psi_i | E_i \quad (4.19)$$

$$\rightarrow \langle \phi_i | e^{\hat{g}} H = \langle \phi_i | e^{\hat{g}} E_i \quad (4.20)$$

$$\rightarrow \langle \phi_i | e^{2\hat{g}} H' = \langle \phi_i | e^{2\hat{g}} E_i, \quad (4.21)$$

$$\langle \tilde{\phi}_i | H' = \langle \tilde{\phi}_i | E_i \quad (4.22)$$

where  $\langle \tilde{\phi}_i | = \langle \phi_i | e^{2\hat{g}}$  is the left-hand eigenvector of  $H'$ . The differing forms of the left-hand and right-hand eigenvectors prove problematic for measuring observables other than the energy. For example, we see that

$$\langle O \rangle = \langle \psi | \hat{O} | \psi \rangle = \langle \phi | e^{\hat{g}} \hat{O} e^{\hat{g}} | \phi \rangle. \quad (4.23)$$

As  $e^{\hat{g}} \hat{O} e^{\hat{g}}$  does not have a terminating Baker-Campbell-Hausdorff expansion, we use the expansion  $\hat{O}' = e^{-\hat{g}} \hat{O} e^{\hat{g}}$  (which does terminate) to write that

$$\langle O \rangle = \langle \phi | e^{2\hat{g}} \hat{O}' | \phi \rangle = \langle \tilde{\phi}_i | \hat{O}' | \phi \rangle. \quad (4.24)$$

As a result, calculating observables other than the energy requires obtaining both the right- and left-hand eigenvectors of the TC Hamiltonian. This challenge has yet to be resolved in studies on the TC method, which have mainly focused on finding the ground state energy of various systems.

Early works on the TC method considered a single Slater determinant together with a Jastrow function, and optimised the form of both parts of the wavefunction using self-consistent equations [211, 212]. The non-Hermitian nature of the TC Hamiltonian limits the applicability of this approach. Subsequent works attempted to instead variationally minimise the variance of the energy [217]. Unfortunately, this approach yields less accurate results than optimising the energy, and can lead to a cost function with  $\mathcal{O}(M^{12})$  terms (the TC Hamiltonian for a molecular system in a Gaussian orbital basis contains  $\mathcal{O}(M^6)$  terms, so measuring the variance requires  $\mathcal{O}(M^{12})$  terms).

More recent work attempted to improve the practicality of the TC method, typically by either: fixing the form of the Jastrow factor and compensating for this error with a larger Slater determinant expansion [218, 219], or developing improved optimisation heuristics [215, 220–223], or by discarding terms to yield a Hermitian approximation to the TC Hamiltonian [224, 225]. The most recent developments in this area use imaginary time-like methods, such as FCIQMC, to optimise the Slater determinant expansion, while freezing the Jastrow factor. This circumvents the non-variational nature of the calculation, and partially ameliorates the difficulties associated with optimising the Jastrow function. Imaginary time-like methods can be used to find the ground state of the TC Hamiltonian, without invoking variational

properties

$$|\psi_0\rangle = \lim_{\tau \rightarrow \infty} e^{-H\tau} |\psi\rangle, \quad (4.25)$$

$$\rightarrow e^{\hat{g}} |\phi_0\rangle = \lim_{\tau \rightarrow \infty} \sum_{k=0}^{\infty} \frac{(e^{\hat{g}} H' e^{-\hat{g}\tau})^k}{k!} e^{\hat{g}} |\phi\rangle, \quad (4.26)$$

$$\rightarrow |\phi_0\rangle = \lim_{\tau \rightarrow \infty} e^{-H'\tau} |\phi\rangle. \quad (4.27)$$

By transferring the dynamic correlation from the right-side wavefunction to the TC Hamiltonian, the TC method also yields a more compact Slater determinant expansion, which is beneficial for the convergence of FCIQMC. This approach has been applied to simulations of plane wave Hamiltonians [226], the Fermi-Hubbard model [216], atomic systems in Gaussian basis sets [214], quantum gases [227], and ultracold atoms [228].

As the TC method is agnostic of the approach used to generate the Slater determinant expansion, it is natural to investigate replacing the classical chemistry methods typically used, with a quantum subroutine. From a quantum perspective, the TC transformation could be beneficial by reducing the number of logical qubits required (through a reduction in the number of orbitals required to account for dynamic correlation) and/or reducing the circuit depth required to prepare a given right-hand eigenstate. From a classical perspective, the use of quantum computing may enable the preparation of Slater determinant expansions that would be computationally expensive to generate using classical methods. However, the non-Hermitian nature of the TC Hamiltonian appears to preclude its use in both the variational quantum eigensolver (which relies on having a variational lower bound when measuring the expectation values of Hermitian operators), or quantum phase estimation (which evolves the system under a unitary function of the Hamiltonian). One possible route would be to

consider approximate, unitary versions of the TC Hamiltonian. This is the approach recently taken in Ref. [229], which used the ‘canonical transcorrelated’ method. However, the severity of these approximations are not yet fully understood by the chemistry community.

Motivated by the success of imaginary time-like methods at dealing with the TC Hamiltonian, I consider applying the aforementioned quantum algorithm for ansatz-based imaginary time evolution to solve this problem in Sec. 4.4.2.

## 4.4 Results

### 4.4.1 Initial benchmarks

In order to test the performance of the algorithm, I first considered the simulation of Hermitian Hamiltonians. I performed classical emulations of applying the algorithm to find the ground states of small molecular systems: the  $H_2$  and LiH molecules in minimal basis sets. Using the approaches discussed in Sec. 2.2, I mapped the second quantised Hamiltonians for these molecules onto qubit systems. For  $H_2$ , I used the Hamiltonian given by Eq. (2.25), which acts on two qubits, and describes the behaviour of two electrons in four spin-orbitals. For LiH, I applied the transforms discussed in Sec. 2.2.3 to obtain a Hamiltonian on 8 qubits, describing two electrons in 8 spin-orbitals. Although these simulations did not explicitly account for global phase (as the necessity of tracking this parameter was not understood at the time), fortunately overparametrization in the ansätze used provided a free parameter to serve as the global phase.

The  $H_2$  simulation was used to verify that the ansatz-based algorithm was able to closely track the true imaginary time evolution. The circuit shown in Fig. 4.4, which can represent any two-qubit state, was used as the ansatz




Figure 4.4: The quantum circuit for preparing the universal two-qubit ansatz applied in the simulation of the  $\text{H}_2$  molecule. Figure reproduced with permission from Ref. [191].




Figure 4.5: The circuit to measure  $A_{2,7} = \Re \left( \frac{\partial \langle \phi(\tau) |}{\partial \theta_2} \frac{\partial | \phi(\tau) \rangle}{\partial \theta_7} \right)$  in the simulation of the  $\text{H}_2$  molecule. In practice, the gates in the dashed box may be omitted. The other terms of  $A$  and  $C$  can be measured using similar circuits. Figure reproduced with permission from Ref. [191].

circuit. This separated the performance of the algorithm from the expressibility of the ansatz. In Fig. 4.5 I show an example circuit used to measure an element of the  $A$  matrix. These simulations (performed using the ProjectQ package) closely emulated the behaviour of an experimentalist wishing to perform the algorithm. I simulated the full circuit, including the ancilla qubit, to obtain all elements of the  $A$  matrix and  $C$  vector. However, because exact wavefunction simulations were performed, I was able to directly measure the expectation values, and so shot noise is not present in these simulations. All 8 parameters were initialised with random values. The timestep was  $\delta\tau = 0.01$ . The results for  $\text{H}_2$  are shown in Fig. 4.6, and confirm that the method closely adheres to the true imaginary time evolution, when the ansatz is sufficiently powerful. Deviation from the true evolution is attributed to the use of an Euler update rule, and finite step




Figure 4.6: Simulation of imaginary time evolution applied to  $H_2$ , at an internuclear separation of  $R = 0.75 \text{ \AA}$ , starting with random initial parameters. The red line is the exact ground state energy. The dashed black line is the exact imaginary time evolution. The blue line is the ansatz-based imaginary time evolution, with the ansatz used shown in Fig. 4.4. The inset plot shows the fidelity between ansatz-based and true imaginary time evolution. The inset plot and main plot share the same  $x$  axis label. Figure reproduced with permission from Ref. [191].

size. These simulations were able to converge to the ground state in all trials.

Simulations were performed on LiH in order to compare the performance of the algorithm with that of gradient descent. These simulations used a hardware efficient ansatz, shown in Fig. 4.7, that aimed to create a flexible trial state using a limited number of gates. For the case of LiH, the ansatz had 137 parameters, approximately one quarter of those needed to describe a general 8 qubit state.

The simulation results are shown in Fig. 4.8. These simulations were performed by Tyson Jones, a collaborator on this project, with input from myself. These simulations (performed using the QuEST package) did not emulate the actions of an experimentalist. Instead, they used direct access to the wavefunction to calculate and store the derivative wavefunctions, and

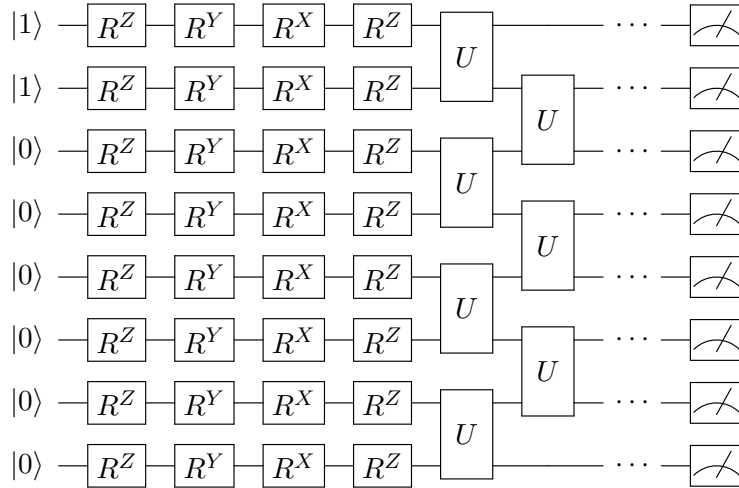


Figure 4.7: The general structure of the ansatz used in the LiH simulations. The circuit structure of the two-qubit rotations block was repeated to depth  $M = 3$ . The form of the  $U$  gate was  $U = e^{i\alpha Y X} e^{i\beta X Y} e^{i\gamma Z Z} e^{i\delta Y Y} e^{i\epsilon X X}$ . In total there were  $(3 \times 5 \times 7) + (4 \times 8) = 137$  parameters. Figure reproduced with permission from Ref. [191].

compute the overlaps required to evaluate the elements in  $A$  and  $C$ . In this way, a single circuit emulation was able to provide the same information as would be obtained from many experimental samples. Tikhonov regularisation was used to stabilise the inversion of  $A$ . We considered two different initial conditions for the simulations: starting from a small perturbation from the Hartree-Fock state, and starting from completely random parameters. The former case closely mimics the approach that could be used if a good initial guess for the parameters is known, while the latter provides a more thorough test of the algorithm. As discussed in Sec. 3.1.3.1, randomizing the parameters will cause the energy gradient to decay exponentially with the circuit depth, which would render the algorithm ineffective. While this 8 qubit simulation was not large enough for this effect to become important, starting from random initial states is not a scalable approach.

For both starting points, the imaginary time method outperformed gradi-

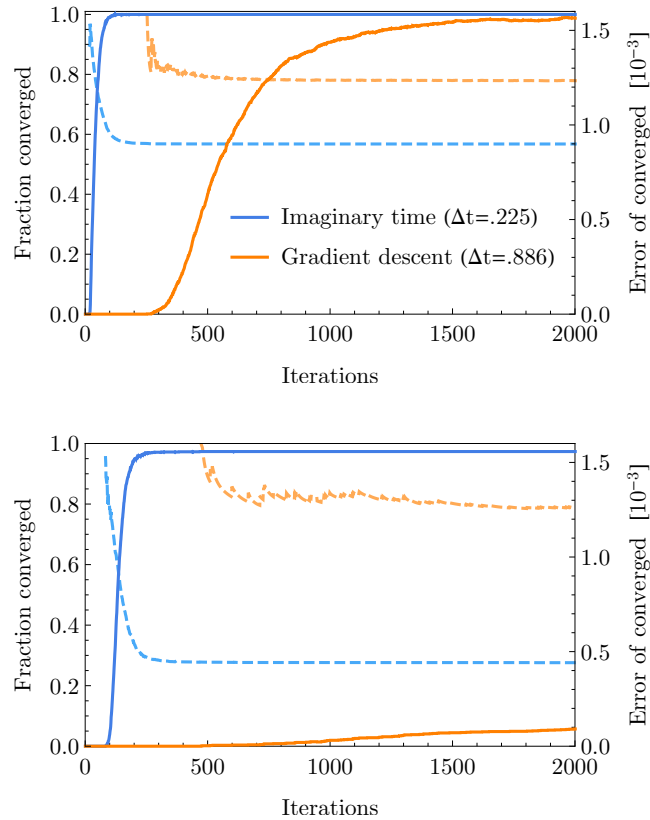


Figure 4.8: Simulations of LiH at an internuclear distance of  $R = 1.45 \text{ \AA}$ . Simulations in the upper plot begin with a small random perturbation (of at most,  $\Delta\theta_j = \pi/50$ ) from the Hartree-Fock state. Simulations in the lower plot have uniformly random initial parameters. The solid lines (against the left axis) indicate the fraction of 1280 simulations which, by the given iteration, have converged to within  $1 \times 10^{-3}$  Hartree of the true ground state. The dashed lines (against the right axis) indicate the average energy error of those simulations within chemical accuracy of the true ground state. The maximum stable stepsize was used for each method such that the energy monotonically decreased in the first 200 iterations. The stable timestep for imaginary time was 0.225, and for gradient descent it was 0.886. These simulations were performed by Tyson Jones, with input from myself. Figure reproduced with permission from Ref. [191].

ent descent. Imaginary time evolution was able to find the ground state more quickly, and accurately. This advantage was more pronounced for random initial states, where the convergence rate was significantly higher than that obtained by gradient descent. Similar results have been subsequently obtained by Refs. [127, 198], which reported improved performance in finding the ground states of spin Hamiltonians when using quantum natural gradient descent (equivalent to the imaginary time algorithm), compared to using vanilla gradient descent, ADAM, and an approximate Hessian method (the BFGS optimiser).

These simulations on small molecular systems (with Hermitian Hamiltonians) confirm that the algorithm captures the desired features of true imaginary time evolution. In the following section, I consider applying the algorithm to non-Hermitian, transcorrelated Hamiltonians.

#### 4.4.2 Application to explicitly correlated Hamiltonians

As discussed in Sec. 4.3.2, classical imaginary time methods have previously been used to treat transcorrelated Hamiltonians. The non-Hermitian nature of the TC Hamiltonian makes it difficult to incorporate into black-box quantum algorithms for finding ground states. As a result, if ansatz-based imaginary time evolution is able to successfully find the ground states of TC Hamiltonians, then it may provide a route to simulate explicitly correlated chemistry using quantum hardware. I tested this by conducting numerical emulations (performed using the Cirq package) of the algorithm applied to the transcorrelated Fermi-Hubbard model. As in the case of the LiH simulations discussed above, I used a finite difference approximation of the gradient, and directly calculated the inner products required to populate  $A$  and  $C$ . The inversion of  $A$  was performed using a generalised inverse of the matrix, obtained from a singular-value decomposition that

included all singular values above  $10^{-6}$ . A timestep of 0.01 was used for the Euler update rule.

As discussed in Sec. 2.1.4, the unmodified Fermi-Hubbard (FH) Hamiltonian is given by

$$H_{\text{FH}} = -t \sum_{\langle i,j \rangle, \sigma} \left( a_{i,\sigma}^\dagger a_{j,\sigma} + a_{j,\sigma}^\dagger a_{i,\sigma} \right) + U \sum_i n_{i,\uparrow} n_{i,\downarrow} \quad (4.28)$$

where  $\langle i, j \rangle$  denotes a sum over nearest-neighbour lattice sites, and  $\sigma$  is a spin-coordinate. I selected interaction strengths of  $t = 1, U = 4$  for the numerical simulations, as this has been observed to be a region of difficulty for classical simulation methods [52].

The TC Hamiltonian was obtained using the Gutzwiller transformation of Refs. [215, 216]. The TC Hamiltonian was given by

$$H'_{\text{FH}} = \left( e^{-J \sum_i n_{i,\uparrow} n_{i,\downarrow}} \right) H_{\text{FH}} \left( e^{J \sum_j n_{j,\uparrow} n_{j,\downarrow}} \right), \quad (4.29)$$

where  $J$  defined the strength of the transformation. This was simplified to [216]

$$H'_{\text{FH}} = H_{\text{FH}} - t \sum_{\langle i,j \rangle, \sigma} \left( a_{i,\sigma}^\dagger a_{j,\sigma} \times \left[ (e^J - 1) n_{j,\bar{\sigma}} + (e^{-J} - 1) n_{i,\bar{\sigma}} - 2(\cosh(J) - 1) n_{i,\bar{\sigma}} n_{j,\bar{\sigma}} \right] \right) \quad (4.30)$$

where  $\bar{\sigma}$  denotes the spin opposite to  $\sigma$ .

Numerical simulations were carried out on  $2 \times 2$  and  $3 \times 2$  lattices, which were mapped onto 8 and 12 qubits (respectively) using the Jordan-Wigner transform. The unmodified and TC Hamiltonians were generated and mapped to qubit operators using the OpenFermion package.

A Hamiltonian variational ansatz [140] was used. This ansatz has proven effective in previous numerical simulations of applying the VQE to the FH model [80, 81, 84, 140]. I used the same ansatz circuit for both the unmodified and TC systems, where the Hamiltonian used in the ansatz was that of the unmodified system. I prepared the quantum register in the lowest eigenstate of the non-interacting ( $U = 0$ ) unmodified Hamiltonian with the same particle number as the ground state of the full Hamiltonian. This state was constructed using a network of Givens rotations [79, 82, 83]. The ansatz was constructed from repeated layers of a Trotterized decomposition of the time evolution operator. I assigned a different parameter to each Pauli rotation in the Trotterized decomposition

$$U = \prod_l \prod_j e^{i\theta_j^l P_j} \quad (4.31)$$

where  $L$  denotes the number of layers in the ansatz, and  $P_j$  are the Pauli strings in the Hamiltonian. I included an additional parameter to represent the global phase of the wavefunction. The initial values of the parameters were perturbed from zero by a random perturbation upper bounded by  $0.02\pi$ , in order to prevent the method from becoming trapped in local minima around the non-interacting initial state.

As a first step, I verified that the algorithm was capable of finding the ground state of the  $2 \times 2$  TC FH Hamiltonian. I set  $J = -0.5$ , similar to the values used in Refs. [215, 216] (as it was not necessary to use the optimal  $J$  value for this trial calculation). In Fig. 4.9, I show a randomly chosen run of the imaginary time algorithm, applied to the  $2 \times 2$  TC Hamiltonian, with a 3-layer ansatz. The energy rapidly converges towards the ground state value, and the imaginary part of the energy decays exponentially as the method moves forwards in imaginary time. The inset of Fig. 4.9 shows




Figure 4.9: A random instance of finding the ground state of the TC Hamiltonian (with  $J = -0.5$ ) for a  $2 \times 2$  Fermi-Hubbard model, with imaginary time evolution. Energies were recorded every 10 timesteps of the algorithm. An ansatz circuit with 3 layers was used. The inset plot shows the evolution of the energy during the final 62.5% of the runtime. Figure reproduced with permission from Ref. [192].

that although the algorithm is no longer strictly variational (the real part of the measured energy no longer upper bounds the ground state value), the real and imaginary parts of the energy converge towards their true values.

I present the average results of ten repetitions of such calculations in Fig. 4.10, for a range of ansatz depths. The upper plot of Fig. 4.10 shows that as the circuit depth was increased, the imaginary time algorithm was able to find better approximations to the right-hand ground state of the TC Hamiltonian. In contrast, the fidelity obtained using gradient descent did not improve as the circuit depth increased, suggesting that algorithms based solely on cost function minimisation may struggle to find the ground state of the TC Hamiltonian. Here, the fidelity was used as a metric, as it can be used to upper bound the error in any possible observable (through its relation to the trace distance [35]), and because it explicitly quantifies the dis-




Figure 4.10: Finding the ground state of the TC Hamiltonian (with  $J = -0.5$ ) for a  $2 \times 2$  Fermi-Hubbard model. Lines are included to guide the eye. The upper plot shows fidelity between the TC right-hand ground state, and the states obtained from imaginary time evolution, or gradient descent, for a given number of ansatz layers. A circuit depth of zero denotes the non-interacting initial state. The lower plot shows the absolute residuals in the real and imaginary parts of the ground state energy value obtained from imaginary time evolution, and gradient descent. The non-interacting initial state gives a TC energy with a negligible imaginary part; this is not plotted, as it would distort the scale. Figure reproduced with permission from Ref. [192].

tance from the desired eigenstate. The optimal metric for a variational simulation may depend on the quantity that we wish to measure. For example, while we might consider using the energy as a metric, this can have known

drawbacks; the classical variational unrestricted Hartree-Fock method can produce lower energies than its restricted counterpart, but at the cost of introducing spin-contamination that leads to the resulting wavefunctions not being spin eigenstates [39]. A good choice of metric will depend on the possibility of enforcing known symmetries, and on the information we wish to extract from the simulation. The number of parameters used in these simulations was  $28L + 1$ , where  $L$  is the number of ansatz layers. This was much smaller than the Hilbert space dimension of  $2^8 = 256$ , for all circuit depths tested. The lower plot in Fig. 4.10 shows the (absolute) real and imaginary components of the energy residual. Once again, while imaginary time evolution was able to find better approximations of the ground state as the circuit depth was increased, gradient descent barely improved upon the non-interacting initial state. It is interesting to compare the gradient descent and imaginary time datapoints for a single layer ansatz. Although the methods achieved similar energy values, imaginary time evolution yielded a much higher fidelity with the true TC ground state. This emphasises how the algorithm attempts to evolve the state in imaginary time, rather than simply optimising the energy.

I also investigated the  $3 \times 2$  FH model, as a more thorough test of using the algorithm to find the ground states of the TC Hamiltonian. I first optimised the  $J$  value of the TC Hamiltonian such that it maximised the fidelity between the right-hand TC eigenvector, and the state produced from an optimised two-layer ansatz (the number of parameters was given by  $46L + 1$ ). The optimal  $J$  value was  $J = -0.6$ . Although optimising the  $J$  value in this way is not efficient, or tractable, for larger system sizes, more efficient classical methods to optimise  $J$  have been developed [215, 216]. Moreover, the expansion of the ground state as a Slater determinant wavefunction can be used to compensate for any shortcomings induced by choosing a

non-optimal  $J$  value. Fig. 4.11 shows the fidelities obtained by targeting the right-hand and left-hand eigenvectors of the TC Hamiltonian, and the ground state of the regular Hamiltonian. The non-interacting initial state (before random perturbations were applied) had a large overlap with the TC right-hand ground state. This fidelity was much higher than the fidelity of 0.81 between the non-interacting initial state and the ground state of the unmodified Hamiltonian. This highlights the additional correlation accounted for when using the TC Hamiltonian.




Figure 4.11: Finding the ground states of the regular and TC Hamiltonians for a  $3 \times 2$  Fermi-Hubbard model. The TC Hamiltonian has  $J = -0.6$ . The figure shows fidelities between the target states specified in the legend, and the states generated by the method given in brackets in the legend. ‘Right (left) TC’ denotes the right (left)-hand lowest energy eigenvector of the TC Hamiltonian. ‘Regular’ denotes the lowest energy eigenvector of the unmodified Hamiltonian. Fidelities were obtained from 10 repetitions of each method (standard error bars are present, but small). Figure reproduced with permission from Ref. [192].

Once again, gradient descent was unable to find the right-hand ground state of the TC Hamiltonian. In contrast, imaginary time evolution was able to successfully find both the right-hand and left-hand eigenstates of the TC

Hamiltonian. For a given circuit depth, I obtained the highest fidelity by using imaginary time evolution to find the TC right-hand eigenvector. A lower fidelity was obtained when using imaginary time evolution to find the ground state of the unmodified Hamiltonian. I obtained even lower fidelities when using imaginary time evolution to find the TC left-hand eigenvector. As noted in Ref. [216], the TC transformation makes the right-hand eigenvector more compact, while making the left-hand eigenvector less compact. This is reflected in my numerical results, where the number of ansatz layers can be likened to the number and degree of fermionic excitations considered. As such, using the TC Hamiltonian may enable us to obtain accurate energies from a quantum simulation, using a lower depth circuit. As the circuit depth of quantum hardware is currently limited by noise, this highlights the potential value of the TC method.

## 4.5 Discussion

In this chapter, I have presented an algorithm for performing ansatz-based imaginary time evolution on quantum computers. This algorithm circumvents the difficulty of constructing a dissipative imaginary time evolution from unitary quantum gates by approximating the true evolution over a short timestep with the closest unitary evolution achievable with a given parameterized circuit. As with classical implementations of imaginary time evolution, this algorithm can be used to find the ground states of quantum systems.

In Sec. 4.4.1 I presented the results of numerical simulations which benchmarked the behaviour of the algorithm. Numerical simulations on the  $H_2$  molecule in a minimal basis set showed that the algorithm closely follows the true imaginary time trajectory. Simulations of the LiH molecule compared the performance of ansatz-based imaginary time evolution and gra-

dient descent, for finding the ground state. Imaginary time evolution was able to converge more rapidly and reliably. In many cases, these algorithms can have similar resource requirements.

I then applied imaginary time evolution to find the ground states of non-Hermitian, explicitly correlated electronic structure Hamiltonians, in Sec. 4.4.2. Unlike gradient descent, the imaginary time algorithm was able to successfully find the ground states of small transcorrelated Fermi-Hubbard Hamiltonians. As expected, using the transcorrelated Hamiltonian enabled the algorithm to find the right-hand eigenstate with a reduced circuit depth (compared to the regular Hamiltonian) – but increased the circuit depth required to approximate the left-hand eigenstate.

Since its proposal in Ref. [191], the imaginary time algorithm has been applied in a number of settings, and by a number of research groups. The algorithm was extended to: find excited states [230], simulate general processes [231], and propagate mixed states [202]. The algorithm has also been applied to: recompiling [232] and discovering [233] quantum circuits, training quantum Boltzmann machines for quantum machine learning [234], simulating models of quantum field theories [235], solving systems of linear equations [236, 237], and pricing financial options [238]. An alternative algorithm was also developed, that approximates imaginary time evolution on a quantum computer without the use of an ansatz [239]. This algorithm introduces a Trotterized approximation to the true imaginary time evolution, and then approximates the imaginary time evolution of the state over a small time interval with a unitary real time evolution. The unitary applied is determined through partial tomography of the state, and solving a resulting system of linear equations. This process remains efficient while the correlations in the system remain local. This algorithm bears similarities to the one discussed herein, but with a relaxation of the circuit structure

enforced by the chosen ansatz circuit.

Several open questions remain to determine the practicality of the algorithm introduced in this chapter. Most notably, there remains the challenge of selecting a suitable ansatz circuit, that is able to well approximate the true imaginary time evolution. While there have been initial steps made in this direction [240], there is not yet a conclusive answer for how well a given ansatz will perform. One route towards this goal may be to take inspiration from ansatz classes that are used in classical simulations of imaginary time evolution [193].

There are a number of avenues for exploring applications of the algorithm to transcorrelated systems. It would be interesting to extend the simulations presented herein to treat transcorrelated molecular Hamiltonians. These simulations are more difficult to perform, due to the large number of resulting Hamiltonian terms. However, the use of the transcorrelated Hamiltonian enables us to work in a smaller basis, thus reducing the number of qubits required for the simulation. Alternatively, one might explore the open question of how to best measure observables other than the energy, given the expected difficulty of obtaining the left-hand eigenvectors. Finally, there are a number of alternative TC transformations [241, 242], which may be even more effective than the methods discussed in this chapter.

Nevertheless, the imaginary time algorithm introduced in this chapter provides a competitive method for finding the ground states of quantum systems on quantum hardware, and is able to circumvent some of the challenges associated with simulating unphysical imaginary time evolution, on a quantum computer. These attributes have led to its use in a number of

follow-on investigations.

## 5 | Error mitigation for quantum simulation

### Contents

---

<a href="#">5.1 Chapter summary</a>	82
<a href="#">5.2 Introduction</a>	82
<a href="#">5.3 Symmetry verification</a>	84
<a href="#">5.3.1 Particle number parity</a>	85
<a href="#">5.3.2 Spin parity</a>	88
<a href="#">5.3.3 Particle number</a>	89
<a href="#">5.3.4 Practical considerations</a>	93
<a href="#">5.4 Results</a>	94
<a href="#">5.5 Discussion</a>	98

---

### 5.1 Chapter summary

This chapter introduces an error mitigation technique that can be used to reduce the impact of noise in quantum simulation algorithms, by monitoring the conservation of symmetries. The work in this chapter was originally published in Ref. [159]:

*Error-Mitigated Digital Quantum Simulation*  
S. McArdle, X. Yuan, S.C. Benjamin  
[Physical Review Letters 122 \(18\), 180501 \(2019\)](#)

My contributions were: developing the idea and implementation (with input on circuit design from Xiao Yuan), and performing all numerical emulations.

### 5.2 Introduction

As discussed in Chapter 3, quantum simulation algorithms can take a number of forms, including: finding the ground states of electronic structure Hamiltonians, evolving spin systems in time, or calculating the transitions between vibrational energy levels. The most thorough investigations into

the resources required to solve classically intractable instances of these problems have all reached similar conclusions. While only 50 – 100 logical qubits may be sufficient, algorithms with provable accuracy bounds yield deep circuits, requiring at least  $10^7$  gates with current approaches [25, 26, 69, 243]. This necessitates the use of quantum error correction, which leads to a large overhead in the number of qubits required. At the time of writing, leading approaches for solving challenging instances of the electronic structure problem require at least  $\mathcal{O}(10^5)$  physical qubits, at  $10^{-3}$  physical error rates [25].

In contrast, the variational quantum eigensolver (VQE) attempts to extract maximum value from the limited coherence time available in non-error-corrected quantum hardware [94]. Nevertheless, while the VQE has been found to possess some resilience to coherent errors [91, 122], its results can still be corrupted by even small amounts of stochastic noise [92, 135]. This is problematic for the viability of quantum algorithms to find the ground states of chemical systems, where high accuracy (at least one part in 1000) is required. Errors can act disastrously in quantum simulation algorithms; creating or destroying particles, or causing spins to relax.

The introduction of quantum error correction or the reduction of physical error rates both come at considerable cost. An actively error corrected qubit has not yet been experimentally demonstrated (at time of writing) due to the difficulties in engineering sub-threshold error rates, measurement feed-forward, and live syndrome decoding. Moreover, while physical error rates have been reduced to  $10^{-3}$  in some systems [147, 149, 150], they have not improved further since 2016. Given the difficulties in these approaches, it is desirable to develop alternative, low cost methods to effectively suppress the effect of errors.

Prior to the development of the work discussed in this chapter, there were three main methods proposed to mitigate the effect of errors in quantum algorithms: error extrapolation [29, 152], the quasiprobability method [29, 30], and the quantum subspace expansion [157, 158]. As discussed in Sec. 3.2.1, these methods effectively lower the error rate of the simulation by taking additional measurements, and using post-processing to extract the noiseless signal from the noisy results. Since the work discussed in this chapter was proposed, a number of additional error mitigation techniques have been developed, as discussed in Sec. 5.5.

In this chapter I discuss a method of mitigating errors in quantum simulation algorithms. This approach uses checks on a suitably constructed trial state to filter errors. It can be used in isolation, or combined with other error mitigation techniques. In Sec. 5.4 I present numerical emulations of applying the technique to an electronic structure calculation on the  $H_2$  molecule. The technique is applicable to calculations of both static properties (such as ground and excited states, and vibrational spectra), and dynamical properties (such as time evolved correlation functions). Herein I focus solely on its application to the ground state problem.

### 5.3 Symmetry verification

Many quantum simulation algorithms exploit the fact that physically relevant states typically reside in a smaller subspace of the full Hilbert space. For example, in VQE simulations, it can be advantageous to begin in a mean-field state, and use a particle-number and spin conserving ansatz [128, 144]. Some ansätze, such as those suggested in Refs. [82, 144, 244] are constructed from individual gates which conserve particle number. If a single bit-flip error occurs, it will create or destroy an electron, radically changing the state. In other number and spin conserving ansätze, like the singlet

unitary coupled cluster (UCC) ansatz (which, in its canonical form [123], is constructed from individual gates which do not necessarily conserve particle number), a single error can spread and degrade the final state even further. In the absence of errors, states are produced with the correct number of: electrons,  $N$ , spin-up electrons,  $N_\uparrow$ , and spin-down electrons,  $N_\downarrow$ . As these quantities are conserved, their relevant parity operators are also conserved;  $\hat{P}_N |\psi\rangle = (-1)^N |\psi\rangle$  and  $\hat{P}_{N_\uparrow/\downarrow} |\psi\rangle = (-1)^{N_{\uparrow/\downarrow}} |\psi\rangle$ . This is similar to the concept of stabiliser states used in quantum error correcting codes [245]. The error mitigation technique introduced in this chapter (later dubbed ‘symmetry verification’ [246]), uses checks on these conserved quantities to detect and filter certain errors.

### 5.3.1 Particle number parity

To detect errors, I introduce an ancilla qubit, and use it to perform measurements of the conserved quantities. The most simple check is of the total electron number parity. This procedure is shown in Fig. 5.1. The circuit enables the detection of any error which creates or destroys an odd number of electrons. In order to derive an error detection rate, I make the following assumptions:

1. The error model is a two-qubit, symmetric depolarising channel, following two-qubit gates.
2. The error rate is low, such that only one error event occurs.
3. The ansatz circuit is built from individual gates which conserve particle number and spin (so errors do not spread to multiple qubits).

