Exploiting Quantum Computers for the Physics Simulations of Many-Body Systems

Hans Hon Sang Chan

The Queen's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2024
Word Count: 38,056

Abstract

Quantum Computers promise to be superior to conventional computers for the simulation and modelling of quantum many-body systems. Their realisation will have important implications in the fields of condensed-matter physics, chemistry, materials science, biochemistry and beyond. As quantum computers are gradually scaled up, one question might be *how exactly would we use quantum computers to model quantum many-body systems that we can't model with conventional computers*? This thesis addresses this question by exploring well-known methodologies and establishing novel quantum computing algorithms, which target the digital simulation of many-body systems on current and future quantum computing technologies, and paves the way for achieving quantum advantage in simulation.

Exploiting Quantum Computers for the Physics Simulations of Many-Body Systems



Hans Hon Sang Chan

The Queen's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2024

Word Count: 38,056

Acknowledgements

Firstly, I would like to express my fondest appreciation to my supervisor, Professor Simon Benjamin. Without his insight, guidance, patience, feedback, challenge and support, my DPhil would not have been the rewarding experience that it was. I also want to thank him for reviewing this thesis. I would also like to whole-heartedly thank my co-supervisor Professor David Tew. His breadth and depth of knowledge and his thoughtfulness have been invaluable throughout my DPhil, but particularly in the first couple of years. I would like to thank the Croucher Foundation and the late Noel Croucher for their financial support through the Oxford-Croucher Scholarship. I would like to acknowledge the University of Oxford Advanced Computing Research (ARC) service, which I used to produce results described in this thesis.

The research in this thesis would not have been possible without key collaborators and domain experts from my research group and beyond. In no particular order, I would like to thank: Richard Meister, Tyson Jones, Matthew Goh, Xiaoning Feng, Adam Siegel, Minjun Jeon, Hamza Jnane, Arthur Rattew, Daniel Marti-Dafcik, Adrian Chapman, Zhenyu Cai, Cica Gustiani, Hugh Burton and Bálint Koczor. These superb scientists have helped me either through direct participation in joint research projects or fruitful discussions.

The opportunity to collaborate with companies in the vibrant quantum computing ecosystem here in the United Kingdom has also enriched my research. I would like to thank Quantinuum Ltd. and Phasecraft Ltd. for hosting me. In particular, I want to thank Alexis Ralli, Nathan Fitzpatrick and David Muñoz Ramo from Quantinuum, and Lana Mineh, Jan Lukas Bosse and Raul Santos from Phasecraft.

I would like to thank all my past mentors in the physical sciences, who have nudged me down this path. These include Mike Bearpark, Javier Segarra Marti, Marcus Medley, Paula Webb, Jane Green, Verity Dutton and many others.

I have enjoyed the brief honour of representing the Oxford University Amateur Boxing Club during my DPhil. I thank the club and particularly its former head coach, David Mace and his team, for some of the most unforgettable experiences during my time at the university. I also want to thank the world-class eye emergency team at St Thomas' Hospital London, for doing a superb job at repairing my retinal detachment as a result of boxing; thanks to their surgeons and care, these injuries healed rapidly and posed negligible interruption to my research.

I would like to thank the team at Matoha Instrumentation Ltd. for their continued efforts in keeping the operations going while some of the leadership hurriedly completes their DPhil.

A humble group of irreplaceable close friends have been a constant source of strength, council, joy and buffoonery for me. I would like to thank, in no particular order: Boris Mang, Ronald Wong, Jasper Wong, Kennard Ng, Jason Kwong, Kevin Wan, Divad Ng, Dexter Chua, Aisha Kamara, Lydia Seeley, Angel Droth-Vicens, Fernando Acero Marchesotti, Felipe Portocarrero, Alejandra Martńez del Navío, Javier Leguina, Erik Valverde, Gigi Pescador, Tobias Drees, Lorenzo Falcao, Emma De Cocker, Martin Holicky, Darlene Sammut, Benjamin Wooley, Estelle Paulus, Johannes Forkel, Håvard Damm-Johnsen, Jeremy Wu, Lorenz Olbrich, Johan Runeson, Lucien Haeller, Amanda Wicki, Ludwig Bauer and Flurin Caviezel. I also want to thank the Acero and Lieberherr families for the generous hospitality they have shown me on many occasions in the last few years.

I would like to thank my parents Kenneth Kam Keung Chan, May Hing Mui Tan, my brother Andre Lok Sang Chan and his partner Hayley Lam. Their unconditional love, care, wisdom and support, expressed in their own beautiful way, has and always will be felt and cherished in my life, even if oceans and continents lie between us. The encouragement from the extended maternal and paternal families, as well as friends of the family, continues to bring much comfort and gratification.

Und s Best chunt z'letscht: I muess de Annina Lieberherr danke. Abgseh vo de offesichtliche Gründ wie ihri edel Charaktereigeschafte, kennt sie alli Gründ warum I sie do nenne will - die Gründ sind jenseits vo Wort und ganz sicher jenseits vo mine Schwizerdütschkenntniss.

Abstract

Quantum Computers promise to be superior to conventional computers for the simulation and modelling of quantum many-body systems. Their realisation will have important implications in the fields of condensed-matter physics, chemistry, materials science, biochemistry and beyond. As quantum computers are gradually scaled up, one question might be *how exactly would we use quantum computers to model quantum many-body systems that we can't model with conventional computers*? This thesis addresses this question by exploring well-known methodologies and establishing novel quantum computing algorithms, which target the digital simulation of many-body systems on current and future quantum computing technologies, and paves the way for achieving quantum advantage in simulation.

Contents

List of Figures **xi**

List of Abbreviations **xv**

1 Introduction **1**

2 Background **5**

2.1 Digital Quantum Computers and their Emulators 5

2.1.1 Emulators 8

2.2 Quantum Simulation of Many-Body Physics 9

2.2.1 Motivation for Quantum Computing 11

2.2.2 Key Concepts for Simulation with Qubits 13

2.2.3 Representation of Quantum States 16

2.2.4 State Preparation 25

2.2.5 Quantum Simulation 35

2.2.6 Measuring Useful Information 45

3 Computing Energy Eigenvalues on Near-Term Quantum Computers **55**

3.1 Introduction 56

3.2 Shortcuts to Adiabaticity for Ground State Preparation 57

3.2.1 Method 60

3.2.2 Case Study: Transverse Field Ising Model 62

3.2.3 Numerical Experiments 66

3.2.4 Beyond the TFIM 70

3.2.5 Discussion 72

3.3 Symmetry-restricted Variational Ansätze for Excited States 73

3.3.1 Methods 74

3.3.2 Numerical Validation 78

3.3.3 Discussion 84

3.4 Concluding Thoughts 85

4 Non-Unitary Evolution on Early Fault-Tolerant Quantum Computers 87
4.1 Introduction . . . 88
4.2 Methods . . . 89
4.2.1 Non-unitary evolution via Quantum Eigenvalue Transform . 89
4.2.2 Key Features and Comparison with Other Methods . . . 91
4.2.3 Applications . . . 93
4.3 Numerical Validation . . . 99
4.3.1 Ground state preparation . . . 101
4.3.2 Gibbs state preparation . . . 105
4.3.3 Open Quantum System Dynamics . . . 107
4.3.4 Hamiltonian scaling: Bounds of extremum eigenvalues . . . 109
4.4 Concluding Thoughts . . . 110

5 Real-Space Grid Physics Simulations on Fault-Tolerant Quantum Computers 113
5.1 Introduction . . . 114
5.2 Methods . . . 119
5.2.1 Split-Operator Hamiltonian Simulation under the Coulomb potential . . . 119
5.2.2 Energy Observable . . . 125
5.2.3 Attenuation and Scattering . . . 126
5.2.4 State Preparation . . . 129
5.2.5 Augmented Split-Operator . . . 136
5.3 Results . . . 138
5.3.1 Trotter Evolution and the Relationship between Spatial and Temporal Resolution . . . 140
5.3.2 Quantum Dynamics Experiments . . . 144
5.3.3 Initial State Preparation . . . 149
5.3.4 Antisymmetrisation of the initial state . . . 152
5.3.5 Augmented Split-Operator . . . 154
5.3.6 Informal Resource Estimate . . . 158
5.3.7 Quantum Hardware and Architecture . . . 165
5.3.8 Additional Analysis and Results . . . 167
5.4 Concluding Thoughts . . . 169

6 Shadow Spectroscopy: Measurement Protocol for Any Quantum Technology **173**
6.1 Introduction . . . 174
6.2 Method . . . 175
6.2.1 Quatum Measurement of Many Signals . . . 175
6.2.2 Classical Post-Processing . . . 178
6.2.3 Summary of Shadow Spectroscopy . . . 182
6.3 Key Features of Shadow Spectroscopy . . . 182
6.3.1 Robustness to shot noise . . . 183
6.3.2 Robustness to algorithmic or gate errors . . . 183
6.3.3 Compatibility with Quantum Chemistry . . . 184
6.3.4 Comparison with other quantum algorithms . . . 188
6.4 Experimental Validation: NISQ Scenario . . . 189
6.4.1 Proof-of-Concept . . . 190
6.4.2 Hardware Demonstration . . . 193
6.5 Experimental Validation: Fault Tolerant Scenario . . . 199
6.5.1 Utility of Local Pauli Operators . . . 199
6.5.2 Methylene Example . . . 200
6.6 Additional Analysis and Experiments . . . 203
6.7 Concluding Thoughts . . . 204

7 Conclusion **209**

References **211**

List of Figures

2.1 Example Pauli gadget quantum circuit primitives that apply the exponential of Pauli string matrix. 16
2.2 The circuit of the Quantum Fourier Transform on a 5-qubit system. 16
2.3 Basis functions in a Discrete Variable Representation (DVR) of wave functions in real space. 24
2.4 Quantum many-body encodings on qubits used in this thesis. 27
2.5 Initialising the HF state on a qubit register in second quantisation. 30
2.6 An early-fault tolerant QSP method; the QET-U circuit and relation to the PITE circuit. 44
2.7 Single Pauli string expectation value measurement on a quantum computer. 47
2.8 Cartoon depicting classical shadows measurement. 50
2.9 Circuit for performing canonical phase estimation. 52

3.1 The Hamiltonian ansatz U^{Ham} versus the Trotterised counter-adiabatic driving ansatz U^{CDTrot} for the Transverse Field Ising Model. 63
3.2 Compilation of the $U_{12}(\phi, \theta)$ unitary into a 2-qubit gate with 2 CNOTs and single qubit rotations. 66
3.3 The recompiled counter-adiabatic driving ansatz U^{CDRecomp} for the Transverse Field Ising Model and the U^{ZYRecomp} ansatz where the gauge potential term is replaced with the YY operator. 66
3.4 Comparing the VQE performance of the Hamiltonian ansatz with the Trotterised counter-adiabatic ansatz for the Transverse Field Ising Model. 68
3.5 Comparing the VQE performance of the recompiled counter-adiabatic ansatz with other ansätze for the Transverse Field Ising Model. . . 69
3.6 Performing VQE using the Hamiltonian ansatz or the recompiled counter-adiabatic ansatz for the 11-qubit Transverse Field Ising Model, applied to a collection of random initial states. 70
3.7 Comparing the Hamiltonian ansatz, the recompiled counter-adiabatic ansatz and an ansatz where the gauge potential operator is replaced with the YY operator for VQE. 71

3.8 Comparing the Hamiltonian ansatz with the recompiled counter-adiabatic ansatz for the TFIM+F system. 72
3.9 Illustration of the ADAPT-VQD approach with symmetry-restriction for computing excited states on NISQ computers. 78
3.10 The exact electronic structure of the lowest-lying eigenstates along the Li-H bond length in the STO-3G basis. 79
3.11 Quantum simulations of the singlet and triplet electronic excited states of LiH along its bond length using the 2-UpCCGSD ansatz. . 81
3.12 Quantum simulations of the singlet and triplet excited states of LiH at 5 bond lengths using adaptively grown ansätze, unique for each geometry and each eigenstate. 83

4.1 Circuit for QET-U and relation to the PITE circuit. 92
4.2 Using single-ancilla quantum eigenvalue transform for Gibbs state preparation 96
4.3 Using single-ancilla quantum eigenvalue transform for non-unitary dynamics simulation. 99
4.4 Quantum eigenvalue transform applied to a scaled Hamiltonian. . . 101
4.5 Ground state convergence properties by performing imaginary-time evolution using QET-U. 104
4.6 Probability of successful state preparation and different initial trial states in the 4-site linear Fermi-Hubbard model. 105
4.7 Gibbs state preparation and partition function calculation using QET-U. 106
4.8 Propagating open quantum system dynamics using QET-U. 109

5.1 Quantum simulation using the SO-QFT approach in the real-space representation. 124
5.2 Un-computing tag ancilla registers used in the preparation of an antisymmetric state in first quantisation. 134
5.3 Three quantum circuit techniques used for real-space quantum simulation. 136
5.4 Relationship between spatial and temporal resolution explored in a SO-QFT simulation of the 2D hydrogen atom. 141
5.5 Real-space simulation of single and double-electron dynamics. . . . 145
5.6 Demonstrating state editing by ancilla phase measurement. 150
5.7 Ground state preparation of the 2D hydrogen atom using the PITE technique. 153
5.8 The performance of the Augmented Split Operator (ASO) technique. 155

5.9 A possible 'multicore' quantum computer architecture for grid-based simulation. 166
5.10 Simulation of a Helium atom in real space. 168

6.1 Flowchart of algorithmic shadow spectroscopy. 175
6.2 Classical post-processing step in shadow spectroscopy. 191
6.3 Scaling analysis of shadow spectroscopy. 193
6.4 Hardware efficient ansatz circuit layers used to approximate spin system evolution, and their measurement on real quantum hardware. 196
6.5 Experimental shadow spectrum of a 6-qubit Heisenberg spin ring obtained using the IBM Jakarta QPU compared with simulation. . 197
6.6 Experimental shadow spectrum of a 6-qubit disordered Heisenberg chain obtained using the 7-qubit IBM Jakarta QPU. 199
6.7 List of the at-most 3-local Pauli strings which give most intense signals for a 12-qubit LiH Hamiltonian. 200
6.8 Discovering the electronic spectrum of methylene with (classically emulated) shadow spectroscopy. 201

List of Abbreviations

1-D, 2-D, 3-D	One-, two- or three-dimensional, referring in this thesis to spatial dimensions
NISQ	Noisy Intermediate-Scale Quantum
CPU	Central Processing Unit, in this thesis used interchangeably with conventional classical computers
GPU	Graphical Processing Unit
HPC	High Performance Computing, dedicated classical resources
HF	Hartree-Fock
CI	Configuration Interaction
CNOT	Controlled-NOT gate
SO-FT	Split Operator Fourier Transform
QFT	Quantum Fourier Transform
JW	the Jordan-Wigner transform
ITE	Imaginary Time Evolution
VQE	Variational Quantum Eigensolver
VQD	Variational Quantum Deflation
U(p)CC	Unitary (paired) Coupled Cluster
QSP	Quantum Signal Processing
QET-U	Quantum Eigenvalue Transform with Unitaries
LCU	Linear Combination of Unitaries

The fundamental laws necessary for the mathematical treatment of a large part of physics and the whole of chemistry are thus completely known,

and the difficulty lies only in the fact that application of these laws leads to equations that are too complex to be solved.

— Paul Dirac [1]

1 Introduction

Motivation

The ability to digitally simulate the behaviour of complex, microscopic quantum systems containing up to thousands of atoms – these might be large molecules or material alloys with a myriad of different elements – allows us to investigate, predict and ultimately design the properties of novel substances in a rapid, cost-effective manner. This is important for lowering the resource needed and shortening the time it takes when exploring new materials and chemicals, which have the potential to solve some of the greatest challenges facing humanity today. Indeed, the modelling of molecular properties was one of the first applications of early computers.

Overwhelming theoretical, algorithmic, software and hardware advances made in the computational sciences over the last century has transformed the way we perform these simulations and how the chemical, biomedical and material sciences are conducted. We can now simulate small molecules on laptop computers, even make predictions about large proteins on high-performance, remote-access computers. At times, sophisticated computer simulations can provide accurate synthetic data about quantum many-body systems *before* experimental verification can be obtained, either because the experimental set-ups are too dangerous or complex to perform, difficult to scale, or the experimenters are limited by measurement technologies

that are too slow or too low-resolution to observe the process. With increasingly sophisticated artificial intelligence and machine learning methodologies, accurate simulations of chemical processes plus experimental data can even be used as valuable inputs to data-driven models, capable of even faster chemical property prediction and design. The importance and relevance of computational simulation in modern chemical and material research cannot be understated.

However, the conventional, 'classical' computer simulation of quantum many-body systems is fundamentally not scalable – this notion is succinctly captured by Dirac's quote in the start of this Chapter. The computer memory and time necessary to store and simulate some of the most accurate digital representations of molecular systems grows exponentially with the number of constituent atoms, and there is not enough computational capacity on planet Earth to even encode some of the largest molecules of key scientific and industrial import. While many modern methods are able to leverage abstractions and simplifications *via* e.g. recognising that not *all* quantum many-body systems have the same degree of particle correlation and quantum entanglement, thereby side-stepping the need for full-quantum representations and treatment in simulations, *any* molecular processes involving a transfer of energy are quantum in nature, and accurate simulations that account for quantum phenomena will be meaningful and useful.

One possible solution to address this scaling bottleneck, first proposed by Feynman, is to model a quantum system of interest by simulating and probing an analogous quantum system which we can control. After all, molecular processes 'simulate' their own physics in real time with no apparent scaling issues. Subsequent researchers suggested that, in theory, a digital computer consisting of bits that are quantum particles, or *qubits*, can substantially outperform classical computers in the speed and scalability of quantum many-body simulations, along with promises of speedups in other tasks including database search and decryption.