While symmetry verification is still applicable under higher noise rates, different noise models, and using other ansätze – as shown by the numerical simulations in Sec. 5.4 – calculating an analytic bound becomes more

difficult without these assumptions.




Figure 5.1: A circuit to check the particle number parity of a trial state. The ancilla should be measured in  $\frac{1}{2}(1 - (-1)^N)$ . If errors occur, and the measured value of the ancilla is not correct, we reinitialise the circuit, rather than measuring the register qubits. The  $R$  gates rotate the qubits into the desired measurement basis, which is not necessarily the computational basis. Reproduced from Ref. [159], with permission.

Under the assumptions above, this parity check can detect 53% of errors. In the assumed noise model, the following errors on qubits  $i, j$  are equally likely

$$\begin{aligned} X_i I_j, \quad I_i X_j, \quad Y_i I_j, \quad I_i Y_j, \quad X_i Z_j, \quad Z_i X_j, \quad Y_i Z_j, \quad Z_i Y_j, \quad (5.1) \\ Z_i I_j, \quad I_i Z_j, \quad X_i X_j, \quad Y_i Y_j, \quad Z_i Z_j, \quad X_i Y_j, \quad Y_i X_j. \end{aligned}$$

All of the errors in the top row change the electron number parity. As a result, this parity check detects  $8/15 \approx 53\%$  of these errors. This is true both if the error occurs during the ansatz circuit (we can detect  $X_i I_j, I_i X_j, Y_i I_j, I_i Y_j, X_i Z_j, Z_i X_j, Y_i Z_j, Z_i Y_j$ ) or during the parity check gate sequence (assuming the use of non-local gates for the parity check, we can detect  $X_r X_a, X_r Y_a, Y_r X_a, Y_r Y_a, Z_r X_a, Z_r Y_a, I_r X_a, I_r Y_a$ , where  $r$  and  $a$  denote register and ancilla qubits, respectively).

This parity check can also be implemented using a different protocol, based on the Hadamard test. This alternative check is shown in Fig. 5.2, and is similar to the circuits used in the imaginary time algorithm discussed in Chapter 4. We can leverage this similarity to perform symmetry verifica-




Figure 5.2: An alternative circuit to measure the particle number parity, based on the Hadamard test. Reproduced from Ref. [159], with permission.

tion of these circuits ‘for free’. The circuits to measure elements of the metric matrix or gradient vector used in the imaginary time algorithm generate the derivative of an ansatz state with respect to one of the parameters. If we focus on the circuit for the gradient elements, then after measuring the ancilla the register is in the state

$$|\phi_{\pm}\rangle = \frac{1}{\sqrt{k_{\pm}}} (|\psi(\vec{\theta})\rangle \pm i\nu H_k \left| \frac{\partial \psi(\vec{\theta})}{\partial \theta_j} \right\rangle) \quad (5.2)$$

where  $|\psi(\vec{\theta})\rangle$  is the ansatz state,  $k_{\pm}$  is a normalisation constant,  $H_k$  is a Hamiltonian Pauli string, and  $\nu$  is a constant for the derivative state vector. For the ansatz circuits and Hamiltonians discussed in this chapter, both  $|\psi(\vec{\theta})\rangle$  and  $H_k \left| \frac{\partial \psi(\vec{\theta})}{\partial \theta_j} \right\rangle$  will have fixed particle number parities (in the absence of errors). As the register qubits are not normally measured in the imaginary time algorithm, we can measure their parity at the same time as measuring the ancilla qubit – providing symmetry verification at no additional cost (beyond an increase in the variance of the estimate). This alternative circuit may also be beneficial when readout errors dominate. If a single readout error occurred in the circuit in Fig. 5.1, it would likely happen on one of the register qubits. Consequently, we would measure the correct parity, and thus ‘accept’ an incorrect energy measurement. In contrast, using the circuit in Fig. 5.2 the error would be most likely to occur on the parity check qubits. This would cause us to incorrectly filter a correct

energy measurement. While this would increase the variance of our result, it would not bias our energy estimate.

### 5.3.2 Spin parity

The circuit shown in Fig. 5.3 can detect additional errors by measuring the parity of both the number of spin-up and spin-down electrons. The spin-parity checks can detect additional two-qubit errors, that change the value of either spin-parity. However, we are still not able to detect all two-qubit bit-flip errors.




Figure 5.3: A circuit to measure the spin-parities. We compute the spin-up parity onto the ancilla, and measure it. We then reset the ancilla to  $|0\rangle$ , and measure the spin-down parity. Reproduced from Ref. [159], with permission.

In the following analysis, I assume that the ansatz circuit is equally likely to apply gates between any two qubits. This may not hold for all ansatz circuits that are particle and number conserving. There are  $\binom{M}{2} = \frac{M(M-1)}{2}$  ways of distributing a two-qubit bit-flip error between the  $M$  spin-orbitals. We are able to detect errors which occur on one orbital of each spin. There are  $\binom{M/2}{1} = M/2$  ways of distributing a single bit-flip error amongst half of the orbitals. As a result, there are  $\frac{M^2}{4}$  such errors that we can detect. When  $M$  is large, we can detect  $\frac{M^2/4}{M(M-1)/2} \approx \frac{M^2/4}{M^2/2} = \frac{1}{2}$  of double bit-flip errors. As there are four possible types of double bit-flip errors, and we can detect half of the occurrences of each of them, this effectively increases the number of detectable errors by  $2/15$ , to  $10/15 \approx 66\%$ . This analysis only holds when the error occurs in the ansatz circuit. However, as the ansatz circuit likely

contains the majority of the gates in the circuit, errors that occur during the parity check should only have a small impact on this detection rate.

### 5.3.3 Particle number

We can also measure the electron number and spin numbers directly, using an iterative procedure. We first write the electron number in binary. We then use the circuit in Fig. 5.4 to measure the first (i.e. rightmost) bit in  $N$ , denoted as  $N_1$ . In the absence of errors, the ancilla is in the state

$$|\phi\rangle = \frac{1}{\sqrt{2}}(|0_a\rangle + e^{N_m\pi i}|1_a\rangle). \quad (5.3)$$

If  $N_m = 0$  we measure the ancilla in  $|+\rangle$ , while if  $N_m = 1$  we measure the ancilla in  $|-\rangle$ . We then repeat the circuit to measure  $N_2$ , using the measurement of  $N_1$  in the rotation  $\omega_2$ . In general, we can generate  $\omega_m$  using the measurements of the  $m - 1$  preceding bits in  $N$ . In total,  $M\log_2(N)$  control gates are needed to measure the electron number, or both spin numbers.




Figure 5.4: The circuit which measures the  $m^{\text{th}}$  bit of the electron number,  $N_m$ . The  $R_m$  gates are given by  $\text{diag}(1, e^{\pi i/2^{m-1}})$ . The gate  $\omega_m$  is given by  $\text{diag}(1, e^{-\text{dec}(N_{m-1}\dots N_1)\pi i/2^{m-1}})$ , where  $\text{dec}(N_{m-1}\dots N_1)$  is the decimal representation of the binary string  $N_{m-1}\dots N_1$ .  $\omega_1$  is taken to be the identity matrix. The ancilla is measured in the  $X$  basis. Reproduced from Ref. [159], with permission.

I provide a worked example for the case of a system with  $N = 3$  electrons in  $M = 6$  orbitals. First apply the circuit shown in Fig. 5.5




Figure 5.5: The circuit which measures the 1<sup>st</sup> bit of the electron number,  $N_1$ . The  $R_1$  gates are given by  $\text{diag}(1, e^{\pi i})$ . Measurement of the ancilla is in the  $X$  basis. Reproduced from Ref. [159], with permission.

After this circuit, if no errors have occurred, the ancilla is in the state

$$|\phi\rangle = \frac{1}{\sqrt{2}}(|0_a\rangle + e^{3\pi i}|1_a\rangle) = |-\rangle. \quad (5.4)$$

Measuring the ancilla in the  $X$  basis yields  $N_1 = 1$ . We then apply the circuit in Fig. 5.6. The  $\omega_2$  gate is given by  $\text{diag}(1, e^{-\text{dec}(N_1)\pi i/2}) = \text{diag}(1, e^{-\pi i/2})$ .

The state of the ancilla is

$$|\phi\rangle = \frac{1}{\sqrt{2}}(|0_a\rangle + e^{\pi i}|1_a\rangle) = |-\rangle. \quad (5.5)$$

As a result, we measure  $N_2 = 1$ . The number of electrons is therefore measured as 3.

It is difficult to determine the detection rate of the particle number / spin number check. I first assume that the wavefunction consists of only one Slater determinant, and that the error event occurs in the ansatz circuit. I also assume an equal number of spin-up and spin-down electrons. We cannot detect any two-qubit bit-flip error which acts on occupied and unoccupied orbitals of the same spin, as this error mimics a spin conserving excitation operator. Considering the spin-up number, then there are  $\frac{N}{2} \times \frac{M-N}{2}$




Figure 5.6: The circuit which measures the 2<sup>nd</sup> bit of the electron number,  $N_2$ . The  $R_2$  gates are given by  $\text{diag}(1, e^{\pi i/2})$ . The  $\omega_2$  gate is given by  $\text{diag}(1, e^{-\text{dec}(N_1)\pi i/2})$ , where  $\text{dec}(N_1)$  is the decimal value of the first bit of  $N$ . Measurement of the ancilla is in the  $X$  basis. Reproduced from Ref. [159], with permission.

such errors. The fraction of undetectable errors,  $f_u$ , is then given by

$$\begin{aligned}
 f_u &= \frac{\frac{N}{2} \times \frac{M-N}{2}}{\frac{M}{4} \left(\frac{M}{2} - 1\right)} \\
 &= \frac{N(M-N)}{M\left(\frac{M}{2} - 1\right)} \\
 &\approx \frac{NM\left(1 - \frac{N}{M}\right)}{\frac{M^2}{2}} \\
 &= 2\eta(1 - \eta),
 \end{aligned} \tag{5.6}$$

where  $\eta = \frac{N}{M}$  is the filling fraction of the system, and I have assumed that  $M \gg 1$ . This expression obtains a maximum value of  $f_u = 0.5$  at  $\eta = 0.5$ . As a result, in the worst case (at half filling) the spin number check can detect half of the two-qubit bit-flip errors occurring on orbitals of the same spin. Overall then, if we were to consider the spin number checks, we would be able to detect:

- All single-qubit bit-flip errors (8/15 errors in the depolarising model).
- All two-qubit bit-flip errors where each happens on orbitals of a dif-

ferent spin (as discussed in Sec. 5.3.2, these errors account for 50% of two-qubit bit-flip errors, so  $1.0 \times 0.5 \times 4/15 = 2$  errors in the depolarising model).

- At least half of the two-qubit bit-flip errors where both errors happen on orbitals of the same spin (these errors account for the other 50% of two-qubit bit-flip errors, so  $0.5 \times 0.5 \times 4/15 = 1$  error in the depolarising model).

As a result, the spin number check can detect at least  $11/15 \approx 73\%$  of depolarising errors, and at most  $12/15 = 80\%$  of depolarising errors (this would correspond to detecting all bit-flip errors).

This analysis only applies to a single Slater determinant, while in reality the trial state will contain multiple determinants. Unfortunately, double bit-flip errors can change the electron and spin numbers of some determinants, while leaving others unchanged. For example, consider the state

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|001011\rangle + |100101\rangle), \quad (5.7)$$

which has an electron number of 3, a spin up number of 2, and a spin down number of 1 (the rightmost 3 orbitals are spin-up, the leftmost 3 orbitals are spin-down). Under the double bit-flip error  $X_0X_1$ , the state becomes

$$|\psi'\rangle = \frac{1}{\sqrt{2}}(|001000\rangle + |100110\rangle). \quad (5.8)$$

This state has the correct spin parities, so the error is undetectable using parity checks. While the first determinant has an incorrect particle number, the second has the correct particle number. Applying the particle number check described above, we measure  $N_1 = 1$ . When measuring the second bit, we measure  $N_2 = 1$  with a 50% probability. If we measure  $N_2 = 0$ , we

correctly discard the state. However, measuring  $N_2 = 1$  would lead us to believe that no errors have occurred. The wavefunction will collapse to

$$|\psi''\rangle = |100110\rangle, \quad (5.9)$$

The fidelity of this state with the true wavefunction is related to the probability of obtaining this measurement outcome. As a result, the detection rate of the particle number check will be below 80% in almost all cases.

While symmetry verification can filter a fraction of the possible errors, the results will not be completely noise free. It is possible to combine symmetry verification with extrapolation by first calculating the expectation value of the observable at the baseline noise rate, while applying symmetry verification. The noise can then be boosted, again applying symmetry verification to mitigate bit-flip errors. These two expectation values obtained at different noise rates can then be used to estimate the noiseless value of the observable.

### 5.3.4 Practical considerations

When experimentally implementing these checks, we must take the limitations of the hardware into account. The circuits are similar to that of a stabiliser evaluation during quantum error correction, which has been investigated for trapped ion and superconducting systems [188, 247]. As can be seen from Fig. 5.1, the optimal implementation requires non-local gates, which are feasible for some trapped ion systems [248]. As the coherence times of trapped ion qubits are considerably longer than their readout times [249], it may be possible to carry out the checks using a single ancilla that is repeatedly measured and reinitialised.

In contrast, superconducting qubits are typically limited to nearest-neighbour

connectivity. We can then realise the parity checks through a ladder of nearest-neighbour CNOT gates, which pass the information along the register to the ancilla (this was observed by Xiao Yuan and in Ref. [246]). This requires at most  $2M$  nearest-neighbour CNOT gates for the parity check in Fig. 5.1. In contrast, the number of gates required for a general UCC ansatz is  $\mathcal{O}(M^3)$  [250]. For calculations on  $M = 50 - 100$  qubits, the ansatz will be the dominant component of the circuit. In superconducting systems the measurement time can be of a comparable order of magnitude to the coherence time [251]. It may therefore be preferable to use multiple ancilla qubits, rather than to repeatedly reinitialise a single ancilla. This modest overhead constitutes two ancilla qubits for the spin-parity check, and  $\log_2 N$  ancilla qubits for the spin number check.

## 5.4 Results

I tested the efficacy of symmetry verification using a VQE calculation on  $\text{H}_2$  in a minimal basis set (two electrons in four spin-orbitals, as discussed in Sec. 2.2.2). I did not consider the parameter update step of the VQE, so as to examine the effect of errors without consideration of a classical optimisation algorithm. Numerical simulations were performed using QuEST, and simulation code can be found at Ref. [252].

I applied a spin-conserving UCC ansatz to the Hartree-Fock state, as discussed in Sec. 3.1.3.1. To detect errors, I performed error-prone checks of both the spin-up and spin-down parities, using the circuit shown in Fig 5.7. This circuit has nearest-neighbour connectivity, and so lower bounds the efficacy of the method. The simulations were designed to mimic the actions of an experimentalist; the energy expectation value was calculated using the Hamiltonian averaging procedure (see Sec. 3.1.3). The number of measurements used is discussed in Appendix A.1.




Figure 5.7: A circuit using a local gateset to measure both the spin-up parity and spin-down parity. The  $R$  gates implement the single-qubit basis rotations required to measure the desired Pauli string.

Initially, I considered energy measurements on a trial state that contained all four possible vectors with two electrons in four spin-orbitals and  $s_z = 0$  (see Sec. 3.1.3.1). A symmetric depolarising noise channel was applied after all single-qubit and two-qubit gates. The two-qubit gate error rate was set 10 times larger than the single-qubit gate error rate. I compared the error in the energy value when applying: no error mitigation, only symmetry verification, only linear extrapolation, and symmetry verification with linear extrapolation. There were 92 single-qubit gates and 56 two-qubit gates in the UCC ansatz circuit – which I treat as being equivalent to 65 two-qubit gates. The parity checks contributed an additional 8 error-prone two-qubit gates.

Fig. 5.8 shows that symmetry verification alone improved the accuracy of the energy measurements, but was less effective than extrapolation. A greater benefit was obtained by combining the two methods. The accuracy of the combined method only worsened slightly as the error rate increased, unlike the two individual methods. The inset plot shows the fraction of detected errors falling approximately linearly with increasing error rate.




Figure 5.8: Comparing methods of mitigating errors in simulations of  $H_2$ . The detection rates shown in the inset were obtained from the numerical simulations. The true energy value was  $-1.1227$  Hartree. The error bars upper bound the standard error in the result. Reproduced from Ref. [159], with permission.

When the error rate was small, symmetry verification detected around 53% of errors. Assuming that symmetry verification could only detect errors which occurred in the ansatz circuit (due to the restricted connectivity considered) then the detection probability can be estimated as the probability of an error happening in the ansatz circuit ( $65/73$ ), multiplied by the probability of detection ( $10/15$ ), which is roughly 59%. I attribute the deviation from this value to the use of a UCC ansatz built from gates which do not individually conserve particle number and spin. This enabled errors to propagate. At higher error rates, multiple errors were likely to occur in the circuit, reducing the fraction of detectable errors to around one third.

I also applied symmetry verification to a calculation of the dissociation curve of  $H_2$ . I compared the energy values obtained when applying: no

error mitigation, just extrapolation, and extrapolation with symmetry verification. The error model consisted of a symmetric depolarising channel (two-qubit error rate of 0.1%, single-qubit error rate of 0.01%, which have previously been experimentally demonstrated [149, 150]), combined with temporally correlated over/under rotations of up to 1%. In the temporally correlated noise model, each discrete gate in the circuit was replaced by a parameterized equivalent. For example, the Hadamard gate  $H$  was replaced with

$$H = \begin{bmatrix} \cos(\epsilon\frac{\pi}{4}) & \sin(\epsilon\frac{\pi}{4}) \\ \sin(\epsilon\frac{\pi}{4}) & -\cos(\epsilon\frac{\pi}{4}) \end{bmatrix} \quad (5.10)$$

Similarly, the CNOT gates were replaced with

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \text{icos}(\epsilon\frac{\pi}{2}) & \text{sin}(\epsilon\frac{\pi}{2}) \\ 0 & 0 & \text{sin}(\epsilon\frac{\pi}{2}) & \text{icos}(\epsilon\frac{\pi}{2}) \end{bmatrix} \quad (5.11)$$

All rotation gates  $R(\theta)$  were replaced with  $R(\epsilon\theta)$  (except for those used to change the measurement basis when measuring a Hamiltonian term). When the parameter  $\epsilon = 1$  the original gates were recovered. On each iteration of the circuit, I generated a random  $\delta$  for each qubit,  $\delta_q$ , uniformly distributed between  $\pm 0.01$ . For all of the single-qubit gates acting on qubit  $q$  in the circuit, I set  $\epsilon_q = 1 - \delta_q/10$ . For the target qubit of two-qubit gates, I set  $\epsilon_q = 1 - \delta_q$ . As such,  $0.99 < \epsilon_q < 1.01$ . In a given iteration of the circuit, each qubit received the same over or under rotation for each of the gates where it was the target qubit.

Fig. 5.9 shows that the combination of symmetry verification and extrapolation was able to obtain chemically accurate energies, even when the results would have otherwise been corrupted by noise. The combined mitigation




Figure 5.9: Comparing methods of mitigating errors in simulations of  $\text{H}_2$ . The inset shows the residual from the true value. The dashed lines mark chemical accuracy ( $\pm 1.6$  mHartree). The error bars show the true standard error in the result, given in Appendix A.1. Reproduced from Ref. [159], with permission.

method yielded a mean absolute residual of 0.2 mHartree. Compared to the unmitigated results, the deviation from the true value was reduced by a median factor of 239 (with a range of 141 to 818, and a mean of 340). Compared to just applying extrapolation, the deviation from the true value was reduced by a median factor of 9.1 (with a range of 6.6 to 35.6, and a mean of 15.0) by combining the two mitigation methods. As energy differences are exponentiated when calculating reaction rates, this improvement may be magnified when performing certain calculations of interest.

## 5.5 Discussion

In this chapter I have introduced symmetry verification; a low cost method for mitigating errors in quantum simulation algorithms. Symmetry verification detects errors using stabiliser check-like circuits which measure the

value of quantities that should be conserved during the algorithm. While it cannot detect all first order errors, it is effective at filtering out bit-flip errors, achieving detection rates of between 53% - 80% (under idealised conditions). I tested the method in noisy emulations of a VQE calculation on the  $H_2$  molecule. By combining symmetry verification with linear extrapolation, essentially noiseless results could be recovered. The error in the energy was reduced by two orders of magnitude compared to the unmitigated results, in some cases.

Since my coauthors and I proposed it in Ref. [159], symmetry verification has been further developed in a number of other publications. Shortly after the release of Ref. [159], a similar proposal was released in Ref. [246]. That work also introduced an ancilla-free version of symmetry verification, which utilises additional measurements and post-processing (and was later experimentally demonstrated [161]). Ref. [253] presented a method that builds upon the alternative parity check circuit in Fig. 5.2. Rather than measuring the particle number parity of the register qubits, that work applies the inverse of the ansatz circuit before measuring the register qubits. Measuring the register qubits in  $|\bar{0}\rangle$  provides a way to ‘verify’ the computation, yielding one of the most effective error mitigation techniques suggested to date.

Other works have combined symmetry verification with other error mitigation techniques. Ref. [151] showed that by optimising the combination of symmetry verification and extrapolation, the method could be effective even when multiple errors were expected to occur in the circuit. Ref. [254] considered combining error detection codes with symmetry verification, and using ancilla-free post-processing to realise the desired post-selection.

Perhaps the most interesting open question from this work (and one which I have spent considerable time thinking about), is whether monitoring conserved quantities, such as the particle number parity, could ever be enough to *correct* errors, rather than just to detect them. Clearly, for a completely general noise model this will not be possible, as the symmetry provides little information about the state. However, if we were to consider a very specialised noise model, perhaps there are cases where this goal becomes achievable. For example, consider a fault tolerant setting, and assume that all logical Clifford operations are noiseless. In this case, the only source of errors in the circuit would be during non-Clifford logical operations, such as arbitrary angle single-qubit rotations. For example, consider the circuit shown in Fig. 5.10, where each qubit is a logical qubit, composed of a number of physical qubits. This circuit implements a logical  $e^{i\theta X_1 Y_2}$  gate, which conserves particle number parity.




Figure 5.10: A circuit on logical qubits that implements the particle number parity conserving logical gate  $e^{i\theta X_1 Y_2}$ . The Clifford gates in this circuit are considered noiseless. I imagine implementing the  $R_z(\theta)$  gate through a CNOT-ladder construction on the physical qubits, which would not be fault tolerant. I assume here for convenience that this leads to a single-qubit depolarising noise channel for the logical qubit. Note that if  $R_z(\theta)$  were implemented by magic state injection, then it would suffer from  $Z$  type errors, which symmetry verification is unable to detect.

Under the noise model discussed above, only the  $R_z(\theta)$  gate may lead to an error. I assume that this noisy gate is followed by single-qubit depolarising noise. In this case, using a particle number parity check we would be able to detect an  $X$  error on the second qubit. If this were the only possibility, and the check was performed immediately before and after this circuit, then

we would be able to correct this error. Unfortunately, a  $Z$  error would be mapped to  $X_1Y_2$  by the circuit, and so would be undetectable. In general, it looks difficult (if not impossible) to find a circuit that is able to provide computational universality for the chemistry problem, while spreading errors in such a way that they would always be detectable using a physical symmetry. In some ways, this is reminiscent of the Eastin-Knill theorem [173], which states that no quantum error correction code can implement a universal set of quantum gates transversally. Whether these limitations can be circumvented in this case, by considering alternative circuit constructions or resource states is an interesting open problem. If a solution were possible, then perhaps this would enable a reduction in the resources required for error correction, or the elimination of costly processes like magic state distillation and injection.

## 6 | Quantum algorithms for simulating molecular vibrations

### Contents

---

<a href="#">6.1 Chapter summary</a> . . . . .	102
<a href="#">6.2 Introduction</a> . . . . .	102
<a href="#">6.3 Mapping vibrations to qubits</a> . . . . .	106
<a href="#">6.4 Simulating vibrational problems</a> . . . . .	113
<a href="#">6.5 Discussion</a> . . . . .	119

---

### 6.1 Chapter summary

This chapter discusses the use of digital quantum algorithms to simulate molecular vibrations. This work was originally published in Ref. [255]:

*Digital quantum simulation of molecular vibrations*  
S. McArdle, A. Mayorov, X. Shan, S.C. Benjamin, X. Yuan  
[Chemical Science, 10, 5725-5735 \(2019\)](#)

The main novel results in this work were: developing a formulation of the vibrational structure problem amenable for simulation on digital quantum computers, introducing a classically inspired ansatz tailored for solving the vibrational structure problem, considering practical approaches to preparing vibrational thermal states, and developing techniques to measure absorption spectra. My contributions were: developing the vibration-to-qubit mappings, developing methods to evaluate Franck-Condon spectra, introducing the unitary vibrational coupled cluster ansatz, and considering how to prepare vibrational thermal states. Numerical emulations of some of these approaches were reported in the original publication [255], but were performed by my collaborator, Xiao Yuan, and so I do not discuss them in this chapter.

### 6.2 Introduction

So far, this thesis has focused on methods to solve the electronic structure problem on quantum hardware. However, finding the low-lying electronic

energy levels of chemical systems is often only the first step in trying to understand and predict their behaviour. For molecules composed of light atoms (i.e. those where relativistic effects can be neglected), the vibrational motion of the nuclei is the next largest contribution to the total energy of the molecule, after the electronic energy. As a result, calculations of static properties (such as energy differences between configurations) must often include the corrections due to vibrational structure in order to obtain accurate results [256, 257].

Vibrations of the nuclei are also important in a number of other chemical scenarios. They determine absorption properties in Raman or infrared spectroscopy [258, 259], as well as governing the behaviour of the molecule in ultrafast laser experiments [260]. Unfortunately, accurately determining the vibrational structure of many molecules can be just as challenging as the electronic structure problem.

In a sense, the existence of the vibrational structure problem is a fallacy, stemming from our use of the Born-Oppenheimer approximation. If this approximation is not made, then we can simulate the electrons and nuclei on an equal footing, and solve for their total energy. Such simulations can be performed most faithfully by using real-space grid-based methods, which are limited to small system sizes by their high costs. In many systems, it is acceptable to introduce the Born-Oppenheimer approximation, separating the electronic and nuclear motions, due to their vastly different timescales. This introduces the concept of potential energy surfaces, where we can consider an electronic energy surface that depends on the positions of the nuclei, which are assumed to be clamped in place.

The vibrational structure problem then assumes that the system is in equilibrium at the minima of these potential energy surfaces. Vibrations of the nuclei consist of small perturbations about this equilibrium geome-

try. These oscillations are governed by the potential experienced by the nuclei, which is determined by the shape of the electronic potential energy surface in the vicinity of the minimum. Evaluating a multidimensional potential energy surface by sampling the electronic energy at a number of nuclear configurations will require a number of measurements scaling exponentially with the number of vibrational modes [261]. As a result, the form of the vibrational potential is typically approximated. One possible approach is to consider sophisticated adaptive sampling [262] or machine learning [263] techniques. Alternatively, one can consider analytic approximations for the potential (e.g. the Morse potential for diatomic systems), or a multivariable Taylor expansion of the potential energy surface [261, 264]. The latter approach is most commonly used in the literature, with the Taylor series often being truncated at second order. This is known as the ‘harmonic approximation’ to the vibrational structure problem, which treats the potential as that of a harmonic oscillator. Improved accuracy can be obtained by including higher-order ‘anharmonic’ terms in the potential. While the harmonic approximation makes the vibrational structure problem more tractable to simulate classically, it breaks down in many cases – for example, the treatment of non-rigid molecules, or when considering highly excited vibrational states [259]. Once the potential has been obtained, approximate wavefunctions can be generated using a number of electronic structure-inspired methods [265], such as: vibrational self-consistent field (VSCF, the vibrational analogue of Hartree-Fock), vibrational coupled cluster (VCC, the analogue of CC methods), and vibrational full configuration interaction (VFCI) – which amounts to exact diagonalisation of the Hamiltonian. As with the corresponding approaches to solving the electronic structure problem, these methods have an accuracy-dependent cost, which limits the system size that they can be applied to. Simulations accurate to VFCI level are constrained to system

sizes with around 4 atoms [266]. Using approximate methods, accurate simulations of some non-rigid molecules have been performed on system sizes of up to 20 atoms [267, 268].

The difficulty of classically solving the vibrational structure problem has led to the development of several quantum algorithms. Prior to the work discussed in this chapter, all of the quantum algorithms specialised to the vibrational structure problem used analog quantum simulators as the simulation platform [269–276]. These approaches exploit the correspondence between the harmonic approximation to the vibrational structure problem, and the Hamiltonian of a bosonic system (such as photons in waveguides, or energy levels in trapped ions). While these proposals provide a natural simulation platform for small system sizes, they have two main limitations. The first is that it can be difficult to move beyond the harmonic approximation in such systems – for example, it is difficult to engineer non-linearities (required to generate anharmonic terms) in photonic networks. Secondly, it is not yet known how to protect analog quantum simulations from errors in a scalable way. As a result, it seems likely that large-scale vibrational structure calculations will need to be performed on digital quantum computers, for which methods of fault tolerant error correction are known.

In this chapter, I discuss approaches for simulating molecular vibrations on digital quantum computers. In Sec. 6.3 I discuss two methods for mapping the vibrational problem onto a system of qubits. In Sec. 6.4 I discuss techniques introduced to calculate vibrational properties of interest, such as energy levels and transition amplitudes.

### 6.3 Mapping vibrations to qubits

As discussed in Chapter 2, the Coulomb Hamiltonian for a molecular system containing  $G$  nuclei, prior to making the Born-Oppenheimer approximation, is given by

$$H = - \sum_i \frac{\nabla_i^2}{2} - \sum_I \frac{\nabla_I^2}{2M_I} - \sum_{i,I} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{I \neq J} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|}. \quad (6.1)$$

where  $M_I$ ,  $\mathbf{R}_I$ , and  $Z_I$  denote the mass, position, and atomic number of the  $I^{\text{th}}$  nucleus, and  $\mathbf{r}_i$  is the position of the  $i^{\text{th}}$  electron. I assume that the wavefunction can be factorized into a tensor product of a nuclear and an electronic wavefunction  $|\phi_N\rangle \otimes |\psi_e\rangle$ . For fixed nuclear positions, one can integrate out the electronic part of the wavefunction by solving the electronic Schrödinger equation. For example, solving for the electronic ground state potential energy surface gives the nuclear Hamiltonian

$$H_v(\mathbf{R}_1, \dots, \mathbf{R}_G) = - \sum_I \frac{\nabla_I^2}{2M_I} + V_e(\mathbf{R}_1, \dots, \mathbf{R}_G) \quad (6.2)$$

where

$$V_e(\mathbf{R}_1, \dots, \mathbf{R}_G) = \min_{\psi_e} \langle \psi_e | \left( - \sum_i \frac{\nabla_i^2}{2} - \sum_{i,I} \frac{Z_I}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{I \neq J} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} \right) | \psi_e \rangle \quad (6.3)$$

is the electronic potential energy surface for the molecule. In general, it is not possible to obtain an analytic expression for the potential energy surface. As mentioned in Sec. 6.2, one of the most simple ways to approximate the potential energy surface is with a multivariable Taylor expansion about the equilibrium positions of the nuclei. I denote the displacement of nucleus  $I$  from its equilibrium coordinate along each axis as  $X_I = (R_I - R_I^0)$ .

The potential can then be expanded as

$$\begin{aligned}
 V_e(\vec{X}_I) \approx & V_e(\vec{0}) + \sum_I \left( \frac{\partial V_e}{\partial X_I} \right) \Big|_{X_I=0} X_I \\
 & + \sum_{I,J} \left( \frac{\partial^2 V_e}{\partial X_I \partial X_J} \right) \Big|_{X_I, X_J=0} X_I X_J + \dots
 \end{aligned} \tag{6.4}$$

The first term in this expansion is a constant energy offset, while the second term is zero because the equilibrium position is a minimum of the potential energy surface. If we truncate the expansion at second order, we can diagonalise the matrix  $\frac{1}{\sqrt{M_I M_J}} \frac{\partial^2 V_e}{\partial X_I \partial X_J} \Big|_{X_I, X_J=0}$  (known as the Hessian matrix for the system) to obtain the mass-weighted normal mode coordinates for the system,  $Q_I$ . This is analogous to the way that we solve for the normal modes of coupled classical harmonic oscillators. For a general molecule, there are  $M = 3G - 6$  normal modes with non-zero values in the diagonalised Hessian ( $M = 3G - 5$  for a linear molecule). These normal modes correspond to the vibrational degrees of freedom. The modes with an eigenvalue of zero correspond to translational or rotational degrees of freedom. I neglect these terms in the Hamiltonian. The vibrational Hamiltonian is then given by

$$H_v = \frac{1}{2} \sum_{m=1}^M (P_m^2 + \omega_m^2 Q_m^2), \tag{6.5}$$

where  $m$  denotes a vibrational mode. Defining the creation and annihilation operators

$$\begin{aligned}
 Q_m &= \frac{1}{\sqrt{2\omega_m}} (a_m + a_m^\dagger) \\
 P_m &= i\sqrt{\frac{\omega_m}{2}} (a_m^\dagger - a_m), \\
 [a_m, a_n^\dagger] &= \delta_{mn}
 \end{aligned} \tag{6.6}$$

where  $P_m$  is the conjugate momentum to the mass-weighted normal mode coordinate  $Q_m$ , enables the Hamiltonian to be expressed as

$$H_v = \sum_{m=1}^M \omega_m \left( a_m^\dagger a_m + \frac{1}{2} \right). \quad (6.7)$$

The transformation to mass-weighted normal mode coordinates uncouples the vibrational modes, enabling us to approximate the motion as a set of independent harmonic oscillators. The eigenfunctions for this Hamiltonian are tensor products of the Gauss-Hermite polynomials. I represent the  $n^{\text{th}}$  Gauss-Hermite function for vibrational mode  $m$  as  $|n\rangle_m$ , such that

$$\begin{aligned} a_m^\dagger a_m |n\rangle_m &= n |n\rangle_m \\ a_m |n\rangle_m &= \sqrt{n} |n-1\rangle_m \\ a_m^\dagger |n\rangle_m &= \sqrt{n+1} |n+1\rangle_m. \end{aligned} \quad (6.8)$$

As discussed above, the harmonic approximation breaks down when considering non-rigid molecules that undergo large oscillations, or when searching for highly excited vibrational eigenstates. These cases necessitate a higher-order expansion of the potential

$$H_v = \frac{1}{2} \sum_{\alpha} (P_{\alpha}^2 + \omega_{\alpha}^2 Q_{\alpha}^2) + \sum_{\alpha, \beta, \gamma} K_{\alpha\beta\gamma} Q_{\alpha} Q_{\beta} Q_{\gamma} + \sum_{\alpha, \beta, \gamma, \delta} K_{\alpha\beta\gamma\delta} Q_{\alpha} Q_{\beta} Q_{\gamma} Q_{\delta} + \dots \quad (6.9)$$

As this Hamiltonian has been obtained from a Taylor series expansion around a minimum of the electronic potential energy surface, it is expected to only give a good description of the vibrational dynamics in the vicinity of the minimum. This Hamiltonian would not be appropriate for describing complex potential energy surfaces, far from the minimum. We can expand each of the operators  $Q$  and  $P$  into a linear combination of creation and annihilation operators. This is equivalent to using the harmonic oscillator eigenstates as basis functions for the vibrational Hilbert space. For an  $l^{\text{th}}$  order

expansion of the potential, the Hamiltonian will contain  $\mathcal{O}(M^l)$  terms. Although the creation and annihilation operators have a bosonic commutation relation, I treat the different vibrational modes as distinguishable particles, and so do not explicitly consider symmetrising the wavefunction.