Thus began the race to build *quantum computers*. Only within the last couple of decades did researchers have all the necessary technology to manufacture the first qubits, then build the first quantum computers – at the time of writing,

these cumbersome machines can take up entire rooms, and use a combination of sophisticated materials engineering plus significant energy to cool, protect and manipulate the few extremely delicate, noise-sensitive qubits at their core. The scene resembles that of computer science in the 1950-60's; with the exception that quantum computer scientists today do not have the luxury of packing more transistors onto a chip by 'shrinking' their size as in conventional computing, since they are already operating at the very limit of atomic or subatomic physics.

Given the immense financial, time and human resources that have been and will be invested into building quantum computers, one might expect a commensurate pay-off beyond the scientific and research value of the endeavour. This is generally expected to start from a demonstration of *quantum advantage* or *quantum utility*, the idea that a quantum computer decisively outperforms a classical computer on a task that is of genuine practical value. The last half a decade saw a number of quantum supremacy claims (quantum computers shown to be faster than a conventional computer on select contrived tasks), even claims of quantum utility and quantum advantage in quantum many-body simulation tasks and beyond. Virtually all of these claims are closely followed by demonstrations which show that these tasks could, in fact, have been solved on a conventional computer with fewer resources than previously anticipated, or that the actual utility and practicality of these tasks are questionable. Most researchers in the field of quantum computing and quantum many-body simulations will therefore likely agree that, as yet, no convincing demonstration of quantum advantage in simulation has been performed.

A natural core question then becomes, "*what will be the best way to use quantum computers for the purpose of quantum many-body simulation, to demonstrate quantum advantage?*" This involves understanding the nature of 'unsolved' problems in the realm of chemical and material simulations, the methodologies that conventional computers struggle with and quantum computers are expected to excel at, the properties and limitations of current or future quantum technologies, and accordingly designing quantum computing algorithms with these insights. This core question is what I seek to address in this thesis.

Contribution

In this thesis, I explore how we might exploit quantum computers from different stages of technological development to simulate the physics of many-body systems, and report a number of contributions I have made to the field of quantum algorithms for this purpose – with particular interest in applications to chemistry and molecular physics.

This includes my work targeting the use of near-term quantum computers: Using spin-restrictions in variational quantum methods to reveal molecular excited states of a particular symmetry [2] and unpublished work I did in an industrial collaboration with the company Phasecraft on using techniques from 'shortcuts to adiabaticity' to improve variational ansatze circuits. It also includes work I did in collaboration with the company Quantinuum on exploiting small but error-corrected quantum computers to perform non-unitary or imaginary time evolution, useful in preparing ground states and thermal states [3]. In addition, a chapter is dedicated to my efforts exploring how we might use future fault-tolerant quantum computers for simulating molecular processes - specifically, the notion of encoding a real-space grid representation of many-body wave functions directly into a quantum computer, and how that might unlock new insights [4]. Finally, I also detail a new method called 'shadow spectroscopy', of which I am a co-creator [5]. It combines a recently-developed, highly-efficient quantum measurement protocol with quantum simulation to reveal energy difference information in a Hamiltonian, and can be used in any era of quantum technology.

If I have seen further it is by standing on the shoulders of Giants.

— Isaac Newton [6]

2 Background

Contents

2.1 Digital Quantum Computers and their Emulators . . . 5
2.1.1 Emulators 8
2.2 Quantum Simulation of Many-Body Physics 9
2.2.1 Motivation for Quantum Computing 11
2.2.2 Key Concepts for Simulation with Qubits 13
2.2.3 Representation of Quantum States 16
2.2.4 State Preparation 25
2.2.5 Quantum Simulation 35
2.2.6 Measuring Useful Information 45

2.1 Digital Quantum Computers and their Emulators

A quantum computer contains arrays of controllable quantum particles – these *qubits* (2-level units) or *qudits* (multi-level units) might be photons, ions, nano-scale particles or subatomic particles embedded in a substrate. *Analog quantum computers* can only apply a limited set of control operations onto the qubits or qudits, and are unable to realise general-purpose computation. These machines promise to be better than conventional computers for a restricted set of applications – the two most prominent being cost-function optimisation via quantum annealing, and the

simulation of some quantum many-body systems – see Ref. [7] for a broad overview of the former and Ref. [8] for the latter. On the other hand, *digital quantum computers* typically use qubits instead of qudits, and well-defined sets of operations (quantum logic *gates*) are combined to perform any transformation on the qubits to arbitrary precision, forming the basis of *universal* quantum computing – see Ref. [9] for a comprehensive reference for quantum computing. Many *quantum algorithms* that run on digital quantum computers have been proposed, some of which promise to be significantly faster than the best equivalent classical method, and a subset of those are meaningful for solving useful, high-value problems – see the following surveys summarising some of these most-researched quantum algorithms [10, 11]. This thesis specifically concerns the potential for digital quantum computers to encode and subsequently simulate *any* quantum many-body system we might be interested in. For the remainder of this thesis, I will only be discussing simulation in the digital quantum computing context.

Digital quantum computing technology is expected to progress in 3 broad stages:

1. Noisy Intermediate Scale Quantum (NISQ) computers: These are quantum computers that have a limited number of qubits with high error rates and can only sustain quantum coherence for a short amount of time. Quantum computers with up to ~ 100 physical qubits will fall into this category – this includes quantum computers that exist at the time of writing. As no error correction can be meaningfully implemented, proposed quantum algorithms and applications that target this hardware seek to circumvent the noise that inevitably results; a variety of quantum resource cost reduction methods and *quantum error mitigation* techniques (see e.g. Ref. [12]) are used for this purpose. See Ref. [13] for a summary of NISQ-targetted algorithms.

2. Early Fault-Tolerant Quantum computers: *Quantum error correction codes* can be used to detect and correct any errors that may have occurred during the execution of qubit operations – these methods typically use multiple physical qubits to encode the information of a single logical, error-corrected

qubit (see the following guide [14] for more on quantum error correction). It is conceivable that quantum technology can reach a level where there are sufficiently many low-noise physical qubits ($\sim$ 1000's of them) to implement robust error correction [15], but the overhead necessary to do so is so high that only a small number of error corrected qubits can be implemented in practice [16]. These computers we refer to as early fault-tolerant quantum computers.

3. Fully Fault-Tolerant Quantum computers: If engineering challenges can be overcome and there are no unforeseen fundamental physical barriers to scaling, we might be able to build quantum computers that will have more than $\sim$10,000's to 100,000's of physical qubits. In this regime, coherence can likely be sustained for extended periods of time, and it will be possible to encode information on a single quantum computer that will exceed the memory capacity of most conventional high performance computing facilities. We are most likely to perform some select computations that will be impossible to replicate on conventional computers with reasonable time and cost using fully fault tolerant quantum computers.

At the time of writing, the largest commercially available digital quantum computers include IBM's Condor processor (1,121 superconducting qubits) [17] and Atom Computing's quantum processor (1,125 neutral atom qubits) [18]. To additionally account for how noise-prone the qubits are, and therefore how complex of a quantum circuit we would be able to perform without noise corruption, some in the industry of fabricating quantum computers have used other metrics to benchmark them. For example, the reader may have come across the quantum volume, a product of the number of qubits in a quantum computer and the maximum circuit depth the computer can execute noise free; the latter is inversely proportional to the number of qubits and the effective error rates of those qubits. While most stakeholders believe that there are no insurmountable barriers that stop us from scaling the number of

qubits and reducing their error rate, there are differing opinions on which types of qubits and architecture will ultimately perform best, and how rapidly we can scale.

As I have alluded to, one of the key milestones to reach fault tolerance is the ability to perform error correction. Significant effort is underway to reduce the resource overhead necessary to do so. Some recent highlights include; research from Google experimentally demonstrated a well-known error correction code called a 'surface code' that was sufficiently scalable to suppress errors, even as the number of qubits increased [19], and the construction of a quantum processor based on logical qubits by researchers at Harvard, built using an array of 280 neutral atoms [20]. See Ref. [21] for more on recent advances in quantum error correction.

2.1.1 Emulators

Qubit emulation allows us to develop and test quantum algorithms on conventional high-performance computers before they are implemented on real quantum hardware (which might have limited availability), or before more powerful and stable quantum computers become available. These emulations can give us detailed control over the properties of the experiment e.g. we might choose to investigate the effects of noise on a quantum algorithm by accounting for various noise models in the emulation, giving us deep insights into how different algorithms might perform in practice, as well as how they might be improved. Recent advances in leveraging Graphical Processing Units (GPU) for accelerated processing of linear algebra has also led to substantial speed-ups of the emulation time. The only limitation of course is that the cost of emulation grows exponentially with the number of qubits, so metrics such as the exact gate costs to perform a given quantum algorithm might have to be extrapolated to large quantum computers.

There are a myriad of qubit emulators. They can largely be divided into two types – state vector/density matrix based emulators, or tensor network based emulators. The former expresses the quantum state of the qubits in either a vector or a density matrix, and can accurately evaluate the action of any arbitrary quantum gate in a circuit on the state via matrix algebra; the memory cost of storing a

state vector limits this to $\sim 30 - 40$ qubits. Tensor network emulators instead represent a quantum circuit in a graphical tensor and structure, and uses tensor network contraction to evaluate observable outcomes – this allows emulation of circuits and systems that display low entanglement with up to 100 qubits. Refer to a recent extensive benchmark which explored the performance of virtually all quantum emulators in the public domain on a number of representative tasks for a detailed survey into their differences [22].

Hitherto, the majority of the literature in quantum algorithms for quantum many-body simulation features, sometimes exclusively, the use of classical emulators of quantum computers, mostly because of their versatility and the lack of sufficiently robust quantum computers at the time of writing. Throughout this thesis, I will also be describing work that mostly uses qubit emulators, particularly when I explored quantum simulation techniques targeting quantum computers that do not yet exist.

2.2 Quantum Simulation of Many-Body Physics

In *ab-initio* quantum many-body physics simulations, we are generally concerned with solving the non-relativistic Time-Dependent Schrödinger Equation:

$$i\hbar\frac{\partial}{\partial t}\psi(\vec{x},t) = \mathcal{H}\psi(\vec{x},t) \tag{2.1}$$

where $\hbar$ is the Planck's constant (in all subsequent text I use atomic units where $\hbar = 1$), $\vec{x} = (\mathbf{x}_1, \mathbf{x}_2 \ldots \mathbf{x}_N)$ are spatial and spin coordinates of every constituent quantum particle in the system, ψ is the normalised, complex, many-body wave function defined by the coordinates of all N particles involved, and $\mathcal{H}$ is the Hamiltonian of the system. For the systems of interest in this thesis the Hamiltonian is a sum of two components:

$$\mathcal{H} = \mathcal{H}_{\text{kin}} + \mathcal{H}_{\text{int}}, \tag{2.2}$$

where

$$\mathcal{H}_{\text{kin}} = -\sum_{p=1}^{P}\frac{1}{2m_p}\nabla_p^2 \tag{2.3}$$

represents the kinetic energy of each of the P particles present. $\mathcal{H}_{\text{int}}$ instead encompasses all interactions – in many cases it is convenient to further resolve according to

$$\mathcal{H}_{\text{int}} = \mathcal{H}_{\text{U}} + \mathcal{H}_{\text{V}}, \tag{2.4}$$

where $\mathcal{H}_{\text{U}}$ represents single-particle interactions with e.g. classical fields, while $\mathcal{H}_{\text{V}}$ represents particle-particle interactions. For charged particles interacting via the Coulomb potential (electron-electron, nucleus-electron, nucleus-nucleus), we would write

$$\mathcal{H}_{\text{V}} = \sum_{p>q}^{P} \frac{q_{p,q}}{|\mathbf{r}_p - \mathbf{r}_q|} \tag{2.5}$$

for suitable charge constants q and spatial coordinates $\mathbf{r}$. In the case of atomic and molecular systems, we could opt to model each nucleus as a full quantum particle in its own right, in which case we may set $\mathcal{H}_{\text{U}} = 0$ unless there are external e.g. electric or magnetic fields which may vary in time; we do not discuss them here.

In most molecular and material simulations, to simplify the problem, the nuclei are not treated explicitly as quantum particles. We might opt to model only the electrons and employ classical fields to represent the M nuclear potentials originating at fixed locations $\mathbf{R}_m$ according to

$$\begin{aligned} \mathcal{H}_{\text{U}} &= \sum_{m=1}^{M} \sum_{p=1}^{P} \frac{Z_{p,m}}{|\mathbf{r}_p - \mathbf{R}_m|} \\ &= \sum_{p=1}^{P} \frac{Z_p}{|\mathbf{r}_p|} \quad \text{for } M = 1 \text{ nucleus at } \mathbf{R}_0 = \mathbf{0} \end{aligned} \tag{2.6}$$

in the single atom case for suitable values of nuclear charges Z_p or $Z_{p,m}$.

Formally, Eq. (2.2) admits the following initial value problem:

$$\psi(\vec{x}, t) = e^{-i\mathcal{H}t}\psi(\vec{x}, 0), \tag{2.7}$$

which falls within the realm of e.g. quantum dynamics, kinetics, scattering, spectroscopy and, if the Hamiltonian included time-changing potential, control and manipulation of a many-body system's dynamics. Alternatively, we might

choose to split the wave function into a stationary and time-dependent element and solve the static eigenvalue problem

$$\mathcal{H}\psi(\vec{x}) = E\psi(\vec{x}), \tag{2.8}$$

which falls within the realm of e.g. ground and excited state state energy determination, static potential energy surfaces and thermodynamics.

2.2.1 Motivation for Quantum Computing

For Hamiltonians with more than two interacting particles in 3D, no analytic solutions of the wave function exist. One has to therefore use numerical techniques, of which a great plethora have been proposed – mostly originating from theoretical chemistry and particle physics. Some solve the quantum many-body dynamics problem (see e.g. Ref. [23] for a good primer) and others the static energy eigenvalue problem (see e.g. Ref. [24]). For the most accurate *ab initio* (first principles) numerical methods that require little to no empirically-motivated inputs, classical computational time and memory grows exponentially with system size (the curse of dimensionality) – limiting them to either modestly sized molecules or heavily simplified abstractions of large systems. Techniques that have improved scalability and shorter absolute compute time find more compact ways of representing the system, but are generally lower in accuracy and require more chemical intuition and user inputs to obtain meaningful results. Physics-informed machine learning techniques and neural network methods have, in recent years, surfaced as promising means of augmenting existing approaches, although lingering hesitations regarding scalability, training data and explainability has not been decisively addressed.

The situation is not dissimilar to well-motivated engineering challenges where the underlying dynamics do not admit analytical solutions and their numerical simulation can become prohibitive; for example, the use of computational fluid dynamics simulations for the design and testing of mechanical wings and air-foils. Simplifying the system under study to lower dimension representations and reliance on expert input are often necessary for computations to begin in the first place, let

alone obtaining useful results. However, computational simulations can often be supplemented with experiments on scaled-down air-foil models placed in controlled wind tunnels. These experiments can provide plenty of insight, and are only limited by the size of the wind tunnel. Analog quantum computers can be thought of as parallels to wind-tunnel experiments; if we could manipulate a system of qubits such that it behaves 'as though' it was another many-body system, we could probe this experiment and obtain properties of the system being modelled, even if accurate computational simulations are out of reach. Digital quantum computing takes this idea further and blurs the line between simulation and experiment; a programmable set of qubits capable of realising any arbitrary unitary operation gives us access to the richness of quantum function evaluation and quantum algorithmic techniques, which can in principle be exploited to represent any quantum system we like, and any operation relevant to its simulation can be implemented.

To this end, the last half a decade saw a significant increase in research exploring the use of quantum computers for the simulation of quantum many-body systems, including both experimental and methodological advances that seek to accelerate the prospects of realising practical quantum advantage. I will refer the reader to the following detailed reviews on quantum simulation [25–27], its application in quantum chemistry and materials [28–32], and specifically methodologies for quantum dynamics or time evolution [33, 34] – we will see later that this plays a critical role in quantum simulation in general. Because quantum-assisted modelling is in essence performing a microscopic experiment, most methodologies can be framed into the following steps:

$$\text{State Preparation} \quad \rightarrow \quad \text{State Driving/Simulation} \quad \rightarrow \quad \text{Measurement.}$$

For the remainder of this Chapter, I will provide a high-level overview of prior concepts involved in these three steps which the research in this thesis builds on. To start, I will briefly summarise some key quantum computing concepts involved in simulation.

2.2.2 Key Concepts for Simulation with Qubits

A qubit is typically visualised as a 2-component complex vector in Bloch sphere, with the poles corresponding to the computational basis states $|0\rangle$ and $|1\rangle$:

$$|\psi\rangle = a_0 |0\rangle + a_1 |1\rangle , \tag{2.9}$$

where a_0 and a_1 are the probability amplitudes associated with each computational basis state – the square of an amplitude equals to the probability that its corresponding basis state is observed when the qubit is measured. Single-qubit gates realise 2×2 unitary matrix transformations of the vector in Bloch sphere – they serve to generate superposition and modify the amplitudes (and thus possible measurement outcomes) of the qubit by creating interference between the two outcomes. Some important single-qubit gates I might mention throughout this thesis include the following:

- The Hadamard gate H, which prepares quantum superpositions of the basis states:
$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$
Specifically, it maps the $|0\rangle$ state to the $|+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$ state, and $|1\rangle$ state to the $|-\rangle = \frac{|0\rangle-|1\rangle}{\sqrt{2}}$ state.

- The Pauli gates $P \in \{X, Y, Z\}$, synonymous with Pauli matrices, which rotates the vector about the x, y, z axes of the Bloch sphere by π radians:
$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$
The X gate is equivalent to the classical NOT gate and applies a bitflip – directly affecting the probability amplitudes between the two outcomes. The Z gate applies a relative phase shift of π radians to the $|1\rangle$ state without changing the probability of the two outcomes. The Y gate in essence combines bit and phase flips.

- The Pauli gates can be generalised to rotations by arbitrary angles about the x, y, z axes of the Bloch sphere:
$$R_x(\theta) = \begin{bmatrix} \cos(\frac{\theta}{2}) & -i\sin(\frac{\theta}{2}) \\ -i\sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}, \quad R_y(\theta) = \begin{bmatrix} \cos(\frac{\theta}{2}) & -\sin(\frac{\theta}{2}) \\ \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}, \quad R_z(\theta) = \begin{bmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{bmatrix}.$$

- The phase or phase shift gate P_φ manifests relative phase shifts to the $|1\rangle$ state by arbitrary angles (generalisation of the Z gate):

$$P_\varphi = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\varphi} \end{bmatrix}.$$

  Some special cases of the phase shift gate include the Z gate ($\varphi = \pi$), the S gate ($\varphi = \pi/2$) and the T gate ($\varphi = \pi/4$).

The entanglement of multiple qubits to create state vectors;

$$|\psi\rangle = a_0 |00\ldots00\rangle + a_1 |00\ldots01\rangle + a_2 |00\ldots10\rangle + \cdots = \sum_{k=0}^{2^N} a_k |k\rangle \tag{2.10}$$

that grow exponentially in the number of qubits N is critical for modelling quantum many-body wave functions, which 'live' in exponentially-scaling Hilbert spaces. This entanglement is typically built up by the controlled-NOT (CNOT) gate, which performs a unitary operation on two qubits specified in the exact same way as a classical CNOT gate. A related multi-qubit gate relevant in simulation in the fault-tolerant regime might be the 3-qubit controlled-controlled-NOT (CCNOT) gate or the Toffoli gate. When combined, these single- and multi-qubit gates form groups of operations that may differ in simulability (if they can be classically emulated efficiently) or universality (if they can be used to express any arbitrary operation [1]). The former is critical in understanding which mechanisms actually make quantum many-body simulations hard beyond the curse of dimensionality. The latter unlocks the versatility of digital quantum computers and lets us program them for simulating any many-body system.

These quantum gates come together to form quantum circuits, which are the foundations of quantum algorithms. Quantum circuits and quantum algorithms are typically characterised in terms of their complexity, which are bounds on how the

[1] A nuanced point not directly relevant to quantum many-body simulation but has real implications on how quantum simulation is performed 'in practice': The ability to realise the continuous set of unitary operations is of course not feasible – this would imply essentially a set of infinite rotation gates. Specific phase gates e.g. the S gate and T gates are important here because we could approximate any (single-qubit) unitary to arbitrary precision efficiently with a finite gate set that uses these types of fixed phase rotations instead of arbitrary rotations – see the Solovay-Kitaev theorem [35]. Attempts to overcome experimental challenges relating to the numerical precision of the rotations include e.g. recent efforts in exploiting interpolation [36].

resource (in terms of e.g. the depth of the quantum circuit, or the number of times we might have to sample the circuit, and thus the runtime) might be expected to scale as a function of the number of qubits N. Throughout this thesis I will use the 'curly-O' nomenclature to represent computational complexity – for example, a circuit or method that scales quadratically with the number of qubits will be $\mathcal{O}(N^2)$. The complexity of quantum circuits and algorithms are instructive bounds that tells us where we might expect quantum advantage and to what extent. However, they do not reflect the expected absolute runtime of an algorithm, and there might be cases where the use of heuristics actually produces algorithms that scale much better than what the computational complexity would suggest – one should therefore be cautious when comparing different quantum algorithms by their complexities. Some quantum circuit primitives or '*motifs*' are often used in the context of evaluating quantum functions relevant in quantum simulation. I give a preface to these presently.

A *Pauli string* or *Pauli term* is a tensor of Pauli matrices operating on different qubits $\bigotimes_i P_i$; it is usually expressed as a list of the letters X, Y, Z with subscripts indicating which operator is applied (non-trivially) to which qubit e.g. $X_0Y_2Z_5$, omitting the identity operator. The number of qubits that a Pauli string acts on non-trivially is the *locality* of the Pauli string e.g. $X_0Y_2Z_5$ is a 3-local Pauli string. Any matrix acting on n-qubits can in principle be expressed as a linear combination of all up-to-n-local Pauli strings; this can be seen when we recognise that the Pauli matrices form a complete basis that span the space of single-qubit operators. Depending on how one encodes a quantum state, physically-relevant operators such as the many-body Hamiltonian can generally be expressed with a linear combination of Pauli strings that grows polynomially in the system size – i.e. $\mathcal{O}(\text{poly}(N))$.

Another group of quantum circuit operators that are relevant in simulation are the exponents of Pauli strings i.e. $e^{i\theta \bigotimes_i P_i}$, which can be thought of as a rotation by angle θ about different axes across multiple qubits. Exponents of Pauli strings can be implemented using a circuit structure called *Pauli gadgets* and will be referred to as such throughout this text (see Fig. 2.1).

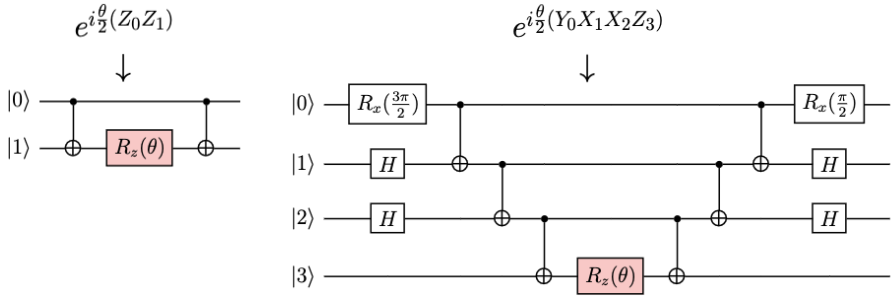


Figure 2.1: Example Pauli gadget quantum circuit primitives that apply the exponential of Pauli string matrix. The 2-qubit gates indicate CNOT gates.

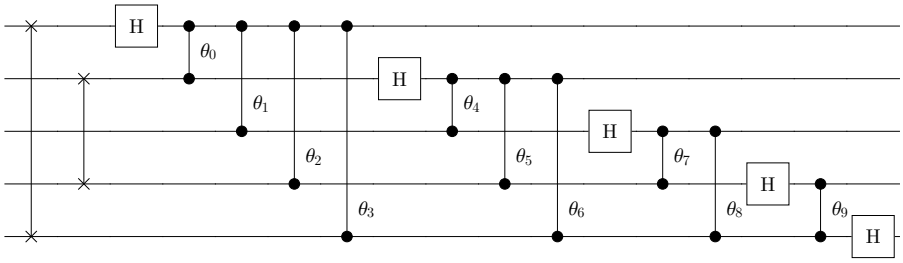


Figure 2.2: The circuit of the Quantum Fourier Transform on a 5-qubit system. The first two 2-qubit gates indicate SWAP gates. The subsequent gates are controlled phase shift gates.

The final type of relevant circuit subroutine is the Quantum Fourier Transform (QFT), which essentially performs a quantum analogue of the discrete Fourier Transform on an N-qubit state with 2^N amplitudes. The QFT can be implemented with the number of gates scaling as $\mathcal{O}(N^2)$, whereas the classical discrete Fourier Transform requires $\mathcal{O}(N2^N)$ – the QFT is therefore said to have an exponential speedup over its classical counterpart.

2.2.3 Representation of Quantum States

In using the aforementioned quantum circuit elements or primitives to perform quantum simulations, one would first need to settle on how to represent a quantum many-body system on a register of qubits. In the literature, the representation of many-body wave functions and the operators that act on them can be divided into two broad categories: first quantisation or second quantisation. First quantisation methods require explicit enforcement of particle symmetry, whereas second quantised methods maintain the correct exchange statistics through the mathematical

operators applied to the wave function [30]. My research spans quantum simulation methodologies that explore both representations.

Second-quantised representations

The second-quantised representation is typically used in conventional molecular and materials modelling. It is the foundation of many sophisticated, modern *ab initio* quantum chemistry methods that have had considerable success in the computation and prediction of molecular and material electronic properties. We start from the expression of the n-particle (typically Fermionic electrons) wave function in a basis of single-particle spin-orbitals or sites χ_i familiar to chemists. An approximate wave function or *ansätze* might be a *Slater determinant*; anti-symmetrised products of 1-electron spin orbitals that capture key aspects of the n-electron wave function such as the correct exchange symmetry and exclusion principle [37–39]:

$$|\psi(\mathbf{x}_1, \mathbf{x}_2, ...\mathbf{x}_n)\rangle = \frac{1}{\sqrt{n!}} \begin{vmatrix} \chi_1(\mathbf{x}_1) & \chi_2(\mathbf{x}_1) & \cdots & \chi_n(\mathbf{x}_1) \\ \chi_1(\mathbf{x}_2) & \chi_2(\mathbf{x}_2) & \cdots & \chi_n(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \chi_1(\mathbf{x}_n) & \chi_2(\mathbf{x}_n) & \cdots & \chi_n(\mathbf{x}_n) \end{vmatrix} \quad \text{(Determinant form)}$$

$$= |\chi(\mathbf{x}_1)\chi(\mathbf{x}_2)...\chi(\mathbf{x}_n)\rangle \quad \text{(Bra-ket notation)}$$

where the set of N single-electron spin orbitals $\{\chi_i(\mathbf{r}_i, \omega_i)\}$ is a product of $N/2$ spatial/molecular orbitals (MOs) $\{\varphi_i(\mathbf{r}_i)\}$ and the spin-up and spin-down functions $\alpha(\omega), \beta(\omega)$:

$$\chi_i(\mathbf{x}_i) = \begin{cases} \varphi_i(\mathbf{r}_i)\alpha(\omega) \\ \varphi_i(\mathbf{r}_i)\beta(\omega) \end{cases} .$$

An equivalent representation of a Slater determinant is a 'tally' of whether a given spin orbital contains within it an electron:

$$|\psi(\mathbf{x}_1, \mathbf{x}_2, ...\mathbf{x}_n)\rangle = |f_{2N-1}, f_{2N-2}, \ldots, f_2, f_1\rangle = |b\rangle \tag{2.11}$$

where $f_i \in \{0, 1\}$ depending on whether the spin orbital i is populated by an electron; we can therefore expect b to be a binary string with N digits and n ones

in it. In the quantum field theory literature this is often referred to as an n-electron Fock state, and the space of all possible number of particles distributed in all possible sites is called the Fock space. Any wave function of suitable symmetry, and indeed any eigenstate of the system Hamiltonian $\mathcal{H}$, can be expressed as a linear combination of Fock state bases:

$$\psi = \sum_{b=0}^{2^N-1} a_b \left|b\right\rangle . \tag{2.12}$$

Although the above Equation accounts for all possible combinations of N-digit binary numbers, for an n-electron system the amplitudes a_b are 0 when a Fock basis state b does not have n number of ones in its binary representation. The form of the MO's are generally computed using the well-established, iterative Hartree-Fock (HF) method. It is a self-consistent approach that finds the lowest-energy individual-electron wave function within the mean-field of all other electrons in the system. The HF Slater determinant is one where all the electrons populate the lowest energy orbitals. In general, the HF Slater determinant would contribute the most to the ground state wave function if it was expanded in the Fock space basis. See Refs. [37, 40] for more details on Hartree-Fock.

Second quantisation captures the essence of the Slater determinant using operators. The occupancy of spin orbitals are described by a set of *creation operators* $\{a_i^\dagger\}$, and elimination of electrons from orbitals by the complementary set of *annihilation operators* $\{a_i\}$. Specifically, a creation operator populates the spin orbital χ_i with an electron:

$$a_i^\dagger \left|\quad\right\rangle = \left|\chi_i\right\rangle , \tag{2.13}$$

and an annihilation operator a_i de-populates an electron from a spin orbital:

$$a_i \left|\chi_i\right\rangle = \left|\quad\right\rangle . \tag{2.14}$$

Therefore any N-electron Slater determinant in the second quantised form can be represented as:

$$a_i^\dagger a_j^\dagger \cdots a_n^\dagger \left|\quad\right\rangle = \left|\chi_i\chi_j\cdots\chi_n\right\rangle . \tag{2.15}$$

Orbital order is important and the operators do not commute. It follows that:

$$a_i^\dagger a_j^\dagger = - a_j^\dagger a_i^\dagger$$

$$a_i a_j = - a_j a_i$$

which obeys the anti-symmetric principle of Fermions, without explicitly retaining antisymmetry in the wave function. It also follows that:

$$a_i a_i = a_i^\dagger a_i^\dagger = 0 \tag{2.16}$$

The first product means that the same electron cannot be annihilated twice. The second product means that no two electron can be created in the same spin orbital. This formalism represents a many-electron wave function which implicitly satisfies Pauli exclusion, and can be manipulated without knowledge of determinant properties.

Operators that act on Slater determinants are also specified by creation and annihilation operators. For example, the spin orbital occupation operator is given by:

$$\hat{n}_i = a_i^\dagger a_i$$

$$\hat{n}_i \left| f_{2N-1}, f_{2N-2}, \dots f_i \dots, f_2, f_1 \right\rangle = f_i \left| f_{2N-1}, f_{2N-2}, \dots f_i \dots, f_2, f_1 \right\rangle$$

which returns the occupation number (either 1 or 0) of the i^{th} spin orbital. The electronic Hamiltonian in second quantisation is:

$$\mathcal{H} = h_{\text{Nu}} + \sum_{p,q}^{N} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{p,q,r,s}^{N} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s \tag{2.17}$$

$$\text{(nuclei repulsion)} \quad h_{\text{Nu}} = \frac{1}{2} \sum_{i \neq j}^{M} \frac{Z_i Z_j}{R_{ij}} \tag{2.18}$$

$$\text{(1-electron integrals)} \quad h_{pq} = \langle p | \hat{h} | q \rangle = \int \chi_p^*(r_1) \left(-\frac{1}{2} \nabla_{r_1}^2 - \sum_{I=1}^{M} \frac{Z_I}{R_{1I}} \right) \chi_q(r_1) dr_1 \tag{2.19}$$

$$\text{(2-electron integrals)} \quad h_{pqrs} = \langle pq | \hat{h} | rs \rangle = \int \chi_p^*(r_1) \chi_q^*(r_2) \frac{1}{r_{12}} \chi_r(r_1) \chi_s(r_2) dr_1 dr_2 \tag{2.20}$$

(Z_i is nuclear charge, M is the number of nuclei, r and R denote electronic and nuclear separation coordinates). Here we invoke the Born-Oppenheimer approximation, decoupling the motion of the nuclei from those of the electrons. The 1-electron terms correspond to individual electron kinetic energies and their attraction to nuclei. The 2-electron terms correspond to inter-electron repulsion.

Second-quantised operators are useful for quantum simulation because there are well-documented transformations that map them to operators that can be implemented on quantum circuits. In the intuitive Jordan-Wigner (JW) encoding, the Fock space of a molecular system is straightforwardly mapped onto the Hilbert space of qubits i.e. the occupation number of the i^{th}-orbital f_i is mapped onto the $|0\rangle$ or $|1\rangle$ state of the i^{th} qubit (the binary string of Eq. (2.11) *is* the binary number represented by qubits). In general, the excitation operators in the JW encoding are mapped to

$$a_p^\dagger \mapsto Z_{:p} A_p^\dagger, \qquad a_p \mapsto Z_{:p} A_p$$
$$a_p^\dagger a_q \mapsto i A_p^\dagger Z_{q:p} A_q$$

where $Z_{:p} := \prod_{k=1}^{p-1} Z_k$, $A = (X + iY)/2$ is the qubit lowering operator and $Z_{q:p} = \prod_{k=p+1}^{q-1} Z_k$. When exploring second quantisation in this thesis, I will be referring exclusively to the JW encoding.

While not relevant for this thesis, alternative encoding methods exist (see e.g. Refs. [41, 42]). In the Parity encoding for example, instead of storing the occupation, the parity information of the particle modes are stored. The Bravyi-Kitaev encoding compromises between Parity and JW encoding, promising to have advantageous scaling and lower quantum circuit resources requirements for larger systems. The number of qubits needed in all of these methods typically scale linearly with the number of basis functions. Second quantised encoding based on real-space grids using finite-difference schemes [43] have also been proposed. Advanced second-quantised encodings that 'taper' the number of qubits or opt to represent only part of the many-body system further promise to unlock quantum simulation on more near-term quantum technologies (see e.g. Ref. [44]).

First-quantised representations

First quantised representations of quantum many-body systems are less frequently encountered in the simulation literature; they typically involve representing the individual particles as functions, and the tensor product of individual particles (after explicitly accounting for particle exchange antisymmetry) is the wave function of the system. It is generally believed that the qubit cost of these types of representation will scale more favourably compared to second-quantised methods. However, the number of qubits needed even to represent small systems in first-quantisation is substantial, and so second-quantisation is generally the representation of choice in the NISQ and early FT era. One specific first-quantised representation of quantum systems that I would like to dedicate some text to preface is the representation of the many-body state on a finite, discretised or tabularised *real-space grid*, which lies at the heart of the research I did for Chapter 5.

Consider a system of P quantum particles in d spatial dimensions, well-localised within a volume of L^d i.e. has negligible wave function contributions beyond the region. I refer to this as a *simulation box*. The wave function of one particle in this box can be discretised into a finite basis of v^d volumetric pixels or *voxels* – where the value of a pixel in the grid is the value of the wave function at the centre of that pixel. The wave function of the entire system is thus a tensor product of every particle, meaning the total number of pixels required to encode the system scales exponentially as v^{dP}. The basis of grids can represent a spectral, momentum plane-wave basis familiar to materials scientists, conventionally referred to as a finite basis representation (FBR), or its dual pseudo-spectral basis in real, physical space, also called the discrete variable representation (DVR). These representations are related by the (discrete) Fourier transform.

In the research I describe in Chapter 5, my collaborators and I explored the first-quantised expression of a P-particle system on a fault-tolerant quantum computer by partitioning the qubits into P registers, and each register is further divided into d spatial sub-registers each with n_r qubits – each particle is discretised into an evenly-spaced grid with 2^{dn_r} pixels or basis functions [45–59]. The coefficients

of the wave function expansion in this grid representation map directly onto the amplitudes of the qubit computational basis. The number of qubits in the register therefore scales linearly as $\mathcal{O}(dPn_r)$ and logarithmically with the number of grid basis functions. The favourable asymptotic scaling is one of the main advantages of first-quantized real-space grid based encoding.