I now consider how to map these vibrational energy levels onto qubit states. The first challenge to overcome is to note that the Hilbert space for this problem is infinite dimensional, as  $|n\rangle$  may take an arbitrarily high value. In practice, one can consider the simulation to take place in a low energy subspace, by truncating at a finite value of  $n$ . This can be considered analogous to the basis set discretisation error for the electronic structure problem. I denote the maximum energy level considered as  $d$ .

I consider two possible mappings: a direct mapping requiring  $\mathcal{O}(Md)$  qubits to store the vibrational wavefunction, and a compact mapping requiring  $\mathcal{O}(M\lceil\log_2(d)\rceil)$  qubits. Due to the relatively large number of qubits required in both mappings, it may be the case that these calculations will be intractable for NISQ processors, and so would have to wait for fault-tolerant processors to become available. The direct mapping was originally introduced in the context of simulating bosonic degrees of freedom on quantum computers in Ref. [277]. This mapping uses  $M$  registers, each containing  $d$  qubits. I represent the Gauss-Hermite polynomial basis state for mode  $m$  as

$$|i\rangle_m = |0\rangle_{m_{d-1}} \otimes \dots \otimes |1\rangle_{m_i} \otimes \dots \otimes |0\rangle_{m_1} \otimes |0\rangle_{m_0} \quad (6.10)$$

where  $m_i$  denotes the  $i^{\text{th}}$  qubit in register  $m$ . The creation and annihilation

operators take a particularly simple form under this mapping

$$\begin{aligned}
a_m &= \sum_{\alpha=0}^{d-2} \sqrt{\alpha+1} (|0\rangle \langle 1|_{m_{\alpha+1}} \otimes |1\rangle \langle 0|_{m_\alpha}) \\
&= \sum_{\alpha=0}^{d-2} \frac{\sqrt{\alpha+1}}{4} ((X+iY)_{m_{\alpha+1}} \otimes (X-iY)_{m_\alpha}), \\
a_m^\dagger &= \sum_{\alpha=0}^{d-2} \sqrt{\alpha+1} (|1\rangle \langle 0|_{m_{\alpha+1}} \otimes |0\rangle \langle 1|_{m_\alpha}) \\
&= \sum_{\alpha=0}^{d-2} \frac{\sqrt{\alpha+1}}{4} ((X-iY)_{m_{\alpha+1}} \otimes (X+iY)_{m_\alpha})
\end{aligned} \tag{6.11}$$

A single creation or annihilation operator is mapped to  $\mathcal{O}(d)$  weight-2 Pauli strings. As a result, a vibrational Hamiltonian expanded to  $l^{\text{th}}$  order will contain  $\mathcal{O}(M^l d^l)$   $2l$ -weight Pauli strings.

The compact mapping achieves an exponential reduction in spatial resources over the direct mapping discussed above, at a cost of more complex expressions for the mode operators. The compact mapping was originally introduced in Ref. [278], where it was used as a mapping for Gaussian nuclear orbitals in ‘nuclear orbital plus molecular orbital’ theory, which attempts to treat the electronic and nuclear degrees of freedom on a similar footing. The compact mapping uses  $M$  registers, each with  $\lceil \log_2(d) \rceil$  qubits. I represent the Gauss-Hermite polynomial basis state for mode  $m$  as

$$|n\rangle_m := |\mathbf{n}\rangle_m = \bigotimes_{i=0}^{\lceil \log_2(d) \rceil - 1} |\mathbf{n}_i\rangle_{m_i} \tag{6.12}$$

where  $\mathbf{n}$  is a binary expansion of integer  $n$ , and  $\mathbf{n}_i$  denotes the  $i^{\text{th}}$  bit of  $\mathbf{n}$  (counting from the right). For example, we can store the state  $|5\rangle := |1\rangle|0\rangle|1\rangle$ , for  $d = 8$ . Under this mapping, the annihilation operators are

given by

$$\begin{aligned}
a_m &= \sum_{s=0}^{d-2} \sqrt{s+1} |s\rangle_m \langle s+1|_m \\
&= \sum_{s=0}^{d-2} \sqrt{s+1} \left( \left( \bigotimes_{i=0}^{\lceil \log_2(d) \rceil - 1} |s_i\rangle_{m_i} \right) \left( \bigotimes_{i=0}^{\lceil \log_2(d) \rceil - 1} \langle (s+1)_i|_{m_i} \right) \right) \\
&= \sum_{s=0}^{d-2} \sqrt{s+1} \left( \bigotimes_{i=0}^{\lceil \log_2(d) \rceil - 1} (|s_i\rangle_{m_i} \langle (s+1)_i|_{m_i}) \right).
\end{aligned} \tag{6.13}$$

For example, for  $d = 4$ ,

$$\begin{aligned}
a_m &= \sqrt{1} |0\rangle_m \langle 1|_m + \sqrt{2} |1\rangle_m \langle 2|_m + \sqrt{3} |2\rangle_m \langle 3|_m \\
&= \sqrt{1} |00\rangle_m \langle 01|_m + \sqrt{2} |01\rangle_m \langle 10|_m + \sqrt{3} |10\rangle_m \langle 11|_m \\
&= \sqrt{1} (|0\rangle \langle 0|_{m_1} \otimes |0\rangle \langle 1|_{m_0}) + \sqrt{2} (|0\rangle \langle 1|_{m_1} \otimes |1\rangle \langle 0|_{m_0}) \\
&\quad + \sqrt{3} (|1\rangle \langle 1|_{m_1} \otimes |0\rangle \langle 1|_{m_0}).
\end{aligned} \tag{6.14}$$

Each term  $|s_i\rangle_{m_i} \langle (s+1)_i|_{m_i}$  is a single-qubit binary projector, which can be mapped to Pauli operators using

$$\begin{aligned}
|0\rangle \langle 1| &= \frac{1}{2}(X + iY) \\
|1\rangle \langle 0| &= \frac{1}{2}(X - iY) \\
|0\rangle \langle 0| &= \frac{1}{2}(I + Z) \\
|1\rangle \langle 1| &= \frac{1}{2}(I - Z)
\end{aligned} \tag{6.15}$$

A single term  $|s\rangle_m \langle s+1|_m$  is thus mapped to  $2^{\lceil \log_2(d) \rceil} = d$  Pauli strings, with weight  $\mathcal{O}(\lceil \log_2(d) \rceil)$ . There are  $d$  such terms for each creation or annihilation operator – so each is mapped to  $\mathcal{O}(d^2)$  Pauli terms. As a result, a vibrational Hamiltonian expanded to  $l^{\text{th}}$  order will contain  $\mathcal{O}(M^l d^{2l})$  Pauli strings with weight  $\mathcal{O}(l \lceil \log_2(d) \rceil)$ .

It is evident that the direct mapping makes less efficient use of the avail-

able Hilbert space, but consequently yields lower weight expressions for vibrational operators, and fewer terms in the expansion. In contrast, the compact mapping makes use of all of the available Hilbert space, but at a cost of leading to an increased number of operators, with higher weight. Consequently, while the direct mapping may appear to require a greater number of qubits than the compact mapping, whether this is actually the case will depend on the context of the simulation performed. For example, if one were performing an error-corrected simulation of time evolution under the qubit-mapped vibrational Hamiltonian, then the depth of the circuit becomes as important a metric as its logical ‘width’. It may be the case that the lower weight, fewer terms, and improved parallelisability of the direct mapping lead to a significant reduction in the circuit depth, compared to that required for the compact mapping. In this case, fewer physical qubits would be required to protect each logical qubit in the direct mapping than in the compact mapping. There may be scenarios in which this saving is significant enough to outweigh the exponential reduction in logical qubits afforded by the compact mapping.

On a similar note, the different mappings may be better suited for different hardware systems, when considering NISQ simulations. For example, superconducting platforms are typically constrained to nearest-neighbour gates. Moreover, such platforms are often coherence time limited, meaning that it is favourable to use wide, shallow circuits. This suggests that the direct mapping may be preferable. In contrast, certain trapped ion systems are able to perform globally entangling operations, such as the Mølmer-Sorenson gate. While they have long coherence times, they typically possess fewer qubits than superconducting systems. As a result, trapped ion systems may be better suited to using the compact mapping, which can require operations on all of the qubits in a register to perform a creation or annihilation operation, and as such is limited to performing some opera-

tions sequentially.

## 6.4 Simulating vibrational problems

The techniques introduced in the previous section can be used to map the vibrational simulation problem onto a quantum computer. As the vibrational Hamiltonian is mapped to a linear combination of Pauli strings, it is straightforward to apply existing quantum simulation algorithms. For example, I could perform a Trotter decomposition of the vibrational Hamiltonian, enabling the implementation of quantum phase estimation, to calculate the vibrational energy levels. In this section I discuss techniques designed specifically for solving problems of molecular vibrations.

The first is a unitary analogue of the vibrational coupled cluster ansatz [265, 279, 280] for approximating low-lying vibrational eigenstates. The ansatz, denoted as UVCC, is designed for use in a VQE calculation, and is given by

$$|\Psi(\vec{\theta})\rangle = \exp(\hat{T}_\theta - \hat{T}_\theta^\dagger) |\Psi_0\rangle, \quad (6.16)$$

where the initial state  $|\Psi_0\rangle$  can be either the ground state of the harmonic Hamiltonian, or a vibrational mean-field or multi-reference state for the full Hamiltonian.  $\hat{T}_\theta$  is the sum of vibrational excitation operators truncated at a specified excitation rank, and  $\vec{\theta}$  are excitation amplitudes. The excitation operators truncated at second order are given by

$$\hat{T}_\theta = \hat{T}_1^\theta + \hat{T}_2^\theta, \quad (6.17)$$

with

$$\begin{aligned}\hat{T}_1^\theta &= \sum_{m=1}^M \sum_{i,j=0}^{d-1} \theta_{i_m j_m} |i\rangle_m \langle j|_m, \\ \hat{T}_2^\theta &= \sum_{m < m'}^M \sum_{i,j,k,l=0}^{d-1} \theta_{i_m j_m k_{m'} l_{m'}} |i\rangle_m \langle j|_m \otimes |k\rangle_{m'} \langle l|_{m'}.\end{aligned}\tag{6.18}$$

The operators can be mapped to qubit operators using either the direct or compact mapping. For example, under the direct mapping the operator  $|3\rangle\langle 1| \rightarrow |1000\rangle\langle 0010|$ , which has the same effect as the operator  $\frac{1}{4}(X_3 - iY_3) \otimes (X_1 + iY_1)$  in the computationally relevant subspace. In the compact mapping, the operator  $|3\rangle\langle 1| \rightarrow |11\rangle\langle 01| = |1\rangle\langle 0|_1 \otimes |1\rangle\langle 1|_0 = \frac{1}{4}(X_1 - iY_1) \otimes (I_0 - Z_0)$ . The UVCC operator can then be Trotterized to split it into a product of exponentiated Pauli strings, each of which can be easily decomposed into single-qubit rotations and CNOT gates. Because of the unitary nature of the ansatz, it obeys the Rayleigh-Ritz variational principle, and the parameters can be minimised directly. This is in contrast to the classically tractable VCC ansatz, which is not variational, and is optimised by solving a number of coupled equations. These equations struggle to converge when used with multireference initial states [266, 281–283], which can be required to describe vibrational resonances in a system. These resonances can arise from near-degeneracies between different vibrational eigenstates caused by the anharmonic terms in the Hamiltonian [284, 285]. The quantum-amenable UVCC ansatz can be applied to multireference initial states, and so provides advantages over the classical VCC ansatz. In this sense, the correspondence between the VCC and the UVCC ansätze is equivalent to that between the CC and UCC ansätze used to solve the electronic structure problem. Although I do not present the results in this chapter, in the original publication of these results (Ref. [255]) the UVCC ansatz was used in a VQE emulation to find the vibrational ground state of H<sub>2</sub>O (numerical simulations performed by Xiao Yuan, a collaborator on

this project).

For many molecular systems, the differences between the vibrational energy levels are of the same order of magnitude as the energy of a room temperature photon. As a result, we may be interested in preparing vibrational thermal states, as well as pure eigenstates. One possible approach to prepare these states is to use the algorithm for ansatz-based imaginary time evolution, discussed in Chapter 4, to prepare a purification of the thermal state (as discussed in Refs. [195, 234]). I double the number of qubits, and initialise the registers in the maximally entangled state

$$|\phi(0)\rangle = \frac{1}{d^{M/2}} \sum_{i_1, i_2, \dots, i_M} (|i_0\rangle_0^A \dots |i_M\rangle_M^A) \otimes (|i_0\rangle_0^B \dots |i_M\rangle_M^B), \quad (6.19)$$

where  $A, B$  denote the two systems, and  $|i_\alpha\rangle_\alpha$  denotes an eigenstate of the number operator for mode  $\alpha$ . It is simple to generate this state when using the compact mapping. To place a mode in an equally weighted superposition of all number states, I apply the Hadamard gate to each qubit. I do this for each mode independently. To generate the equivalent state under the direct mapping, I create a  $d$ -qubit W state for each vibrational mode register. This can be achieved using the algorithm presented in Ref. [286], which I discuss in more detail for the purpose of preparing Dicke states in Chapter 7. Having prepared these states, I then entangle system  $A$  with system  $B$  by performing a CNOT gate between each qubit in system  $A$  and the corresponding qubit in system  $B$ . By writing  $|i_0\rangle_0 \dots |i_M\rangle_M$  as  $|k\rangle$ , I denote the initial state as

$$|\phi(0)\rangle = \frac{1}{\sqrt{d^M}} \sum_k |k\rangle^A \otimes |k\rangle^B. \quad (6.20)$$

After preparing this state, I apply the imaginary time algorithm introduced

in Chapter 4, evolving under the Hamiltonian  $H_v^A \otimes I^B$  for imaginary time  $\tau = \frac{1}{2T}$ , where  $T$  is the desired temperature of the thermal state,

$$\begin{aligned}
|\phi(\tau)\rangle &= \frac{e^{-H_v^A \otimes I^B \tau} |\phi(0)\rangle}{\sqrt{\langle \phi(0) | e^{-2H_v^A \otimes I^B \tau} | \phi(0) \rangle}} \\
&= \frac{\frac{1}{\sqrt{d^M}} \sum_k e^{-H_v^A \tau} |k\rangle^A \otimes |k\rangle^B}{\sqrt{\frac{1}{d^M} \sum_{i,j} \langle i |^A \langle i |^B e^{-2H_v^A \otimes I^B \tau} |j\rangle^A |j\rangle^B}} \\
&= \frac{\sum_k e^{-H_v^A \tau} |k\rangle^A \otimes |k\rangle^B}{\sqrt{\sum_j \langle j |^A e^{-2H_v^A \tau} |j\rangle^A}} \\
&= \frac{\sum_k e^{-H_v^A \tau} |k\rangle^A \otimes |k\rangle^B}{\sqrt{\text{Tr}(e^{-2H_v^A \tau})}}.
\end{aligned} \tag{6.21}$$

Tracing out register  $B$  then prepares the following mixed state

$$\begin{aligned}
\text{Tr}_B(|\phi(\tau)\rangle \langle \phi(\tau)|) &= \frac{\sum_k e^{-H_v^A \tau} |k\rangle^A \langle k|^A e^{-H_v^A \tau}}{\text{Tr}(e^{-2H_v^A \tau})} \\
&= \frac{e^{-2H_v^A \tau}}{\text{Tr}(e^{-2H_v^A \tau})} \\
&= \frac{e^{-\frac{H_v^A}{T}}}{\text{Tr}\left(e^{-\frac{H_v^A}{T}}\right)},
\end{aligned} \tag{6.22}$$

which corresponds to the thermal state for the system at temperature  $T$ . Because I am using the ansatz-based imaginary time algorithm, the non-unitary imaginary time evolution of the initial state is implemented by a unitary ansatz circuit. This ansatz circuit should act on both registers, creating entanglement between them. This can be distinguished from preparing the maximally entangled state, evolving just system  $A$  in imaginary time, and tracing out register  $B$  (which is equivalent to removing register  $B$  from the calculation, and sampling over input states  $|k\rangle$  in register  $A$ ), because the imaginary time evolution implemented by the ansatz-based approach is normalised, and so would lead to an incorrectly weighted state.

In vibrational simulations, it is often of interest to calculate the spectra arising from transitions between different vibrational and electronic states. I assume that the system is initially in the state  $|\psi_i\rangle_e |\phi_i\rangle_N$ , and seek the probability of transition to the state  $|\psi_f\rangle_e |\phi_f\rangle_N$  upon absorption of a photon. Using Fermi's Golden rule, the transition probability is given by

$$\begin{aligned} \mathcal{P} &= |\langle \psi_f | \langle \phi_f | \mu | \phi_i \rangle_N | \psi_i \rangle_e|^2 \\ &= |\langle \psi_f | \mu_e | \psi_i \rangle_e \langle \phi_f | \phi_i \rangle_N + \langle \psi_f | \psi_i \rangle_e \langle \phi_f | \mu_N | \phi_i \rangle_N|^2. \end{aligned} \quad (6.23)$$

where  $\mu_e/\mu_N$  are the electronic and nuclear dipole moments, respectively. The assumption that  $\mu_e$  does not depend on the coordinates of the nuclei is known as the Franck-Condon approximation [287]. In reality however, the positions of the electrons are influenced by the locations of the nuclei, and so we can consider a Taylor expansion of the electronic transition dipole moment in terms of the nuclear coordinates (known as the Herzberg-Teller expansion [288]). I do not consider these non-Condon effects in this chapter. For a transition between vibrational levels belonging to the same potential energy surface, the first term is zero, and  $\mathcal{P} = |\langle \phi_f | \mu_N | \phi_i \rangle_N|^2$ . For vibronic transitions (where an electron is excited to a higher potential energy surface) the second term vanishes, so  $\mathcal{P} = |\langle \psi_f | \mu_e | \psi_i \rangle_e \langle \phi_f | \phi_i \rangle_N|^2$ . I focus on the latter case, where the goal of the quantum calculation is to evaluate  $|\langle \phi_f | \phi_i \rangle_N|^2$ . This is complicated by the fact that the coordinate systems for the harmonic oscillator basis sets of these two states differ. The excited electronic potential energy surface will likely be displaced, rotated and stretched relative to the initial potential energy surface. There are two possible approaches one can consider to solve this problem. The first is to apply a unitary basis change directly onto the state vector. This is known as the Doktorov unitary [289]. Unfortunately, this matrix is obtained by

exponentiating sums of creation and annihilation operators, which means that Trotterization (or another, related technique, such as the linear combination of unitaries method) is required in order to decompose the operator into single- and two-qubit gates. A more straightforward approach is to express the excited state vibrational eigenstates in the basis of the initial state vibrational eigenstates. The coordinates of the two systems can be related using the Duschinsky transform  $\vec{Q}_f = \mathbf{S}\vec{Q}_i + \vec{\delta}$ , where  $\mathbf{S}$  is the Duschinsky mode mixing matrix, and  $\vec{\delta}$  is a displacement vector [270, 290]. It may be difficult to exactly relate the two coordinate systems, and in practice a number of approximations are used (see Ref. [259] Section 2.2 for further details). Assuming that an accurate transformation can be found between  $\vec{Q}_i$  and  $\vec{Q}_f$ , these can be used to obtain  $H_v^i$  and  $H_v^f$  in the same harmonic oscillator basis set. I assume the ability to prepare the eigenstates  $|\phi_i\rangle$  and  $|\phi_f'\rangle$  (where the prime denotes the use of the same harmonic oscillator basis set describing the initial state  $|\phi_i\rangle$ ) with unitaries  $V_i$  and  $V_f'$ . The vibrational transition strength can then be calculated using a swap test circuit, as shown in Fig. 6.1.




Figure 6.1: A swap test circuit to measure the squared-overlap between the two vibrational states  $|\phi_i\rangle = V_i |\bar{0}\rangle$  and  $|\phi_f'\rangle = V_f' |\bar{0}\rangle$ .

An alternative approach is to prepare the state  $V_f'^{\dagger} V_i |\bar{0}\rangle$ , and calculate the probability of measuring all qubits in the  $|0\rangle$  state, to yield  $|\langle \bar{0} | V_f'^{\dagger} V_i |\bar{0}\rangle|^2 = |\langle \phi_f' | \phi_i \rangle|^2$ .

## 6.5 Discussion

In this chapter, I have discussed how to simulate molecular vibrations on digital quantum computers. By truncating the Hilbert space to a low energy subspace, I was able to map a canonical form of the anharmonic vibrational Hamiltonian onto a qubit Hamiltonian that can be expressed as a linear combination of Pauli strings. This form of the Hamiltonian is amenable to use in generic quantum simulation algorithms, such as phase estimation and the VQE. To target specific vibrational phenomena, I introduced the UVCC ansatz as a quantum analogue of the classically tractable VCC ansatz, for approximating the low energy vibrational eigenstates. I have also presented methods to prepare thermal states of the vibrational Hamiltonian, using the imaginary time algorithm discussed in Chapter 4. Finally, I showed how these techniques can be used in conjunction with a SWAP-test style circuit, to calculate absorption spectra in molecules.

This work offers a number of advantages over existing analog quantum algorithms for simulating molecular vibrations. In particular, it is easy to incorporate anharmonic terms in the Hamiltonian using the methods introduced herein, whereas this can be difficult in some analog simulation platforms. Moreover, considering digital quantum computation enables the use of quantum error correction to protect simulations against noise. One limitation compared to analog simulators is the necessity of truncating the energy level of the harmonic oscillator. However, this truncation is often enforced indirectly in analog calculations by experimental limitations. For example, deficiencies in number resolving photon detectors mean that it can be difficult to consider arbitrarily high harmonic oscillator eigenstates.

Since the publication of this work, there have been a number of follow-up investigations by other research groups, which have built on the work

discussed herein. In Ref. [291], the compact and direct mappings were compared to alternative approaches; the Gray code, and block unary encodings. The authors of that work investigated the resources required to implement vibrational operators using these mappings. Similar work in Ref. [292] examined how these resources were affected by constrained connectivity in the underlying hardware. In Ref. [293], an algorithm was presented to calculate an entire Franck-Condon spectra, rather than just the intensity of a given transition. This algorithm requires the preparation of relevant initial states, and so can use the techniques presented herein as subroutines. More recently, Ref. [294] considered the use of alternative, more general basis sets for solving the vibrational structure problem. The authors of that work made use of the direct mapping and UVCC ansatz to calculate the ground states of small molecular systems, including performing a small calculation on a superconducting quantum processor. Finally, a recent review [295] investigated how to include non-Condon effects in vibrational simulations, using methods similar to those discussed here. Those authors claimed that quantum advantage may be demonstrated for vibrational structure problems before it is achieved for electronic structure calculations, due to the difficulties in classically simulating small non-rigid molecules. These systems may require few logical qubits to simulate, but will require a high-order expansion of the potential, leading to a Hamiltonian with a large number of terms.

The work introduced in this chapter, and the publications described above, have laid the groundwork for simulations of vibrations on digital quantum hardware. The next challenge that needs to be overcome is ascertaining and optimising the cost of these calculations. In particular, it would be interesting to consider a fault-tolerant resource estimate for a vibrational calculation, analogous to those performed for the electronic structure prob-

lem [25–27, 68, 69]. This will help to determine whether the lower gatecount expected from the direct mapping is able to outweigh the spatial efficiency of the compact mapping. One limitation of algorithms for the vibrational structure problem is that the Hamiltonian may contain a large number of terms, which may translate to a large gatecount. There are at least two interesting possibilities for mitigating this issue. The first is to use a tensor factorization of the vibrational potential. Similar factorizations of the Coulomb operator have recently been used to reduce the cost of quantum algorithms for solving the electronic structure problem [27, 69, 72, 250]. While tensor factorizations have recently been incorporated into classical vibrational calculations, these have predominantly been used to approximate the wavefunction itself (as in tensor network methods) rather than to compress the vibrational Hamiltonian [296–298]. Alternatively, we could consider using a Hamiltonian simulation algorithm whose cost does not depend explicitly on the number of terms in the Hamiltonian. For example, the cost of the qDRIFT protocol introduced in Sec. 3.1.1 scales with the 1-norm of the Hamiltonian, rather than with the number of terms. As the vibrational potential is obtained through a Taylor expansion it is likely that the strengths of terms in the vibrational Hamiltonian will decay rapidly. Similarly, one can consider locality constraints, which may further reduce the magnitude of interactions between more distant vibrations. As a result, qDRIFT and related methods [106] may be able to improve upon the gatecounts obtained from naive Trotterization of the vibrational Hamiltonian.

Developing methods to simulate molecular vibrations expands the range of chemical phenomena that we will one day be able to explore using fault tolerant digital quantum computers. It is hoped that these simulations will enable us to predict, via simulation, how a given system may interact with light, or react with another substance. Being able to perform such simu-

lations may reduce the cost of developing new compounds and materials, and open up new avenues in autonomous design. The work discussed in this chapter constitutes a step in this direction, and provides another setting in which to use some of the techniques discussed in this thesis.

# 7 | Quantum simulation of muon spectroscopy experiments

## Contents

---

<b>7.1 Chapter summary</b> . . . . .	<b>123</b>
<b>7.2 Introduction</b> . . . . .	<b>124</b>
7.2.1 Muon spectroscopy . . . . .	125
7.2.2 Muon polarisation functions . . . . .	128
<b>7.3 Quantum algorithm for muon spectroscopy analysis</b> . .	<b>134</b>
7.3.1 Mapping the spin system to qubits . . . . .	136
7.3.2 Preparing the initial state of the system . . . . .	138
7.3.3 Evolving the state in time . . . . .	141
7.3.4 Measuring the polarisation . . . . .	142
<b>7.4 Results</b> . . . . .	<b>144</b>
7.4.1 Noiseless simulations . . . . .	147
7.4.2 Noisy simulations . . . . .	155
7.4.3 Resource estimates . . . . .	161
<b>7.5 Discussion</b> . . . . .	<b>167</b>

---

## 7.1 Chapter summary

This chapter discusses how quantum computers can be used to analyse the data arising from muon spectroscopy experiments. This work has been released as Ref. [299]:

*Learning from physics experiments, with quantum computers: Applications in muon spectroscopy*  
S. McArdle  
[arXiv:2012.06602 \(2020\)](https://arxiv.org/abs/2012.06602)

The experimental data analysed as part of this work was obtained by John Wilkinson and Stephen Blundell [300], who kindly provided me with access to their data.

## 7.2 Introduction

As discussed throughout this thesis, the majority of work to date on quantum simulation algorithms has focused on solving the electronic structure problem. Current estimates for the resources required to solve classically intractable instances of this problem require on the order of  $10^5 - 10^6$  physical qubits [25, 27, 69, 72] – well beyond our current capabilities. In this final research chapter, I consider the quantum resources required to solve a new class of simulation problem; the analysis of muon spectroscopy experiments. Muon spectroscopy is a technique for quantitatively studying the interactions present in a sample of interest, by studying the time evolution of spin-polarised (anti)-muons implanted into the sample. Accurate analysis of these experiments can require a fully quantum treatment of the system dynamics. In some cases, this problem appears challenging to solve on classical computers, admitting few simplifications, and often requiring exact diagonalisation of the system Hamiltonian [300–302].

In Sec. 7.3 I present an efficient quantum algorithm to solve this problem. In Sec. 7.4 I report classical numerical emulations of this algorithm, to investigate its scaling behaviour and noise robustness. I used these emulations to analyse muon spectroscopy data produced by a real experiment, finding good agreement with recent state-of-the-art classical analysis [300]. In Sec. 7.4.3 I discuss both the near-term and error corrected resources required to run the algorithm for problem sizes of interest.

Before presenting the quantum algorithm for muon spectroscopy analysis, I provide an introduction to muon spectroscopy experiments in Sec. 7.2.1 and Sec. 7.2.2. I refer the interested reader to Refs. [303, 304] for further information on muon spectroscopy.

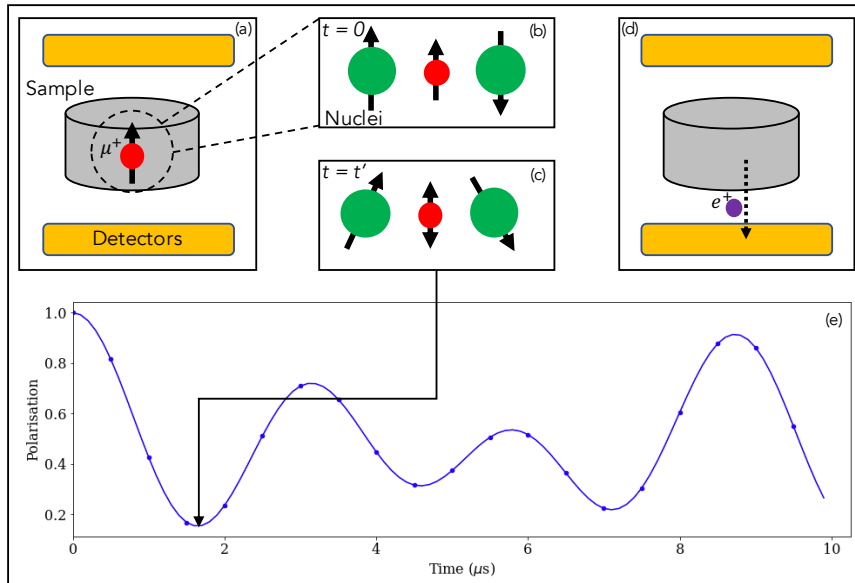


Figure 7.1: A schematic of a muon spectroscopy experiment, as described in the main text. Spin polarised muons are implanted into a sample of interest (a), whereupon they interact with nearby nuclear spins (b). This causes an oscillation of the muon spin polarisation (c). When the muon decays, the positron is preferentially emitted in the direction of the muon spin at the time that it decayed (d). The normalised asymmetry between positron counts registered by the two detectors gives the time evolution of the muon polarisation (e).

### 7.2.1 Muon spectroscopy

Muon spectroscopy, more commonly known as muon spin rotation, relaxation and resonance ( $\mu^+\text{SR}$ ), is an experimental technique closely related to other spin-based methods for probing magnetic interactions, such as nuclear magnetic resonance or electron spin resonance. The technique uses spin-polarised beams of positive (anti)-muons (hereafter referred to as muons, as the ordinary muon is rarely used in muon spectroscopy experiments) to probe the interactions in a sample of interest. A typical  $\mu^+\text{SR}$  experiment proceeds as follows, and as shown in Fig. 7.1:

- The  $\mu^+$  is a positively charged, spin- $\frac{1}{2}$  particle, with a mass approximately  $1/9^{\text{th}}$  that of the proton. It can be produced at particle ac-

celerators, from the decay of pions. As a consequence of the weak interaction (through which this decay proceeds), the muons are spin-polarised. Following convention, I choose the initial spin-polarisation of the muon beam to be along the positive  $Z$  axis, such that the muons are in the state  $|0\rangle \langle 0|_{\mu}$ .

- The muons are directed into a sample of interest, and are brought to rest by electrostatic interactions, which do not depolarise the beam [303]. Muons can be implanted one-by-one (continuous wave beams) or in a pulse containing hundreds of other muons (pulsed beams). Continuous wave beams have better time resolution than pulsed beams (1 ns vs 0.1  $\mu$ s) but cannot record spectra to the same maximum duration (10  $\mu$ s vs 30  $\mu$ s). As a result, the two beams are better suited for investigating fast and slow dynamics, respectively.
- The muons interact with nearby nuclear spins, and with the magnetic fields present at their rest site. This causes the muon spin polarisation to evolve according to the Schrödinger equation.
- With a half-life of 2.2  $\mu$ s, the muon decays into a positron, electron neutrino and muon antineutrino. Due to the weak interaction, the positron is emitted preferentially in the direction of the muon spin polarisation.
- The emitted positrons are detected by detectors placed along the  $Z$  axis. The normalised difference between the forward and backwards detector counts ( $N_{+/-}$ , respectively) yields an asymmetry function

$$A(t) = \frac{N_+ - \alpha N_-}{N_+ + \alpha N_-} \quad (7.1)$$

where  $\alpha$  accounts for detection inefficiencies and imbalances. This

value is converted to a polarisation function, using

$$P(t) = \frac{A(t) - A_{\text{bg}}}{A_0} \quad (7.2)$$

where  $A_{\text{bg}}$  gives the background positron count, and  $A_0 = A(0) - A_{\text{bg}}$ . The polarisation function is normalised to between  $\pm 1$ , and corresponds to the spin-polarisation of the muon beam at time  $t$

$$\begin{aligned} P(t) &= \text{Tr}(Z_\mu \rho_s(t)) \\ &= \text{Tr}(Z_\mu e^{-iH_s t} [ |0\rangle \langle 0|_\mu \otimes \rho_e(0) ] e^{iH_s t}), \end{aligned} \quad (7.3)$$

where  $Z_\mu$  is the Pauli  $Z$  matrix acting on the muon,  $H_s$  is the system Hamiltonian,  $\rho_s(t)$  is the state of the muon-environment system at time  $t$ , and  $\rho_e(0)$  is the initial state of the environment.

- Measuring the time evolution of the polarisation function enables us to infer the interactions felt by the muon at its rest site.

Experiments can be carried out at high pressures ( $\sim 1.5$  GPa), low temperatures (less than 1 K), with high strength transverse or longitudinal magnetic fields ( $\sim 8$  T), or with the application of time-dependent radio-frequency pulses [304].  $\mu^+$ SR experiments have been used to investigate a range of physical phenomena. For example, we can measure the temperature dependence of oscillations in the polarisation function (or the lack thereof) in order to investigate phase transitions in magnetic materials, such as low dimensional spin chains [305]. Other experiments have measured the polarisation function of muons in semiconductors at a range of temperatures, in order to investigate the diffusion of muons within the sample. The muon acts as a light proton, so these experiments can be used to examine the effects of hydrogenic defect diffusion in semiconductors [306]. Similar  $\mu^+$ SR experiments have examined  $\text{Li}^+$  ion diffusion in battery materials [307–

309]. In superconducting systems,  $\mu^+$ SR experiments have been used to: measure the superconducting electron density, determine phase diagrams, and characterise vortex lattices [304].  $\mu^+$ SR experiments have also been applied to biological and chemical systems; for example, to investigate oxygen dependent effects in the radiation treatment of cancer [310].

Muon spectroscopy is a versatile technique, that has provided insights on a range of physical systems. The technique is still undergoing active development, including the introduction of lower energy muon beams which can be used to probe surface effects [304]. However, there are also theoretical challenges for  $\mu^+$ SR that are desirable to address. While some systems can be analysed using a mean-field, or semi-classical approach, others appear to require a fully quantum treatment [301, 302, 311, 312]. In the following section, I discuss the simulation and analysis of muon polarisation functions in more detail.

### 7.2.2 Muon polarisation functions

In order to analyse the polarisation function arising from a given  $\mu^+$ SR experiment, we can compare it to a theoretical polarisation function obtained from a physical model of the studied system. The accuracy of these theoretical polarisation functions is determined by both the level of detail included in the model, and the method used to simulate the model. This can be likened to the field of computational chemistry, where the accuracy of a simulation is determined by both the physical effects included in the model (e.g. the Born-Oppenheimer approximation, or relativistic effects) and the method used to simulate the system (e.g. mean-field approaches, or exact diagonalisation). I focus first on the different methods used to obtain the theoretical polarisation function, before considering other physical effects

that can be incorporated into the model.

On a most basic level, one can use a semiclassical, mean-field approach, which considers how the muon polarisation evolves, given a model for the surrounding magnetic field distribution. This approach can yield analytically derivable formulae [304, 313] that are widely applied within the  $\mu^+$ SR field, and are acceptably accurate in many circumstances. If greater accuracy is required, techniques have been developed which treat the muon interactions with its local spin environment semi-classically [314, 315]. In some circumstances, these methods can yield high accuracy at a modest computational cost. However, they cannot always account for strong spin-spin interactions, or the effect of quadrupole interactions [304]. Finally, the highest level of accuracy for a given model can be obtained by using a quantum mechanical analysis. The cost of these calculations is believed to scale exponentially with the size of the system simulated, due to the computational complexity of storing highly entangled quantum states.

In this chapter, I focus on systems which require a fully quantum treatment in order to obtain accurate polarisation functions. Two techniques have been developed by the muon community for exactly simulating the polarisation function. The first relies on exact diagonalisation of the muon-environment Hamiltonian. Because the thermal energies encountered in  $\mu^+$ SR experiments are typically much larger than the nuclear energy levels, the environment is normally assumed to be in the maximally mixed state  $\rho_e(0) = I_e/D_e$  where  $D_e$  is the Hilbert space dimension of the environment.

The polarisation function is given by

$$\begin{aligned}
P(t) &= \frac{1}{D_s} \text{Tr}([Z_\mu \otimes I_e] e^{-iH_s t} [|0\rangle \langle 0|_\mu \otimes I_e] e^{iH_s t}) \\
&= \frac{1}{2D_s} \text{Tr}([Z_\mu \otimes I_e] e^{-iH_s t} [(I_\mu + Z_\mu) \otimes I_e] e^{iH_s t}) \\
&= \frac{1}{2D_s} \text{Tr}([Z_\mu \otimes I_e] e^{-iH_s t} [Z_\mu \otimes I_e] e^{iH_s t}).
\end{aligned} \tag{7.4}$$

I use a resolution of the identity in terms of eigenstates of the system Hamiltonian;  $|m\rangle, |n\rangle$

$$\begin{aligned}
P(t) &= \frac{1}{2D_s} \sum_m \langle m| Z_\mu e^{-iH_s t} Z_\mu e^{iH_s t} |m\rangle \\
&= \frac{1}{2D_s} \sum_m \sum_n \langle m| Z_\mu e^{-iH_s t} |n\rangle \langle n| Z_\mu e^{iH_s t} |m\rangle \\
&= \frac{1}{2D_s} \sum_m \sum_n |\langle m| Z_\mu \otimes I_e |n\rangle|^2 e^{i(E_m - E_n)t},
\end{aligned} \tag{7.5}$$

where  $E_m$  is the eigenvalue of eigenstate  $|m\rangle$ . Performing an exact diagonalisation of the Hamiltonian can thus be used to calculate the polarisation function. However, the dimension of the Hamiltonian will scale exponentially with the number of spins considered, which limits the size of these calculations. These exact calculations have been performed for a range of systems [300–302, 304, 312, 316–318], with Hilbert space dimensions up to around 2048 [300, 302].

An alternative method, introduced in Ref. [319], reduces the memory cost of the calculation, at the expense of introducing statistical uncertainty into the measured result, which can be reduced through sampling. The technique makes use of a random-phase-approximation, whereby the wavefunction at time  $t$  is set to

$$|\psi_s(t)\rangle = U_1(t) \left[ \frac{1}{\sqrt{D_e}} \sum_j e^{i\theta_j} |0\rangle_\mu |j\rangle_e \right], \tag{7.6}$$

where  $U_1(t)$  is a first-order Trotter approximation to the time evolution operator,  $\theta_j$  are randomly chosen on each sample of the algorithm, and  $|j\rangle_e$  denote unentangled basis states of the environment. In this case, an approximation for the polarisation function can be obtained from

$$\begin{aligned}
 P(t) &\approx \langle \psi_s(t) | Z_\mu \otimes I_e | \psi_s(t) \rangle \\
 &= \sum_j \frac{1}{D_e} \langle 0 |_\mu \langle j |_e U_1(t)^\dagger Z_\mu U_1(t) | 0 \rangle_\mu | j \rangle_e \\
 &\quad + \sum_{j,k} \frac{1}{D_e} e^{i(\theta_k - \theta_j)} \langle 0 |_\mu \langle j |_e U_1(t)^\dagger Z_\mu U_1(t) | 0 \rangle_\mu | k \rangle_e.
 \end{aligned} \tag{7.7}$$

The first term is equal to the polarisation function that we wish to measure (up to an error induced by the Trotterization of the time evolution operator, which I discuss in more detail in Sec. 7.3). As the phases in the second term are chosen randomly, many of these terms cancel, leading to a small error on the polarisation function. This error can be reduced by repeating the method with independently randomly generated values of  $\theta_j$ . The error also decreases as the Hilbert space dimension of the system increases. While exact diagonalisation requires manipulating matrices of dimension  $D_s \times D_s$ , this product-formula based approach only requires storing the wavefunction, which has dimension  $D_s$ . This has enabled much larger calculations to be performed using this method, including simulations with Hilbert space dimension  $\sim 2^{17}$  [320] and  $2^{26}$  [321]. This latter calculation is, to the best of my knowledge, the previous largest  $\mu^+$ SR calculation performed to date with an exact method.

Exact methods for calculating the polarisation function have predominantly been used for two purposes: locating the muon rest site, and studying muon diffusion. Determining the muon rest site is an important challenge in  $\mu^+$ SR experiments [322]. Accurate simulation can be employed in con-

junction with density functional theory [323–327] and experimental methods [328] to give greater certainty on the muon location. Exact simulation has repeatedly been employed when studying fluorinated materials, due to the strong dipolar interaction between the spin- $\frac{1}{2}$  fluorine nuclei and the muon, and the large electronegativity of the fluorine ion which ‘traps’ the muon [329]. This effect has been used to locate the muon rest site in fluorinated polymers [318, 330], molecular magnets [317], and ionic crystals [300, 331]. Similar calculations were used to identify the muon rest sites in the high temperature superconductor  $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$  [321]. Previous calculations have often included heuristic terms in the polarisation function, to compensate for the limited system sizes that can be simulated using these costly techniques.

Polarisation functions calculated from quantum models have also been used to study muon diffusion. I discuss this topic in more detail in Ref [299]. However, it suffices to note here that using approximate methods can lead to inaccurate estimates of the diffusion rate [320, 332, 333].

Having established the scenarios where exact calculations are often utilised in muon spectroscopy analysis, I now consider the interactions typically included in our models for systems of interest. These include: dipolar interactions between spins, quadrupole interactions between nuclear spins and electric field gradients in the sample, and the coupling of spins to time-dependent or independent magnetic fields. The dipolar contribution is given by

$$H_D = \frac{1}{2} \sum_{i,j} \frac{\hbar^2 \mu_0 \gamma_i \gamma_j}{4\pi r_{ij}^3} \left[ \vec{S}_i \cdot \vec{S}_j - 3(\vec{S}_i \cdot \hat{r}_{ij})(\vec{S}_j \cdot \hat{r}_{ij}) \right], \quad (7.8)$$

where  $\mu_0$  is the permeability of free space,  $\gamma_i$  is the gyromagnetic ratio of spin  $i$ ,  $\vec{r}_{ij}$  is the vector connecting spins  $i$  and  $j$ , and  $\vec{S}_i$  are the vectorized generalised spin matrices for a particle with spin quantum number  $s$ . I

have moved the factor of  $\hbar$  from the spin matrices to the coefficient of the sum. For example, for a spin- $\frac{1}{2}$  particle  $\vec{S}_i = \frac{1}{2}(X_i, Y_i, Z_i)$ , where  $X_i, Y_i, Z_i$  are the Pauli matrices acting on spin  $i$ . These can be generalised to spins of higher dimension, which I discuss in more detail in Sec. 7.3. This Hamiltonian contains  $\mathcal{O}(N^2)$  terms, where  $N$  is the number of spins considered in the simulation. Some previous calculations have neglected the dipolar interactions between the nuclear spins, as they are often much smaller than those between the muon and the nuclear spins [319, 321]. Neglecting these interactions reduces the number of terms in the Hamiltonian to  $\mathcal{O}(N)$ .

There is also an interaction between the quadrupole moment of nuclei with  $s > \frac{1}{2}$ , and any non-zero electric field gradients in the sample. Such electric field gradients can often be induced by the presence of the muon. The quadrupole interaction Hamiltonian is given by [300]

$$H_Q = \sum_{i \in \mathcal{Q}} \frac{\hbar e Q_i (1 + \Gamma_i)}{2s_i(2s_i - 1)} \left( \vec{S}_i^T \cdot \mathbf{G}(\vec{r}_i) \cdot \vec{S}_i \right), \quad (7.9)$$

where  $\mathcal{Q}$  is the set of nuclei in the simulation with quadrupole moments,  $e$  is the electron charge,  $Q_i$  and  $\Gamma_i$  are the quadrupole coupling factor and anti-shielding factor (respectively) of the  $i^{\text{th}}$  spin, and  $\mathbf{G}(\vec{r}_i)$  is the electric field gradient tensor at position  $\vec{r}_i$ . The elements of  $\mathbf{G}(\vec{r}_i)$  are  $G_{\alpha\beta}(\vec{r}_i) = \frac{\partial^2 V(\vec{r}_i)}{\partial r_\alpha \partial r_\beta}$ , where  $V(\vec{r}_i)$  is the Coulomb potential at position  $\vec{r}_i$ . This Hamiltonian contains  $\mathcal{O}(N_Q)$  terms, each acting on a single spin (where  $N_Q$  is the number of nuclei with quadrupole moments).

The Zeeman interactions of each spin with an applied magnetic field are given by

$$H_M(t) = \sum_i \hbar \gamma_i \vec{S}_i \cdot \vec{B}(t), \quad (7.10)$$

where  $\vec{B}(t)$  is the magnetic field.

In this section, I have discussed theoretical methods to generate simulated polarisation functions. The accuracy of these polarisation functions depends on both the model for the system (i.e. which interactions and effects are included), and the method used to solve the model. The most accurate calculations require a fully quantum treatment. Unfortunately, the exponentially scaling cost of such methods has restricted simulations to small systems, with Hilbert space dimensions of less than  $2^{30}$ . In Sec. 7.3 I show how this exponential cost can be circumvented by using quantum hardware as the simulation platform.

### 7.3 Quantum algorithm for muon spectroscopy analysis

In this section, I discuss how quantum computers can be used to analyse muon spectroscopy data. The algorithm is illustrated in Fig. 7.2, and can be summarised as follows:

1. Map the spin system of interest to qubits.
2. Prepare the quantum registers in the desired initial state.
3. Evolve the system in time for the desired duration,  $t$ .
4. Measure the muon  $Z$  expectation value,  $P(t)$ .

We can repeat steps 2 – 4 of this process at many different  $t$  values in order to obtain a simulated version of the polarisation function for the given system. Below, I discuss in more detail how each of these steps can be implemented.

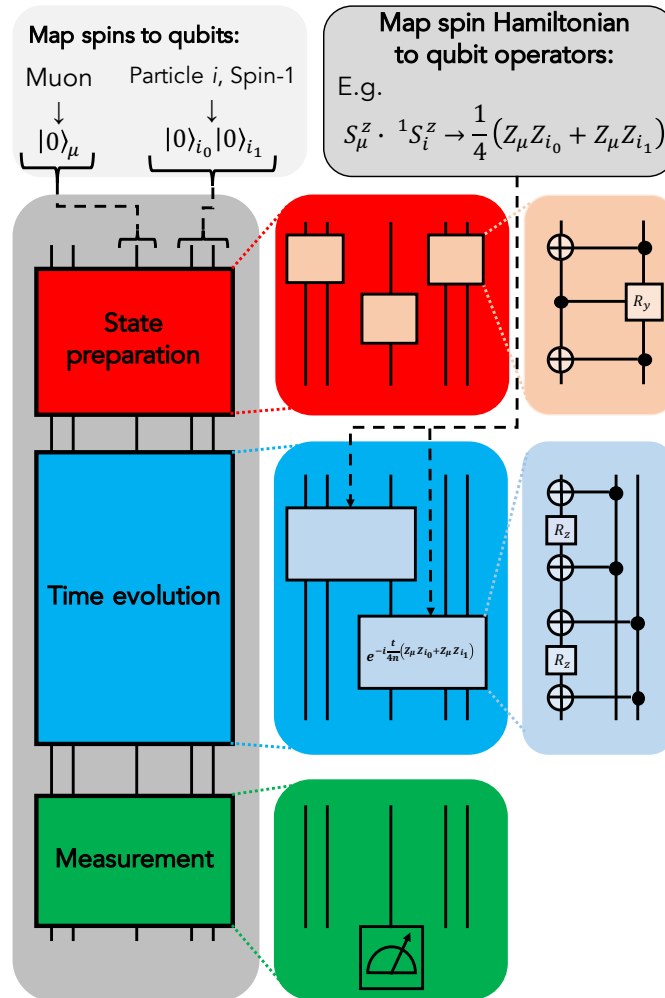


Figure 7.2: An outline of the quantum algorithm for simulating muon polarisation functions. I first map the spins onto qubits via the mapping in Sec. 7.3.1, which yields the qubit operators corresponding to the spin Hamiltonian of the system. I prepare the desired initial state on these qubit registers using the techniques described in Sec. 7.3.2. This state can be propagated in time using the methods in Sec. 7.3.3. Finally, the polarisation value of the muon qubit can be read out, as described in Sec. 7.3.4. This diagram shows the most simple form of the algorithm. This circuit must be repeated many times in order to estimate the  $Z$  expectation value of the muon at a single time value. Repeating this calculation at a number of different time values will generate a simulated polarisation function.

### 7.3.1 Mapping the spin system to qubits

I first consider how to map the spin system onto the quantum computer. For spin- $\frac{1}{2}$  particles, this mapping takes a simple form

$$\begin{aligned} \left| \frac{\mathbf{1}}{2} \right\rangle &:= |0\rangle \\ \left| -\frac{\mathbf{1}}{2} \right\rangle &:= |1\rangle. \end{aligned} \tag{7.11}$$

This leads to a natural mapping of the spin operators of particle  $i$

$$\begin{aligned} S_i^x &= \frac{1}{2} X_i \\ S_i^y &= \frac{1}{2} Y_i \\ S_i^z &= \frac{1}{2} Z_i. \end{aligned} \tag{7.12}$$

When considering particles with higher spin  $s$ , there are a number of possible mappings. One approach is to store the spin of the particle in a register with  $\log_2(2s + 1)$  qubits [291] (essentially this approach uses the ‘compact’ mapping introduced for simulating vibrations in Chapter 6). However, this compact mapping can lead to more complicated expressions for the spin operators, and may be problematic if the spin multiplicity is not a power of 2. In this chapter, I employ a spin-to-qubit mapping based on symmetric quantum states of a given Hamming weight (known as Dicke states) [334]. For example, a spin-1 particle is stored as

$$\begin{aligned} |\mathbf{1}\rangle &:= |00\rangle \\ |\mathbf{0}\rangle &:= \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle) \\ |-\mathbf{1}\rangle &:= |11\rangle. \end{aligned} \tag{7.13}$$

This mapping can be derived by considering the joint Hilbert space of two spin- $\frac{1}{2}$  particles, which can be partitioned into a spin-1 triplet space, and a

spin-0 singlet space. This mapping is an embedding of the 3-dimensional triplet space into the 4-dimensional two-qubit Hilbert space. The spin operators for particle  $i$  are given by

$$\begin{aligned} S_i^x &= \frac{1}{2}(X_{i_0} + X_{i_1}) \\ S_i^y &= \frac{1}{2}(Y_{i_0} + Y_{i_1}) \\ S_i^z &= \frac{1}{2}(Z_{i_0} + Z_{i_1}). \end{aligned} \tag{7.14}$$

This mapping can be generalised to the case of a spin- $s$  particle by assigning each of the states  $|j\rangle, j \in [-s, \dots, s]$  to the symmetric superposition of states with Hamming weight  $h$  such that  $h = s - j$ . For example, a spin- $\frac{3}{2}$  particle can be represented by

$$\begin{aligned} \left| \frac{3}{2} \right\rangle &:= |000\rangle \\ \left| \frac{1}{2} \right\rangle &:= \frac{1}{\sqrt{3}}(|001\rangle + |010\rangle + |100\rangle) \\ \left| -\frac{1}{2} \right\rangle &:= \frac{1}{\sqrt{3}}(|011\rangle + |101\rangle + |110\rangle) \\ \left| -\frac{3}{2} \right\rangle &:= |111\rangle. \end{aligned} \tag{7.15}$$

The spin operator for a spin- $s$  particle  $i$  is given by

$${}^s S_i^\alpha = \frac{1}{2} \sum_{j=0}^{2s-1} P_{i_j}^\alpha, \tag{7.16}$$

where  $\alpha \in [x, y, z]$  denotes which of the Pauli matrices  $P^\alpha \in [X, Y, Z]$  are used, and  $j$  denotes the qubits in register  $i$ . For example, the  $S^x$  operator on particle  $i = 1$  with spin  $s = \frac{3}{2}$  is given by

$$\frac{3}{2} S_1^x = \frac{1}{2}(X_{1_0} + X_{1_1} + X_{1_2}). \tag{7.17}$$

The total number of qubits required is given by  $\sum_{i=1}^N 2s_i$ , where  $s_i$  is the spin of the  $i^{\text{th}}$  particle, and  $N$  is the number of particles in the system. This mapping can be used to obtain qubit representations of the system Hamiltonians given in Eqs. (7.8, 7.9, 7.10). The dipole interaction is the dominant contribution to the number of terms in the Hamiltonian. Each term  $s_i S_i^\alpha \otimes s'_j S_j^\beta$  is mapped to  $4s_i s'_j$  two-qubit Pauli terms. As a result, the Hamiltonian contains up to  $\mathcal{O}(N^2 s_{\max}^2)$  two-qubit Pauli terms, where  $s_{\max}$  is the largest spin value in the system.

### 7.3.2 Preparing the initial state of the system

I consider two approaches to prepare the system in an appropriate initial state. As discussed in Sec. 7.2, this is typically taken to be  $\rho_0 = |0\rangle \langle 0|_\mu \otimes \frac{I_e}{D_e}$ . The first approach emulates Nature; I prepare the environment register in a state  $|k\rangle_e = \bigotimes_{i \in e} |j_i\rangle_i$  that is the tensor product of each environment spin in a randomly selected  $s_z$  basis state. Each state  $|k\rangle_e$  is chosen with probability  $p_k = \frac{1}{D_e}$ . Repeating the simulation many times gives

$$\begin{aligned} \sum_k p_k \text{Tr} \left( (Z_\mu \otimes I_e) e^{-iHt} (|0\rangle \langle 0|_\mu \otimes |k\rangle \langle k|_e) e^{iHt} \right) \\ = \text{Tr} \left( (Z_\mu \otimes I_e) e^{-iHt} (|0\rangle \langle 0|_\mu \otimes \frac{I_e}{D_e}) e^{iHt} \right), \end{aligned} \quad (7.18)$$

as required. This method converges as  $1/\sqrt{\omega}$ , where  $\omega$  is the number of samples taken.

The second approach to effectively sample from the time evolved mixed state is a modified version of the random-phase-approximation method [319] discussed in Eqs. (7.6–7.7). I first initialise the environment register in an equal superposition of all possible  $|k\rangle_e$  states. For the case of spin- $\frac{1}{2}$  particles, this can be accomplished by applying a Hadamard gate to each of

the qubits. For the case of higher-spin particles, I discuss below how to construct this superposition. These steps will prepare the state

$$|\psi\rangle = \frac{1}{\sqrt{D_e}} \sum_k |0\rangle_\mu |k\rangle_e, \quad (7.19)$$

which can be compared with the state given in Eq. (7.6). In order to generate the desired random phase for each basis state  $|k\rangle_e$ , I apply random single particle  $R_z(\theta)$  rotations to each particle. This procedure will not generate a state with completely independent phases for each basis state. If it is necessary to further randomise the state, one could apply layers of  $R_z$  and controlled- $R_z$  rotations between the different particles. Averaging the results of several simulations (each with a different set of randomly chosen  $\theta_i$ ) yields the polarisation at the desired time, as shown by Eq. (7.7). While the asymptotic convergence properties of the two methods are the same, the latter method yields a smaller error for a given number of samples.

For both of the methods discussed above, it is necessary to construct the states  $|k\rangle_e$ , either alone, or in superposition. These states are tensor products of the environment spins each in an arbitrary state  $|j\rangle$ , that is an  $s_z$  eigenstate. As discussed in Sec. 7.3.1, I have mapped these spin states onto qubit Dicke states. Efficiently constructing Dicke states (or superpositions of Dicke states) on quantum computers has remained challenging for a number of years [335–338], but has recently been made possible with the elegant inductive solution of Ref. [286]. This approach can be summarised as follows, and is explained in more detail in Appendix. A.2.1. I first define a Dicke state with Hamming weight  $h$ , on  $n$  qubits, as  $|D_h^n\rangle$ . For example, the  $|-1/2\rangle$  state in Eq. (7.15) equates to  $|D_2^3\rangle$ . Then observe that

$$|D_h^n\rangle = \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \quad (7.20)$$

Assume the existence of a unitary operator  $U_{n,k}$  such that  $U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} = |D_h^n\rangle$  for all  $h \leq k$ . Through induction, it can be shown that this unitary operator exists, and how to construct it from typical single- and two-qubit gates. For example, note that

$$|D_h^n\rangle = U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \quad (7.21)$$

and

$$\begin{aligned} |D_h^n\rangle &= \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \\ &= U_{n-1,k} \left[ \sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} \right. \\ &\quad \left. + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle \right] \\ &= U_{n-1,k} V_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \end{aligned} \quad (7.22)$$

implying that

$$U_{n,k} = U_{n-1,k} V_{n,k}. \quad (7.23)$$

As shown in Appendix. A.2.1, these relations can be repeated recursively, to obtain a circuit purely in terms of the  $V_{n,k}$  type gates. Ref. [286] show how these gates can be constructed from CNOT gates and single-qubit rotations (see also Appendix. A.2.1). The entire state preparation circuit has a depth of  $\mathcal{O}(n)$ , and requires  $\mathcal{O}(kn)$  gates, even on a 1D linear chain of qubits. The algorithm also requires no additional ancillary qubits. Because the unitary  $U_{n,k}$  that creates the Dicke states was defined to work for all input states with  $h \leq k$ , the unitary  $U_{n,n}$  can be used to construct all Dicke states  $|D_{h \leq n}^n\rangle$ . As a result, inputting the superposition

$$\sqrt{\frac{1}{n+1}} \sum_{k=0}^n |0\rangle^{\otimes n-k} |1\rangle^{\otimes k} \quad (7.24)$$

outputs an equally weighted superposition of all of the possible Dicke states. Applying this circuit to each environment spin generates the state shown

in Eq. (7.19). It is straightforward to generate the state in Eq. (7.24), using a ladder of control- $R_y$  gates, as shown in Ref. [286]. As a result, a similar number of gates are required for both of the initial state preparation methods described in this section.

### 7.3.3 Evolving the state in time

As discussed in Sec. 3.1.1, there are a number of quantum algorithms for performing time evolution. In this chapter, I focus predominantly on product-formula methods (Trotterization), which have previously been observed to lead to lower gate counts than some asymptotically more efficient algorithms when considering the time evolution of spin systems [243]. In both this section and Sec. 7.4, I make use of the Trotter error bounds discussed in Sec. 3.1.1. For a second-order Trotterization of the time evolution operator, the error in the unitary is upper-bounded by [96]

$$\|U(t) - U_2(t)\| := \epsilon \sim \mathcal{O}\left(\frac{(L\Lambda t)^3}{n^2}\right), \quad (7.25)$$

where  $U(t)$  is the true time evolution operator,  $\| \cdot \|$  denotes the spectral norm,  $L$  is the number of terms in  $H$ ,  $n$  is the number of Trotter steps used, and  $\Lambda = \max_{\alpha} \|h_{\alpha}\|$ . As discussed above, the dipolar Hamiltonian can be decomposed into up to  $L \sim \mathcal{O}(N^2 s_{\max}^2)$  two-qubit Pauli terms, where  $s_{\max}$  is the largest spin value in the system. Considering a second-order product formula, the number of Trotter steps required to obtain Trotter-error  $\epsilon$  is upper bounded by

$$n \sim \frac{N^3 s_{\max}^3 \Lambda^{\frac{3}{2}} t^{\frac{3}{2}}}{\epsilon^{\frac{1}{2}}}. \quad (7.26)$$

Each second-order Trotter step requires  $\mathcal{O}(L)$  gates to implement, resulting in a total gate count that scales at worst as

$$G \sim \frac{N^5 s_{\max}^5 \Lambda^{\frac{3}{2}} t^{\frac{3}{2}}}{\epsilon^{\frac{1}{2}}}. \quad (7.27)$$

This estimate was obtained by considering a loose bound on the second-order Trotter error, which does not take into account commutativity between terms. As a result, it is likely that significantly better scaling could be obtained by using tighter error bounds, and considering favourable compilations of the Trotter ordering. Similar progress has been seen in the context of quantum chemistry calculations. While initial estimates suggested a scaling of  $\mathcal{O}(N^{11})$  [115], more careful analysis reduced the asymptotic scaling to  $\mathcal{O}(N^{5.5})$  [116]. Furthermore, the aim of the simulation algorithm is to obtain an accurate estimate for the value of the muon polarisation function at a given time. As a result, we are not interested in the Trotter error  $\epsilon$  directly, but in the error it induces on  $\text{Tr}(Z_\mu \rho(t))$ . It is likely that the Trotter error will provide a loose upper bound for this error, as has been seen previously [96, 339]. It is also interesting to consider if randomized compilations of the Trotter circuit [97], or related algorithms, such as qDRIFT [105], could be used to further reduce the gatecount required. As discussed in Sec. 3.1.1, the scaling of qDRIFT does not depend explicitly on the number of terms in the Hamiltonian. This makes it an interesting candidate for systems with a large number of weakly interacting terms – such as the dipolar Hamiltonian considered in muon spectroscopy, which exhibits rapid power-law decay.

### 7.3.4 Measuring the polarisation

The final stage of the algorithm consists of measuring the  $Z$  expectation value of the muon qubit. The most straightforward way to measure this is via preparation of the state at time  $t$ , followed by a projective measurement in the computational basis. This process must be repeated  $\mathcal{O}(1/\epsilon^2)$  times, where  $\epsilon$  is the desired precision. The repetition rate of a quantum processor can depend on a number of factors, including: the depth of the circuit, the

speed of executing quantum gates, and whether error correction is used.

Alternative quantum algorithms can reduce the time cost to  $\mathcal{O}(1/\epsilon)$ . These techniques rely on a combination of quantum amplitude amplification and phase estimation [340, 341]. These approaches use a constant number of samples, but require a circuit depth of  $\mathcal{O}(d_U/\epsilon)$ , where  $d_U$  is the circuit depth required to implement the time evolution circuit  $U(t)$ . This is achieved by repeatedly evolving the register under a controlled version of the operator

$$\begin{aligned}\Omega &:= \Omega_1 \Omega_2, \\ \Omega_1 &= U(t) [I_s - 2 |\Phi\rangle \langle \Phi|] U^\dagger(t), \\ \Omega_2 &= Z_\mu \Omega_1 Z_\mu,\end{aligned}\tag{7.28}$$

where  $|\Phi\rangle$  is the initial state of the system. This conditional evolution is controlled by the state of an auxiliary register. Controlling the evolution on an auxiliary register enables us to perform quantum phase estimation on the unitary  $\Omega$ , which we can use to extract the value of  $\bar{Z}_\mu$  to the desired precision. A more detailed discussion of this approach is given in Refs. [340, 341].

The steps outlined in this section can be used to obtain a simulated polarisation function for the system of interest. In order to use this function to analyse experimental data, the quantum simulation routine can be incorporated into an optimisation loop. In Sec. 7.4, I consider the problem of locating the muon rest site. I first specify the positions of each particle in the system, and use this information to generate the spin and qubit Hamiltonians. I can then use the quantum simulation routine outlined above to calculate the simulated polarisation function. An optimisation loop can be used to minimise the value of a loss function between the generated and experimental data, by updating the positions of the particles in the system.

## 7.4 Results

In order to investigate the practicality of the algorithm discussed in Sec. 7.3, I performed numerical simulations of systems with up to 29 qubits. I considered a muon interacting with spin- $\frac{1}{2}$  fluorine nuclei in a sample of calcium fluoride ( $\text{CaF}_2$ ) via the dipolar Hamiltonian in Eq. (7.29). The calcium nuclei in the sample were neglected because they have spin-0 with an abundance of  $\sim 99.9\%$  [300]. The experimental data analysed was obtained in Ref. [300]<sup>1</sup>, which then used state-of-the-art classical simulations to determine the impact of the muon on the system geometry. Those simulations used exact diagonalisation of an 11 spin Hamiltonian – but scaled the interactions between the muon and more distant nuclei to act as a proxy for the remaining nuclei in the sample.

The geometry studied consisted of a simple cubic lattice of  $\text{F}^-$  ions, with a lattice constant of 2.72 Å. The muon implantation site was taken to be between two adjacent fluorines, as shown in Fig. 7.3. As the system was composed of only spin- $\frac{1}{2}$  particles, each particle was mapped to a single qubit. The Hamiltonian was given by

$$H_D = \frac{1}{2} \sum_{i,j} \frac{\hbar^2 \mu_0 \gamma_i \gamma_j}{4\pi r_{ij}^2} \left[ \vec{S}_i \cdot \vec{S}_j - 3(\vec{S}_i \cdot \hat{r}_{ij})(\vec{S}_j \cdot \hat{r}_{ij}) \right], \quad (7.29)$$

where  $\mu_0$  is the permeability of free space,  $\gamma_i$  is the gyromagnetic ratio of spin  $i$  ( $\gamma_\mu = 2\pi \times 1.355 \times 10^8 \text{ Hz}\cdot\text{T}^{-1}$ ,  $\gamma_F = 2\pi \times 4.006 \times 10^7 \text{ Hz}\cdot\text{T}^{-1}$ ),  $\vec{r}_{ij}$  is the vector connecting spins  $i$  and  $j$ , and  $\vec{S}_i = \frac{1}{2}(X_i, Y_i, Z_i)$ . The polarisation

---

<sup>1</sup>The data was taken using the pulsed muon beam at the ISIS Facility, Rutherford Appleton Laboratory, UK. Experiments were performed in zero magnetic field, at a temperature of 50 K. Around 181 million muon decay events are included in the dataset.