For the purposes of this thesis, the FBR is the momentum plane wave basis:

$$\phi_k(x) = L^{-\frac{1}{2}} \exp\left(\frac{i2\pi kx}{L}\right) \quad \leftrightarrow \quad |k\rangle\,. \tag{2.21}$$

Here k is an integer indexing the possible momentum a particle could be in, and $|k\rangle$ refers to the computational basis state which, regarded as a binary string, corresponds to k. Defining $\rho = 2^{n_r-1}$ and noting that we have 2ρ basis states in our computer's sub-register, a natural choice for the allowed k is to run from $-\rho \leq 0 \leq \rho - 1$. Therefore a 1D single-particle state ψ would be represented by our sub-register $|\psi\rangle_x$ according to

$$\psi = L^{-\frac{1}{2}} \sum_{k=-\rho}^{\rho-1} a_k e^{i2\pi kx/L} \quad \leftrightarrow \quad |\psi\rangle_x^{\mathrm{KS}} = \sum_{k=-\rho}^{\rho-1} a_k\,|k\rangle\,. \tag{2.22}$$

The superscript KS denotes the k-space representation.

The dual of the spectral FBR space i.e. the DVR space can be generated on a quantum computer by applying, to each sub-register, a quantum Fourier transform (QFT) denoted by U_{QFT}. The sub-register as a whole will be transformed as

$$|\psi\rangle_x^{\mathrm{RS}} = U_{\mathrm{QFT}}\,|\psi\rangle_x^{\mathrm{KS}} = \sum_{n=-\rho}^{\rho-1} b_n\,|n\rangle\,, \tag{2.23}$$

where

$$b_n = \frac{1}{\sqrt{2\rho}} \sum_{k=-\rho}^{\rho-1} \exp\left(i\frac{n\pi}{\rho}k\right) a_k. \tag{2.24}$$

The superscript RS indicates the real-space representation. When we wish to return to the original, k-space representation we can simply use the inverse QFT:

$$|\psi\rangle_x^{\mathrm{KS}} = U_{\mathrm{QFT}}^{\dagger}\,|\psi\rangle_x^{\mathrm{RS}} = \sum_{k=-\rho}^{\rho-1} a_k\,|k\rangle\,, \tag{2.25}$$

where:

$$a_k = \frac{1}{\sqrt{2\rho}} \sum_{n=-\rho}^{\rho-1} \exp\left(-i\frac{k\pi}{\rho}n\right) b_n. \tag{2.26}$$

The basis functions represented by each computational basis $|n\rangle$ appearing in Eq. (2.23) are peaked at (but not strictly localised around) the spatial point $x_n = \frac{n}{\rho}\frac{L}{2}$. The upper left panel of Fig. 2.3 shows a plot for the case that $n_r = 6$, $n = 0$. Specifically, the inferred mapping is

$$\phi_n(x) = P^n_{n_r}(x) \quad \leftrightarrow \quad |n\rangle \tag{2.27}$$

with

$$P^n_{n_r}(x) = \exp\left(-\frac{i\pi x'}{L}\right)\sqrt{\frac{2}{\rho L}}\sum_{j=1}^{\rho}\cos\frac{\pi(2j-1)x'}{L} \tag{2.28}$$

where $x' = x - x_n$ with $x_n = \frac{nL}{2\rho}$, and $\rho = 2^{n_r} - 1$. The function $P^n_{n_r}(x)$, which I will now informally refer to as a *pixel function*, serves the role of an approximation or "smear" of the Dirac delta function $\Delta(x - x_n)$, the sharpness increasing as 2^{n_r} tends to infinity. The separation between the peaks of adjacent pixel functions, e.g. $x_2 - x_1$, is $\delta r = L/2^{n_r}$, and I now define the *spatial resolution* as the reciprocal of this quantity, i.e. the number of pixel functions per unit distance:

$$\delta r^{-1} = 2^{n_r}/L. \tag{2.29}$$

A suitable decomposition to represent any 1D wave function $\Psi(x)$ in a quantum register is therefore very intuitive:

$$\Psi(x) \approx \frac{C}{\sqrt{\delta r}}\sum_{n=-\rho}^{\rho-1}\Psi(x_n)P^n_{n_r}(x) \leftrightarrow \frac{C}{\sqrt{\delta r}}\sum_{n=-\rho}^{\rho-1}\Psi(x_n)\,|n\rangle\,. \tag{2.30}$$

That is, the amplitude of $|n\rangle$ is found simply by sampling the target wave function at the point x_n. Here C is a normalisation constant that will be close to unity providing that (a) the target wave function has negligible amplitude outside of the simulation box and (b) the target wave function varies slowly with respect to δr. Intuitively, one can think of these spatial basis functions $P^n_{n_r}(x)$ analogously to pixels as used in conventional digital photographs: the greater the number of spatial pixels or grid points, the more features of the wave function are adequately captured.

Whereas the k-space representation maps qubit basis states to wave functions with sharp values of k, in the dual representation the mapping is to wave functions that are not perfectly sharp around points in real-space. Nonetheless, because of

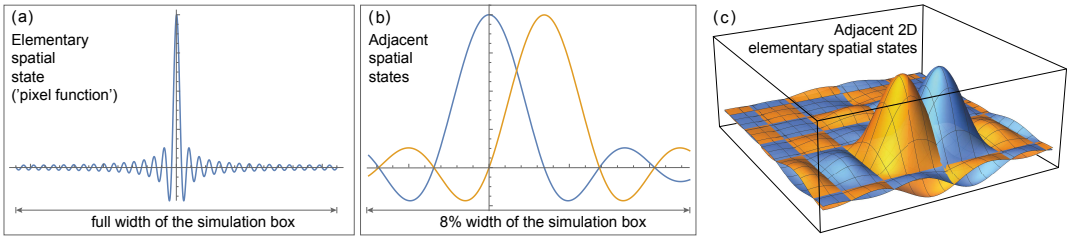


Figure 2.3: Basis functions in a Discrete Variable Representation (DVR) of wave functions in real space. **(a)** A 'pixel function' $P^n_{n_r}(x)$ at the spatial index $n = 0$, with $n_r = 6$ qubits (neglecting its complex phase, see Eq. (2.28)). **(b)** Adjacent pixel functions $n = 0$ and $n = 1$; where a function has its primary peak, all others are zero. **(c)** 2D pixel functions formed from products of 1D basis functions. Figure reproduced from Ref.[4] – created by Simon Benjamin and me.

their Dirac-like nature, one can analyse techniques and protocols *as if* they were Dirac delta basis functions, secure in the knowledge that in the high-resolution limit this becomes exact. This gives the visually intuitive picture that a particle wave function *is* a (pixelated) wave function in real-space, supported by a basis of sharp, evenly-spaced functions.

The real and k-space Cartesian grid dual representations are elegant options for quantum simulation. Firstly, this is because in the k-space representation, momentum operators, and therefore the kinetic parts of a many-body Hamiltonian $\mathcal{H}_{\text{kin}}$ are separable and exactly local (diagonal); at the same time, in the dual real-space representation, spatial operators and thus the potential parts of the Hamiltonian $\mathcal{H}_{\text{int}}$ are also separable and approximately diagonal. This reduces the complexity of implementing these operators, especially their time evolution, in these respective bases. Secondly, it is very efficient to switch between these two representations on a quantum computer using the quantum Fourier transform (QFT) U_{QFT}. The appealing conceptual simplicity over second quantisation can be regarded as another merit of first-quantised grid representation. The duality between the momentum- and real-space representations under the U_{QFT} has been explored in multiple studies that perform quantum simulation in first quantised grid representations, including, but not limited to, Refs. [47, 49, 52, 54, 55].

The generalisation to 2D or 3D is a natural one: The sub-registers tensor together to form the complete representation of a given particle. The 3D analogue

of Eq. (2.30)) is

$$\begin{aligned}\Psi(\mathbf{r}) &\approx C\delta r^{-\frac{3}{2}}\sum_{n,m,l=-\rho}^{\rho-1}\Psi(x_n,y_m,z_l)P_{n_r}^n(x)P_{n_r}^m(y)P_{n_r}^l(z)\\ &\leftrightarrow C\delta r^{-\frac{3}{2}}\sum_{n,m,l=-\rho}^{\rho-1}\Psi(x_n,y_m,z_l)\left|n\right\rangle\left|m\right\rangle\left|l\right\rangle. \end{aligned} \tag{2.31}$$

To represent a P-particle wave function $\Psi(\mathbf{r}_1,...,\mathbf{r}_P)$ we need only extend in the natural fashion,

$$\Psi(\mathbf{r}_1,...,\mathbf{r}_P) = C\delta r^{-\frac{3P}{2}}\sum_{\{n_1..l_P\}=-\rho}^{\rho-1}\Psi(x_{n_1},y_{m_1},...,z_{l_P})\left|n_1\right\rangle\left|m_1\right\rangle...\left|l_P\right\rangle. \tag{2.32}$$

2.2.4 State Preparation

After deciding how to represent a many-body system on a register of qubits, the first step to its simulation might be to prime the qubits and prepare a relevant initial state.

State Loading

To start, we might be interested in preparing the state of an n-qubit register (its probability amplitudes) according to a predefined analytic function. This is most relevant in e.g. the real-space grid simulation scenario, where the analytic functions represent single particle functions in real space. While I never directly used these techniques in my research (the use of emulated quantum computers means that I could simply 'load' a target state vector into the emulator's memory), I will briefly review some of the literature that describes methods and costs for real quantum processors.

In 2009, Ward *et al.* [60] extended the work of Zalka [46] to develop such a method for generating physically relevant states in grid representations, which I summarise here. We begin with the task of loading suitable initial states into the particle quantum register. The objective is to distribute a single particle state ϕ centred about (x_0,y_0,z_0) within a simulation box of width L across the amplitudes of an N-qubit register

$$\left|0\right\rangle\left|0\right\rangle\left|0\right\rangle \Rightarrow \phi(\mathbf{r})\left|n\right\rangle\left|l\right\rangle\left|m\right\rangle \tag{2.33}$$

where $\mathbf{r} = (n\delta r - x_0, l\delta r - y_0, m\delta r - z_0)$ and $\delta r = L/2^{n_r}$. L must be sufficiently large to capture the significant parts of the wave function and the sum of all the amplitudes squared should be unity. In 1998, Zalka proposed a constructive algorithm for loading single particle distributions onto a register [46], which was subsequently either rediscovered or generalised, most notably by Grover and Rudolph in 2002 [61], in a number of contexts [62–65]. The method splits the norm of the single particle function N times across the N qubits by performing qubit rotations, the angles for which are determined by computing exponentially many integrals. Complex-valued functions will require additional procedures for including the phase in each amplitude. There are conflicting opinions on the efficiency of the method; some argue that the method has an exponential $\mathcal{O}(2^N)$ overhead [64, 66], while others claim that it can be efficient [67], perhaps using quantum computers to solve the integrals [60]. In recent years, more novel approaches for initial state loading techniques have been proposed. Ollitrault *et al.* prepared Gaussian amplitudes using hybrid variational algorithms [52]. Holmes and Matsura used matrix product states to represent functions, which maps to state preparation circuits with linear-depth [66]. In particular, a noise-resistant technique for efficient preparation of distributions from Rattew *et al.* [68, 69] shows promise for generalisation to the single particle functions required. Recently, Matsushita *et al.* exploited the relationship between Slater and Lorentz functions to perform probabilistic encoding of the desired analytic amplitudes efficiently [70].

I will conclude this Section with a graphic summarising the second-quantised, Fock-space representation encoding and the first-quantised, real-space representation encoding used in this work (see Fig. 2.4).

Symmetrisation

Assuming we have multiple qubit registers that each correspond to individual particles in the first-quantised setting, we would want to encode the correct particle symmetry to the registers. Protocols for creating an antisymmetric superposition on a quantum computer have been explored in several publications dating back

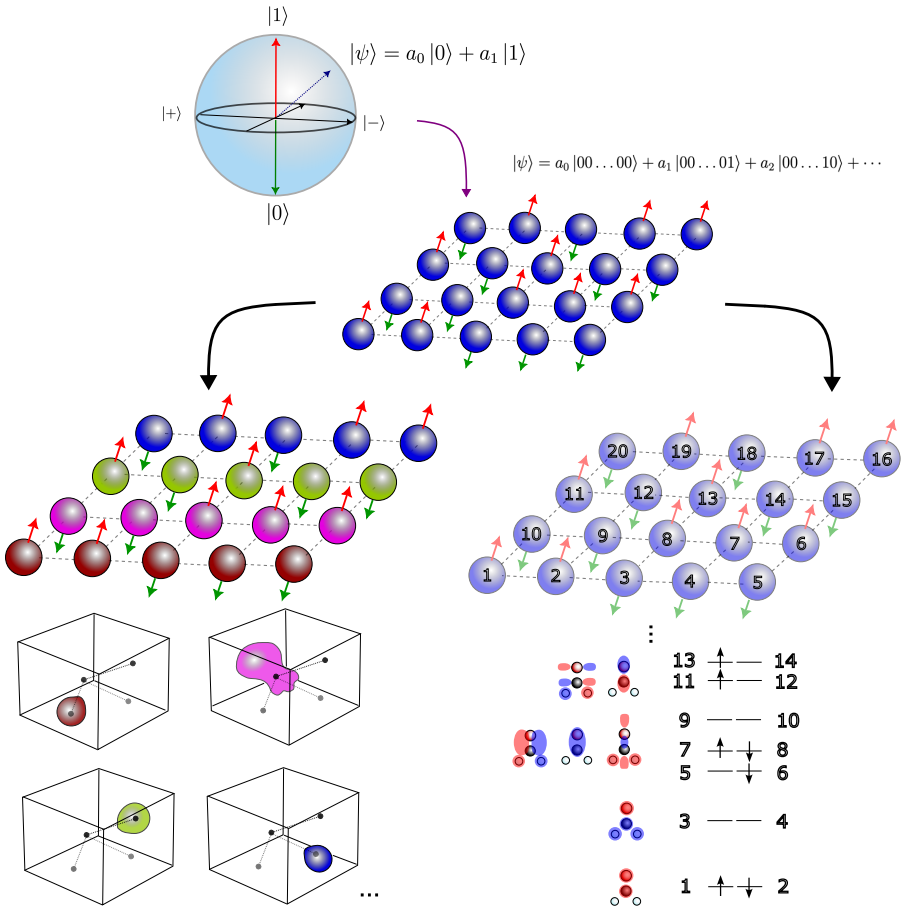


Figure 2.4: Quantum many-body encodings on qubits used in this thesis. On the left, subregisters of qubits are used to represent particle functions in real space in first quantisation. On the right, the individual qubits encode the electron count of molecular spin orbitals in second quantisation. Original graphic created by me.

at least as far as Abrams and Lloyd in 1997 [71]. The relatively-recent 2018 paper of Berry *et al.* [72] describes a deterministic, poly-logarithmic algorithm to create a state of the form,

$$|\text{prm}\rangle \equiv \frac{1}{\sqrt{P!}} \sum_{i}^{P!} \text{perm}_i(|0\rangle\,|1\rangle\,...\,|P-1\rangle). \tag{2.34}$$

Here we have P particles thus P sets of qubits, each of size $s = \lceil \ln P \rceil$, which represent the binary numbers 0 to $P-1$. The notation $\text{perm}_i()$ indicates the i^{th} permutation, and is understood to include the sign appropriate to antisymmetry (that is, $(-1)^k$ if there are k pair-wise swaps needed to permute from the canonical ascending order).

The state $|\text{prm}\rangle$ is a superposition of binary labels, rather than a superposition of states in the grid-based representation as we require. In principle we can obtain the latter once we have the former. Specifically, we prepare $|\text{prm}\rangle$, and distribute each of the P sets of qubits as follows: Each will occupy the first s qubits of one of our P particle-representing registers, while the remaining qubits within each particle register are in state zero, $|\underline{0}\rangle$. We now apply to each register an operator that maps these integer numbers to the desired grid-based eigenstates as follows:

$$M_{\text{unpack}}\left(|j\rangle\,|\underline{0}\rangle\right) = |\psi_j\rangle\,,$$

where ψ_j with $j = 0...(P-1)$ are the single-particle basis states from which we are building the initial state. While this is conceptually simple, the operator M_{unpack} may be challenging to implement since it must encode knowledge of every ψ_j and moreover it should be (near-)deterministic if it is to succeed over all P particle registers.

Instead one might wish to begin by preparing a simple product state across our particle-registers, i.e.

$$|\psi_0\rangle\,|\psi_1\rangle\,...\,|\psi_{P-1}\rangle\,, \tag{2.35}$$

and then find some means to antisymmetrise this, thus generating a Slater determinant [24]. The preparation of each state $|\psi_j\rangle$ would be relatively easy, performed independently for each particle register and optionally involving non-deterministic approaches with a high success probability per-attempt. This route is clearly attractive, but relies upon an efficient means of subsequently antisymmetrising the product state in Eq. (2.35).

Second Quantised Slater Determinants as Initial States

When a register of qubits instead is representing the Fock state i.e. the population of molecular spin orbitals as in the JW encoding in second quantisation, a Slater determinant is trivial to prepare on a quantum computer - it is simply a binary string representing which orbital is occupied, and is generated by applying the bit flip gate (Pauli X gate) on the corresponding qubit. The HF Slater determinant

of an n-electron system is therefore the binary number representing the number $2^n - 1$. Slater determinants corresponding to excited states (i.e. placing electrons at higher energy orbitals) can likewise be trivially generated by placing X gates on other qubits. In this thesis, I regularly use this technique to prepare the initial state of a second-quantised simulation. Explicit antisymmetrisation is not necesary as that is imposed by the second-quantised operators.