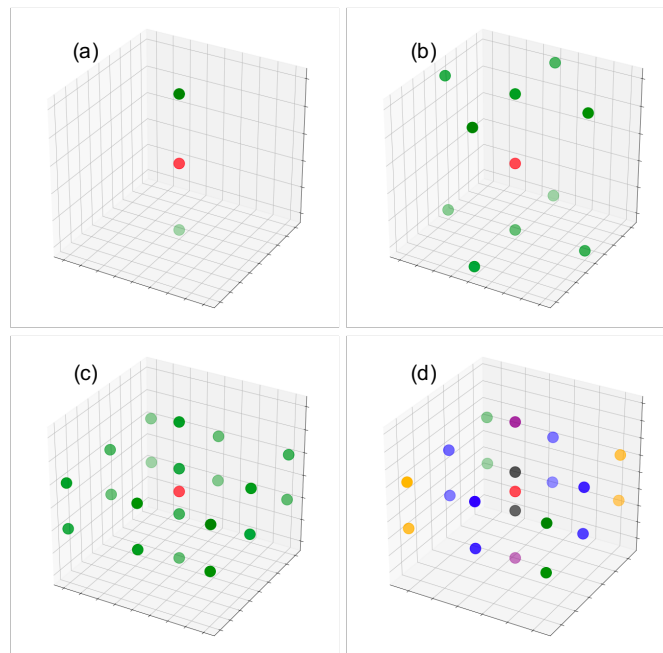


Figure 7.3: A selection of geometries investigated for the  $\text{CaF}_2 + \mu^+$  system. The calcium ions are spin-0 with an abundance of around 99.9 %, so I only included the fluorine ions (green in a - c) and the muon (red) in my simulations. a) The three spin  $\text{F}-\mu^+-\text{F}$  system. b) The 11 spin  $\text{F}-\mu^+-\text{F}$  system with 8 next-nearest-neighbour fluorines. c) The 21 spin system, which included the next-nearest 10 fluorines to the muon. d) The 21 spin system, with fluorines grouped by colour. The distances between the muon and each fluorine in a given group were the same. When fitting simulated data for the 21 qubit system to experimental data, I separately parameterised the  $\text{F}-\mu$  distances according to these groups.

function was appropriately averaged for this polycrystalline system

$$\begin{aligned} \langle P(t) \rangle = \frac{1}{3} & \left[ \text{Tr}(Z_\mu e^{-iH_s t} [|0\rangle \langle 0|_\mu \otimes \rho_e(0)] e^{iH_s t}) \right. \\ & + \text{Tr}(X_\mu e^{-iH_s t} [|+\rangle \langle +|_\mu \otimes \rho_e(0)] e^{iH_s t}) \\ & \left. + \text{Tr}(Y_\mu e^{-iH_s t} [|+Y\rangle \langle +Y|_\mu \otimes \rho_e(0)] e^{iH_s t}) \right] \end{aligned} \quad (7.30)$$

where  $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ ,  $|+Y\rangle = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$ . I considered simulations involving an increasing number of fluorine nuclei. The smallest system considered required three qubits, representing the muon and its two nearest-neighbour (nn) fluorines. Additional fluorine nuclei were then added in ‘shells’ determined by their distance from the muon. In the ‘next-nearest-neighbour’ (n-nn) shell, there were an additional 8 nuclei (thus 11 qubits total). The ‘next-next-nearest-neighbour’ (nn-nn) shell contributed an additional 10 nuclei (thus 21 qubits total). The nnn-nn shell added 8 nuclei (thus 29 qubits total).

I utilised a range of simulation techniques, in order to investigate a number of different properties of the proposed algorithm. Simulations of the random-phase-approximation based approach were carried out using the QuEST package. These simulations manually generated the initial state shown in Eq. (7.6), and then carried out quantum circuit emulations of first and second-order Trotter decompositions of the time evolution operator. Because these simulations used the random-phase-approximation approach, sampling noise was present in the results.

I also performed quantum circuit-level simulations of running the algorithm on a quantum processor, using the density matrix simulator included in Cirq. Due to the increased computational cost of storing and manipulating the density matrix, these calculations were restricted to systems of up

to 11 qubits. However, these simulations enabled the initialisation of the environment in a maximally mixed state, eliminating the sampling error present in wavefunction-based approaches. These density matrix simulations also provided a more efficient way to investigate the effects of circuit-level noise on the algorithm.

I also performed exact diagonalisation of Hamiltonians of systems with up to 11 qubits. This provided exact results which were used to quantify the error introduced by Trotterizing the time evolution operator.

### 7.4.1 Noiseless simulations

The random-phase-approximation method contains two sources of algorithmic error; error introduced by Trotterization, and sampling errors arising from working with a random-phase augmented wavefunction, rather than with an actual mixed state. In this section, I present the results of numerical simulations that quantify the magnitude and scaling of these sources of error. Given the similarities between the random-phase-approximation method and the quantum algorithm introduced in Sec. 7.3, these simulations help to quantify the quantum resources required to run the algorithm.

#### 7.4.1.1 Quantifying errors

I first investigated the Trotter error for the 3 and 11 qubit systems. I isolated the Trotter error from the sampling error using the density matrix simulator. I used both first and second-order Trotter formulae, where Hamiltonian terms with the largest magnitude were placed first in the product formula sequence. These results were compared to those obtained from exact diagonalisation. I compared the numerically simulated Trotter error to two different bounds for the Trotter error, given in Sec. 3.1.1. The first bound is given by Eq. (3.2) for a first order Trotterization, and Eq. (3.5) for second-order Trotter, and represents a loose bound on the error between

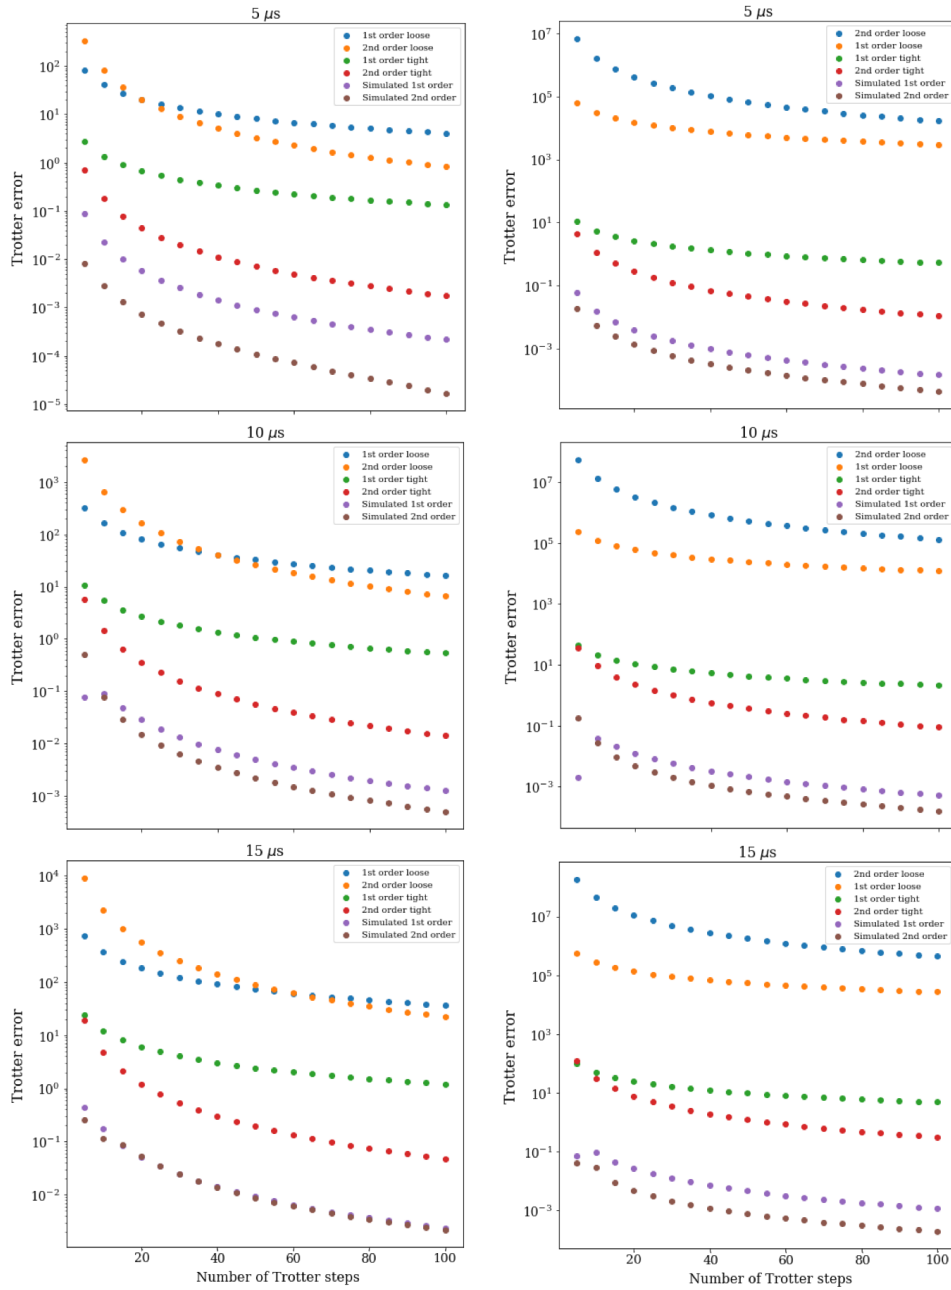


Figure 7.4: The Trotter error in the 3 (left) and 11 (right) qubit systems. The error bounds used are described in the main text. Hamiltonian terms with the largest coefficients were placed first in the Trotter formula.

the exact time evolution operator and the unitary specified by the product formula. The second, tighter bound, which considers commutativity

between Hamiltonian terms, is given by Eq. (3.3) for a first order Trotterization, and Eq. (3.6) for second-order Trotter (but with the norm moved inside of the sums in both cases, which loosens the bound). These quantities upper bound the error in any observable measured after time evolution, and so were strictly larger than the error in the polarisation value obtained from the numerical simulations. The Trotter errors for the 3 and 11 qubit systems are shown in Fig. 7.4. I plot the Trotter errors obtained at time values of 5, 10, and 15  $\mu\text{s}$ .

In all cases considered, the ‘loose’ error bounds were orders of magnitude larger than the ‘tight’ error bounds, which in turn were at least an order of magnitude larger than the numerically observed Trotter errors. As discussed above, this can be partially explained by the fact that the analytic bounds give errors in the unitary evolution, while the numerical results give the error in a specific observable. In addition, despite the improved bounds given by the ‘tight’ formulae, they are still known to not be completely tight to numerical results [96]. However, the fact that the tight bounds were still at least an order of magnitude larger than the numerical results highlights the value in work to bound the error in specific observables, rather than existing worst-case bounds. While initial steps have been taken in this direction [96, 339], it would be interesting to consider if tighter bounds are possible for the case of muon spectroscopy, given the simple form of the observable measured.

I also investigated the error present in the random-phase-approximation method, which arises from sampling pure state wavefunctions

$$|\psi_s(t)\rangle = U(t) \left[ \frac{1}{\sqrt{D_e}} \sum_j e^{i\theta_j} |0\rangle_\mu |j\rangle_e \right]. \quad (7.31)$$

I considered two possible randomization schemes. The first scheme, referred to as the ‘random-phase-approximation (RPA) method’ was exactly




Figure 7.5: The sampling error present in the random-phase-approximation (RPA) and dephasing methods described in the main text, for the 11 qubit system with 40 second-order Trotter steps. Each point is the average of 20 values taken in the interval  $[0, 10)\mu\text{s}$ .

the same as the classical method of Ref. [319], and considered  $\theta_j$  chosen uniformly at random in the range  $[0, 2\pi)$ . The second, referred to as the ‘dephasing method’ (as it bears resemblance to passing an equal superposition state of qubits through an  $N$ -qubit dephasing channel) considered  $\theta_j$  chosen at random from the discrete set  $\{0, \pi\}$ . The sampling error was isolated by comparing the results from the wavefunction-based approaches to those obtained from density matrix simulations – using Trotterized time evolution for both methods. Fig. 7.5 shows that the sampling error arising from the RPA method was around an order of magnitude smaller than that of the dephasing method, due to the increased randomization of the former.

In Appendix A.2.2 I present the results of additional simulations used to test the convergence of both the Trotter and sampling errors for systems too large to study with exact diagonalisation.

### 7.4.1.2 Analysing $\text{CaF}_2 + \mu^+$ spectra

In Fig. 7.6 I plot the angular averaged polarisation functions obtained for the 3, 11, 21, and 29 qubit  $\text{CaF}_2 + \mu^+$  systems for the first 9.5  $\mu\text{s}$  of evolution. The simulated data points for the 3 and 11 qubit systems were obtained from a density matrix simulation, using 30 second-order Trotter steps, which sufficed to measure the polarisation function to an accuracy of less than  $10^{-2}$ . The polarisation functions obtained from exact diagonalisation are plotted for comparison. The simulated data points for the 21 and 29 qubit systems were obtained using the RPA method, with 48 samples used for the 21 qubit system, and a single sample used for the 29 qubit system. Both of these simulations also used 30 second-order Trotter steps. Fig. 7.6 shows that introducing additional nuclei caused a damping effect on the polarisation function, which appears well converged by 29 qubits. This observation supports the numerical results of Ref. [300], which showed that additional environmental spins act as a source of decoherence, as polarisation leaks from the muon to the environment.

Having confirmed that the algorithm could produce the qualitative results expected, I then used it to locate the muon rest site in  $\text{CaF}_2$ . This was achieved by parametrizing the muon-fluorine distances, and generating a polarisation function for a given geometry. I used the Nelder-Mead algorithm to minimise the reduced- $\chi^2$  value of the simulated data, by optimising the fitting parameters. It was necessary to use a derivative free optimisation method because of the sampling noise present in the RPA method, which was larger than the finite-difference steps used to calculate the gradient in many black-box optimisation algorithms. I used this approach to optimise the geometry of the 21 qubit system. Five geometric fitting parameters were used, which I describe via the colours used in Fig. 7.3d:

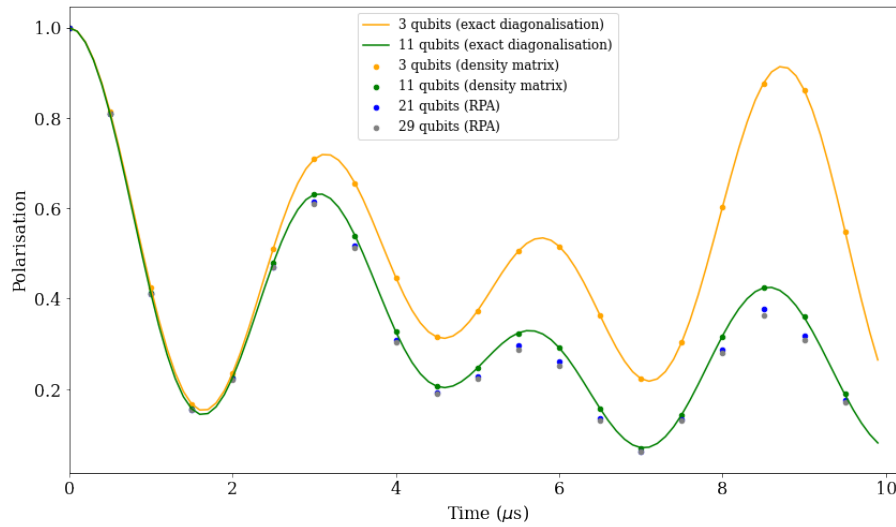


Figure 7.6: Polarisation functions for the unperturbed 3, 11, 21, and 29 qubit  $\text{CaF}_2 + \mu^+$  systems. All simulations used 30 second-order Trotter steps to simulate time evolution. The 3 and 11 qubit results were obtained using a density matrix simulation, that eliminates sampling error. The 21 and 29 qubit simulations were performed using the random-phase-approximation method, with 48 and 1 samples, respectively.

1. The distance between the muon (red) and the nearest-neighbour fluorines (both coloured black).
2. The distance between the muon and the next-nearest-neighbour fluorines (all coloured blue).
3. The distance between the muon and the green next-next-nearest-neighbour fluorines.
4. The distance between the muon and the purple next-next-nearest-neighbour fluorines.
5. The distance between the muon and the orange next-next-nearest-neighbour fluorines.

together with two parameters to fit the asymmetry

$$A(t) = A_0 P(\theta_1, \dots, \theta_5; t) + A_{\text{bg}} \quad (7.32)$$

I used the RPA method to generate the polarisation function, with 40 second-order Trotter steps, and a single sample for each data point.

The fitted data, shown in the upper plot of Fig. 7.7, is in qualitative agreement with the experimental data, particularly at early times. The fit obtained a reduced  $\chi^2$  value of 2.13, which suggests an incomplete fit to the data. I attribute this to a combination of limited experimental data at times beyond 10  $\mu\text{s}$ , as well as the ineffective nature of the Nelder-Mead optimisation algorithm, which was liable to becoming trapped in local minima. The fitting procedure caused the nearest-neighbour fluorines to move towards the muon by 0.188 Å. This is in excellent agreement with the calculation of Ref. [300], which yielded a value of 0.190(1) Å for the same quantity. The results of Ref. [300] were obtained by fitting an 11 spin system to the experimental data, and considering two physical parameters; the muon – nearest-neighbour fluorines distance, and a factor that scaled the strength of the next-nearest-neighbour interactions to act as a proxy for more distant nuclei. My fit was obtained by moving the next-nearest-neighbour fluorines towards the muon by 0.206 Å. While Ref. [300] did not explicitly consider the effect of perturbing the positions of the more distant nuclei, that work carried out density functional theory (DFT) calculations suggesting that the next-nearest-neighbour fluorines only move towards the muon by around 0.03 Å. I carried out additional numerical simulations with this geometry for the 21 spin system, shown in the lower plot of Fig. 7.7. This geometry resulted in a worse fit for the experimental data, with a  $\chi^2_\nu$  value of 4.44. This implies that either:

1. The system size of 21 spins was still not large enough to fully capture

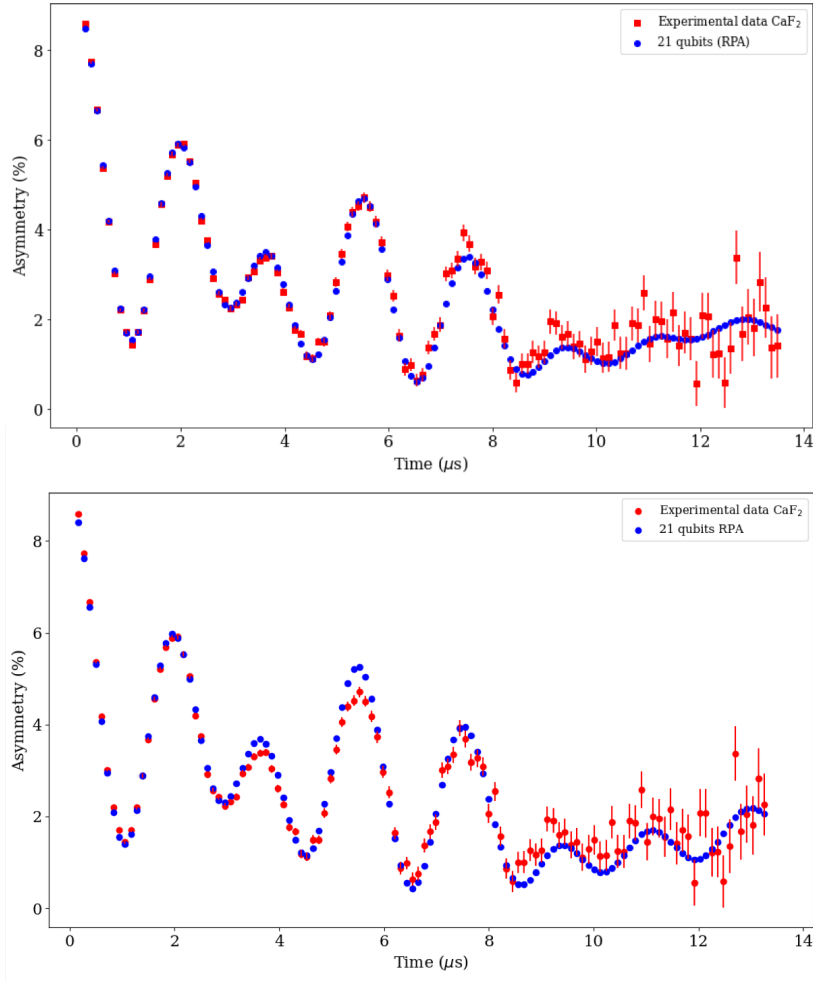


Figure 7.7: Fits of the 21 qubit simulated data (using the RPA method and 40 second-order Trotter steps) to the experimental data obtained in Ref. [300]. Upper: The fitting parameters were the 5 muon-fluorine distance parameters described in the main text, and a scaling factor and offset to convert the simulated polarisation value to an asymmetry value. The fit was performed using the gradient-free Nelder-Mead algorithm. 1 sample of the RPA method was used per point. Lower: The geometry was fixed; the nn and n-nn fluorines were moved towards the muon by 0.19 Å and 0.027 Å, respectively. The nn-nn fluorines were at their equilibrium positions. The polarisation function was fitted to the asymmetry data using  $A(t) = A_0P(t) + A_{\text{bg}}$ , where  $A_0$  and  $A_{\text{bg}}$  were fitting parameters. The fit was performed using the Levenberg-Marquardt algorithm. 48 samples of the RPA method were used per data point.

the continuum extrapolation of Ref. [300], and larger simulations are required (which are impractical to perform on classical hardware).

2. The fit obtained by the simulation may suggest an inaccuracy in the DFT results. However, from a physical perspective, it would be surprising if the more distant nuclei were more strongly attracted towards the muon than the nearest-neighbours.

The most likely explanation may be a combination of these factors, together with the fact that a number of similarly good local minima were present in the  $\chi^2$  surface for the parameter space. This highlights the value in utilising complementary techniques to analyse muon spectroscopy data, as well as the benefit provided by having access to as large a simulation of the system as is possible.

#### 7.4.2 Noisy simulations

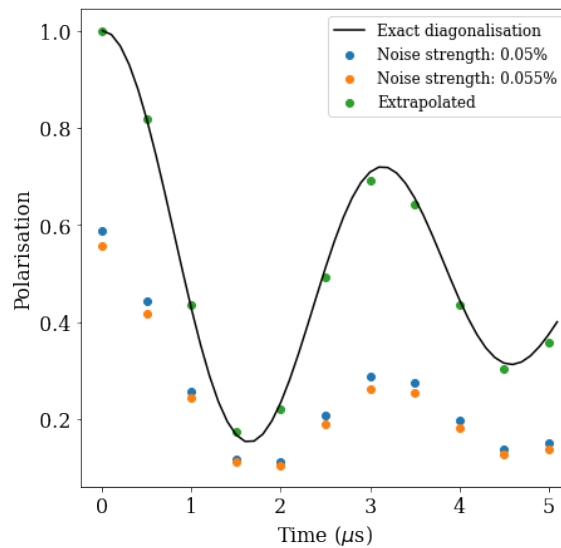


Figure 7.8: Using exponential extrapolation of data obtained at two different noise rates to infer the noise free polarisation function of the 3 qubit system. The extrapolation was performed using Eq. (7.33). The expected number of errors in the circuit was 0.79.

As discussed throughout this thesis, physical noise can corrupt the results of calculations. To investigate the effects of noise on the algorithm introduced in this chapter, I carried out noisy density matrix simulations of the 3 qubit  $F\text{-}\mu^+\text{-}F$  system. The circuit used 20 second-order Trotter steps, which sufficed to obtain an accuracy of less than  $10^{-3}$  in the polarisation function at all values within the first  $5\ \mu\text{s}$  (in the noiseless case). I merged adjacent single-qubit gates together, to yield a circuit with 900 single-qubit gates, and 680 two-qubit gates. The noise model was a single-qubit depolarising channel applied to each qubit involved in a gate, following every gate in the circuit. The physical error rate was set at  $p = 5 \times 10^{-4}$ , a factor of two better than the best error rates demonstrated to date [147, 149, 150]. The same circuit depth was used for all of the points calculated. This meant that the Trotter error was likely smaller in points taken at earlier times than at later times, while the effective physical noise rate of the circuit was similar at all points. In reality, it would be preferable to fix the Trotter error along the polarisation function, which would enable the use of a shorter depth circuit to simulate points at earlier times – thus reducing the physical noise in those datapoints.

To combat the effects of noise, I utilised the exponential extrapolation technique introduced in Sec. 3.2.1. The stretch factor was heuristically set to  $\lambda = 1.1$ . The extrapolated expectation value was given by

$$P_0 = \left( \frac{P_\epsilon^\lambda}{P_{\lambda\epsilon}} \right)^{\frac{1}{\lambda-1}}, \quad (7.33)$$

where  $P_\epsilon$  was the polarisation value calculated with the baseline noise rate, and  $P_{\lambda\epsilon}$  was the polarisation value calculated with the boosted noise rate. We see from Fig. 7.8 that exponential extrapolation was able to recover almost noiseless results, despite the large damping of the polarisation function in the unmitigated case. The noise strength of  $p = 5 \times 10^{-4}$  corre-

sponded to an expected number of errors per circuit of 0.79. The mean absolute error in the polarisation function after exponential extrapolation was 0.011.

In order to probe the noise robustness of the application of the algorithm to data analysis, I investigated locating the muon rest site using noisy simulated data for the 11 qubit system. Rather than considering the circuit-level noise model discussed above, I effectively engineered noise in the polarisation function through undersampling in the RPA/dephasing method. As shown in Fig. 7.5, it is possible to generate mean errors of 0.0243, 0.0068, and 0.0022 in the polarisation function by using the dephasing method with 10 samples, the RPA method with 1 sample, and the RPA method with 10 samples, respectively. There are two important assumptions present in this noise model. Firstly, it presumes that physical noise would cause the data points to become normally distributed around the noiseless value, as happens for the sampling noise. As is evident from Fig. 7.8, this is not necessarily the case; many sources of noise will simply cause a decay in the polarisation function. However, we can also see from Fig. 7.8 that performing extrapolation can change how the data points are distributed around the noiseless result. When shot noise in the quantum algorithm is taken into account, this may enable us to approach a normal distribution. Another path to recover normally distributed results could be to use the quasiprobability method of error mitigation described in Sec. 3.2.1. This approach has previously been observed to give results that are approximately normally distributed around the true values [30]. Secondly, the noise model assumes that the magnitude of the noise is the same for all data points. In reality, as discussed above, it would be better to fix the Trotter error along the polarisation function, and thus consider shorter depth circuits at earlier times. This would reduce the noise in earlier data points, mimicking the error bars

present in the experimental results. My simulations used 40 second-order Trotter steps, and considered two geometric fitting parameters:

1. The muon – nearest-neighbour fluorines distance.
2. The muon – next-nearest-neighbour fluorines distance.

along with two parameters to fit the asymmetry

$$A(t) = A_0 P(\theta_1, \theta_2; t) + A_{bg} \quad (7.34)$$

Once again, I used the Nelder-Mead algorithm to minimise the error-weighted least-squares residuals between the experimental data and the simulated data.

In Fig. 7.9 I plot examples of the fitted data obtained at the three noise values listed above. Even at high noise rates of around 2.5% in each data point, the fitting procedure was able to approximately capture the overall shape of the polarisation function. As the strength of the noise was reduced, increasingly better fits to the experimental data were obtained.

In order to better quantify the degree of noise robustness present in the algorithm, I plot the error that noise induced on the two geometric fitting parameters in Fig. 7.10. The noiseless values of the fitted parameters were obtained by fitting the 11 qubit polarisation function, generated by exact diagonalisation, to the experimental data, using the Nelder-Mead algorithm. The upper plot in Fig. 7.10 shows that as the average noise in each individual data point was reduced, an improved estimate of the noiseless muon – nearest-neighbour fluorine distance was obtained. Interestingly, the fractional error in the parameter was around an order of magnitude smaller than minimum fractional error in the polarisation value (the maximum value of the polarisation function is 1, so the minimum fractional error is equivalent to the noise strength of the simulation). I attribute this

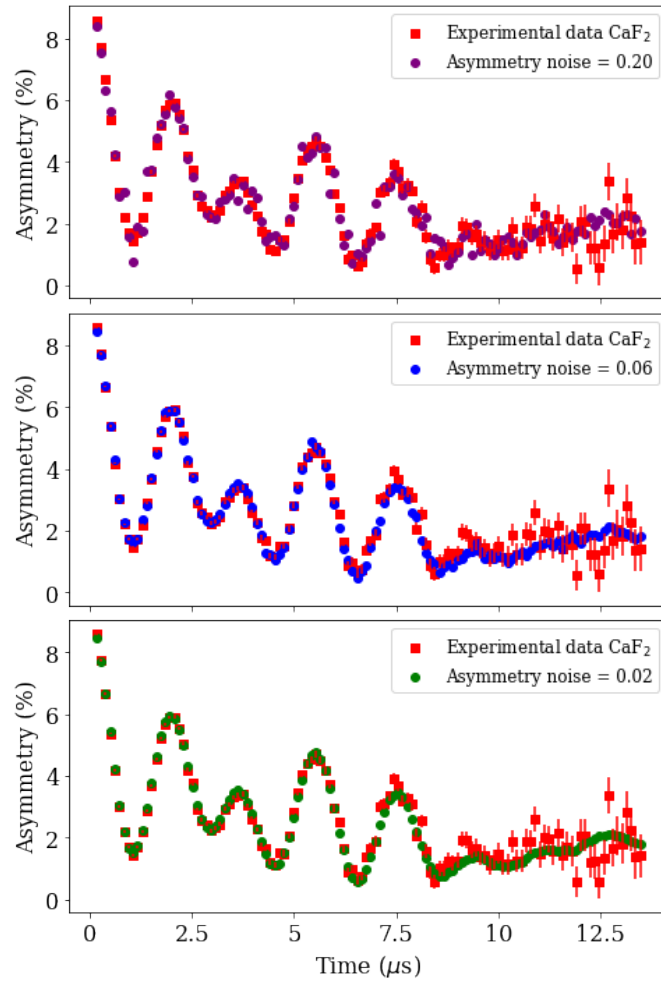


Figure 7.9: Fits of the 11 qubit simulated data with 40 second-order Trotter steps to the experimental data obtained in Ref. [300]. The noise in each simulated datapoint was varied by controlling the number of samples used. The fit was performed using the gradient-free Nelder-Mead algorithm. The fitting parameters were the muon-fluorine distances described in the main text, and a scaling factor and offset to convert the simulated polarisation value to an asymmetry value.

noise resilience to the fact that the fitting procedure was extracting a global property of the data (the muon – fluorine distance) rather than the value of any individual data point. The lower plot in Fig. 7.10 suggests that the muon – next-nearest-neighbour fluorine distance was less resilient to noise, as the fractional error in the parameter was approximately the same order




Figure 7.10: The fractional error in the geometric fitting parameters used to fit the simulated muon spectroscopy data to the experimental data in Fig. 7.9. The noiseless values of the parameters were obtained through exact diagonalisation of the 11 qubit Hamiltonian to obtain a polarisation function, which was fitted using the Nelder-Mead algorithm.

of magnitude as (but larger than) the magnitude of the noise in the polarisation function. Nevertheless, I am optimistic that the algorithm could be made more noise resilient by: 1) Reducing the noise strength at earlier times by adapting the number of Trotter steps used (as discussed above), and 2) Using an optimisation routine known to be more resilient to noise than the Nelder-Mead algorithm, such as Bayesian optimisation.

Linking back to the earlier simulations of circuit-level noise, I note the following:

1. The fitting routine obtained accurate parameter values when the error

in each simulated polarisation data point was less than around 0.01 (for the 11 qubit system).

2. Exponential extrapolation provided a mean absolute error in the polarisation function of 0.011 when the expected number of errors in the circuit was around 0.8 (for the 3 qubit system).

Assuming that the circuit-level result holds for larger system sizes (in reality, the results may improve in larger system sizes, as when errors do occur they will be less likely to occur on the crucial muon qubit), I infer that the algorithm is able to tolerate an expected error rate of around 0.8 errors per circuit, on average. In the 11 qubit case, this would enable us to learn the muon – nearest-neighbour fluorine distance to an accuracy better than 0.1%, and the muon – next-nearest-neighbour fluorine distance to an accuracy of around 5%.

### 7.4.3 Resource estimates

In this section I discuss the quantum resources required to perform the calculations described above. For NISQ resource estimates, I consider the number of two-qubit gates as the cost metric, as they provide the largest contribution to the circuit error rate in many hardware systems. For error corrected simulations, I take the approach discussed in Sec. 3.2.2. I considered a calculation performed using the surface code, and sought to minimise the number of T gates. I used the rule-of-thumb that around 100 T gates can be used to generate an arbitrary angle single-qubit rotation, which is appropriate for the gate count estimates in this chapter. In these resource estimates, I focus in particular on the 11 qubit system. This system is too small to be a viable target for beyond-classical simulation (as it can be solved with exact diagonalisation of the system Hamiltonian) – however, this means that it is possible to get tight bounds on the Trotter error

for this system, which makes it possible to obtain more accurate resource estimates.

### 7.4.3.1 NISQ resource estimates

As shown in Fig. 7.4, approximately 40 second-order Trotter steps were required to reduce the Trotter error in the observable to  $10^{-3}$  when calculating  $P(t=15\ \mu\text{s})$  for the 11 qubit system. I performed a basic compilation of the circuit, by merging adjacent single-qubit gates. This resulted in a gate count of  $5 \times 10^4$  single-qubit gates, and  $3.9 \times 10^4$  two-qubit gates. Assuming that single-qubit gates have an error rate ten times lower than two-qubit gates, then this circuit is roughly equivalent to implementing  $4.4 \times 10^4$  two-qubit gates. As discussed in Sec. 7.4.2, I observed that for the 3 qubit system, exponential extrapolation was able to recover acceptably accurate results when there were around 0.8 expected errors per circuit. Achieving this circuit error rate for the 11 qubit system with  $4.4 \times 10^4$  two-qubit gates, would require a two-qubit gate error rate of  $p = 2 \times 10^{-5}$ . This is approximately two orders of magnitude lower than the current lowest two-qubit gate error rates [147, 149, 150]. Simulations of larger, classically challenging system sizes would require an even larger number of gates. The most promising avenue for realising these calculations before the advent of quantum error correction is to significantly reduce the gate count required. One possible route towards this goal would be to consider Trotter orderings that reduce the Trotter error, or that enable a larger number of gates to be cancelled. We could alternatively consider neglecting the dipolar interactions between the nuclei, retaining only the muon-fluorine dipolar interactions (as has previously been done in classical simulations [321]). Even if these optimisations can be incorporated, it appears challenging for NISQ devices to surpass classical  $\mu^+$ SR simulation capabilities. This is not necessarily reflective of the performance of the algorithm described in this

chapter, and instead serves to highlight both the quality of classical simulations of quantum systems, and the challenges inherent in many NISQ algorithms, resulting from high noise rates in existing quantum processors.

### 7.4.3.2 Error corrected resource estimates

I followed the approach to surface code resource estimation taken in Ref. [185], which considers surface code operations implemented using lattice surgery, complemented with magic state distillation for implementing T gates. I initially sought to minimise the spatial resources required. I used a ‘compact’ data block of physical qubits to store each logical qubit. For a system with  $Q$  logical qubits, using a compact block results in needing  $\lceil 1.5Q + 3 \rceil$  surface code tiles. For a distance  $d$  surface code, each surface code tile consists of  $2d^2$  physical qubits ( $d^2$  data qubits and  $d^2$  syndrome qubits).

The circuit for the 11 qubit system, with 40 second-order Trotter steps, contained  $1.96 \times 10^4$  non-Clifford single-qubit rotations, which I assumed could be synthesised from  $1.96 \times 10^6$  T gates. I used a 15-1 distillation block for magic state distillation, which produces output magic states with an error rate of at most  $35p^3$  (where  $p$  is the physical error rate), and requires an additional 11 surface code tiles. For the 11 qubit system, this resulted in  $B = 31$  surface code tiles in total. With this setup, a magic state could be consumed every  $11d$  surface code cycles. The calculation therefore needed to ‘survive’ for  $T = 11d \times (1.96 \times 10^6)$  code cycles. The size of the surface code required to perform the calculation was determined using [185]

$$B \cdot T \cdot 0.1(100p)^{(d+1)/2} < \epsilon \quad (7.35)$$

where  $p$  is the physical error rate, and  $\epsilon$  is the target circuit error rate. The fault-tolerant resources required are shown in Table. 7.1, for a number of

$p$	$\epsilon$	$d$	Physical qubits	Circuit time (s)
$10^{-3}$	0.01	22	30,008	474
$10^{-4}$	0.01	10	6200	216
$10^{-3}$	0.8	18	20,088	388
$10^{-4}$	0.8	8	3968	172

Table 7.1: Error corrected resource estimates for simulations of the 11 qubit  $\text{CaF}_2 + \mu^+$  system with 40 second-order Trotter steps. Here,  $p$  is the physical error rate,  $\epsilon$  is the expected number of errors per circuit, and  $d$  is the surface code distance used.

scenarios. I considered resource estimates with a realistic two-qubit gate error rate of  $p = 10^{-3}$ , as well as a more optimistic error rate of  $p = 10^{-4}$ . I also considered two target noise suppression levels. In the first case, I considered suppressing the error rate such that there would be  $\epsilon = 0.01$  errors in the circuit, on average. In the second case, I assumed that the algorithm was able to tolerate a noise rate of  $\epsilon = 0.8$  errors in the circuit, on average. This choice of target circuit error rate was motivated by the success of exponential extrapolation in obtaining accurate results for the 3 qubit system in the presence of noise of this magnitude. I assumed a surface code decoding cycle time of  $1 \mu\text{s}$ . We can see from Table 7.1 that while the error robustness of the algorithm reduced the resources required by a factor of around 1.5, this was overshadowed by the exponential improvements arising from reductions in the physical error rate.

I also estimated the resources required to simulate the 29 qubit system. I assumed that 50 second-order Trotter steps would be needed, which led to a circuit with  $2.3 \times 10^7$  T gates. I used the same surface code setup described above, and considered a physical noise rate of  $p = 10^{-3}$ , and a target circuit error rate of  $\epsilon = 0.8$ . I found that a distance 21 surface code was needed, requiring around 51,000 physical qubits. The calculation would take 5234 seconds to run. On first inspection, this is a modest number of

physical qubits, compared to existing resource estimates for solving classically challenging chemistry calculations with quantum hardware, which have previously been estimated to require around  $10^5$  [25] to  $10^6$  [27, 69, 72] physical qubits at  $p = 10^{-3}$  physical error rates (but admittedly, are solving more challenging problem instances than the 29 qubit simulations discussed here).

Unfortunately, there are limitations to the approach taken in my resource estimates, and it is necessary to think more carefully about how these calculations would be performed on error corrected quantum computers. As seen in Table. 7.1, performing a single iteration of even the 11 qubit simulation would take several minutes, assuming a surface code decoding cycle time of  $1 \mu\text{s}$  (which may be optimistic for slower systems, such as trapped ion quantum computers [342]). Estimating the polarisation value to a precision of  $10^{-2}$  would take around 10,000 repetitions using direct sampling. The calculation must then be repeated at a number of simulated time values to perform a single step of the optimisation subroutine. In turn, the optimisation routine may require hundreds of iterations to converge. As a result, if the methods discussed in this chapter are to prove useful for analysing experimental data, their runtime must be reduced.

As discussed in Sec. 7.3.4, amplitude amplification can be used to reduce the time cost of measuring an observable to  $\mathcal{O}(\frac{1}{\epsilon_m})$ , where  $\epsilon_m$  is the desired precision. This comes at a cost of increasing the circuit depth required by a factor of  $1/\epsilon_m$ . While a full resource estimate of using this technique is beyond the scope of this work, I performed a rough estimate of how this approach would compare to that discussed above. I assumed a value of  $\epsilon_m = 10^{-2}$ , increasing the circuit depth required by a factor of 100. For the 29 qubit system, this increased the T count estimate to  $2.3 \times 10^9$  T gates. I neglected the resources required for magic state distillation, and focused on the resources required to reduce the expected number of errors in the

logical qubits to less than  $\epsilon = 0.01$ . Assuming that distilling a magic state still takes  $11d$  surface code cycles, this was achieved when

$$\lceil (1.5 \times 29) + 3 \rceil \cdot (11 \times 2.3 \times 10^9) \cdot 0.1(100p)^{(d+1)/2} < 0.01. \quad (7.36)$$

This expression is satisfied for  $d \geq 29$ . The compact data block for the 29 logical qubits would correspond to around 98,000 physical qubits. With this minimal setup, the calculation would take at least 8 days to complete. As discussed in Ref. [185], it is possible to reduce the runtime of the calculation by exchanging spatial resources for time resources. Adding more qubits enables faster distillation and consumption of magic states, to the point where a magic state can be consumed every  $d$  surface code clock cycles. The addition of further qubits enables parallelisation of T gates using quantum teleportation. If the T gates could be fully parallelised into around  $8 \times 10^7$  layers, and each layer could be implemented in one surface code cycle, a calculation of the polarisation value along a single axis, at a single point in time, would finish in 80 seconds. Building a quantum computer with these capabilities would require billions of qubits [185] – considerably larger than any machines that are planned within the coming decades.

Given the challenges associated with building even a small error corrected quantum computer, it is natural to ask whether there are other ways to reduce the runtime of the algorithm. The best way to achieve this is with algorithmic improvements. As in the NISQ case, better compilation routines, or the use of more efficient time evolution algorithms, could reduce the number of gates required to execute this proposal. In particular, we could consider alternative Hamiltonian decompositions that reduce the Trotter error by grouping commuting terms [343], or that exploit the locality of power law interactions [344]. Similarly, we could utilise techniques that ex-

exploit symmetries in the Hamiltonian coefficients to reduce the total number of T gates required for the circuit [25, 345].

## 7.5 Discussion

In this chapter, I have introduced and investigated a quantum algorithm for the analysis of data arising from muon spectroscopy experiments, which can be exponentially costly to perform using classical methods. I carried out numerical emulations of applying the algorithm to the  $\text{CaF}_2 + \mu^+$  system, and obtained results in good agreement with recent start-of-the-art classical analysis [300]. My simulations are the largest to date in the muon field, with a Hilbert space size of  $2^{29}$ .

By considering the impact of noise on the algorithm, I was able to estimate the quantum resources required to perform classically challenging instances of muon spectroscopy analysis. The algorithm appeared to be relatively noise robust, stemming from the aim of extracting a global parameter from the fitted polarisation function, rather than targeting any individual data point. These results may find use beyond this work, and suggest that analysing noisy experimental data may be a good target for future quantum computers.

Nevertheless, the resource estimates presented in this section highlight the challenges faced by both this algorithm, and many other quantum simulation algorithms. In the context of NISQ simulations, the gate counts produced by the algorithm are likely too large to simulate on devices with realistic noise rates, even if error mitigation techniques are utilised. As discussed in Chapter 5, this is also a key challenge for NISQ algorithms targeting the electronic structure problem. While noise can ultimately be overcome using quantum error correction, challenges also remain in this

arena. At first glance, analysing muon spectroscopy data appeared to require fewer fault tolerant resources than solving challenging instances of the electronic structure problem. However, the large number of repetitions required resulted in an impractically long runtime. While this runtime can be reduced by increasing the size of the quantum computer, or parallelising the calculation across multiple quantum machines, these both come at significant cost. I would argue that this is not necessarily a limitation of the algorithm presented here, but a challenge facing many quantum simulation algorithms. For example, while Ref. [243] performed a more rigorous analysis of the gate counts required to perform classically intractable simulations of spin- $\frac{1}{2}$  systems, it stopped short of estimating the resources required to estimate a given observable. Those gate counts, on the order of  $10^8$  T gates, are similar to those presented in Sec. 7.4.3.2. This will lead to the same problem discussed herein, when taking into account the resources required to estimate a given observable, or to sample a range of conditions. Similarly, while existing estimates for solving the electronic structure problem on small fault tolerant quantum processors [25–27, 68, 69, 72] show how to obtain the energy of a classically intractable system in hours or days, these works typically stop short of considering what problem that actually solves. In order to optimise a molecular geometry, estimate other observables on the ground state, or elucidate a phase diagram, these calculations will likely have to be repeated many times – resulting in what may be a prohibitively long calculation time. As such, I stress that it is important to consider the total resources required to solve a given problem with a quantum computer, and not just the quantum resources required to run the corresponding circuit once. The embarrassingly parallel nature of many classical algorithms (including the classical emulations used in this chapter) and the low cost and ubiquity of classical hardware, will place stringent requirements on the performance of future quantum algorithms. Sim-

ilar arguments were recently made by Ref. [346], in the context of whether quadratic quantum speedups will be sufficient to show quantum advantage on realistic problem instances.

There are a number of optimisations that could be introduced to make quantum algorithms more practical for muon data analysis. In particular, the use of more efficient time evolution algorithms, or improved compilations of the time evolution primitives, could dramatically reduce the number of gates required to achieve a given accuracy. Liaising with the muon community might also enable the introduction of problem-specific optimisations of the algorithm. For example, one could consider incorporating the ideas of Ref. [300], which scaled the interaction strengths of more distant nuclei to act as a proxy for the rest of the sample.

A related question is whether the quantum algorithm introduced in this chapter could be simulated efficiently using a classical computer. The environment being in a mixed initial state may motivate the belief that approximate classical methods may be able to efficiently simulate this problem. While answering this question is beyond the scope of this work, the algorithm introduced here has clear links to the ‘one clean qubit’ (DQC1) model of quantum computing [347], which has so far resisted efficient classical simulation [348–351].

Finally, it is interesting to ask if other, more complex quantum algorithms could be applied to analysing muon spectroscopy data. One possibility could be to use quantum read-only memory (QROM) [26] to load the experimental data values into the quantum computer. We could then attempt to compute all of the simulated data in superposition (the state would resemble  $\sum_t |t\rangle |P(t)\rangle$ , where  $t$  denotes the simulated time), and use existing

quantum machine learning algorithms to extract fitting parameters of interest. While the details of this approach will likely be more complicated than outlined here, I note that this is a slightly different approach than is typically considered in quantum machine learning algorithms. We would be exploiting the exponential speedup in calculating  $P(t)$  (and potentially polynomial speedups in parameter fitting), and do not mind that there is no speedup for loading in the data. We would seek to load in a number of datapoints that is generally constant for a given type of muon spectroscopy experiment (as it is defined by the characteristics of the beam type used). Given the numerous possible avenues for exploration, I believe that both muon spectroscopy, and the analysis of data arising from other experiments underpinned by quantum mechanics, could be promising targets for future quantum computers.

## 8 | Conclusion

Quantum simulation appears to be one of the most exciting and promising applications of quantum computing. If we are able to tame the spectre of noise, then we may soon be able to carry out calculations that would be intractable to perform on classical hardware. Nevertheless, the high overheads imposed by current methods of error correction lead to dauntingly high resources for such calculations. As a result, it is crucial to develop new ways to tackle quantum simulation problems.

One possible approach is to consider alternative simulation algorithms. In Chapter 4 I introduced and discussed an algorithm for ansatz-based imaginary time evolution, that could be used to find the ground states of chemical systems. The possibility of applying this algorithm to explicitly correlated (non-Hermitian) chemical Hamiltonians appears particularly exciting, as it may provide a new way to account for dynamic correlation in quantum simulations.

Another option is to develop lower cost methods for noise mitigation. I discussed one such approach, known as symmetry verification, in Chapter 5. By exploiting our knowledge of symmetries that should be respected by the simulation, the method is able to filter out a large number of errors, using modest additional resources. Such error mitigation techniques will be invaluable as we seek to perform classically intractable calculations using NISQ algorithms.

An alternative approach to reducing the costs of classically intractable simulations is to ‘move the goalposts’. Rather than trying to improve the efficiency of solving the electronic structure problem, we can try to identify alternative, more promising simulation targets. This is the approach taken in Chapters 6 & 7 of this thesis. In Chapter 6 I discussed how molecular

vibrations could be simulated on quantum hardware. Many existing quantum simulation tools can be re-purposed to target vibrational quantities of interest, where they face new opportunities and challenges. In Chapter 7 I introduced and investigated a quantum algorithm for analysing data from muon spectroscopy experiments. The algorithm yielded T gate counts of approximately  $2 \times 10^7$  (when minimising the footprint of the quantum computer) for classically challenging simulations of 29 qubit systems. This can be compared to estimates of around  $10^7 - 10^8$  T gates for simulating classically challenging instances of time evolution of the Heisenberg model (30 – 50 logical qubits) [243], around  $4 \times 10^6$  T gates for phase estimation of  $8 \times 8$  Fermi-Hubbard lattices (162 logical qubits) [352], and around  $3 \times 10^{10}$  Toffoli gates (2196 logical qubits) for phase estimation of the FeMoco molecule [27] discussed in Sec. 2.1.4. This investigation showed the value of examining new systems through a quantum lens – but also highlighted the importance of considering the resources required to solve a complete simulation problem (rather than just to run a single circuit instance). When this factor is taken into account, these resource estimates may increase significantly. Consequently, we may need to think carefully about whether the simulations listed above will be able to provide useful scientific insights, with fewer than  $10^6$  physical qubits, in a reasonable timeframe. Based on provable resource estimates, a sufficiently large error corrected quantum computer will be able to carry out simulations that appear intractable for classical computers. When exactly such a machine will become available is difficult to predict, and many challenges remain [353, 354]. Whether the NISQ machines developed en route to fault-tolerance are able to provide scientific insights for certain systems of interest remains to be seen. Achieving this goal will likely hinge on choosing simulation targets that maximise the strengths of NISQ hardware – as is done in the field of analog quantum simulation.

The last three years have seen rapid change in the field of quantum simulation, including: repeated drops in fault tolerant resource estimates for simulation problems [25–27, 69, 72], tightening of Trotter error bounds [96, 97, 243, 344, 355, 356], the development of a number of error mitigation techniques [139, 159, 160, 162, 253, 357, 358], an explosion in the number of quantum circuit emulation packages available for use [359–361], the increasing availability of cloud quantum processors [362–364], and the achievement of quantum supremacy [147].

To sustain the excitement over quantum simulation in the coming years, we will need to continue to explore ways to lower the resources required for challenging simulations. The work presented in this thesis has made contributions towards this goal in a number of ways. Nevertheless, there is (and likely always will be) much work left to be done. I hope to continue to contribute to this thriving research area – and look forward to the day when quantum simulations of quantum systems are commonplace.

## References

- [1] P. Dirac, “Quantum mechanics of many-electron systems,” *Proceedings of the Royal Society of London*, vol. 123, Apr 1929.
- [2] R. P. Feynman, “Simulating physics with computers,” *International journal of theoretical physics*, vol. 21, no. 6, pp. 467–488, 1982.
- [3] T. H. Johnson, S. R. Clark, and D. Jaksch, “What is a quantum simulator?,” *EPJ Quantum Technology*, vol. 1, p. 10, Jul 2014.
- [4] P. Hauke, F. M. Cucchietti, L. Tagliacozzo, I. Deutsch, and M. Lewenstein, “Can one trust quantum simulators?,” *Reports on Progress in Physics*, vol. 75, no. 8, p. 082401, 2012.
- [5] I. H. Deutsch, “Harnessing the power of the second quantum revolution,” *PRX Quantum*, vol. 1, p. 020101, Nov 2020.
- [6] I. Buluta and F. Nori, “Quantum simulators,” *Science*, vol. 326, no. 5949, pp. 108–111, 2009.
- [7] I. M. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation,” *Rev. Mod. Phys.*, vol. 86, pp. 153–185, Mar 2014.
- [8] A. Aspuru-Guzik and P. Walther, “Photonic quantum simulators,” *Nature physics*, vol. 8, no. 4, pp. 285–291, 2012.
- [9] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, “Cold bosonic atoms in optical lattices,” *Phys. Rev. Lett.*, vol. 81, pp. 3108–3111, Oct 1998.
- [10] I. Bloch, J. Dalibard, and S. Nascimbene, “Quantum simulations with ultracold quantum gases,” *Nature Physics*, vol. 8, no. 4, pp. 267–276, 2012.

- 
- [11] R. Blatt and C. F. Roos, “Quantum simulations with trapped ions,” *Nature Physics*, vol. 8, no. 4, pp. 277–284, 2012.
- [12] A. A. Houck, H. E. Türeci, and J. Koch, “On-chip quantum simulation with superconducting circuits,” *Nature Physics*, vol. 8, no. 4, pp. 292–299, 2012.
- [13] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch, “Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional bose gas,” *Nature physics*, vol. 8, no. 4, pp. 325–330, 2012.
- [14] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, “Exploring the many-body localization transition in two dimensions,” *Science*, vol. 352, no. 6293, pp. 1547–1552, 2016.
- [15] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator,” *Nature*, vol. 551, no. 7682, pp. 601–604, 2017.
- [16] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, *et al.*, “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature*, vol. 551, no. 7682, pp. 579–584, 2017.
- [17] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, *et al.*, “Quantum phases of matter on a 256-atom programmable quantum simulator,” *arXiv preprint arXiv:2012.12281*, 2020.
- [18] T. S. Cubitt, A. Montanaro, and S. Piddock, “Universal quantum

- hamiltonians," *Proceedings of the National Academy of Sciences*, vol. 115, no. 38, pp. 9497–9502, 2018.
- [19] T. Kohler, S. Piddock, J. Bausch, and T. Cubitt, "Translationally-invariant universal quantum hamiltonians in 1d," *arXiv preprint arXiv:2003.13753*, 2020.
- [20] P. M. Poggi, N. K. Lysne, K. W. Kuper, I. H. Deutsch, and P. S. Jessen, "Quantifying the sensitivity to errors in analog quantum simulation," *PRX Quantum*, vol. 1, p. 020308, Nov 2020.
- [21] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.
- [22] K. R. Brown, "Energy protection arguments fail in the interaction picture," *Phys. Rev. A*, vol. 76, p. 022327, Aug 2007.
- [23] S. Lloyd, "Universal quantum simulators," *Science*, vol. 273, no. 5278, pp. 1073–1078, 1996.
- [24] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, "Simulated quantum computation of molecular energies," *Science*, vol. 309, no. 5741, pp. 1704–1707, 2005.
- [25] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, "Improved Fault-Tolerant Quantum Simulation of Condensed-Phase Correlated Electrons via Trotterization," *Quantum*, vol. 4, p. 296, July 2020.
- [26] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, "Encoding electronic spectra in quantum circuits with linear t complexity," *Phys. Rev. X*, vol. 8, p. 041015, Oct 2018.

- [27] J. Lee, D. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, “Even more efficient quantum computations of chemistry through tensor hypercontraction,” *arXiv preprint arXiv:2011.03494*, 2020.
- [28] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, and J. I. Cirac, “Analogue quantum chemistry simulation,” *Nature*, vol. 574, no. 7777, pp. 215–218, 2019.
- [29] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Physical review letters*, vol. 119, no. 18, p. 180509, 2017.
- [30] S. Endo, S. C. Benjamin, and Y. Li, “Practical quantum error mitigation for near-future applications,” *Physical Review X*, vol. 8, no. 3, p. 031027, 2018.
- [31] T. Jones, A. Brown, I. Bush, and S. C. Benjamin, “Quest and high performance simulation of quantum computers,” *Scientific reports*, vol. 9, no. 1, pp. 1–11, 2019.
- [32] Google, “Cirq.” <https://github.com/quantumlib/Cirq>, 2020.
- [33] D. S. Steiger, T. Häner, and M. Troyer, “ProjectQ: an open source software framework for quantum computing,” *Quantum*, vol. 2, p. 49, Jan. 2018.
- [34] J. McClean, N. Rubin, K. Sung, I. D. Kivlichan, X. Bonet-Monroig, Y. Cao, C. Dai, E. S. Fried, C. Gidney, B. Gimby, *et al.*, “Openfermion: the electronic structure package for quantum computers,” *Quantum Science and Technology*, 2020.
- [35] M. A. Nielsen and I. Chuang, “Quantum computation and quantum information,” 2002.

- [36] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, "Quantum computational chemistry," *Rev. Mod. Phys.*, vol. 92, p. 015003, Mar 2020.
- [37] H. Eyring, "The activated complex in chemical reactions," *The Journal of Chemical Physics*, vol. 3, no. 2, pp. 107–115, 1935.
- [38] M. G. Evans and M. Polanyi, "Some applications of the transition state method to the calculation of reaction velocities, especially in solution," *Trans. Faraday Soc.*, vol. 31, pp. 875–894, 1935.
- [39] T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular electronic-structure theory*. John Wiley & Sons, 2014.
- [40] A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.
- [41] C. Hattig, W. Klopper, A. Köhn, and D. P. Tew, "Explicitly correlated electrons in molecules," *Chemical Reviews*, vol. 112, no. 1, pp. 4–74, 2012. PMID: 22206503.
- [42] K. D. Vogiatzis, M. V. Polynski, J. K. Kirkland, J. Townsend, A. Hashemi, C. Liu, and E. A. Pidko, "Computational approach to molecular catalysis by 3d transition metals: Challenges and opportunities," *Chemical Reviews*, vol. 119, no. 4, pp. 2453–2523, 2019.
- [43] G. D. Purvis and R. J. Bartlett, "A full coupled-cluster singles and doubles model: The inclusion of disconnected triples," *The Journal of Chemical Physics*, vol. 76, no. 4, pp. 1910–1918, 1982.
- [44] M. Schutz and H.-J. Werner, "Local perturbative triples correction (t) with linear cost scaling," *Chemical Physics Letters*, vol. 318, no. 4, pp. 370 – 378, 2000.
- [45] M. Schutz, "Low-order scaling local electron correlation methods. iii.

- linear scaling local perturbative triples correction (t)," *The Journal of Chemical Physics*, vol. 113, no. 22, pp. 9986–10001, 2000.
- [46] H.-J. Werner and M. Schutz, "An efficient local coupled cluster method for accurate thermochemistry of large systems," *The Journal of Chemical Physics*, vol. 135, no. 14, p. 144116, 2011.
- [47] R. J. Bartlett and M. Musiał, "Coupled-cluster theory in quantum chemistry," *Rev. Mod. Phys.*, vol. 79, pp. 291–352, Feb 2007.
- [48] P. Hobza and J. Spöner, "Toward true dna base-stacking energies:mp2, ccSD(t), and complete basis set calculations," *Journal of the American Chemical Society*, vol. 124, no. 39, pp. 11802–11808, 2002. PMID: 12296748.
- [49] T. Yamazaki, S. Matsuura, A. Narimani, A. Saidmuradov, and A. Zaribafiyani, "Towards the practical application of near-term quantum computers in quantum chemistry simulations: A problem decomposition approach," *arXiv:1806.01305*, 2018.
- [50] D. I. Lyakh, M. Musiał, V. F. Lotrich, and R. J. Bartlett, "Multireference nature of chemistry: The coupled-cluster view," *Chemical Reviews*, vol. 112, no. 1, pp. 182–243, 2012. PMID: 22220988.
- [51] M. A. Watson and G. K.-L. Chan, "Excited states of butadiene to chemical accuracy: Reconciling theory and experiment," *Journal of Chemical Theory and Computation*, vol. 8, no. 11, pp. 4013–4018, 2012. PMID: 26605568.
- [52] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, and E. Gull, "Solutions of

- the two-dimensional hubbard model: Benchmarks and results from a wide range of numerical algorithms," *Phys. Rev. X*, vol. 5, p. 041041, Dec 2015.
- [53] B. O. Roos, P. R. Taylor, and P. E. Siagbahn, "A complete active space scf method (casscf) using a density matrix formulated super-ci approach," *Chemical Physics*, vol. 48, no. 2, pp. 157 – 173, 1980.
- [54] H. Lischka, D. Nachtigallova, A. J. A. Aquino, P. G. Szalay, F. Plasser, F. B. C. Machado, and M. Barbatti, "Multireference approaches for excited states of molecules," *Chemical Reviews*, vol. 118, no. 15, pp. 7293–7361, 2018. PMID: 30040389.
- [55] U. Schollwöck, "The density-matrix renormalization group in the age of matrix product states," *Annals of physics*, vol. 326, no. 1, pp. 96–192, 2011.
- [56] S. Sharma, K. Sivalingam, F. Neese, and G. K.-L. Chan, "Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics," *Nature chemistry*, vol. 6, no. 10, p. 927, 2014.
- [57] Y. Kurashige, G. K.-L. Chan, and T. Yanai, "Entangled quantum electronic wavefunctions of the mn 4 cao 5 cluster in photosystem ii," *Nature chemistry*, vol. 5, no. 8, p. 660, 2013.
- [58] R. Olivares-Amaya, W. Hu, N. Nakatani, S. Sharma, J. Yang, and G. K.-L. Chan, "The ab-initio density matrix renormalization group in practice," *The Journal of Chemical Physics*, vol. 142, no. 3, p. 034102, 2015.
- [59] S. Szalay, M. Pfeiffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider, and A. Legeza, "Tensor product methods and entanglement optimization for ab initio quantum chemistry," *International Journal of Quantum Chemistry*, vol. 115, no. 19, pp. 1342–1391, 2015.

- [60] G. H. Booth, A. J. W. Thom, and A. Alavi, "Fermion monte carlo without fixed nodes: A game of life, death, and annihilation in slater determinant space," *The Journal of Chemical Physics*, vol. 131, no. 5, p. 054106, 2009.
- [61] N. M. Tubman, J. Lee, T. Y. Takeshita, M. Head-Gordon, and K. B. Whaley, "A deterministic alternative to the full configuration interaction quantum monte carlo method," *The Journal of Chemical Physics*, vol. 145, no. 4, p. 044112, 2016.
- [62] G. H. Booth and A. Alavi, "Approaching chemical accuracy using full configuration-interaction quantum monte carlo: A study of ionization potentials," *The Journal of Chemical Physics*, vol. 132, no. 17, p. 174104, 2010.
- [63] B. M. Austin, D. Y. Zubarev, and W. A. Lester, "Quantum monte carlo and related approaches," *Chemical Reviews*, vol. 112, no. 1, pp. 263–288, 2012. PMID: 22196085.
- [64] J. Spencer, N. Blunt, and W. Foulkes, "The sign problem and population dynamics in the full configuration interaction quantum monte carlo method," *The Journal of chemical physics*, vol. 136, no. 5, p. 054110, 2012.
- [65] J. J. Shepherd, G. E. Scuseria, and J. S. Spencer, "Sign problem in full configuration interaction quantum monte carlo: Linear and sublinear representation regimes for the exact wave function," *Phys. Rev. B*, vol. 90, p. 155130, Oct 2014.
- [66] K. Ghanem, A. Y. Lozovoi, and A. Alavi, "Unbiasing the initiator approximation in full configuration interaction quantum monte carlo," *The Journal of Chemical Physics*, vol. 151, no. 22, p. 224108, 2019.
- [67] M. Podewitz, M. T. Stiebritz, and M. Reiher, "An enquiry into the-

- oretical bioinorganic chemistry: How heuristic is the character of present-day quantum chemical methods?," *Faraday Discuss.*, vol. 148, pp. 119–135, 2011.
- [68] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, "Elucidating reaction mechanisms on quantum computers," *Proceedings of the National Academy of Sciences*, vol. 114, no. 29, pp. 7555–7560, 2017.
- [69] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, "Qubitization of Arbitrary Basis Quantum Chemistry Leveraging Sparsity and Low Rank Factorization," *Quantum*, vol. 3, p. 208, Dec. 2019.
- [70] B. K. Burgess and D. J. Lowe, "Mechanism of molybdenum nitrogenase," *Chemical Reviews*, vol. 96, no. 7, pp. 2983–3012, 1996. PMID: 11848849.
- [71] B. M. Hoffman, D. Lukoyanov, Z.-Y. Yang, D. R. Dean, and L. C. Seefeldt, "Mechanism of nitrogen fixation by nitrogenase: The next stage," *Chemical Reviews*, vol. 114, no. 8, pp. 4041–4062, 2014. PMID: 24467365.
- [72] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, "Quantum computing enhanced computational catalysis," *arXiv preprint arXiv:2007.14460*, 2020.
- [73] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov, "How will quantum computers provide an industrially relevant computational advantage in quantum chemistry?," *arXiv preprint arXiv:2009.12472*, 2020.
- [74] E. Dagotto, "Correlated electrons in high-temperature superconductors," *Rev. Mod. Phys.*, vol. 66, pp. 763–840, Jul 1994.
- [75] J. Hubbard, "Electron correlations in narrow energy bands," *Proceed-*

- ings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 276, no. 1365, pp. 238–257, 1963.
- [76] P. W. Anderson, “Superconductivity in high  $t_c$  cuprates: The cause is no longer a mystery,” *Physica Scripta*, vol. T102, no. 1, p. 10, 2002.
- [77] P. A. Lee, N. Nagaosa, and X.-G. Wen, “Doping a mott insulator: Physics of high-temperature superconductivity,” *Rev. Mod. Phys.*, vol. 78, pp. 17–85, Jan 2006.
- [78] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, “Colloquium: Theory of intertwined orders in high temperature superconductors,” *Rev. Mod. Phys.*, vol. 87, pp. 457–482, May 2015.
- [79] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, “Solving strongly correlated electron models on a quantum computer,” *Phys. Rev. A*, vol. 92, p. 062318, Dec 2015.
- [80] Z. Cai, “Resource estimation for quantum variational simulations of the hubbard model,” *Phys. Rev. Applied*, vol. 14, p. 014059, Jul 2020.
- [81] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, “Strategies for solving the fermi-hubbard model on near-term quantum computers,” *Phys. Rev. B*, vol. 102, p. 235122, Dec 2020.
- [82] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, “Quantum Simulation of Electronic Structure with Linear Depth and Connectivity,” *Phys. Rev. Lett.*, vol. 120, p. 110501, Mar 2018.
- [83] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, “Quantum algorithms to simulate many-body physics of correlated fermions,” *Phys. Rev. Applied*, vol. 9, p. 044036, Apr 2018.
- [84] J.-M. Reiner, F. Wilhelm-Mauch, G. Schön, and M. Marthaler, “Find-

- ing the ground state of the hubbard model by variational methods on a quantum computer with gate errors," *Quantum Science and Technology*, vol. 4, p. 035005, May 2019.
- [85] P. Jordan and E. Wigner, "Über das paulische äquivalenzverbot," *Zeitschrift für Physik*, vol. 47, pp. 631–651, Sep 1928.
- [86] S. B. Bravyi and A. Y. Kitaev, "Fermionic quantum computation," *Annals of Physics*, vol. 298, no. 1, pp. 210 – 226, 2002.
- [87] J. T. Seeley, M. J. Richard, and P. J. Love, "The bravyi-kitaev transformation for quantum computation of electronic structure," *The Journal of Chemical Physics*, vol. 137, no. 22, p. 224109, 2012.
- [88] A. Tranter, S. Sofia, J. Seeley, M. Kaicher, J. McClean, R. Babbush, P. V. Coveney, F. Mintert, F. Wilhelm, and P. J. Love, "The bravyi-kitaev transformation: Properties and applications," *International Journal of Quantum Chemistry*, vol. 115, no. 19, pp. 1431–1441, 2015.
- [89] V. Havlicek, M. Troyer, and J. D. Whitfield, "Operator locality in the quantum simulation of fermionic models," *Phys. Rev. A*, vol. 95, p. 032332, Mar 2017.
- [90] A. Tranter, P. J. Love, F. Mintert, and P. V. Coveney, "A comparison of the bravyi-kitaev and jordan-wigner transformations for the quantum simulation of quantum chemistry," *Journal of Chemical Theory and Computation*, vol. 14, no. 11, pp. 5617–5630, 2018. PMID: 30189144.
- [91] P. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, *et al.*, "Scalable quantum simulation of molecular energies," *Physical Review X*, vol. 6, no. 3, p. 031007, 2016.
- [92] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M.

- Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [93] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, "Tapering off qubits to simulate fermionic hamiltonians," *arXiv preprint arXiv:1701.08213*, 2017.
- [94] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. Oâbrien, "A variational eigenvalue solver on a photonic quantum processor," *Nature communications*, vol. 5, p. 4213, 2014.
- [95] D. S. Abrams and S. Lloyd, "Simulation of many-body fermi systems on a universal quantum computer," *Phys. Rev. Lett.*, vol. 79, pp. 2586–2589, Sep 1997.
- [96] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, "A theory of trotter error," 2019.
- [97] A. M. Childs, A. Ostrander, and Y. Su, "Faster quantum simulation by randomization," *Quantum*, vol. 3, p. 182, Sept. 2019.
- [98] D. W. Berry, A. M. Childs, and R. Kothari, "Hamiltonian simulation with nearly optimal dependence on all parameters," in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 792–809, Oct 2015.
- [99] D. W. Berry and A. M. Childs, "Black-box hamiltonian simulation and unitary implementation," *Quantum Info. Comput.*, vol. 12, pp. 29–62, Jan. 2012.
- [100] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, "Simulating hamiltonian dynamics with a truncated taylor series," *Phys. Rev. Lett.*, vol. 114, p. 090502, Mar 2015.

- [101] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, "Exponential improvement in precision for simulating sparse hamiltonians," in *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*, STOC '14, (New York, NY, USA), pp. 283–292, ACM, 2014.
- [102] G. H. Low and I. L. Chuang, "Hamiltonian Simulation by Qubitization," *Quantum*, vol. 3, p. 163, July 2019.
- [103] G. H. Low, T. J. Yoder, and I. L. Chuang, "Methodology of resonant equiangular composite quantum gates," *Phys. Rev. X*, vol. 6, p. 041067, Dec 2016.
- [104] G. H. Low and I. L. Chuang, "Optimal hamiltonian simulation by quantum signal processing," *Phys. Rev. Lett.*, vol. 118, p. 010501, Jan 2017.
- [105] E. Campbell, "Random compiler for fast hamiltonian simulation," *Phys. Rev. Lett.*, vol. 123, p. 070503, Aug 2019.
- [106] Y. Ouyang, D. R. White, and E. T. Campbell, "Compilation by stochastic Hamiltonian sparsification," *Quantum*, vol. 4, p. 235, Feb. 2020.
- [107] A. Y. Kitaev, "Quantum measurements and the abelian stabilizer problem," *Preprint at <http://arxiv.org/abs/quant-ph/9511026>*, 1995.
- [108] D. S. Abrams and S. Lloyd, "Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors," *Phys. Rev. Lett.*, vol. 83, pp. 5162–5165, Dec 1999.
- [109] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, *et al.*, "Towards quantum chemistry on a quantum computer," *Nature chemistry*, vol. 2, no. 2, p. 106, 2010.

- [110] Y. Wang, F. Dolde, J. Biamonte, R. Babbush, V. Bergholm, S. Yang, I. Jakobi, P. Neumann, A. Aspuru-Guzik, J. D. Whitfield, *et al.*, “Quantum simulation of helium hydride cation in a solid-state spin register,” *ACS nano*, vol. 9, no. 8, pp. 7769–7774, 2015.
- [111] R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O’Brien, and M. G. Thompson, “Witnessing eigenstates for quantum simulation of hamiltonian spectra,” *Science Advances*, vol. 4, no. 1, 2018.
- [112] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O’Brien, and M. G. Thompson, “Experimental bayesian quantum phase estimation on a silicon photonic chip,” *Phys. Rev. Lett.*, vol. 118, p. 100503, Mar 2017.
- [113] J. Du, N. Xu, X. Peng, P. Wang, S. Wu, and D. Lu, “Nmr implementation of a molecular hydrogen quantum simulation with adiabatic state preparation,” *Phys. Rev. Lett.*, vol. 104, p. 030502, Jan 2010.
- [114] Z. Li, M.-H. Yung, H. Chen, D. Lu, J. D. Whitfield, X. Peng, A. Aspuru-Guzik, and J. Du, “Solving quantum ground-state problems with nuclear magnetic resonance,” *Scientific reports*, vol. 1, p. 88, 2011.
- [115] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, “Gate-count estimates for performing quantum chemistry on small quantum computers,” *Phys. Rev. A*, vol. 90, p. 022305, Aug 2014.
- [116] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, “Chemical basis of trotter-suzuki errors in quantum chemistry simulation,” *Phys. Rev. A*, vol. 91, p. 022311, Feb 2015.
- [117] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and

- A. Aspuru-Guzik, "Exponentially more precise quantum simulation of fermions in second quantization," *New Journal of Physics*, vol. 18, no. 3, p. 033032, 2016.
- [118] R. Babbush, D. W. Berry, Y. R. Sanders, I. D. Kivlichan, A. Scherer, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, "Exponentially more precise quantum simulation of fermions in the configuration interaction representation," *Quantum Science and Technology*, vol. 3, no. 1, p. 015006, 2017.
- [119] R. Meister, S. C. Benjamin, and E. T. Campbell, "Tailoring term truncations for electronic structure calculations using a linear combination of unitaries," *arXiv preprint arXiv:2007.11624*, 2020.
- [120] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, and R. Babbush, "Improved techniques for preparing eigenstates of fermionic hamiltonians," *npj Quantum Information*, vol. 4, no. 1, p. 22, 2018.
- [121] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, "Quantum algorithm for spectral measurement with a lower gate count," *Phys. Rev. Lett.*, vol. 121, p. 010501, Jul 2018.
- [122] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New Journal of Physics*, vol. 18, no. 2, p. 023023, 2016.
- [123] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. J. Love, and A. Aspuru-Guzik, "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz," *Quantum Science and Technology*, vol. 4, no. 1, p. 014008, 2018.
- [124] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya, *et al.*,

- “Quantum chemistry in the age of quantum computing,” *Chemical reviews*, vol. 119, no. 19, pp. 10856–10915, 2019.
- [125] J. R. McClean, R. Babbush, P. J. Love, and A. Aspuru-Guzik, “Exploiting locality in quantum computation for quantum chemistry,” *The Journal of Physical Chemistry Letters*, vol. 5, no. 24, pp. 4368–4380, 2014. PMID: 26273989.
- [126] K. J. Sung, J. Yao, M. P. Harrigan, N. C. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, “Using models to improve optimizers for variational quantum algorithms,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044008, 2020.
- [127] D. Wierichs, C. Gogolin, and M. Kastoryano, “Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer,” *Phys. Rev. Research*, vol. 2, p. 043246, Nov 2020.
- [128] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature communications*, vol. 9, no. 1, pp. 1–6, 2018.
- [129] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, “Variational quantum algorithms,” *arXiv preprint arXiv:2012.09265*, 2020.
- [130] F. G. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, “For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances,” *arXiv preprint arXiv:1812.04170*, 2018.
- [131] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, “An initialization strategy for addressing barren plateaus in parametrized quantum circuits,” *Quantum*, vol. 3, p. 214, 2019.
- [132] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib,

- “Layerwise learning for quantum neural networks,” *Quantum Machine Intelligence*, vol. 3, no. 1, pp. 1–11, 2021.
- [133] T. Volkoff and P. J. Coles, “Large gradients via correlation in random parameterized quantum circuits,” *Quantum Science and Technology*, vol. 6, no. 2, p. 025008, 2021.
- [134] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, “Self-verifying variational quantum simulation of lattice models,” *Nature*, vol. 569, no. 7756, pp. 355–360, 2019.
- [135] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, *et al.*, “Quantum chemistry calculations on a trapped-ion quantum simulator,” *Physical Review X*, vol. 8, no. 3, p. 031022, 2018.
- [136] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, “Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure,” *Phys. Rev. A*, vol. 95, p. 020501, Feb 2017.
- [137] Y. Nam, J.-S. Chen, N. C. Pienti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, *et al.*, “Ground-state energy estimation of the water molecule on a trapped-ion quantum computer,” *npj Quantum Information*, vol. 6, no. 1, pp. 1–6, 2020.
- [138] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, “Error mitigation extends the computational reach of a noisy quantum processor,” *Nature*, vol. 567, no. 7749, pp. 491–495, 2019.
- [139] , F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett,

- N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. J. Huggins, L. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, H. Neven, M. Y. Niu, T. E. O'Brien, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, D. Strain, K. J. Sung, M. Szalay, T. Y. Takeshita, A. Vainsencher, T. White, N. Wiebe, Z. J. Yao, P. Yeh, and A. Zalcman, "Hartree-fock on a superconducting qubit quantum computer," *Science*, vol. 369, no. 6507, pp. 1084–1089, 2020.
- [140] D. Wecker, M. B. Hastings, and M. Troyer, "Progress towards practical quantum variational algorithms," *Phys. Rev. A*, vol. 92, p. 042303, Oct 2015.
- [141] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [142] M. R. Hoffmann and J. Simons, "A unitary multiconfigurational coupled-cluster method: Theory and applications," *The Journal of Chemical Physics*, vol. 88, no. 2, pp. 993–1002, 1988.
- [143] R. J. Bartlett, S. A. Kucharski, and J. Noga, "Alternative coupled-cluster ansätze ii. the unitary coupled-cluster method," *Chemical Physics Letters*, vol. 155, no. 1, pp. 133 – 140, 1989.
- [144] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, "Quantum algorithms for electronic structure calcula-

- tions: Particle-hole hamiltonian and optimized wave-function expansions," *Phys. Rev. A*, vol. 98, p. 022322, Aug 2018.
- [145] H. R. Grimsley, D. Claudino, S. E. Economou, E. Barnes, and N. J. Mayhall, "Is the trotterized uccsd ansatz chemically well-defined?," *Journal of Chemical Theory and Computation*, vol. 16, no. 1, pp. 1–6, 2019.
- [146] M. Y. Niu, V. Smelyanskyi, P. Klimov, S. Boixo, R. Barends, J. Kelly, Y. Chen, K. Arya, B. Burkett, D. Bacon, *et al.*, "Learning non-markovian quantum noise from moire-enhanced swap spectroscopy with deep evolutionary algorithm," *arXiv preprint arXiv:1912.04368*, 2019.
- [147] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [148] IBM, "Quantum takes flight: Moving from laboratory demonstrations to building systems." <https://qiskit.org/>, 2018.
- [149] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, "High-fidelity quantum logic gates using trapped-ion hyperfine qubits," *Phys. Rev. Lett.*, vol. 117, p. 060504, Aug 2016.
- [150] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, *et al.*, "High-fidelity universal gate set for be 9+ ion qubits," *Physical review letters*, vol. 117, no. 6, p. 060505, 2016.
- [151] Z. Cai, "Multi-exponential error extrapolation and combining error mitigation techniques for nisq applications," *arXiv preprint arXiv:2007.01265*, 2020.
- [152] Y. Li and S. C. Benjamin, "Efficient variational quantum simulator in-

- corporating active error minimization," *Phys. Rev. X*, vol. 7, p. 021050, Jun 2017.
- [153] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, "Digital zero noise extrapolation for quantum error mitigation," 2020.
- [154] O. Shehab, K. Landsman, Y. Nam, D. Zhu, N. M. Linke, M. Keesan, R. C. Pooser, and C. Monroe, "Toward convergence of effective-field-theory simulations on digital quantum computers," *Phys. Rev. A*, vol. 100, p. 062319, Dec 2019.
- [155] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, "Learning-based quantum error mitigation," 2020.
- [156] M. Huo and Y. Li, "Self-consistent tomography of temporally correlated errors," *arXiv*, pp. arXiv-1811, 2018.
- [157] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, "Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states," *Phys. Rev. A*, vol. 95, p. 042308, Apr 2017.
- [158] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, "Computation of molecular spectra on a quantum processor with an error-resilient algorithm," *Phys. Rev. X*, vol. 8, p. 011021, Feb 2018.
- [159] S. McArdle, X. Yuan, and S. Benjamin, "Error-mitigated digital quantum simulation," *Phys. Rev. Lett.*, vol. 122, p. 180501, May 2019.
- [160] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, "Low-cost error mitigation by symmetry verification," *Phys. Rev. A*, vol. 98, p. 062339, Dec 2018.
- [161] R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. Bultink,

- X. Fu, C. Price, V. Ostroukh, N. Muthusubramanian, A. Bruno, *et al.*, “Experimental error mitigation via symmetry verification in a variational quantum eigensolver,” *Physical Review A*, vol. 100, no. 1, p. 010302, 2019.
- [162] N. C. Rubin, R. Babbush, and J. McClean, “Application of fermionic marginal constraints to hybrid quantum algorithms,” *New Journal of Physics*, vol. 20, no. 5, p. 053020, 2018.
- [163] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, “Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers,” *arXiv preprint arXiv:1907.13117*, 2019.
- [164] B. M. Terhal, “Quantum error correction for quantum memories,” *Rev. Mod. Phys.*, vol. 87, pp. 307–346, Apr 2015.
- [165] S. J. Devitt, W. J. Munro, and K. Nemoto, “Quantum error correction for beginners,” *Reports on Progress in Physics*, vol. 76, no. 7, p. 076001, 2013.
- [166] R. Raussendorf, “Key ideas in quantum error correction,” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1975, pp. 4541–4565, 2012.
- [167] D. A. Lidar and T. A. Brun, *Quantum Error Correction*. Cambridge University Press, 2013.
- [168] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error,” in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pp. 176–188, ACM, 1997.
- [169] D. Gottesman, “Theory of fault-tolerant quantum computation,” *Physical Review A*, vol. 57, no. 1, p. 127, 1998.

- [170] P. W. Shor, "Fault-tolerant quantum computation," in *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on*, pp. 56–65, IEEE, 1996.
- [171] E. Knill, R. Laflamme, and W. Zurek, "Accuracy threshold for quantum computation," 1996.
- [172] A. M. Steane, "Error correcting codes in quantum theory," *Physical Review Letters*, vol. 77, no. 5, p. 793, 1996.
- [173] B. Eastin and E. Knill, "Restrictions on transversal encoded quantum gate sets," *Phys. Rev. Lett.*, vol. 102, p. 110502, Mar 2009.
- [174] E. T. Campbell, B. M. Terhal, and C. Vuillot, "Roads towards fault-tolerant universal quantum computation.," *Nature*, vol. 549, no. 7671, p. 172, 2017.
- [175] P. Webster, M. Vasmer, T. R. Scruby, and S. D. Bartlett, "Universal fault-tolerant quantum computing with stabiliser codes," *arXiv preprint arXiv:2012.05260*, 2020.
- [176] A. Y. Kitaev, "Quantum computations: algorithms and error correction," *Russian Mathematical Surveys*, vol. 52, no. 6, pp. 1191–1249, 1997.
- [177] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, "Surface code quantum computing with error rates over 1%," *Phys. Rev. A*, vol. 83, p. 020302, Feb 2011.
- [178] A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg, "Towards practical classical processing for the surface code," *Phys. Rev. Lett.*, vol. 108, p. 180501, May 2012.
- [179] A. M. Stephens, "Fault-tolerant thresholds for quantum error correction with the surface code," *Phys. Rev. A*, vol. 89, p. 022321, Feb 2014.

- [180] B. J. Brown, K. Laubscher, M. S. Kesselring, and J. R. Wootton, "Poking holes and cutting corners to achieve clifford gates with the surface code," *Physical Review X*, vol. 7, no. 2, p. 021029, 2017.
- [181] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, "Surface code quantum computing by lattice surgery," *New Journal of Physics*, vol. 14, no. 12, p. 123011, 2012.
- [182] S. Bravyi and A. Kitaev, "Universal quantum computation with ideal clifford gates and noisy ancillas," *Phys. Rev. A*, vol. 71, p. 022316, Feb 2005.
- [183] N. J. Ross and P. Selinger, "Optimal ancilla-free clifford+t approximation of z-rotations," *Quantum Info. Comput.*, vol. 16, p. 901â953, Sept. 2016.
- [184] S. Pallister, "A jordan-wigner gadget that reduces t count by more than 6x for quantum chemistry applications," 2020.
- [185] D. Litinski, "A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery," *Quantum*, vol. 3, p. 128, Mar. 2019.
- [186] D. Litinski, "Magic State Distillation: Not as Costly as You Think," *Quantum*, vol. 3, p. 205, Dec. 2019.
- [187] H. Bombín, "Single-shot fault-tolerant quantum error correction," *Phys. Rev. X*, vol. 5, p. 031043, Sep 2015.
- [188] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, p. 032324, Sep 2012.
- [189] J. O’Gorman and E. T. Campbell, "Quantum computation with realistic magic-state factories," *Phys. Rev. A*, vol. 95, p. 032338, Mar 2017.

- [190] C. Gidney and M. Ekerdt, “How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits,” 2019.
- [191] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, “Variational ansatz-based quantum simulation of imaginary time evolution,” *npj Quantum Information*, vol. 5, Sept. 2019.
- [192] S. McArdle and D. P. Tew, “Improving the accuracy of quantum computational chemistry using the transcorrelated method,” 2020.
- [193] L. Hackl, T. Guaita, T. Shi, J. Haegeman, E. Demler, and I. Cirac, “Geometry of variational methods: dynamics of closed quantum systems,” 2020.
- [194] G. Vidal, “Efficient simulation of one-dimensional quantum many-body systems,” *Phys. Rev. Lett.*, vol. 93, p. 040502, Jul 2004.
- [195] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, “Theory of variational quantum simulation,” *Quantum*, vol. 3, p. 191, Oct. 2019.
- [196] A. McLachlan, “A variational solution of the time-dependent schrodinger equation,” *Molecular Physics*, vol. 8, no. 1, pp. 39–44, 1964.
- [197] R. Cheng, “Quantum geometric tensor (fubini-study metric) in simple quantum system: A pedagogical introduction,” 2010.
- [198] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum Natural Gradient,” *Quantum*, vol. 4, p. 269, May 2020.
- [199] K. Mitarai and K. Fujii, “Methodology for replacing indirect measurements with direct measurements,” *Phys. Rev. Research*, vol. 1, p. 013006, Aug 2019.
- [200] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, “Evalu-

- ating analytic gradients on quantum hardware," *Phys. Rev. A*, vol. 99, p. 032331, Mar 2019.
- [201] B. van Straaten and B. Koczor, "Measurement cost of metric-aware variational quantum algorithms," 2020.
- [202] B. Koczor and S. C. Benjamin, "Quantum natural gradient generalised to non-unitary circuits," 2019.
- [203] E. v. d. Berg, Z. K. Mineev, and K. Temme, "Model-free readout-error mitigation for quantum expectation values," *arXiv preprint arXiv:2012.09738*, 2020.
- [204] T. Kato, "On the eigenfunctions of many-particle systems in quantum mechanics," *Communications on Pure and Applied Mathematics*, vol. 10, no. 2, pp. 151–177, 1957.
- [205] M. Nooijen and R. J. Bartlett, "Elimination of coulombic infinities through transformation of the hamiltonian," *The Journal of Chemical Physics*, vol. 109, no. 19, pp. 8232–8240, 1998.
- [206] A. Grüneis, S. Hirata, Y.-y. Ohnishi, and S. Ten-no, "Perspective: Explicitly correlated electronic structure theory for complex systems," *The Journal of Chemical Physics*, vol. 146, no. 8, p. 080901, 2017.
- [207] L. Kong, F. A. Bischoff, and E. F. Valeev, "Explicitly correlated r12/f12 methods for electronic structure," *Chemical Reviews*, vol. 112, pp. 75–107, Jan 2012.
- [208] E. A. Hylleraas, "Neue berechnung der energie des heliums im grundzustande, sowie des tiefsten terms von ortho-helium," *Zeitschrift für Physik*, vol. 54, pp. 347–366, May 1929.
- [209] J. C. Slater, "Central fields and rydberg formulas in wave mechanics," *Phys. Rev.*, vol. 31, pp. 333–343, Mar 1928.

- [210] W. Kutzelnigg, "r<sup>12</sup>-dependent terms in the wave function as closed sums of partial wave amplitudes for large l," *Theoretica chimica acta*, vol. 68, pp. 445–469, Dec 1985.
- [211] S. F. Boys, N. C. Handy, and J. W. Linnett, "A condition to remove the indeterminacy in interelectronic correlation functions," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 309, no. 1497, pp. 209–220, 1969.
- [212] S. F. Boys, N. C. Handy, and J. W. Linnett, "The determination of energies and wavefunctions with full electronic correlation," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 310, no. 1500, pp. 43–61, 1969.
- [213] J. O. Hirschfelder, "Removal of electron—electron poles from many-electron hamiltonians," *The Journal of Chemical Physics*, vol. 39, pp. 3145–3146, Dec. 1963.
- [214] A. J. Cohen, H. Luo, K. Guthrie, W. Dobrautz, D. P. Tew, and A. Alavi, "Similarity transformation of the electronic schrödinger equation via jastrow factorization," *The Journal of Chemical Physics*, vol. 151, no. 6, p. 061101, 2019.
- [215] S. Tsuneyuki, "Transcorrelated Method: Another Possible Way towards Electronic Structure Calculation of Solids," *Progress of Theoretical Physics Supplement*, vol. 176, pp. 134–142, 06 2008.
- [216] W. Dobrautz, H. Luo, and A. Alavi, "Compact numerical solutions to the two-dimensional repulsive hubbard model obtained via nonunitary similarity transformations," *Phys. Rev. B*, vol. 99, p. 075119, Feb 2019.
- [217] N. Handy, "On the minimization of the variance of the transcor-

- related hamiltonian," *Molecular Physics*, vol. 21, no. 5, pp. 817–828, 1971.
- [218] S. Ten-no, "A feasible transcorrelated method for treating electronic cusps using a frozen gaussian geminal," *Chemical Physics Letters*, vol. 330, no. 1, pp. 169 – 174, 2000.
- [219] O. Hino, Y. Tanimura, and S. Ten-no, "Application of the transcorrelated hamiltonian to the linearized coupled cluster singles and doubles model," *Chemical Physics Letters*, vol. 353, no. 3, pp. 317 – 323, 2002.
- [220] N. Umezawa and S. Tsuneyuki, "Transcorrelated method for electronic systems coupled with variational monte carlo calculation," *The Journal of Chemical Physics*, vol. 119, no. 19, pp. 10015–10031, 2003.
- [221] N. Umezawa and S. Tsuneyuki, "Excited electronic state calculations by the transcorrelated variational monte carlo method: Application to a helium atom," *The Journal of Chemical Physics*, vol. 121, no. 15, pp. 7070–7075, 2004.
- [222] N. Umezawa, S. Tsuneyuki, T. Ohno, K. Shiraishi, and T. Chikyow, "A practical treatment for the three-body interactions in the transcorrelated variational monte carlo method: Application to atoms from lithium to neon," *The Journal of Chemical Physics*, vol. 122, no. 22, p. 224101, 2005.
- [223] M. Ochi, K. Sodeyama, R. Sakuma, and S. Tsuneyuki, "Efficient algorithm of the transcorrelated method for periodic systems," *The Journal of Chemical Physics*, vol. 136, no. 9, p. 094108, 2012.
- [224] H. Luo, "Variational transcorrelated method," *The Journal of Chemical Physics*, vol. 133, no. 15, p. 154109, 2010.
- [225] H. Luo, "Complete optimisation of multi-configuration jastrow wave

- functions by variational transcorrelated method," *The Journal of Chemical Physics*, vol. 135, no. 2, p. 024109, 2011.
- [226] H. Luo and A. Alavi, "Combining the transcorrelated method with full configuration interaction quantum monte carlo: Application to the homogeneous electron gas," *Journal of Chemical Theory and Computation*, vol. 14, no. 3, pp. 1403–1411, 2018. PMID: 29431996.
- [227] P. Jeszenszki, H. Luo, A. Alavi, and J. Brand, "Accelerating the convergence of exact diagonalization with the transcorrelated method: Quantum gas in one dimension with contact interactions," *Phys. Rev. A*, vol. 98, p. 053627, Nov 2018.
- [228] P. Jeszenszki, U. Ebling, H. Luo, A. Alavi, and J. Brand, "Eliminating the wave function singularity for ultracold atoms by similarity transformation," 2020.
- [229] M. Motta, T. P. Gujarati, J. E. Rice, A. Kumar, C. Masteran, J. A. Latone, E. Lee, E. F. Valeev, and T. Y. Takeshita, "Quantum simulation of electronic structure with transcorrelated hamiltonian: increasing accuracy without extra quantum resources," 2020.
- [230] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, "Variational quantum algorithms for discovering hamiltonian spectra," *Phys. Rev. A*, vol. 99, p. 062304, Jun 2019.
- [231] S. Endo, J. Sun, Y. Li, S. C. Benjamin, and X. Yuan, "Variational quantum simulation of general processes," *Phys. Rev. Lett.*, vol. 125, p. 010501, Jun 2020.
- [232] T. Jones and S. C. Benjamin, "Quantum compilation and circuit optimisation via energy dissipation," 2018.
- [233] X. Xu, S. C. Benjamin, and X. Yuan, "Variational circuit compiler for quantum error correction," 2019.

- [234] C. Zoufal, A. Lucchi, and S. Woerner, "Variational quantum boltzmann machines," 2020.
- [235] J. Liu and Y. Xin, "Quantum simulation of quantum field theories as quantum chemistry," 2020.
- [236] X. Xu, J. Sun, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, "Variational algorithms for linear algebra," 2019.
- [237] H.-Y. Huang, K. Bharti, and P. Rebentrost, "Near-term quantum algorithms for linear systems of equations," 2019.
- [238] F. Fontanela, A. Jacquier, and M. Oumgari, "A quantum algorithm for linear pdes arising in finance," 2019.
- [239] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, "Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution," *Nature Physics*, vol. 16, pp. 205–210, Feb 2020.
- [240] G. Salis and N. Moll, "Short-depth trial-wavefunctions for the variational quantum eigensolver based on the problem hamiltonian," 2019.
- [241] E. Neuscamman, H. Changlani, J. Kinder, and G. K.-L. Chan, "Non-stochastic algorithms for jastrow-slater and correlator product state wave functions," *Phys. Rev. B*, vol. 84, p. 205132, Nov 2011.
- [242] J. M. Wahlen-Strothman, C. A. Jiménez-Hoyos, T. M. Henderson, and G. E. Scuseria, "Lie algebraic similarity transformed hamiltonians for lattice model systems," *Phys. Rev. B*, vol. 91, p. 041114, Jan 2015.
- [243] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, "Toward the first quantum simulation with quantum speedup," *Proceedings of the National Academy of Sciences*, vol. 115, no. 38, pp. 9456–9461, 2018.

- [244] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, "Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm," *npj Quantum Information*, vol. 6, no. 1, pp. 1–9, 2020.
- [245] D. Gottesman, "Stabilizer codes and quantum error correction," *arXiv:9705.052*, 1997.
- [246] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. O'Brien, "Low-cost error mitigation by symmetry verification," *Physical Review A*, vol. 98, no. 6, p. 062339, 2018.
- [247] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller, "Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation," *Phys. Rev. X*, vol. 7, p. 041061, Dec 2017.
- [248] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Experimental comparison of two quantum computing architectures," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3305–3310, 2017.
- [249] R. Noek, G. Vrijsen, D. Gaultney, E. Mount, T. Kim, P. Maunz, and J. Kim, "High speed, high fidelity detection of an atomic hyperfine qubit," *Opt. Lett.*, vol. 38, pp. 4735–4738, Nov 2013.
- [250] M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, and G. K. Chan, "Low rank representations for quantum simulation of electronic structure," *arXiv preprint arXiv:1808.02625*, 2018.
- [251] G. Wendin, "Quantum information processing with superconducting circuits: a review," *Reports on Progress in Physics*, 2017.

- [252] S. McArdle, "Github repository." <https://github.com/sammcardle30>, 2018.
- [253] T. E. O'Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, "Error mitigation via verified phase estimation," *arXiv preprint arXiv:2010.02538*, 2020.
- [254] J. R. McClean, Z. Jiang, N. C. Rubin, R. Babbush, and H. Neven, "Decoding quantum errors with subspace expansions," *Nature Communications*, vol. 11, no. 1, pp. 1–9, 2020.
- [255] S. McArdle, A. Mayorov, X. Shan, S. Benjamin, and X. Yuan, "Digital quantum simulation of molecular vibrations," *Chemical science*, vol. 10, no. 22, pp. 5725–5735, 2019.
- [256] R. P. de Tudela, F. J. Aoiz, Y. V. Suleimanov, and D. E. Manolopoulos, "Chemical reaction rates from ring polymer molecular dynamics: Zero point energy conservation in  $\mu\text{-H}_2 \rightarrow \text{MuH-H}$ ," *J. Phys. Chem. Lett.*, vol. 3, pp. 493–497, feb 2012.
- [257] A. Gross and M. Scheffler, "Role of zero-point effects in catalytic reactions involving hydrogen," *J. Vac. Sci. Technol. A: Vacuum, Surfaces, and Films*, vol. 15, pp. 1624–1629, may 1997.
- [258] C. Zhu, K. K. Liang, M. Hayashi, and S. H. Lin, "Theoretical treatment of anharmonic effect on molecular absorption, fluorescence spectra, and electron transfer," *Chem. Phys.*, vol. 358, pp. 137–146, mar 2009.
- [259] J. Huh, "Unified description of vibronic transitions with coherent states," *PhD Thesis*, 2011.
- [260] T. Seideman, *Forming Superposition States, in Computational Molecular Spectroscopy*. Chichester : Wiley, 2000.

- [261] O. Christiansen, "Selected new developments in vibrational structure theory: potential construction and vibrational wave function calculations," *Physical Chemistry Chemical Physics*, vol. 14, no. 19, pp. 6672–6687, 2012.
- [262] D. G. Artiukhin, E. L. Klinting, C. König, and O. Christiansen, "Adaptive density-guided approach to double incremental potential energy surface construction," *The Journal of Chemical Physics*, vol. 152, no. 19, p. 194105, 2020.
- [263] G. Schmitz, I. H. Godtlielsen, and O. Christiansen, "Machine learning for potential energy surfaces: An extensive database and assessment of methods," *The Journal of chemical physics*, vol. 150, no. 24, p. 244113, 2019.
- [264] O. Christiansen, "Vibrational structure theory: new vibrational wave function methods for calculation of anharmonic vibrational energies and vibrational contributions to molecular properties," *Physical Chemistry Chemical Physics*, vol. 9, no. 23, pp. 2942–2953, 2007.
- [265] O. Christiansen, "Vibrational coupled cluster theory," *The Journal of chemical physics*, vol. 120, no. 5, pp. 2149–2159, 2004.
- [266] T. Carrington Jr, "Perspective: Computing (ro-) vibrational spectra of molecules with more than four atoms," *The Journal of Chemical Physics*, vol. 146, no. 12, p. 120902, 2017.
- [267] A. Baiardi, C. J. Stein, V. Barone, and M. Reiher, "Vibrational density matrix renormalization group," *Journal of Chemical Theory and Computation*, vol. 13, no. 8, pp. 3764–3777, 2017.
- [268] P. S. Thomas, T. Carrington Jr, J. Agarwal, and H. F. Schaefer III, "Using an iterative eigensolver and intertwined rank reduction to compute vibrational spectra of molecules with more than a dozen atoms:

- Uracil and naphthalene," *The Journal of chemical physics*, vol. 149, no. 6, p. 064108, 2018.
- [269] S. Joshi, A. Shukla, H. Katiyar, A. Hazra, and T. S. Mahesh, "Estimating franck-condon factors using an nmr quantum processor," *Phys. Rev. A*, vol. 90, p. 022303, Aug 2014.
- [270] J. Huh, G. G. Guerreschi, B. Peropadre, J. R. McClean, and A. Aspuru-Guzik, "Boson sampling for molecular vibronic spectra," *Nat Photonics*, vol. 9, no. 9, p. 615, 2015.
- [271] J. Huh and M.-H. Yung, "Vibronic boson sampling: Generalized gaussian boson sampling for molecular vibronic spectra at finite temperature," *Scientific reports*, vol. 7, no. 1, p. 7462, 2017.
- [272] W. R. Clements, J. J. Renema, A. Eckstein, A. A. Valido, A. Lita, T. Gerrits, S. W. Nam, W. S. Kolthammer, J. Huh, and I. A. Walmsley, "Approximating vibronic spectroscopy with imperfect quantum optics," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 51, no. 24, p. 245503, 2018.
- [273] Y. Shen, Y. Lu, K. Zhang, J. Zhang, S. Zhang, J. Huh, and K. Kim, "Quantum optical emulation of molecular vibronic spectroscopy using a trapped-ion device," *Chem. Sci.*, vol. 9, pp. 836–840, 2018.
- [274] C. Sparrow, E. Martín-López, N. Maraviglia, A. Neville, C. Harrold, J. Carolan, Y. N. Joglekar, T. Hashimoto, N. Matsuda, J. L. OâBrien, *et al.*, "Simulating the vibrational quantum dynamics of molecules using photonics," *Nature*, vol. 557, no. 7707, p. 660, 2018.
- [275] S. Chin and J. Huh, "Quantum computing for molecular vibronic spectra and gaussian boson sampling," *Journal of Physics: Conference Series*, vol. 1071, p. 012009, aug 2018.
- [276] L. Hu, Y.-C. Ma, Y. Xu, W.-T. Wang, Y.-W. Ma, K. Liu, H.-Y. Wang, Y.-

- P. Song, M.-H. Yung, and L.-Y. Sun, "Simulation of molecular spectroscopy with circuit quantum electrodynamics," *Sci. Bull.*, vol. 63, no. 5, pp. 293–299, 2018.
- [277] R. D. Somma, G. Ortiz, E. H. Knill, and J. Gubernatis, "Quantum simulations of physics problems," *Proc.SPIE*, vol. 5105, pp. 5105 – 5105 – 8, 2003.
- [278] L. Veis, J. ViÅk, H. Nishizawa, H. Nakai, and J. Pittner, "Quantum chemistry beyond bornâoppenheimer approximation on a quantum computer: A simulated phase estimation study," *Int. J. Quantum Chem.*, vol. 116, no. 18, pp. 1328–1336, 2016.
- [279] O. Christiansen, "A second quantization formulation of multimode dynamics," *J. Chem. Phys.*, vol. 120, no. 5, pp. 2140–2148, 2004.
- [280] V. Nagalakshmi, V. Lakshminarayana, G. Sumithra, and M. D. Prasad, "Coupled cluster description of anharmonic molecular vibrations. application to o3 and so2," *Chemical physics letters*, vol. 217, no. 3, pp. 279–282, 1994.
- [281] W. Mizukami and D. P. Tew, "A second-order multi-reference perturbation method for molecular vibrations," *The Journal of chemical physics*, vol. 139, no. 19, p. 194108, 2013.
- [282] S. Heislbetz and G. Rauhut, "Vibrational multiconfiguration self-consistent field theory: Implementation and test calculations," *The Journal of chemical physics*, vol. 132, no. 12, p. 124102, 2010.
- [283] F. Pfeiffer and G. Rauhut, "Multi-reference vibration correlation methods," *The Journal of chemical physics*, vol. 140, no. 6, p. 064110, 2014.
- [284] E. Fermi, "Über den ramaneffekt des kohlendioxyds," *Zeitschrift für Physik*, vol. 71, no. 3-4, pp. 250–259, 1931.