In many scenarios, a single Slater determinant, specifically the HF state, is the dominant contribution and an adequate starting point that captures the essence of a many-body system's elusive ground state. This is not always the case; the HF reference alone is not a good ansatz for systems that have e.g. strong electron correlation, and multiple Fock state configurations would contribute if the quantum simulation task was an eigenvalue problem e.g. to prepare the ground state. It is worth mentioning at this point that some researchers [73] have highlighted unresolved ongoing challenges and uncertainties associated with the preparation of suitable initial states specifically in the task of ground state preparation on quantum computers, which threatens the promised quantum advantage in many-body quantum simulation.

It has been proposed that using chemical insight and intuition to construct good initial states beyond a single Slater determinant will help. This might involve e.g. restricting the initial state to respect certain symmetries such as the total angular momentum or spin of the particles in the system; spin is characterised by the $\hat{S}^2$ and $\hat{S}_z$ operators, and any single Slater determinant will be an eigenfunction of $\hat{S}^2$ but not the $\hat{S}_z$ operator [74]. It is possible to construct initial states that are superpositions of Slater determinants that give spin-adapted or spin-coupled initial states, which are better starting states especially for eigenvalue problems. I refer the reader to recent work that explores the preparation of initial states for chemical systems where a single Slater determinant might not be adequate [75]. In Chapter 6, I also reference related techniques for preparing superpositions of Slater determinants.

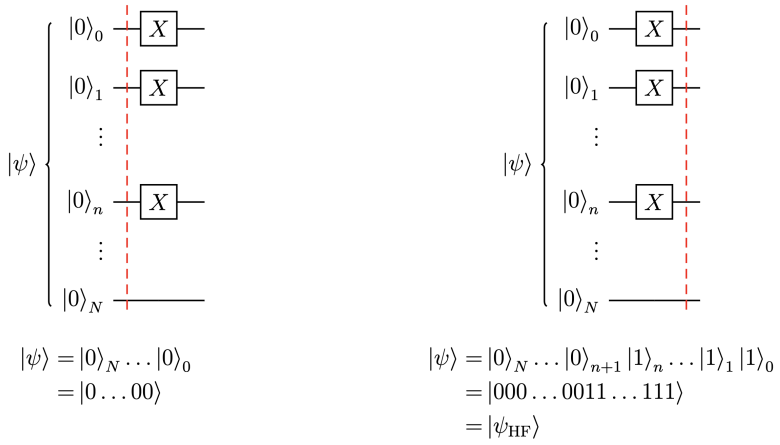


Figure 2.5: Initialising the HF state on a qubit register in second quantisation. The X gate is the classical equivalent of the NOT gate, and flips the qubit value corresponding to occupied orbitals from 0 to 1.

Variational Ansätze

In many-body simulation, we often want to drive the initial state we have prepared on a quantum computer towards another state by application of a quantum circuit. This is relevant in e.g. the final preparation steps of a state, or in the special case of the preparation of a time evolved state (see next Section). Indeed, the preparation of analytic functions (earlier in this Section) on a register of qubits can be thought of as the driving of an initial state vector towards a desired state vector.

One of the most widely studied techniques in the current era of NISQ computing is the Variational Quantum Eigensolver (VQE). They are a class of quantum-classical hybrid algorithms that promises to bridge the gap between available NISQ computers and ground state preparation of challenging quantum many-body systems [76]. These methods approximate wave functions, or *ansätze*, with parameterised, unitary operations $U(\vec{\theta})$:

$$|\psi(\vec{\theta})\rangle = U(\vec{\theta})\,|\psi_{\text{ref}}\rangle \tag{2.36}$$

where the parameters $\{\theta_i\} \in [0, 2\pi)$ and $|\psi_{\text{ref}}\rangle$ might be an appropriate initial reference state; in second-quantised simulations, this might be the Hartree-Fock state (see Section 2.2.4). The ansatz unitary is then transformed into a parameterised quantum circuit (or ansatz circuit), which encodes the unitary operations onto

a register of qubits. The quantum circuit expression for a Hermitian operator corresponding to an observable e.g. a Hamiltonian $\mathcal{H}$, is applied onto the ansatz, and the expectation value of its corresponding observable E is extracted by repeated measurements of the qubits' values (I will explain how qubit binary measurements map to an expectation value in Section 2.2.6). An approximate ground state of $\mathcal{H}$ is obtained by invoking the variational principle, where the quantum circuit parameters are optimised until E is minimised:

$$E(\vec{\theta}) = \min_{\vec{\theta}} \left(\langle \psi(\vec{\theta}) | \, \mathcal{H} \, | \psi(\vec{\theta}) \rangle \right). \tag{2.37}$$

The update of $\vec{\theta}$ (given E values measured from a quantum computer) is performed on conventional computers running typical optimisation routines. Beyond solving static eigenvalue problems, variational ansätze can also be used to prepare time-evolved states (see variational quantum simulation [77]); I will talk about its use in my research in Chapter 6.

While enthusiasm has declined recently, some researchers still believe that VQE can potentially demonstrate quantum advantage on near-term NISQ hardware. This claim is supported by a number of notable VQE demonstrations on real quantum computers available today [78–85]. Yet a number of significant blockers that hinder the performance of VQE remains largely unresolved. VQE ansatz circuits that have shallow gate depths tend to use many circuit parameters (see e.g. the Hardware Efficient Ansatz [30]) – these ansätze are known to suffer from large optimisations costs due to relatively 'flat' optimisation landscapes (see barren plateaus [86, 87]). On the other hand, ansatz circuits that have more problem-informed structure tend to use fewer parameters, thus making the optimisation task easier; however, they may also involve deeper circuits, and even small systems can require a number of gates that existing quantum hardware cannot execute without major corruption by noise. Most contemporary research into VQE involves addressing these blockers. I talk about my research contributions in this direction in Chapter 3.

One of the first ansatz structures proposed is the variational Hamiltonian ansatz [88, 89]. It is motivated by the time-dependent Schrödinger Equation,

specifically application of the time-evolution unitary $e^{-i\mathcal{H}t}$. Given that a many-body system's Hamiltonian matrix can be decomposed into a sum of Pauli strings $\mathcal{H} = \sum_k a_k \bigotimes_j P_j^k$, we can construct a physically-motivated ansatz by exponentiation of each Pauli string, then replacing the time variable with ansatz parameters:

$$U(\vec{\theta}) = \prod_k e^{-i\theta_k \bigotimes_j P_j^k} \tag{2.38}$$

I will be using this ansatz in research described in the first half of Chapter 3.

In quantum-enabled electronic structure calculations, a popular, physically-motivated class of VQE ansatz is the Unitary Coupled Cluster (UCC). It is a variational and more robust alternative to the traditional coupled cluster method [90, 91]. Although prohibitively inefficient on a classical computer [92], the parameterised excitation and relaxation operators in the Trotterised UCCSD ansatz can be mapped onto Pauli gadgets which combine to prepare the wave function approximation efficiently on available qubits.

The UCC ansatz is given by Equation 2.39:

$$U(\vec{\theta}) = e^{\hat{T}(\vec{\theta}) - \hat{T}^\dagger(\vec{\theta})} \tag{2.39}$$

where the excitation operators are given by Equation 2.40:

$$\begin{aligned} \hat{T}(\vec{\theta}) - \hat{T}^\dagger(\vec{\theta}) \;=\; & \sum_{i,a} \theta_i^a (a_i^\dagger a_a - a_a^\dagger a_i) \\ & + \sum_{i,j,a,b} \theta_{ij}^{ab} (a_i^\dagger a_j^\dagger a_a a_b - a_a^\dagger a_b^\dagger a_i a_j) \\ & + \dots \end{aligned} \tag{2.40}$$

and $\{i, j\}$ indexes occupied orbitals and $\{a, b\}$ indexes unoccupied orbitals. The above unitary $U(\vec{\theta})$ containing a single exponent is not executable on a quantum computer. Instead, the Trotter formula [93] is invoked:

$$e^{A+B} = \lim_{\rho \to \infty} (e^{\frac{A}{\rho}} e^{\frac{B}{\rho}})^\rho \tag{2.41}$$

to approximate Eq. (2.39) by taking a Trotter step $\rho = 1$, sometimes referred to as the disentangled (factorised) UCC:

$$U_{t1}(\vec{\theta}) = \prod_m e^{\theta_m(\tau_m - \tau_m^\dagger)} \tag{2.42}$$

where m indexes all possible excitations, $\theta_m \in \{\theta_i^a, \theta_{ij}^{ab}, \dots\}$ and $\tau_m \in \{a_a^\dagger a_i, a_a^\dagger a_b^\dagger a_i a_j \dots\}$. It has been shown that specific orderings of operators in this product of unitaries can in principle exactly parameterise any state [91].

In practice, the excitations are truncated to only include single and double excitations. This UCCSD ansatz has been popular in the VQE for quantum chemistry literature and is typically the benchmark for more cost effective methods [93]. There are variants of this ansatz which include generalised excitations in $U(\vec{\theta}) = e^{\hat{T}(\vec{\theta}) - \hat{T}^\dagger(\vec{\theta})}$. For example, in the k-UpCCGSD ansatz, the unitary is replaced by a sequence of generalised single excitations and only paired double excitations (pairwise excitations from one molecular orbital to another):

$$\hat{T}(\vec{\theta}) - \hat{T}^\dagger(\vec{\theta}) = \sum_{p,r} \theta_p^r (a_p^\dagger a_r - a_r^\dagger a_p) + \sum_{p,r} \theta_{p_\alpha p_\beta}^{r_\alpha r_\beta} (a_{p_\alpha}^\dagger a_{p_\beta}^\dagger a_{r_\alpha} a_{r_\beta} - a_{r_\alpha}^\dagger a_{r_\beta}^\dagger a_{p_\alpha} a_{p_\beta}) \tag{2.43}$$

and k products of the unitary operator are taken to increase the flexibility of the ansatz:

$$U(\vec{\theta_1}, \vec{\theta_2}, ..., \vec{\theta_k}) = e^{\hat{T}(\vec{\theta_1}) - \hat{T}^\dagger(\vec{\theta_1})} e^{\hat{T}(\vec{\theta_2}) - \hat{T}^\dagger(\vec{\theta_2})} ... e^{\hat{T}(\vec{\theta_k}) - \hat{T}^\dagger(\vec{\theta_k})}. \tag{2.44}$$

The sequence of unitary operators are then translated into Pauli gadgets via the JW mapping of second quantised operators.

A lot of research has been dedicated to improving upon these ansatz unitaries [94], either to introduce more structure and reduce the number of parameters, or to reduce the number of gates needed. Some noteworthy ansatz construction schemes involve selecting operators to append to an ansatz based on predefined criteria, so that the ansatz circuit is kept as short as possible; I used one such approach, called ADAPT-VQE [95], to reveal excited states (see latter half of Chapter 3). See [96] for a review of some of the recent advances in ansatz construction.

Preparing excited states

An extension of the VQE called the Variational Quantum Deflation (VQD) uses a VQE-like mechanism for driving initial states towards excited states [97]. VQD starts with approximating the ground state $\left|\tilde{\psi}_g\right\rangle$ of an n-qubit system using a VQE

computation. It then optimises an ansatz circuit $\left|\psi(\vec{\theta})\right\rangle$ on another set of n qubits to find the next eigenstate using the method of Lagrangian multipliers; concurrently minimising the overlap between the ground state in addition to minimising the variational energy expectation:

$$\min_{\vec{\theta}} \left(\langle\psi(\vec{\theta})| \mathcal{H} |\psi(\vec{\theta})\rangle + \beta| \langle\tilde{\psi}_g|\psi(\vec{\theta})\rangle |^2 \right) \tag{2.45}$$

where β is a Lagrangian multiplier larger than the maximum energy gap that is to be calculated. To find additional eigenstates, the Lagrangian is the sum of overlaps between the currently optimimised ansatz and any previously converged ansätze:

$$\sum_k \beta_k| \langle\tilde{\psi}_k|\psi(\vec{\theta})\rangle |^2 \tag{2.46}$$

Most of the VQD schemes proposed hitherto use identical parameterised ansatz circuits $|\Psi(\vec{\theta})\rangle$ for preparing both ground and excited states[90, 97, 98]. To ensure these fixed-length ansätze have sufficient variational flexibility for preparing different eigenstates, symmetry breaking excitation operators are also needed beyond the UCCSD set to explore the relevant parts of the Hilbert space. This is achieved by inclusion of generalised excitation and relaxation operators in the ansätze, such as UCC Generalised Singles and Doubles (UCCGSD) [97] and k-UCC Generalised Singles and paired Doubles (k-UpCCGSD) [90]. The increase in the number of operators leads to a commensurate increase of Pauli gadgets in the ansatz preparation circuit, which would be impossible to prepare noise-free on NISQ quantum computers. Estimation of the overlap between two eigenstates on a quantum computer further amplifies this problem: Overlap calculations on quantum computers use either the vacuum test, which doubles the circuit depth, or the destructive SWAP test (see later in this Chapter), which doubles the number of qubits and has a large entangling CNOT gate overhead [97] (entangling gates require qubits to interact, and are particularly prone to noise). Other variational quantum algorithms for preparing excited states on a register of qubits use similar projective strategies (see e.g. Ref. [99]).

2.2.5 Quantum Simulation

The discussion so far has focused on the preparation of primarily 'static' initial states on a quantum computer. An important family of quantum algorithms involve the preparation of a quantum state that has experienced a finite period of time evolution – equivalent to performing time propagation on some initial state that has been prepared on a quantum computer. This is effected via application of quantum circuits that approximate the unitary $e^{-i\mathcal{H}t}$ onto the initial state $|\psi(0)\rangle$:

$$|\psi(t)\rangle = e^{-i\mathcal{H}t}\,|\psi(0)\rangle \tag{2.47}$$

($\mathcal{H}$ can be time-dependent; in this thesis I do not consider these scenarios). The time and memory cost of performing Eq. (2.47) on a classical computer grows exponentially with system size; thus, the ability to efficiently perform time evolution or Hamiltonian simulation links directly to the potential predictive power of simulating many-body dynamics on quantum computers – however, I will also explain later on in this Chapter why Hamiltonian simulation in fact forms the basis of many other quantum algorithms, including the quantum simulation of static properties.

In seeking to use a digital quantum computer as our simulator, one may adapt the rich set of numerical methods for solving the time-dependent Schrödinger equation on classical machines. The reader can refer to the following references for wavepacket methods [100–106], more compact techniques where single particle wavepackets represent multiple degrees of freedom e.g. the Multi-Configuration Time Dependent Hartree (MCTDH) method [107–110] and its many flavours [111, 112], as well as path integral methods [113], and trajectory-based approaches [114] which are getting the most attention in the computational chemistry literature currently. These different techniques have been applied to great effect in studying phenomena and experiments where quantum effects play a significant role in the underlying physics of many-body systems, including; proton-coupled electron transfer [115], isotopic effects [23], femto- and attosecond laser chemistry [116], vibronic spectroscopy [40], plasmonic chemistry [117] and quantum control[118].

Wavepacket time-propagation methods have been used for simulation since at least the 1980's. They are the perhaps the 'purest' forms of numerically integrating the time-dependent Schrödinger Equation, seeking systematic approximations to the time evolution unitary operator $e^{-i\mathcal{H}t}$ without necessarily invoking simplified representations of the state $|\psi\rangle$. While these methods suffer from the most prohibitive scaling in the classical computing setting, they can provide a structured framework for understanding the different ways the time evolution operator $e^{-i\mathcal{H}t}$ is approximated in the quantum computing setting. Refer to Ref. [34] for a detailed review of these techniques; I will only describe three general approaches here.

Trotterisation

To approximate the time evolution operator $e^{-i\mathcal{H}t}$, the total time t can be discretised into shorter time intervals δt such that

$$e^{-i\mathcal{H}t} = \underbrace{e^{-i\mathcal{H}\delta t}e^{-i\mathcal{H}\delta t}\dots e^{-i\mathcal{H}\delta t}}_{N \text{ times}} = U^N(\delta t). \tag{2.48}$$

The short time evolution operator is then approximated by invoking the Lie-Trotter-Suzuki product formula (or Trotterisation): Given that $\mathcal{H}$ can be broken into a sequence of individual terms:

$$\mathcal{H} = \sum_{j=1} a_j H_j \tag{2.49}$$

The Trotter-Suzuki Formula approximates the time evolution operator as:

$$e^{-i\mathcal{H}\delta t} \approx \prod_j e^{-ia_j H_j \delta t} + \mathcal{O}(\delta t^2) \tag{2.50}$$

Assuming that each $e^{-ia_j H_j t}$ can be efficiently compiled into a quantum circuit, the intuition is to simulate the complete operator by switching rapidly between the terms, thereby effectively having all the terms active at once [119]. The form of the circuit depends on the Hamiltonian and the representation of the system. The error is a function of the sum of the commutators of the individual terms H_j, thus Trotter-Suzuki methods are most useful for systems where Hamiltonian

terms (approximately) commute. I report the use of Trotterisation to time evolve Hamiltonians in the Pauli basis in Chapter 4 and Chapter 6.

A Trotterisation time evolution scheme of particular interest in this thesis, specifically regarding real-space simulation research that I report in Chapter 5, is the dynamic Fourier or split-operator Fourier transform (SO-FT) method, which dates back at least as far as the work of Feit and Fleck in 1976 [120–122]. The scheme is popular and often employed in small-molecule applications, predominantly in propagating nuclear wavepackets on pre-calculated static or time-dependent electronic potential energy surfaces. It has found uses in reactive scattering [123, 124] and excited-state chemistry [125–129], as well as modelling quantum processes like tunnelling [130, 131] and superfluid turbulence [132], and indeed quantum control for designing quantum devices [133, 134]. Application to the relativistic Klein-Gordon equation [135] and non-linear Schrödinger equation [136–138] extend the method's scope. Its success is partly attributed to the simplicity of its construction, as it only involves updating the vector describing the state, and partly to the low computational price paid for obtaining efficiently converged results [139].

Building on the works of Zalka [46] and Wiesner [45] from the mid 1990's, it was observed just over a decade ago by Kassal *et al.* [47] that it will be advantageous to store high-dimensional spatial grid representations of molecular wave functions on digital quantum computers using qubits that scale only linearly with the degrees of freedom. They also showed that, by exploiting the remarkably efficient quantum Fourier transform (QFT), a quantum version of the SO method (the SO-QFT) will be able to simulate the dynamics of a chemical Hamiltonian in polynomial time. Indeed, the authors challenged the *status quo* of limiting the split-operator method to propagating nuclei on electronic potential energy surfaces.

Although the SO-QFT method is oriented towards future fault-tolerant quantum computing, it is nonetheless gaining increasing interest. The method has been investigated in the context of propagating one-dimensional Gaussian wavepackets [48], quantum harmonic oscillator [50], non-adiabatic dynamics across coupled, pre-computed potential energy surfaces [52], imaginary-time evolution for molecular

geometry optimisation [59] as well as ground state, Gibbs state and partition function determination [54]. I want to highlight works geared towards understanding the scaling of the SO-QFT at mid-to-large scale, specifically that of Jones *et al.* which looked at the resource necessary for a fault-tolerant implementation [49], the extensive investigation into first quantized basis-set simulations for chemistry of Su *et al.*, which also included block-encoding the SO-QFT Hamiltonian for simulation using qubitization and interaction picture techniques [53], and analysis by Childs *et al.* which provided concrete upper bounds for the gate complexity of SO-QFT to arbitrary order [55].

Linear Combination of Unitaries

Linear Combination of Unitaries (LCU) is a general quantum algorithmic technique. Quantum computing operations are naturally unitary, but the sum of unitary matrices is not necessarily unitary. To implement this linear combination, we can first use an approach called block-encoding to embed a sum of unitary matrices into the subspace of a larger matrix – in general, this would rely on ancillary qubits that are entangled to the unitary matrices in the LCU. The action of the LCU on a register of qubits is realised with a non-unitary probabilistic measurement – this typically involves performing a 'weak' measurement on the ancillary qubits used in the block encoding. LCU is useful for time evolution because the time evolution operator can be Taylor expanded [140]:

$$e^{-i\mathcal{H}\delta t} = 1 - i\mathcal{H}\delta t - \frac{\mathcal{H}^2\delta t^2}{2} + \ldots, \tag{2.51}$$

meaning we can approximate time evolution if we can essentially apply powers of the Hamiltonian (non-unitary operations) via LCU [141].

Time evolution via LCU has a circuit complexity which scales linearly with the propagation time. With the Truncated Dyson Series [142], LCU has also been shown to be very efficient at simulating time-dependent Hamiltonians. Interaction Picture Simulation techniques [143, 144] exploit this by transforming time-independent Hamiltonians into time-dependent Hamiltonians in the interaction frame. However,

LCU require the aforementioned block-encoding approach and thus ancillary qubits are generally required, which can become cumbersome. In this thesis, I explore single-ancilla LCU techniques that are useful for embedding Hamiltonian matrices in non-unitary quantum simulation techniques (see Chapter 4 and Chapter 5).

Quantum Signal Processing

Another group of quantum algorithms that uses weak ancilla qubit measurements is Quantum Signal Processing (QSP). It was observed that QSP provides a framework for unifying all quantum algorithms with known speedups [145]; in particular, it was pointed out that QSP promises to be the optimal approach in the long term for simulating Hamiltonian time evolution on quantum computers [146]. This is perhaps not surprising; the technique of expanding both real- and imaginary-time evolution operators in a basis of Chebyshev polynomial transforms, which is the central tenet in QSP methods, has been known since as early as the 90's to provide some of the most accurate quantum wavepacket dynamics simulations (see e.g. the Chebyshev scheme [102]) in chemical physics. However, the quantum resource cost for performing QSP can be substantial. For example, QSP typically uses an LCU subroutine to encode the Hamiltonian matrix [145], requring an entire register of error-corrected ancilla qubits with a size proportional to log of the number of expansion terms in the Hamiltonian and many associated entangling gates. This places the prospect of performing QSP well into the mature fault-tolerant quantum computing regime.

Specifically, QSP prepares matrix functions on quantum computers. Consider a sparse n-qubit ($2^n \times 2^n$) Hamiltonian $\mathcal{H} = \sum_k a_k \bigotimes_j P_j^k$, where $\bigotimes_j P_j^k$ are a polynomial number of Pauli strings. Given an n-qubit initial state $|\psi_0\rangle$ and an input model for embedding $\mathcal{H}$, QSP manifests a d-degree matrix polynomial transformation $\mathcal{P}$,

$$|\psi_0\rangle \rightarrow P(\mathcal{H})\,|\psi_0\rangle \tag{2.52}$$

up to some normalisation. It achieves this with a repeating circuit structure (a *QSP sequence*) that interleaves (1) a fixed *signal operator* – a large unitary matrix $W_x(\mathcal{H})$

which embeds $\mathcal{H}$ into its top left block, with (2) a sequence of *signal processing operators* – single-qubit rotations with classically pre-computed rotation angles or *phase factors* $\{\varphi_k\}$ that act on the embedded $\mathcal{H}$ and are responsible for realising the desired $\mathcal{P}(\mathcal{H})$. It can be shown that the polynomial transform is in fact applied directly onto the eigenvalues of $\mathcal{H}$ that have been embeded, hence it is also somtimes referred to as the Quantum Eigenvalue Transform (QET) for when $\mathcal{H}$ is Hermitian. The unitary of the QSP sequence is thus:

$$U_x(\mathcal{H}) = e^{i\varphi_0 Z} \prod_{k=1}^{d} W_x(\mathcal{H}) e^{i\varphi_k Z} = \begin{bmatrix} \mathcal{P}(\mathcal{H}) & * \\ * & * \end{bmatrix},$$
$$\text{where } W_x(\mathcal{H}) = \begin{bmatrix} \mathcal{H} & i\sqrt{1-\mathcal{H}^2} \\ i\sqrt{1-\mathcal{H}^2} & \mathcal{H} \end{bmatrix}.$$

Upon application of the QSP sequence circuit to $|\psi_0\rangle$, the signal processing ancilla qubit is measured in the computational basis. If it is in the $|0\rangle$ state, the post-selected n-qubit state should correspond to achieving the transformation of Eq. (2.52).

The versatility of QSP boils down to the fact that any analytic function of the Hamiltonian $f(\mathcal{H})$ can be arbitrarily well-approximated on a quantum computer, through the classical optimisation of phase factors $\{\varphi_k\}$ that generate a polynomial approximation of the transform $\mathcal{P}(\mathcal{H}) \approx f(\mathcal{H})$. The query complexity for performing sparse Hamiltonian simulation using QSP was found to be optimal [146].

In practice, additional non-unitary quantum operations might be necessary to implement block encoding of the Hamiltonian, which would use extra ancilla qubits that increase logarithmically with the number of terms in the Hamiltonian. The desired transformation of the Hamiltonian eigenvalues is thus only achieved if every ancilla qubit for LCU are measured to be in the $|0\rangle$ state – incurring an associated measurement or shots overhead. Some recent efforts seek to reduces this cost, using amplitude amplification techniques to boost the success of measuring the desired ancilla outcome (see e.g. Ref. [141]).

The number of phase factors in the QSP sequence is also linear in d; in general, higher-order approximations, thus deeper circuits, are necessary to perform the desired transforms accurately. Thanks to recent innovations, such as using a

symmetric set of phase factors [147], the cost for classical pre-computation of phase factors for high-degree polynomials is no longer an inhibiting step.

The action of the QSP sequence can be interpreted as a set of orthogonal rotations in the Bloch sphere of a single qubit [145]. We define a signal operator which is an R_z rotation by a fixed angle θ, and a set of signal processing operators which are R_x rotations each with different phase factors $\{\varphi_i\}$. The resulting amplitude of the $|0\rangle$ state upon application of this QSP sequence $\langle 0| U(\theta) |0\rangle$ is a linear combination of $\cos(2k\theta)$ (top left block of $U(\theta)$). Rewriting this expression using the definition of Chebyshev polynomials of the first kind $\cos(k\theta) = T_k(\cos\theta)$ shows that it is in fact an even polynomial $\mathcal{P}(\theta)$ in the basis of Chebyshev polynomials:

$$\mathcal{P}(\theta) = \sum_{k=0}^{d/2} c_k T_{2k}(\cos\theta), \tag{2.53}$$

where the coefficients c_k are determined by the phase factors, and the degree d is determined by the number of phase factors in the QSP sequence.

QSP algorithms which use few or even only a single additional ancilla qubit [148–152] have been proposed. These techniques circumvent the expensive LCU block-encoding model, and opt instead to use e.g. ancilla-controlled unitary Hamiltonian simulation (or the real-time evolution oracle) as the Hamiltonian input model. Its important to clarify here that although these techniques can be used to perform polynomial transforms of Hamiltonians, they would *not* be used to approximate Hamiltonian evolution especially when the time evolution oracle itself is used to embed the Hamiltonian. Using controlled unitary evolution this way is employed in many other early fault-tolerant state preparation algorithms, from standard routines such as the Hadamard test and Kitaev's iterative phase estimation, the Probabilistic Imaginary Time Evolution approach [153] (see next Section) and the rodeo algorithm [154, 155].

The recently-proposed Quantum Eigenvalue Transform with Unitary matrices (QET-U) approach is able to apply a desired polynomial transform of a Hermitian Hamiltonian matrix using only a single extra ancilla qubit [148]. It uses the

alternative W_z-convention of the QSP sequence, where the signal operator is diagonal, and instead R_x rotations are performed on the ancilla qubit:

$$U_z(\mathcal{H}) = e^{i\varphi_0 X} \prod_{k=1}^{d} W_z(\mathcal{H}) e^{i\varphi_k X}. \tag{2.54}$$

Replacing the block-encoding signal operator with the forwards-backwards real-time evolution oracle $e^{\pm i\mathcal{H}t}$ [153]:

$$W_z(\mathcal{H}) = \begin{bmatrix} e^{-i\mathcal{H}t} & 0 \\ 0 & e^{i\mathcal{H}t} \end{bmatrix} \tag{2.55}$$

enables the embedding of Hamiltonian eigenvalues for QSP simply by performing Trotter evolution; t is divided into regular time steps $t = N\delta t$ such that $e^{-i\mathcal{H}t} = (e^{-i\mathcal{H}\delta t})^N$, and each time step propagator is then partitioned into a product of propagators

$$e^{-i\mathcal{H}\delta t} = \prod_k e^{-i\delta t a_k \bigotimes_i P_i^k} + \mathcal{O}(\delta t^2), \tag{2.56}$$

for every term in the Hamiltonian decomposition. This avoids the additional qubit overhead for performing LCU block-encoding. Specifically, the transformation achieved after application of the QSP sequence is:

$$|\psi_0\rangle \to \frac{P(\mathcal{H})}{\sqrt{\langle\psi_0| \mathcal{P}(\mathcal{H})^2 |\psi_0\rangle}} |\psi_0\rangle = \frac{1}{\sqrt{\mathcal{N}}} \mathcal{P}(\mathcal{H}) |\psi_0\rangle , \tag{2.57}$$

where $1/\sqrt{\mathcal{N}}$ is a normalisation factor and $\mathcal{N}$ is the probability of success. The Hamiltonian function

$$\mathcal{P}(\mathcal{H}) = \sum_{k=0}^{d/2} c_k T_{2k}(\cos \mathcal{H}) \tag{2.58}$$

is a d-degree polynomial, which is a superposition of Chebyshev polynomials of $\cos \mathcal{H}$, and $\mathcal{P}^\dagger(\mathcal{H}) = \mathcal{P}(\mathcal{H})$.

The unitary nature of the real-time evolution oracle does imply that eigenvalue transform performed with QET-U is limited to Hermitian input matrices. Using Trotterisation to implement the time evolution oracle, while resource-friendly, introduces an additional Trotter error atop the polynomial approximation of the target function, but can be controlled with a combination of fine time-stepping

and high-order product formulae. Moreover, the probability of successful ancilla measurement $\mathcal{N}$ is sensitive to the initial state $|\psi_0\rangle$ and the desired transform. Post-selection implies that a low probability of success can incur a high shots overhead.

Finally, QSP requires that the embedded Hamiltonian is scaled by its operator norm; e.g. in LCU, every Hamiltonian term is normalised by the operator norm. In QET-U, this manifests as the restriction that any eigenvalue transform $\mathcal{P}(\mathcal{H})$ must have a range $[-1, 1]$, for a Hamiltonian with eigenvalues between the domain $[\eta, \frac{\pi}{2} - \eta]$, where η is chosen to be a small constant for improved optimisation of the phase factors. As a result, one must scale and shift the energy spectrum of the Hamiltonian $\mathcal{H}^{\text{scale}}$ such that:

$$E_0^{\text{scale}} = \eta \tag{2.59}$$

$$E_{N-1}^{\text{scale}} = \frac{\pi}{2} - \eta. \tag{2.60}$$

Optionally, one might also scale the range of $f(\mathcal{H})$. The authors of QET-U propose pre-computing (at least) the lower bounds of the ground state energy $\tilde{E}_0$ and maximum energy eigenvalue $\tilde{E}_{N-1}$, and using those to embed $\mathcal{H}^{\text{scale}}$. This is achieved with an additional global phase σ and scaling of the evolution time t of the time evolution oracle:

$$e^{-i\mathcal{H}^{\text{scale}}} = e^{-i(\mathcal{H}t+\sigma)}, \quad \text{where} \tag{2.61}$$

$$t = \frac{\pi - \eta}{\tilde{E}_{N-1} - \tilde{E}_0} \tag{2.62}$$

$$\sigma = \eta - \tilde{E}_0 t. \tag{2.63}$$

In Chapter 4, I will be exploring the use of single-ancilla QET-U for imaginary time or non-unitary evolution.

Non-unitary Evolution

The discussion so far has focused on approximating the real time evolution unitary operator $e^{-i\mathcal{H}t}$, the time and memory cost of which grows exponentially with system size when performed on classical computers. A similar challenge affects the preparation of a state that has been subjected to the non-unitary imaginary time

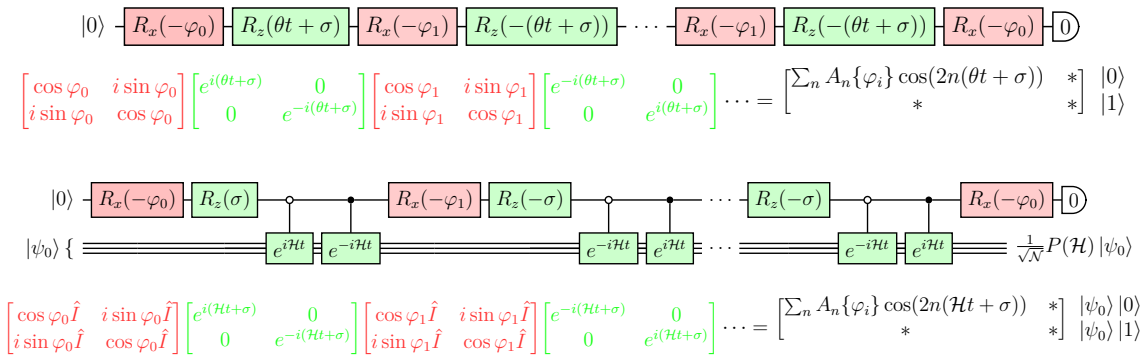


Figure 2.6: An early-fault tolerant QSP method; the QET-U circuit and relation to the PITE circuit. It can be seen that the operations of the single qubit rotation sequence (TOP) and the QET-U circuit (BOTTOM) manifest identical transform, except the former performs an eigenvalue transform on the variable θ and the latter a transform on the matrix $\mathcal{H}$. Figure reproduced from [3] – I was its original creator.

evolution (ITE) operator $e^{-\mathcal{H}\tau}$. Specifically, this involves the numerical integration of the normalised Wick-rotated Schrödinger Equation:

$$|\psi(\tau)\rangle = A(\tau)e^{-\mathcal{H}\tau}|\psi_0\rangle\,, \tag{2.64}$$

where $\tau = it$ and $A(\tau) = 1/\sqrt{\langle\psi_0|\,e^{-2\mathcal{H}\tau}\,|\psi_0\rangle}$ is a normalisation constant. It can be shown that, by noting any state can be expressed as a linear combination of the eigenstates of $\mathcal{H}$, that application of long-periods of ITE is equivalent to exponentially attenuating higher energy contributions to an initial state, eventually distilling down to only the ground state. I will speak more about the other uses of non-unitary evolution in Chapter 4.

In recent years, many quantum computing algorithms for approximate-ITE have been proposed, which promise to overcome these prohibitive scaling costs. These methods use different techniques to circumvent the fact that the ITE operator is non-unitary and cannot be directly implemented using quantum circuit gate operations alone: The five main classes of quantum imaginary time algorithms are variational approaches (VITE) [156–161], tensor network compiled circuits for imaginary time evolution [162], the quantum imaginary time evolution (QITE) method [163], the probabilistic imaginary time evolution (PITE) [153], and QSP-type methods [150, 164]. The key metrics in comparing the utility of these methods,

beyond circuit complexity, are the probability of successfully preparing the desired state and the necessary classical pre- and post-processing resource.

The QITE algorithm is perhaps the most widely discussed in the literature. It uses Pauli tomography to approximate non-unitary imaginary-time propagation with Trotterized unitary evolution. Many variations and improvements of the QITE have since been proposed. The method is not probabilistic, but requires classical computing resources for tomography. The VITE is a general class of hybrid quantum-classical approaches, where a parameterised quantum circuit ansatz is optimised classically such that the ansatz convergence mimics the action of ITE. This method is also not probabilistic, but requires classical computing resources for variational optimisation. PITE achieves non-unitary evolution by entangling real time evolution of an n-qubit system to an ancilla qubit in an LCU-type approach, and triggering a partial collapse of the n-qubit state by measuring the ancilla qubit. A successful outcome prepares a function proportional to $\cos(I - H\delta t)$, thereby approximating to first order the short-ITE operator for a small time step. I will be discussing this method in greater detail in Chapter 4 and Chapter 5.