- [285] E. L. Sibert III, "Modeling vibrational anharmonicity in infrared spectra of high frequency vibrations of polyatomic molecules," *The Journal of chemical physics*, vol. 150, no. 9, p. 090901, 2019.
- [286] A. Bärttschi and S. Eidenbenz, "Deterministic preparation of dicke states," in *International Symposium on Fundamentals of Computation Theory*, pp. 126–139, Springer, 2019.
- [287] R. T. Birge, "The band spectra of carbon monoxide," *Phys. Rev.*, vol. 28, pp. 1157–1181, Dec 1926.
- [288] G. Herzberg and E. Teller, "Schwingungsstruktur der elektronenübergänge bei mehratomigen molekülen," *Zeitschrift für Physikalische Chemie*, vol. 21B, no. 1, pp. 410 – 446, 01 Feb. 1933.
- [289] E. Doktorov, I. Malkin, and V. Man'ko, "Dynamical symmetry of vibronic transitions in polyatomic molecules and the franck-condon principle," *Journal of Molecular Spectroscopy*, vol. 64, no. 2, pp. 302 – 326, 1977.
- [290] F. Duschinsky, "The importance of the electron spectrum in multi atomic molecules. concerning the franck-condon principle," *Acta Physicochim. URSS*, vol. 7, pp. 551–566, 1937.
- [291] N. P. Sawaya, T. Menke, T. H. Kyaw, S. Johri, A. Aspuru-Guzik, and G. G. Guerreschi, "Resource-efficient digital quantum simulation of d-level systems for photonic, vibrational, and spin-s hamiltonians," *npj Quantum Information*, vol. 6, no. 1, pp. 1–13, 2020.
- [292] N. P. Sawaya, G. G. Guerreschi, and A. Holmes, "On connectivity-dependent resource requirements for digital quantum simulation of d-level particles," *arXiv preprint arXiv:2005.13070*, 2020.
- [293] N. P. Sawaya and J. Huh, "Quantum algorithm for calculating molec-

- ular vibronic spectra," *The journal of physical chemistry letters*, vol. 10, no. 13, pp. 3586–3591, 2019.
- [294] P. J. Ollitrault, A. Baiardi, M. Reiher, and I. Tavernelli, "Hardware efficient quantum algorithms for vibrational structure calculations," *Chemical Science*, 2020.
- [295] N. P. Sawaya, F. Paesani, and D. P. Tabor, "Near-and long-term quantum algorithmic approaches for vibrational spectroscopy," *arXiv preprint arXiv:2009.05066*, 2020.
- [296] N. K. Madsen, I. H. Godtlielsen, S. A. Losilla, and O. Christiansen, "Tensor-decomposed vibrational coupled-cluster theory: Enabling large-scale, highly accurate vibrational-structure calculations," *The Journal of Chemical Physics*, vol. 148, no. 2, p. 024103, 2018.
- [297] I. H. Godtlielsen, B. Thomsen, and O. Christiansen, "Tensor decomposition and vibrational coupled cluster theory," *The Journal of Physical Chemistry A*, vol. 117, no. 32, pp. 7267–7279, 2013. PMID: 23662994.
- [298] I. H. Godtlielsen, M. B. Hansen, and O. Christiansen, "Tensor decomposition techniques in the solution of vibrational coupled cluster response theory eigenvalue equations," *The Journal of chemical physics*, vol. 142, no. 2, p. 024105, 2015.
- [299] S. McArdle, "Learning from physics experiments, with quantum computers: Applications in muon spectroscopy," *arXiv preprint arXiv:2012.06602*, 2020.
- [300] J. M. Wilkinson and S. J. Blundell, "Information and decoherence in a muon-fluorine coupled system," *Phys. Rev. Lett.*, vol. 125, p. 087201, Aug 2020.
- [301] M. Celio and P. F. Meier, "Exact calculation of the muon polarization function," *Hyperfine Interactions*, vol. 18, no. 1, pp. 435–440, 1984.

- [302] J. Lord, S. Cottrell, and W. Williams, "Muon spin relaxation in strongly coupled systems," *Physica B: Condensed Matter*, vol. 289-290, pp. 495 – 498, 2000.
- [303] S. J. Blundell, "Spin-polarized muons in condensed matter physics," *Contemporary Physics*, vol. 40, no. 3, pp. 175–192, 1999.
- [304] A. Yaouanc and P. D. De Reotier, *Muon spin rotation, relaxation, and resonance: applications to condensed matter*, vol. 147. Oxford University Press, 2011.
- [305] T. Lancaster, B. M. Huddart, R. C. Williams, F. Xiao, K. J. A. Franke, P. J. Baker, F. L. Pratt, S. J. Blundell, J. A. Schlueter, M. B. Mills, A. C. Maahs, and K. E. Preuss, "Probing magnetic order and disorder in the one-dimensional molecular spin chains  $\text{CuF}_2(\text{pyz})$  and  $[\text{In}(\text{hfac})_3(\text{boaDTDA})]_n$  ( $\text{In} = \text{sm}, \text{la}$ ) using implanted muons," *Journal of Physics: Condensed Matter*, vol. 31, p. 394002, jul 2019.
- [306] V. G. Storchak and N. V. Prokof'ev, "Quantum diffusion of muons and muonium atoms in solids," *Rev. Mod. Phys.*, vol. 70, pp. 929–978, Jul 1998.
- [307] J. Sugiyama, K. Mukai, Y. Ikedo, H. Nozaki, M. Månsson, and I. Watanabe, "Li diffusion in  $\text{Li}_x\text{COO}_2$  probed by muon-spin spectroscopy," *Phys. Rev. Lett.*, vol. 103, p. 147601, Sep 2009.
- [308] J. Sugiyama, H. Nozaki, I. Umegaki, Y. Higuchi, S. P. Cottrell, S. Hori, R. Kanno, and M. Månsson,  *$\mu^+$ SR Study on Li Ionic Conductors*.
- [309] I. Umegaki, S. Kawauchi, H. Nozaki, H. Sawada, H. Nakano, M. Harada, S. P. Cottrell, F. C. Coomer, M. Telling, and J. Sugiyama, *Observation of Li Diffusion in Cathode Sheets of Li-ion Battery by  $\mu^+$ SR*.

- [310] A. D. Pant, K. Nagamine, I. Shiraki, E. Torikai, K. Shimomura, F. L. Pratt, H. Ariga, K. Ishida, and J. S. Schultz, "Muonium response to oxygen content in biological aqueous solutions for cancer research," *Journal of Physics: Conference Series*, vol. 551, p. 012043, dec 2014.
- [311] M. Celio, "Corrections to the strong collision model," *Hyperfine Interactions*, vol. 31, pp. 153–155, Dec 1986.
- [312] E. Holzschuh and P. F. Meier, "Zero-field spin relaxation of positive muons," *Phys. Rev. B*, vol. 29, pp. 1129–1134, Feb 1984.
- [313] R. Kubo and T. Toyabe, *Magnetic Resonance and Relaxation*, ed. R. Blinc. North-Holland, Amsterdam, 1966.
- [314] P. D. de Reotier, A. Yaouanc, and S. Meshkov, "Effect of nuclear spin dynamics on the  $^{139}\text{La}$  zero field depolarisation function," *Physics Letters A*, vol. 162, no. 2, pp. 206 – 212, 1992.
- [315] P. D. de Reotier and A. Yaouanc, "Quantum calculation of the muon depolarization function: effect of spin dynamics in nuclear dipole systems," *Journal of Physics: Condensed Matter*, vol. 4, pp. 4533–4556, may 1992.
- [316] M. Celio and P. F. Meier, "Spin relaxation of positive muons due to dipolar interactions," *Phys. Rev. B*, vol. 27, pp. 1908–1910, Feb 1983.
- [317] T. Lancaster, S. J. Blundell, P. J. Baker, M. L. Brooks, W. Hayes, F. L. Pratt, J. L. Manson, M. M. Conner, and J. A. Schlueter, "Muon-fluorine entangled states in molecular magnets," *Phys. Rev. Lett.*, vol. 99, p. 267601, Dec 2007.
- [318] T. Lancaster, F. L. Pratt, S. J. Blundell, I. McKenzie, and H. E. Assender, "Muon–fluorine entanglement in fluoropolymers," *Journal of Physics: Condensed Matter*, vol. 21, p. 346004, aug 2009.

- [319] M. Celio, “New method to calculate the muon polarization function,” *Phys. Rev. Lett.*, vol. 56, pp. 2720–2723, Jun 1986.
- [320] P. Dalmas de Réotier, A. Yaouanc, and J. P. Boucher, “Model dependence of the longitudinal  $\mu$ sr depolarisation function in nuclear dipole systems,” *Hyperfine Interactions*, vol. 65, pp. 1121–1127, Feb 1991.
- [321] W. Huang, V. Pacradouni, M. P. Kennett, S. Komiya, and J. E. Sonier, “Precision search for magnetic order in the pseudogap regime of  $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$  by muon spin relaxation,” *Phys. Rev. B*, vol. 85, p. 104527, Mar 2012.
- [322] S. Blundell, A. Steele, T. Lancaster, J. Wright, and F. Pratt, “A bayesian approach to magnetic moment determination using  $\hat{\mu}_{\text{sr}}$ ,” *Physics Procedia*, vol. 30, pp. 113 – 116, 2012. 12th International Conference on Muon Spin Rotation, Relaxation and Resonance ( $\hat{\mu}_{\text{SR}}2011$ ).
- [323] P. Bonfà and R. De Renzi, “Toward the computational prediction of muon sites and interaction parameters,” *Journal of the Physical Society of Japan*, vol. 85, no. 9, p. 091014, 2016.
- [324] F. Bernardini, P. Bonfà, S. Massidda, and R. De Renzi, “Ab initio strategy for muon site assignment in wide band gap fluorides,” *Phys. Rev. B*, vol. 87, p. 115148, Mar 2013.
- [325] I. J. Onuorah, P. Bonfà, R. De Renzi, L. Monacelli, F. Mauri, M. Candra, and I. Errea, “Quantum effects in muon spin spectroscopy within the stochastic self-consistent harmonic approximation,” *Phys. Rev. Materials*, vol. 3, p. 073804, Jul 2019.
- [326] J. S. Möller, P. Bonfà, D. Ceresoli, F. Bernardini, S. J. Blundell, T. Lancaster, R. D. Renzi, N. Marzari, I. Watanabe, S. Sulaiman, and M. I. Mohamed-Ibrahim, “Playing quantum hide-and-seek with

- the muon: localizing muon stopping sites," *Physica Scripta*, vol. 88, p. 068510, dec 2013.
- [327] J. S. Möller, D. Ceresoli, T. Lancaster, N. Marzari, and S. J. Blundell, "Quantum states of muons in fluorides," *Phys. Rev. B*, vol. 87, p. 121108, Mar 2013.
- [328] G. M. Luke, J. H. Brewer, S. R. Kreitzman, D. R. Noakes, M. Celio, R. Kadono, and E. J. Ansaldo, "Muon diffusion and level crossing in copper," *Hyperfine Interactions*, vol. 64, pp. 721–727, Feb 1991.
- [329] J. H. Brewer, S. R. Kreitzman, D. R. Noakes, E. J. Ansaldo, D. R. Harshman, and R. Keitel, "Observation of muon-fluorine "hydrogen bonding" in ionic crystals," *Phys. Rev. B*, vol. 33, pp. 7813–7816, Jun 1986.
- [330] K. Nishiyama, S. Nishiyama, and W. Higemoto, "Asymmetric  $f \mu f$  interaction of the muon in polyfluorocarbons," *Physica B: Condensed Matter*, vol. 326, no. 1, pp. 41 – 45, 2003.
- [331] D. Noakes, E. Ansaldo, S. Kreitzman, and G. Luke, "The  $(f \mu f)$  ion in solid fluorides," *Journal of Physics and Chemistry of Solids*, vol. 54, no. 7, pp. 785 – 792, 1993.
- [332] R. Kadono, J. Imazato, T. Matsuzaki, K. Nishiyama, K. Nagamine, T. Yamazaki, D. Richter, and J.-M. Welter, "Quantum diffusion of positive muons in copper," *Phys. Rev. B*, vol. 39, pp. 23–41, Jan 1989.
- [333] G. M. Luke, J. H. Brewer, S. R. Kreitzman, D. R. Noakes, M. Celio, R. Kadono, and E. J. Ansaldo, "Muon diffusion and spin dynamics in copper," *Phys. Rev. B*, vol. 43, pp. 3284–3297, Feb 1991.
- [334] B. Koczor, "On phase-space representations of spin systems and their relations to infinite-dimensional quantum states," *Ph.D. Thesis*, 2019.

- [335] D. Bacon, I. L. Chuang, and A. W. Harrow, "The quantum schur transform: I. efficient qudit circuits," 2005.
- [336] D. Bacon, I. L. Chuang, and A. W. Harrow, "Efficient quantum circuits for schur and clebsch-gordan transforms," *Phys. Rev. Lett.*, vol. 97, p. 170502, Oct 2006.
- [337] W. M. Kirby and F. W. Strauch, "A practical quantum algorithm for the schur transform," *Quantum Info. Comput.*, vol. 18, p. 721â742, Aug. 2018.
- [338] H. Krovi, "An efficient high dimensional quantum Schur transform," *Quantum*, vol. 3, p. 122, Feb. 2019.
- [339] M. Heyl, P. Hauke, and P. Zoller, "Quantum localization bounds trotter errors in digital quantum simulation," *Science advances*, vol. 5, no. 4, p. eaau8342, 2019.
- [340] D. Wang, O. Higgott, and S. Brierley, "Accelerated variational quantum eigensolver," *Phys. Rev. Lett.*, vol. 122, p. 140504, Apr 2019.
- [341] E. Knill, G. Ortiz, and R. D. Somma, "Optimal quantum measurements of expectation values of observables," *Phys. Rev. A*, vol. 75, p. 012328, Jan 2007.
- [342] V. Schäfer, C. Ballance, K. Thirumalai, L. Stephenson, T. Ballance, A. Steane, and D. Lucas, "Fast quantum logic gates with trapped-ion qubits," *Nature*, vol. 555, no. 7694, pp. 75–78, 2018.
- [343] E. van den Berg and K. Temme, "Circuit optimization of Hamiltonian simulation by simultaneous diagonalization of Pauli clusters," *Quantum*, vol. 4, p. 322, Sept. 2020.
- [344] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, "Locality and digital quantum

- simulation of power-law interactions," *Phys. Rev. X*, vol. 9, p. 031006, Jul 2019.
- [345] C. Gidney, "Halving the cost of quantum addition," *Quantum*, vol. 2, p. 74, June 2018.
- [346] R. Babbush, J. McClean, C. Gidney, S. Boixo, and H. Neven, "Focus beyond quadratic speedups for error-corrected quantum advantage," 2020.
- [347] E. Knill and R. Laflamme, "Power of one bit of quantum information," *Phys. Rev. Lett.*, vol. 81, pp. 5672–5675, Dec 1998.
- [348] K. Fujii, H. Kobayashi, T. Morimae, H. Nishimura, S. Tamate, and S. Tani, "Impossibility of classically simulating one-clean-qubit model with multiplicative error," *Phys. Rev. Lett.*, vol. 120, p. 200502, May 2018.
- [349] A. Datta and G. Vidal, "Role of entanglement and correlations in mixed-state quantum computation," *Phys. Rev. A*, vol. 75, p. 042310, Apr 2007.
- [350] T. Morimae, K. Fujii, and J. F. Fitzsimons, "Hardness of classically simulating the one-clean-qubit model," *Phys. Rev. Lett.*, vol. 112, p. 130502, Apr 2014.
- [351] T. Morimae, "Hardness of classically sampling the one-clean-qubit model with constant total variation distance error," *Phys. Rev. A*, vol. 96, p. 040302, Oct 2017.
- [352] E. T. Campbell, "Early fault-tolerant simulations of the hubbard model," *arXiv preprint arXiv:2012.09238*, 2020.
- [353] J. M. Gambetta, J. M. Chow, and M. Steffen, "Building logical qubits

- in a superconducting quantum computing system," *npj Quantum Information*, vol. 3, no. 1, p. 2, 2017.
- [354] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, "Quantum computers," *Nature*, vol. 464, pp. 45 EP –, Mar 2010. Review Article.
- [355] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, "Faster digital quantum simulation by symmetry protection," *arXiv preprint arXiv:2006.16248*, 2020.
- [356] Y. Su, H.-Y. Huang, and E. T. Campbell, "Nearly tight trotterization of interacting electrons," *arXiv preprint arXiv:2012.09194*, 2020.
- [357] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, "Virtual distillation for quantum error mitigation," *arXiv preprint arXiv:2011.07064*, 2020.
- [358] B. Koczor, "Exponential error suppression for near-term quantum devices," *arXiv preprint arXiv:2011.05942*, 2020.
- [359] R. LaRose, "Overview and Comparison of Gate Level Quantum Software Platforms," *Quantum*, vol. 3, p. 130, Mar. 2019.
- [360] Q. O.-S. Foundation, "List of open quantum projects." [https://qosf.org/project\\_list/](https://qosf.org/project_list/), 2020.
- [361] Quantiki, "List of qc simulators." <https://www.quantiki.org/wiki/list-qc-simulators>, 2020.
- [362] IBM, "Ibm quantum experience." <https://www.ibm.com/quantum-computing/experience>, 2020.
- [363] C. Rigetti, "Introducing rigetti quantum cloud services." <https://medium.com/rigetti/>

[introducing-rigetti-quantum-cloud-services-c6005729768c](#),  
2018.

[364] A. Quantum, “Amazon braket.” <https://aws.amazon.com/braket/>, 2020.

## A | Appendices

### A.1 Appendix: Error mitigation for quantum simulation

#### A.1.1 Number of measurements and error analysis

The numerical simulations in Chapter 5, Sec. 5.4 were designed to mimic the actions of an experimentalist; the expectation value of each term in the Hamiltonian was found by repeating the circuit and measurement procedure many times.

The number of measurements required was set by the desired precision. Measurements were distributed optimally among the different Hamiltonian terms [122, 123], such that the number of measurements for each term was proportional to its strength. The standard error in each measurement is then given as follows.

The standard error in the mean is given by

$$\alpha = \frac{\sigma}{\sqrt{N}} \quad (\text{A.1})$$

where  $N$  is the number of measurements performed, and  $\sigma$  is the standard deviation of the result. The standard deviation of a measurement of one of the Pauli strings in the Hamiltonian,  $H_j$ , is given by

$$\sigma_{H_j} = \sqrt{\langle \psi | H_j^2 | \psi \rangle - \langle \psi | H_j | \psi \rangle^2} = \sqrt{1 - \langle \psi | H_j | \psi \rangle^2} \leq 1. \quad (\text{A.2})$$

The standard error in the energy measurement is then upper bounded by

$$\alpha_E = \sqrt{\sum_i \alpha_{h_i}^2} = \sqrt{\sum_i |h_i|^2 \frac{\sigma_{H_j}^2}{N_i}} \leq \sqrt{\sum_i \frac{|h_i|^2}{N_i}}, \quad (\text{A.3})$$

where  $N_i$  is the number of measurements used for each term in the Hamiltonian, and  $h_i$  is the coefficient of term  $i$  in the Hamiltonian. Measurements were distributed optimally [122, 123], setting

$$N_i = \frac{|h_i|}{h_{\max}} k \quad (\text{A.4})$$

where  $h_{\max}$  is the largest coefficient in the Hamiltonian, and  $k$  is the number of measurements allocated to the largest term in the Hamiltonian. Substituting this expression into Eq. A.3,

$$\alpha_E \leq \sqrt{\sum_i \frac{|h_i| h_{\max}}{k}} \quad (\text{A.5})$$

Solving for  $k$ , and substituting back into Eq. A.4 provides an expression for the number of measurements required per term, as a function of the standard error

$$N_i = \frac{|h_i| \sum_j |h_j|}{\alpha_E^2}. \quad (\text{A.6})$$

The total number of measurements is then

$$N = \frac{(\sum_i |h_i|)^2}{\alpha_E^2}. \quad (\text{A.7})$$

When extrapolation is performed, the standard error is increased by a factor that depends on the ‘stretch-factor’,  $\lambda$ , used in the extrapolation. When performing extrapolation, I used the same total number of samples for each expectation value, which were divided equally between two points for a linear extrapolation. The extrapolated value is given by

$$O_{\text{extrap}} = \frac{\lambda O(\epsilon) - O(\lambda\epsilon)}{\lambda - 1}. \quad (\text{A.8})$$

As such, the standard error was increased by a factor of

$$\omega = \frac{\sqrt{2(\lambda^2 + 1)}}{(\lambda - 1)}. \quad (\text{A.9})$$

The stretch factor used in the simulations shown in Fig. 5.8 depended on the error rate,  $\epsilon$ , as

$$\lambda = 1 + \frac{0.001}{\epsilon}. \quad (\text{A.10})$$

The stretch factor used in the simulations shown in Fig. 5.9 was

$$\lambda = 1.5. \quad (\text{A.11})$$

For the results presented in Fig. 5.9, I calculated the standard error using the expectation values  $\sqrt{1 - \langle \psi | H_j | \psi \rangle^2}$ , rather than the loose upper bound. The resulting standard errors in the results were: Detection & Extrapolation = [0.05, 0.07, 0.04, 0.04, 0.04] mHartree, Extrapolation = [0.05, 0.07, 0.04, 0.05, 0.04] mHartree.

For the results presented in Fig. 5.8, I used the loose upper bound for the standard error, obtained as described above.

## A.2 Appendix: Quantum simulation of muon spectroscopy experiments

### A.2.1 Preparing Dicke states

In this Appendix, I discuss the method developed by Ref. [286] to prepare Dicke states on quantum computers. As discussed in Chapter 7, to prepare a Dicke state with Hamming weight  $h$ , acting on  $n$  qubits, the algorithm requires  $\mathcal{O}(hn)$  gates,  $\mathcal{O}(n)$  depth, and  $n$  qubits. Here, I present a slightly less rigorous, but more pedagogical overview of the algorithm introduced

in Ref. [286], and refer the reader to the original reference for more information.

The state preparation algorithm proceeds recursively, making use of the following expression for Dicke states

$$|D_h^n\rangle = \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \quad (\text{A.12})$$

Assume the existence of a unitary operator  $U_{n,k}$  such that  $U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} = |D_h^n\rangle$  for all  $h \leq k$ . As I show below, this operator exists, and can be constructed from typical single- and two-qubit gates.

As a first step, note that

$$|D_h^n\rangle = U_{n,k} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \quad (\text{A.13})$$

and

$$\begin{aligned} |D_h^n\rangle &= \sqrt{\frac{h}{n}} |D_{h-1}^{n-1}\rangle \otimes |1\rangle + \sqrt{\frac{n-h}{n}} |D_h^{n-1}\rangle \otimes |0\rangle. \\ &= \left( U_{n-1,k} \otimes I \right) \left[ \sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle \right] \\ &= \left( U_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right) |0\rangle^{\otimes n-h} |1\rangle^{\otimes h}, \end{aligned} \quad (\text{A.14})$$

with

$$\left( I_{n-k-1} \otimes V_{n,k} \right) |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} = \sqrt{\frac{h}{n}} |0\rangle^{\otimes n-h} |1\rangle^{\otimes h} + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes n-1-h} |1\rangle^{\otimes h} |0\rangle. \quad (\text{A.15})$$

These relations imply that

$$U_{n,k} = \left( U_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right). \quad (\text{A.16})$$

Here,  $U_{n-1,k}$  acts on the leftmost  $n-1$  qubits.  $V_{n,k}$  is defined on  $n$  qubits,

but acts trivially on the first  $n - k - 1$  qubits, so can be considered to only act on the final  $k + 1$  qubits. We can then recurse this relationship

$$U_{n-1,k} = \left( U_{n-2,k} \otimes I_2 \right) \cdot \left( I_{n-k-2} \otimes V_{n-1,k} \right), \quad (\text{A.17})$$

where  $V_{n-1,k}$  acts on the leftmost  $k + 1$  of the final  $k + 2$  qubits in the state.

This enables us to write that

$$U_{n,k} = \left( U_{k+1,k} \otimes I_{n-k-1} \right) \dots \left( I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right). \quad (\text{A.18})$$

Repeating the analysis above for  $U_{k+1,k}$ , and considering just the  $k + 1$  qubits acted upon gives

$$\begin{aligned} U_{k+1,k} |0\rangle |1\rangle^{\otimes k} &= |D_k^{k+1}\rangle \\ &= \sqrt{\frac{k}{k+1}} |D_{k-1}^k\rangle \otimes |1\rangle + \sqrt{\frac{1}{k+1}} |1\rangle^{\otimes k} \otimes |0\rangle \\ &= \left( U_{k,k-1} \otimes I \right) \left[ \sqrt{\frac{k}{k+1}} |0\rangle |1\rangle^{\otimes k-1} \otimes |1\rangle + \sqrt{\frac{1}{k+1}} |1\rangle^{\otimes k} \otimes |0\rangle \right] \\ &= \left( U_{k,k-1} \otimes I \right) \cdot \left( V_{k+1,k} \right) |0\rangle |1\rangle^{\otimes k-1} \otimes |1\rangle, \end{aligned} \quad (\text{A.19})$$

where  $V_{k+1,k}$  acts on all of the  $k + 1$  qubits. Continuing the recursion above

$$\begin{aligned} U_{n,k} &= \left( U_{k+1,k} \otimes I_{n-k-1} \right) \dots \left( I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right) \\ &= \left( U_{k,k-1} \otimes I_{n-k} \right) \cdot \left( V_{k+1,k} \otimes I_{n-k-1} \right) \dots \left( I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right) \\ &= \dots \\ &= \left( V_{2,1} \otimes I_{n-2} \right) \cdot \left( V_{3,2} \otimes I_{n-3} \right) \cdot \dots \cdot \left( V_{k,k-1} \otimes I_{n-k} \right) \cdot \left( V_{k+1,k} \otimes I_{n-k-1} \right) \cdot \\ &\quad \left( I \otimes V_{k+2,k} \otimes I_{n-k-2} \right) \cdot \dots \cdot \left( I_{n-k-2} \otimes V_{n-1,k} \otimes I \right) \cdot \left( I_{n-k-1} \otimes V_{n,k} \right). \end{aligned} \quad (\text{A.20})$$

This can be expressed this more concisely via the expression in Lemma 2 of

Ref. [286]:

$$U_{n,k} = \prod_{l=2}^k \left( V_{l,l-1} \otimes I_{n-l} \right) \cdot \prod_{l=k+1}^n \left( I_{l-k-1} \otimes V_{l,k} \otimes I_{n-l} \right). \quad (\text{A.21})$$

This confirms the existence of the gate  $U_{n,k}$ , if we are able to construct the unitary  $V_{n,k}$  for arbitrary  $n, k$ . As discussed in Ref. [286],  $V_{n,k}$  can be built from standard single-, two- and three-qubit gates as follows.

We have that

$$\left( V_{n,k} \right) |0\rangle^{\otimes k+1-h} |1\rangle^{\otimes h} = \sqrt{\frac{h}{n}} |0\rangle^{\otimes k+1-h} |1\rangle^{\otimes h} + \sqrt{\frac{n-h}{n}} |0\rangle^{\otimes k-h} |1\rangle^{\otimes h} |0\rangle. \quad (\text{A.22})$$




Figure A.1: A circuit which yields the unitary operator given by Eq. (A.23).

This unitary only changes the value of the zeroth qubit from the right, and the  $h^{\text{th}}$  qubit from the right. First, consider the circuit shown in Fig. A.1. Setting  $\theta = 2\cos^{-1}(\sqrt{1/n})$  gives the unitary

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{\frac{1}{n}} & -\sqrt{\frac{n-1}{n}} & 0 \\ 0 & \sqrt{\frac{n-1}{n}} & \sqrt{\frac{1}{n}} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (\text{A.23})$$

Similarly, the circuit in Fig. A.2 with  $\theta = 2\cos^{-1}(\sqrt{\alpha/n})$  acts trivially on all




Figure A.2: A circuit which yields the unitary operator which I refer to as  $W_n^\alpha$ . This operator acts trivially on all input states, except those noted in Eq. (A.24).

input states, except for

$$\begin{aligned} |011\rangle &\rightarrow \sqrt{\frac{\alpha}{n}} |011\rangle + \sqrt{\frac{n-\alpha}{n}} |110\rangle, \\ |110\rangle &\rightarrow -\sqrt{\frac{n-\alpha}{n}} |011\rangle + \sqrt{\frac{\alpha}{n}} |110\rangle. \end{aligned} \quad (\text{A.24})$$

We can now use these building blocks to construct the unitary  $V_{n,k}$ . First apply the gate in Fig. A.1 to the final two qubits in the register, and then repeatedly apply the gate  $W_n^\alpha$  shown in Fig. A.2, incrementing the value of  $\alpha$  each time. As we are applying these gates on the state  $|0\rangle^{\otimes n-h} |1\rangle^{\otimes h}$ , the first  $h$  qubits encountered are all in the  $|1\rangle$  state, so these gates act trivially. After these  $h$  qubits, the three qubit  $W_n^h$  gate encounters the state  $|0\rangle_{n-h} |1\rangle_{n-h+1} |1\rangle_n$ , which it transforms into

$$\sqrt{\frac{h}{n}} |0\rangle_{n-h} |1\rangle_{n-h+1} |1\rangle_n + \sqrt{\frac{n-h}{n}} |1\rangle_{n-h} |1\rangle_{n-h+1} |0\rangle_n. \quad (\text{A.25})$$

The subsequent  $W_n^{i>h}$  gates act trivially on both branches of this superposition. Taken together, these gates carry out the transform in Eq. (A.22), for all  $h \leq k$ . It is shown in Ref. [286] how the three-qubit  $W$  gate can be decomposed into single- and two-qubit gates, and how the circuit can be implemented with  $\mathcal{O}(kn)$  gates,  $\mathcal{O}(n)$  depth, and  $n$  qubits on a linear array of qubits, with nearest-neighbour connectivity.

## A.2.2 Additional numerical simulations of Trotter and sampling error

In this Appendix, I present the results of additional simulations to test the convergence of sampling error and Trotter error, for systems too large to study with exact diagonalisation.




Figure A.3: Comparing the sampling error obtained using both the random-phase-approximation (RPA) method and the dephasing method, at a given number of samples, to that obtained at 9600 samples. The RPA method rapidly converged to the value obtained at 9600 samples. While this does not bound the sampling error at 9600 samples, we see from Fig. 7.5 that the RPA method can obtain an accuracy of less than  $10^{-3}$  for the 11 qubit system with 100 samples. Given that the sampling error decreases with both system size, and the number of samples, I infer that the error is likely small for even a modest number of samples with the RPA method for the 21 qubit system.

As shown in Eq. (7.7), the error in the sampling-based methods decreases as the size of the simulated system increases. I investigated this by considering the sampling error in a 21 qubit simulation. This system was too large to carry out density matrix simulations for comparison. In lieu of this, I investigated the convergence of the obtained polarisation function as the

number of samples was increased. In Fig. A.3 I plot this behaviour for both the RPA and dephasing methods. The RPA method rapidly converged, and exhibited small fluctuations around its converged value. In particular, when using the RPA method only 20 samples sufficed to converge the polarisation function to within  $10^{-4}$  of the value obtained with 9600 samples.

I used a similar approach to investigate the convergence of Trotter errors for systems too large to be exactly simulated. In Fig. A.4 I plot how the polarisation function converged as the number of second-order Trotter steps was increased, for both the 21 and 29 qubit systems. These simulations were performed using the RPA method, with 48 samples for the 21 qubit simulation, and a single sample for the 29 qubit simulation. I compared the polarisation function at a range of Trotter steps to that obtained with 30 Trotter steps. This metric does not provide a bound on the Trotter error at 30 Trotter steps. However, given that the convergence error in both cases is monotonically decreasing, and less than  $10^{-2}$  by 20 Trotter steps, this may be taken as an indication that the polarisation function is rapidly converging as the number of Trotter steps is increased.




Figure A.4: Convergence of the second-order Trotter error for both the 21 and 29 qubit systems, as a function of the number of Trotter steps used. I plot how the polarisation converged towards the value at 30 second-order Trotter steps (which may still be far from the true value).