2.2.6 Measuring Useful Information

The final steps to a quantum-assisted many-body simulation routine would be to measure the qubits and extract useful information – in quantum simulation, the observable that we are most interested in would typically be the energy of the system. Measurement of the qubits would 'collapse' the superposition of basis states that has been manipulated by quantum circuit operations into a singular binary outcome. Because of the probabilistic nature of quantum measurement, we would generally want to repeat the same circuit and measurement sequence a number of times (performing *shots*) to obtain accurate simulation results. It is important at this point to ensure that the 'nature' of information that we want to extract from a register of qubits can be efficiently extracted using a number of shots that doesn't scale prohibitively with system size, since otherwise this would wash out any advantages gained during the quantum state preparation and simulation

stages. As I described in the previous Chapter, a partial, weak measurement of ancillary qubits can also be useful in performing non-unitary operations and produce some useful insight. I will now summarise key measurement techniques relevant to quantum simulation of many-body systems.

Single Expectation Value Estimation

In a VQE-type calculation with JW encoding for example, the total energy expectation of the ansatz is the sum of energy expectation of each Pauli string in the qubit Hamiltonian; given $\mathcal{H} = \sum_j h_j \bigotimes_i P_i^j$:

$$\langle E \rangle = \left\langle \psi(\vec{\theta}) \right| \mathcal{H} \left| \psi(\vec{\theta}) \right\rangle \tag{2.65}$$

$$= \sum_j h_j \left\langle \psi(\vec{\theta}) \right| \bigotimes_i P_i^j \left| \psi(\vec{\theta}) \right\rangle \tag{2.66}$$

Each of these terms are estimated by sampling the qubits in N_s shots (the larger N_s is, the more accurate the estimation) and observing the parity of the binary outcomes. These map to the eigenvalues of the Pauli string which are then used to inform the contribution of each Pauli term to the overall expectation value. This means that, during a single VQE ansatz optimisation step, for a Hamiltonian with N_o non-trivial Pauli terms, we will have to independently prepare the same ansatz circuit $N_s \times N_o$ times to produce a single estimation of the energy expectation value (N_o in general would scale polynomially with system size). Techniques exist that leverage the grouping of commuting Pauli terms together, allowing multiple terms to be measured simultaneously and reducing the absolute cost [165]; however, if the optimisation landscape of the ansatz is challenging, more optimisation steps will likely be required, and the number of shots can become prohibitive. Fig. 2.7 elaborates on the statistical nature of these measurements.

Another statistically sampled single-data-point measurement that is relevant to this thesis is the measurement of overlap between two sets of qubits – this is useful in e.g. variational routines for converging to excited states, where the overlap between ansatz circuits representing eigenstates are also minimised. The overlap

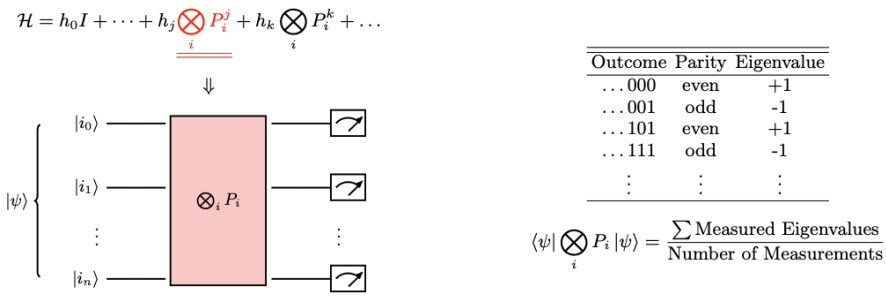


Figure 2.7: Single Pauli string expectation value measurement on a quantum computer. (LEFT) Circuit for measuring one Pauli term. Upon measuring, each qubit collapses to either 0 or 1. The is repeated multiple times and the results (shots) are recorded. (RIGHT) A shot table mapping each measurement outcome to an eigenvalue (parity is whether the number of 1's in the measurement is odd or even, which corresponds to an eigenvalue of -1 and +1 respectively). The energy expectation of the Pauli term P_i is thus the average eigenvalue over multiple shots multiplied by h_i (h_i computed from Hartree-Fock mean-field calculations. Figure reproduced from [2] – I was its original creator.

between two unitaries is calculated using the vacuum test:

$$\text{since} \quad |\Phi_i\rangle = V_i |0\ldots00\rangle \qquad \text{and} \qquad \left|\Psi(\vec{\theta})\right\rangle = U |0\ldots00\rangle$$

where V_i and U are unitary reference and state preparation operators, then the overlap between the two states is the expectation of measuring $|0\ldots00\rangle$ when the $V_i^\dagger U$ circuit is applied

$$\left\langle \Phi_i \middle| \Psi(\vec{\theta}) \right\rangle = \langle 0\ldots00| V_i^\dagger U |0\ldots00\rangle .$$

This technique appends one ansatz circuit onto the other and thus in the naïve case doubles the depth of the circuit.

Measuring Excited State Energies

In most cases, energy measurement such as that explained above would be used to extract the ground state energy from a quantum simulation. Low-lying excited states are of course also useful when quantum many-body systems undergo energy perturbations. Excited state methods in quantum modelling include e.g. the aforementioned variational deflation [97] approach which combines expectation value and overlap measurement to prepare excited states on qubits with variational ansätze. Techniques such as equations-of-motion [166, 167], VITE [168] and subspace-based

[169] have also been proposed. I will now briefly mention specifically those that use a quantum computer to generate a subspace for computing excited states, which has some overlap with my research in Chapter 6.

In subspace expansion [170], one would start from preparing an estimated ground state $|\tilde{\psi}_G\rangle$ on a quantum computer and apply low-weight operators as excitations to this initial state. Typically Pauli operators are used to produce a new set of states $|\psi_k\rangle = \mathcal{P}_k \left|\tilde{\psi}_G\right\rangle$, from which the elements $\mathcal{H}_{kj} = \langle\psi_k| \mathcal{H} |\psi_j\rangle$ and $S_{kj} = \langle\psi_k|\psi_j\rangle$ are measured using aforementioned measurement techniques. Finally, diagonalising $\mathcal{H}_{kj}$ on a classical computer reveals an improved spectrum and ground-state energy guess than that of $|\tilde{\psi}_G\rangle$. A related group of methods are Krylov [171] and Krylov-like subspaces [163, 172]. These algorithms generate a subspace by applying to a supplied initial state either (a) powers of the Hamiltonian, or (b) real- or imaginary-time Hamiltonian evolution steps which approximate powers of the Hamiltonian in the small time-step limit (as per the Taylor expansion of the matrix exponential). Computing the subspace Hamiltonian and overlap matrix elements on a quantum computer and diagonalising it classically gives approximations to eigenvalues of the modelled Hamiltonian.

Classical Shadows

Instead of measuring only a single observable at any given time, a recent advance is the classical shadow procedure [173–178]. In it, we perform randomised measurements to several copies of an unknown quantum state, the outcomes of which can then be used to reconstruct a very large number of observable expectation values with provable sample complexities. While many variants of classical shadows exist [174–178], for the purpose of review I will recapitulate the near-term friendly variant that can be used to reconstruct expected values of local Pauli operators:

- We apply a random unitary U to rotate a copy of an unknown quantum state. For local Pauli strings, the unitaries are chosen randomly from single qubit Clifford gates on each qubit and the procedure is thus equivalent to randomly

selecting to measure in the X, Y or Z bases – we measure each qubit to obtain N-bit measurement outcomes $\left|\hat{b}_i\right\rangle \in \{0,1\}^N$.

- We then generate the classical shadows by applying the inverse of the measurement channel $\mathcal{M}$, which can be done efficiently as the channel chosen is a distribution over Clifford circuits. The classical snapshots are generated as $\hat{\rho}_i = \mathcal{M}^{-1}(U^\dagger \left|\hat{b}_i\right\rangle \left\langle\hat{b}_i\right| U)$, the classical shadows are collections of these N_s snapshots $S(\rho; N_s) = [\hat{\rho}_1, ..., \hat{\rho}_{N_s}]$.

- From these classical shadows we can construct K estimators of ρ from our N_s snapshots as
  $\hat{\rho}_{(k)} = \frac{1}{r}\sum_{i=(k-1)r+1}^{kr} \hat{\rho}_i$ with $r = \lfloor N_s/K \rfloor$ and classically calculate estimators of the Pauli expectation values $\hat{o}_i(N_s, K) = \text{median}\{\text{tr}\Big(O_i\hat{\rho}_{(1)}\Big), ..., \text{tr}\Big(O_i\hat{\rho}_{(K)}\Big)\}$.

- The sample complexity of obtaining these estimators of N_o Pauli operators of locality q to error ϵ is $\mathcal{O}[3^q \log(N_o)/\epsilon^2]$.

In this context, the classical shadow approach is very NISQ friendly as it only requires applying random single-qubit rotations immediately prior to measurement in the standard basis. The approach also has a provable sample efficiency: Recall that q-local Pauli strings are tensor products of single-qubit Pauli operators $P^{(q)} \in \{\text{Id}, X, Y, Z\}^{\otimes N}$ such that $P^{(q)}$ acts non-trivailly (with X, Y or Z) to q qubits and trivially (with Id) on all other qubits. To reconstruct N_o at most q-local Pauli strings, one needs to collect a number of samples $N_s = \mathcal{O}(\log N_o)$ that grows only logarithmically as we increase N_o. Furthermore, the classical computational resources are linear as $\mathcal{O}(N_o)$ and fast, optimised, and hardware agnostic code for quantum emulation can be found (see e.g. QuEST [179]). By applying this procedure, we can determine N_o time-dependent signals with a similarly advantageous sample complexity. Fig. 2.8 summarises the ideas behind classical shadows.

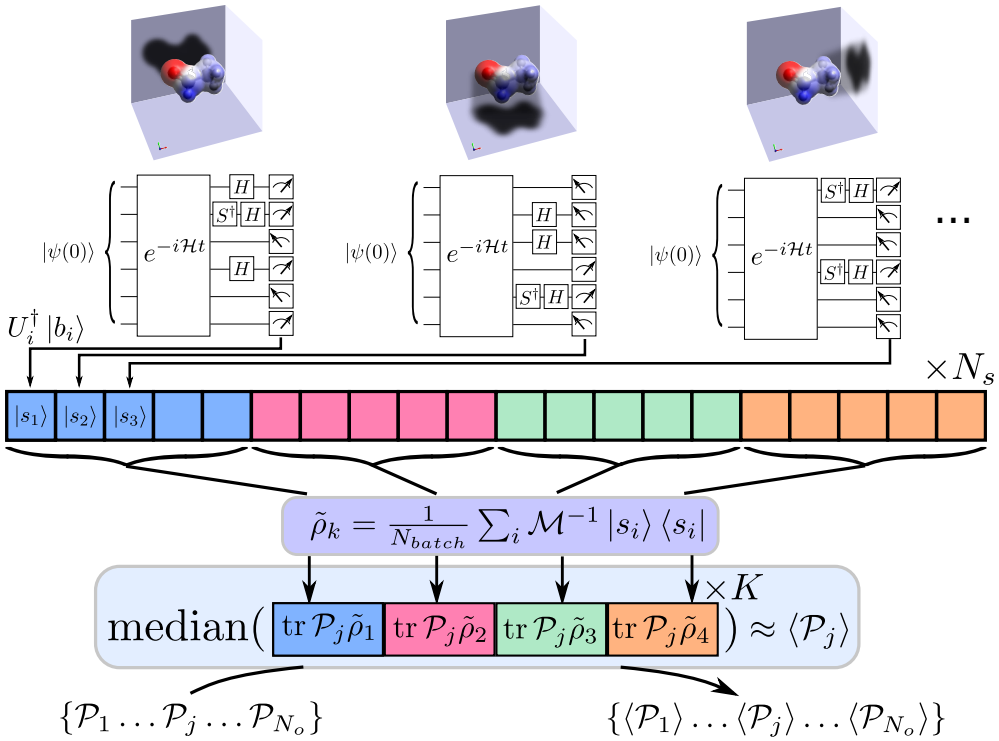


Figure 2.8: Cartoon depicting classical shadows measurement. The same unitary (in this case a time evolution circuit) is used to prepare N_s copies of the same state. The qubits are then measured in a random basis – we can conceptualise this as casting 'shadows' of a quantum many-body system in randomised directions. The binary outcomes are then captured and a means-of-means method is used to reconstruct a set of Pauli observables. Original graphic created by me.

Phase Estimation

Phase estimation through the use of ancilla qubits [180] is one of the first proposed and most fundamental qubit measurement techniques, employed in many quantum computing applicaitons. I will close this Chapter by summarising the standard multi-qubit phase estimation method and recent developments.

The logarithmically-scaling phase estimation algorithm [180] uses significant ancilla qubit overhead to measure the phase that a system acquires as it evolves forward in time – this involves controlling time evolution or Hamiltonian simulation circuits with ancillary qubits, and subjecting an initial state to this machinery. The phase, which is equivalent to what chemists would refer to as the frequencies of the autocorrelation function signal $\langle\psi(0)|\psi(t)\rangle = \langle\psi(0)|\, e^{-i\mathcal{H}t}\, |\psi(0)\rangle$, are encoded in the ancillary qubits and can be extracted extremely efficiently; in suitable cases,

even in a single shot. However, the circuit depths and number of qubits required for a high quality canonical phase estimation experiment well-exceeds the current (and most likely near-future) capabilities of quantum hardware.

Consider an ancillary register with S qubits and a working register containing the state of interest. We first conditionally apply $2^{S-1}N$ time evolution steps to the state, controlled by the most significant ancillary qubit. We then repeat by conditionally applying $2^{S-2}N$ time evolution steps, now controlled by the next most significant ancillary qubit, and repeat again until every ancillary qubit has been used and the least significant qubit controls only N time steps. At this point, the quantum computer is in the state

$$\frac{1}{\sqrt{2^S}} \sum_{m=0}^{2^S-1} |m\rangle \, U^{mN}(\delta t) \, |\psi\rangle \, . \tag{2.67}$$

If $|\psi\rangle$ is an eigenstate $|\Psi_n\rangle$ of $\mathcal{H}$ with energy E_n, the state is thus

$$\begin{aligned} &\frac{1}{\sqrt{2^S}} \, |\Psi_n\rangle \sum_{m=0}^{2^S-1} e^{-iE_n mN\delta t} \, |m\rangle \\ =&\frac{1}{\sqrt{2^S}} \, |\Psi_n\rangle \sum_{m=0}^{2^S-1} e^{-i\phi m} \, |m\rangle \end{aligned} \tag{2.68}$$

where $\phi = E_n N \delta t$. The phase information ϕ is now encoded in the ancillary qubit register, and can be efficiently extracted using an inverse QFT on the ancillary register. Indeed, if $\phi = 2^k$ where $k \in \mathbb{N}$, then a QFT of the ancillary register should yield the binary representation of k upon measurement. The circuit for performing phase estimation is captured in Fig. 2.9.

The phase estimation element is analogous to performing a classical Fourier analysis of the autocorrelation signal to reveal the power spectrum, where resonance peaks correspond to the stationary eigenstates of the system. Specifically, the power spectrum is peaked at the eigenvalues of $\mathcal{H}$ with heights corresponding to $|c_n|^2$

$$\frac{1}{\sqrt{2\pi}} \int \langle \psi(t) | \psi(0) \rangle \, e^{i\omega t} dt = \sqrt{2\pi} \sum_n |c_n|^2 \delta(\omega - E_n), \tag{2.69}$$

where c_n is the contribution of each eigenstate to $|\psi\rangle = \sum_n c_n \, |\Psi_n\rangle$. In practice, the discrete and finite time integration limits the resolution of the power

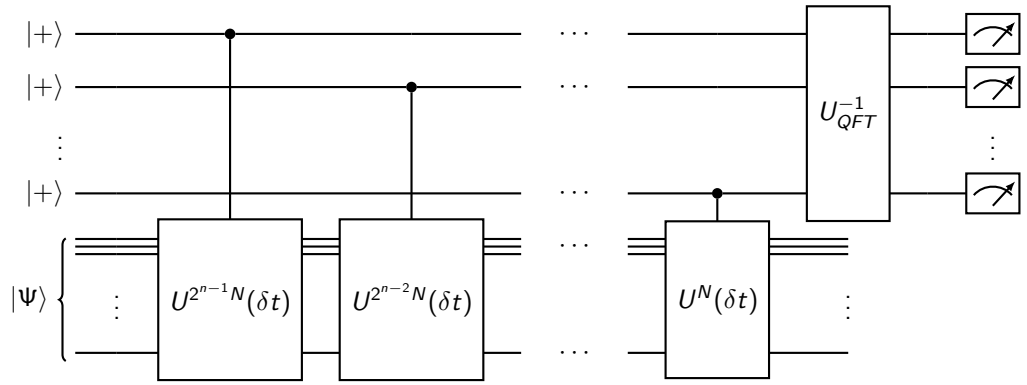


Figure 2.9: Circuit for performing canonical phase estimation. using an n-qubit ancilla register. The i^{th} most significant ancilla qubit controls the propagation of $2^{n-i}N$ time evolution steps in the main register. Figure reproduced from [4] – I was its original creator.

spectrum, given by

$$\delta\omega = \frac{2\pi}{M_{\text{tot}} N \delta t} \tag{2.70}$$

where M_{tot} is the total number of time samples.

If $|\psi\rangle$ is instead a mixture of eigenstates, the probability of observing a given binary outcome during ancillary qubit measurement would reflect the $|c_n^2|$ contributions of constituent eigenstates to $|\psi\rangle$. This would be useful in e.g. a quantum chemistry task for ground state energy determination, where the input state might be Slater determinants with significant ground-state contributions. Phase estimation can also be leveraged as a state-preparation method: Assuming that the initial state $|\psi\rangle$ is dominated by e.g. the ground state of a many-body $\mathcal{H}$, with some contributions from excited states, a phase estimation measurement is likely to result in observing the phase of the ground state in the ancillary qubits, and the resulting state on the main register of qubits will be that of the ground state, with higher energy contributions projected out. The success of this state-preparation approach is therefore almost entirely predicated on an initial state that has substantial overlap with the 'true' ground state of $\mathcal{H}$ – which some believe is in itself a non-trivial challenge (see e.g. Refs. [73, 181, 182]).

Many techniques have been proposed to measure the autocorrelation signal more efficiently. For example, in the Iterative Phase Estimation approach (which I will

spend more time talking about in Chapter 5), the autocorrelation signal is estimated by controlling the time evolution unitary with only a single ancilla qubit, measuring the outcome of the ancilla, and repeating that procedure for multiple shots and at different points during the evolution. The probability of obtaining the $|0\rangle$ state *is* the autocorrelation signal. While the approach has an ultra-low qubit overhead, it does need many measurements to mitigate shot noise, and require sampling at many time steps for high resolution. In the last two years, a group of methods called statistical phase estimation has been proposed which use sophisticated classical signal processing techniques to extract as much information as possible from the single ancilla signal, thereby bringing down the sampling cost and circuit depth for generating said signal on a quantum computer [183–185] – see Ref. [186] for recent experimental demonstrations of statistical phase estimation. Bespoke architectures have also been developed [187] to reduce quantum resources required for performing ancilla-controlled time evolution. Problem-specific techniques have also been developed for quantum chemistry applications. For example, phase estimation for investigating high energy excitations [188] have been proposed. Refs. [189, 190] use Hadamard-test circuits to estimate expected values of excitation operators – and ancilla-free variants are also discussed. Ref. [191] proposes to create a superposition of two eigenstates and performs a Ramsey-type experiment by applying a time evolution. The approach alleviates the use of an ancilla-controlled evolution but requires a Prep and a $\text{Prep}^\dagger$ oracle that prepares a superposition of two eigenstates. Another ancilla-free approach might be Ref. [192], which suggests first applying a Haar-random circuit to the $|0\rangle$ state, applying a time evolution operator, applying the inverse of the random circuit and finally measuring the survival probability as probability of the $|0\rangle$ state.

Do what you can, with what you've got,
where you are.

— Theodore Roosevelt [193]

3

Computing Energy Eigenvalues on Near-Term Quantum Computers

Contents

3.1 Introduction 56
3.2 Shortcuts to Adiabaticity for Ground State Preparation 57
3.2.1 Method 60
3.2.2 Case Study: Transverse Field Ising Model 62
3.2.3 Numerical Experiments 66
3.2.4 Beyond the TFIM 70
3.2.5 Discussion 72
3.3 Symmetry-restricted Variational Ansätze for Excited States 73
3.3.1 Methods 74
3.3.2 Numerical Validation 78
3.3.3 Discussion 84
3.4 Concluding Thoughts 85

In the first part of this Chapter, I include work done in collaboration with Raul Santos at the University of Oxford and the company Phasecraft Limited. I performed all of the numerical simulations reported, based on initial ideas from Raul Santos. The work is currently in preparation for publication. In the second part of this chapter, I use text from the Physical Chemistry Chemical Physics paper "Molecular excited state calculations with adaptive wavefunctions on a quantum eigensolver emulation:

reducing circuit depth and separating spin states" [2] verbatim – I was the first and original author of the text. The work was done in collaboration with coauthors Nathan Fitzpatrick, Javier Segarra Marti, Michael Bearpark and David Tew at the University of Oxford. My coauthors and I developed the methods therein as a group, based on my initial ideas. I performed all of the numerical simulations reported.

3.1 Introduction

In the current era of Noisy Intermediate Scale Quantum (NISQ) computers, enthusiasm in attempting to demonstrate quantum advantage on existing or near-term quantum computers persists. These quantum computers will likely continue to have a limited number of physical qubits, suffer from high error rates, and have a brief coherence time. Therefore, for a NISQ computer to be useful for quantum simulation, a simulation algorithm should execute relevant quantum circuits before errors and noise corrupt the logic. The use of resources like the number of qubits and the depth of quantum gates should also be kept low and reasonably scalable with system size. In addition, if we intend to demonstrate any potential quantum speedup it would be prudent to keep the overall number of circuit measurement repeats manageable.

Many hybrid quantum-classical simulation algorithms, which perform part of the computation on quantum computers and the rest on conventional high-performance computers, have been proposed which promise to overcome these challenges of NISQ computing. In the last three years, use of the Variational Quantum Eigensolver (VQE) hybrid algorithm [78–85] to estimate the ground and low energy states of a quantum many-body system has received the most attention as a viable means of demonstrating quantum advantage. While no VQE experiments indicating meaningful quantum advantage has been reported to date, efforts to improve VQE are still ongoing. Even if it comes to pass that quantum error correction and fault-tolerance will be strictly necessary to achieve any quantum advantage for simulation of many-body systems, VQE will likely remain relevant for the initial

state preparation steps before simulation, and any improvements to VQE explored now will no doubt be useful then.

Recall that, within the variational eigensolver paradigm, a quantum computer is used to prepare a parameterised *ansatz* wavefunction of a many-body state:

$$\left|\psi(\vec{\theta})\right\rangle = U(\vec{\theta})\left|\psi_0\right\rangle \tag{3.1}$$

and the parameters are optimised such that the energy expectation value of the ansatz is minimised;

$$\min_{\vec{\theta}} E(\vec{\theta}) = \left\langle \psi(\vec{\theta}) \right| \mathcal{H} \left| \psi(\vec{\theta}) \right\rangle, \tag{3.2}$$

thereby providing an estimate of the ground state of the Hamiltonian $\mathcal{H}$. most of the efforts in VQE improvements therefore include finding suitable initial states $|\psi_0\rangle$ that are close to the target (often ground) state, reducing the circuit depth of $U(\vec{\theta})$, while ensuring that the total number of shot measurements to measure and optimise $E(\vec{\theta})$ remains manageable (see e.g. Refs. [194, 195]. This is supplemented by other innovations that include e.g. finding representations of the many-body problem that are qubit-frugal (see e.g. Ref. [44]), as well as understanding the effects of noise.

In this Chapter, I discuss two projects that I undertook which sought to address some of these aforementioned pain-points of VQE. One focused on the preparation of ground states and computing the energy, and the other for excited states.

3.2 Shortcuts to Adiabaticity for Ground State Preparation

Consider a bipartite Hamiltonian typical in many-body systems, given by

$$\mathcal{H}_\lambda = (1-\lambda)\mathcal{H}_0 + \lambda\mathcal{H}_1, \tag{3.3}$$

where $0 \leq \lambda \leq 1$ is a variable quantifying the relative contributions of each Hamiltonian. One natural approach for finding the ground state of a Hamiltonian $\mathcal{H}(\lambda_{\text{targ}})$ with a specific value of λ_{targ}, in the sense of both computational studies and physical experiments, is to invoke the adiabatic theorem of Born and Fock and

perform *adiabatic state preparation* – a technique that has been explored extensively in the literature for quantum algorithms. The adiabatic theorem says that; when an initial eigenstate is subjected to sufficiently slow perturbation, it would remain an eigenstate of the instantaneous Hamiltonian, assuming the gap between the eigenstate and the instantaneous Hamiltonian's spectrum does not close up [196]. In the context of state preparation, one could start from the case where $\lambda = 0$ and prepare the ground state of $\mathcal{H} = \mathcal{H}_0$ (which might be easily achievable), and slowly time evolving the state while 'turning on' the contributions of $\mathcal{H}_1$ such that $\lambda \to \lambda_{\text{targ}}$ in long times. At the end of this adiabatic evolution, the state should become the ground state of $\mathcal{H}(\lambda_{\text{targ}})$ [197].

At an abstract level, this process resembles the VQE optimisation of a *Hamiltonian ansatz* [88, 198]. Recall that the Hamiltonian ansatz takes inspiration from the Trotterised time evolution operator $\prod_{j=1}^{N} e^{i\mathcal{H}\delta t}$, where each time step δt and the Hamiltonian coefficients $\lambda, (1-\lambda)$ are replaced with free-variable circuit parameters, forming a unitary product ansatz U^{Ham} of the following form;

$$U^{\text{Ham}}(\vec{\theta}, \vec{\phi}) = \prod_{j=1}^{N} e^{i\theta_j \mathcal{H}_0} e^{i\phi_j \mathcal{H}_1}; \tag{3.4}$$

N in this case is the number of ansatz 'layers' in the circuit, and each ansatz layer indexed j has a unique parameter set θ_j, ϕ_j. This would then be applied on an appropriate initial state e.g. the ground state of $\mathcal{H}_0$ on a quantum computer $\left|\psi(\vec{\theta}, \vec{\phi})\right\rangle = U^{\text{Ham}}(\vec{\theta}, \vec{\phi}) \left|\psi_0\right\rangle$ and classical optimisation of the parameters performed. Such an optimisation of the Hamiltonian ansatz's parameters in $U^{\text{Ham}}(\vec{\theta}, \vec{\phi})$ so as to drive $|\psi_0\rangle$ towards the ground state of $\mathcal{H}$ is reminiscent of adiabatic state preparation [198]. Given this superficial resemblance, one naturally wonders if innovations that makes adiabatic state preparation more efficient could inspire techniques that improves the quantum resource use of VQE within the Hamiltonian ansatz framework.

One such innovation is a group of techniques called the shortcuts to adiabaticity. Shortcuts to adiabaticity are methods that address a key requirement in adiabatic state preparation which can in some cases become prohibitive; that the adiabatic

evolution should be sufficiently slow to prevent the system from exciting into higher energy states. Shorcuts to adiabaticity seek to address this with the introduction of a *gauge potential* term $\mathcal{A}_\lambda$ to the Hamiltonian to create a new counter-adiabatic driving Hamiltonian $\mathcal{H}_{\text{CD}}$:

$$\mathcal{H}_{\text{CD}} = \mathcal{H}_\lambda + \frac{d\lambda}{dt}\mathcal{A}_\lambda. \tag{3.5}$$

This gauge potential is intended to modify the Hamiltonian, so that transitions between the instantaneous ground state and higher energy eigenstates are symmetry restricted – allowing the adiabatic evolution to proceed much faster than otherwise. In practice, designing the ideal gauge potential for many-body systems requires understanding how the spectrum is influenced as a function of λ and thus the diagonalisation of $\mathcal{H}_\lambda$, impossible in practice for large, interesting molecules. Claeys *et al.* proposed a general l-order approximate of the gauge potential $\mathcal{A}^{(l)}_{\lambda(t)}$ [199], given by the following nested commutator:

$$\mathcal{A}^{(l)}_{\lambda(t)} = i\sum_{k=1}^{l} a_k \underbrace{[\mathcal{H}_\lambda, [\mathcal{H}_\lambda, \dots [\mathcal{H}_\lambda}_{2k-1}, \frac{\partial}{\partial\lambda}\mathcal{H}_\lambda]]]. \tag{3.6}$$

This is motivated by the following: We can invoke the Baker-Campbell-Hausdorff (BCH) expansion to write

$$e^{-i\mathcal{H}_\lambda t}\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda e^{i\mathcal{H}_\lambda t} = \sum_{k=0}^{\infty}\frac{(-it)^k}{k!}\underbrace{[\mathcal{H}_\lambda, [\mathcal{H}_\lambda, \dots [\mathcal{H}_\lambda}_{k}, \frac{\partial}{\partial\lambda}\mathcal{H}_\lambda]]]; \tag{3.7}$$

in Ref. [199], the authors suggest that the gauge potential can be expressed as:

$$\mathcal{A}_\lambda = \lim_{\epsilon\to 0^+}\int_0^\infty e^{-\epsilon t}\left(e^{-i\mathcal{H}_\lambda t}\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda e^{i\mathcal{H}_\lambda t} - \sum_n |n\rangle\langle n|\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda|n\rangle\langle n|\right)dt,$$

where the term $\sum_n |n\rangle\langle n|\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda|n\rangle\langle n|$ cancels the diagonal elements ($|n\rangle$ are instantaneous eigenstates of $\mathcal{H}_\lambda$), and the off diagonal elements must satisfy:

$$\begin{aligned}\langle m|\mathcal{A}_\lambda|n\rangle &= \lim_{\epsilon\to 0^+}\int_0^\infty e^{-\epsilon t}\langle m|e^{-i\epsilon_m t}\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda e^{i\epsilon_n t}|n\rangle\, dt \\ &= \lim_{\epsilon\to 0^+}\frac{1}{\epsilon + i(\epsilon_m - \epsilon_n)}\langle m|\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda|n\rangle.\end{aligned}$$

Substitution of Eq. (3.7) into the above gives us an infinite nested commutator expression of the gauge potential. Similarly, we can make the same observation

for a 'gapped' or regularised gauge potential that is exact for $(\epsilon_m - \epsilon_n) > \Delta$ and vanishing for $(\epsilon_m - \epsilon_n) < \Delta$:

$$\mathcal{A}_\lambda^\Delta = i\sum_{k=1}^{\infty}\frac{(-1)^k}{k!}\Delta^{-2k}\underbrace{[\mathcal{H}_\lambda,[\mathcal{H}_\lambda,\ldots[\mathcal{H}_\lambda}_{2k-1},\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda]]], \tag{3.8}$$

where likewise the off diagonal elements must satisfy

$$\begin{aligned}\langle m|\,\mathcal{A}_\lambda^\Delta\,|n\rangle =& i\sum_{k=0}^{\infty}\frac{(-1)^k}{k!}\Delta^{-2k}(\epsilon_m-\epsilon_n)^{2k-1}\,\langle m|\,\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda\,|n\rangle \\ =& -i\frac{(1-e^{-(\epsilon_m-\epsilon_n)^2/\Delta^2})}{(\epsilon_m-\epsilon_n)}\,\langle m|\,\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda\,|n\rangle\,.\end{aligned}$$

This regularised gauge potential suppresses excitations, or energy gaps, above Δ. In practice, the series of nested commutators is truncated as per Eq. (3.6). Taking the first order ($l = 1$) approximation for example;

$$\begin{aligned}\mathcal{A}^{(1)}_{\lambda(t)} &= ia_1[\mathcal{H},\frac{\partial}{\partial\lambda}\mathcal{H}_\lambda] \\ &= ia_1[\mathcal{H},\frac{\partial}{\partial\lambda}\{(1-\lambda)\mathcal{H}_0+\lambda\mathcal{H}_1\}] \\ &= ia_1[\mathcal{H},\mathcal{H}_1-\mathcal{H}_0] \\ &= ia_1\{[(1-\lambda)\mathcal{H}_0+\lambda\mathcal{H}_1](\mathcal{H}_1-\mathcal{H}_0)- \\ &\qquad (\mathcal{H}_1-\mathcal{H}_0)[(1-\lambda)\mathcal{H}_0+\lambda\mathcal{H}_1]\} \\ &= ia_1(\mathcal{H}_0\mathcal{H}_1-\mathcal{H}_1\mathcal{H}_0) \\ &= ia_1[\mathcal{H}_0,\mathcal{H}_1]\end{aligned}$$

we end up with the anti-Hermitian commutator $[\mathcal{H}_0, \mathcal{H}_1]$, which is relatively easy to compute and implement.

3.2.1 Method

With my collaborator Raul Santos, I explored whether a Hamiltonian ansatz constructed from a counter-adiabatic driving Hamiltonian $\mathcal{H}_{\mathrm{CD}}$, where the gauge potential is approximated with Eq. (3.6), might outperform a naïve Trotterised Hamiltonian ansatz U_{Ham} (Eq. (3.4)) in VQE. This insight has been explored in the context of the quantum approximate optimisation algorithm (QAOA) for solving

combinatorial optimisation problems on near-term quantum computers [200, 201] – which performs a circuit parameter optimisation routine virtually identical to VQE, except that the target state represents the solution of the combinatorial optimisation problem instead of a many-body system's ground state.

Taking the first order approximate of the counter-adiabatic driving Hamiltonian

$$\mathcal{H}_{\mathrm{CD}}^{(1)}(t) = \mathcal{H}(t) + i\frac{d\lambda(t)}{dt}a_1[\mathcal{H}_0, \mathcal{H}_1], \tag{3.9}$$

the Hamiltonian ansatz constructed from this Hamiltonian is

$$U\left(\vec{\alpha}, \vec{\phi}, \vec{\theta}\right) = \prod_{j=1}^{N} e^{i(\alpha_j \mathcal{H}_0 + \phi_j \mathcal{H}_1 + i\theta_j[\mathcal{H}_0, \mathcal{H}_1])}. \tag{3.10}$$

Note that the commutator is anti-Hermitian – in other words, the eigenvalues of $i[\mathcal{H}_0, \mathcal{H}_1]$ are real. We can straightforwardly Trotterise this ansatz to:

$$U^{\mathrm{CDTrot}}\left(\vec{\alpha}, \vec{\phi}, \vec{\theta}\right) = \prod_{j=1}^{N} e^{i\alpha_j \mathcal{H}_0} e^{i\phi_j \mathcal{H}_1} e^{i^2 \theta_j[\mathcal{H}_0, \mathcal{H}_1]}. \tag{3.11}$$

In my experiments, I also explored the possibility of compiling the ansatz such that two-body terms are not Trotterised and are instead implemented together. As I will show presently, this ansatz might look like:

$$U^{\mathrm{CDRecomp}}\left(\vec{\alpha}, \vec{\phi}, \vec{\theta}\right) = \prod_{j=1}^{N} e^{i\alpha_j \mathcal{H}_0} e^{i(\phi_j \mathcal{H}_1 + i\theta_j[\mathcal{H}_0, \mathcal{H}_1])} \tag{3.12}$$

All of these ansätze are optimised with respect to the target Hamiltonian $\mathcal{H}(\lambda_{\mathrm{targ}})$ – that is, the objective function minimised is

$$\min_{\vec{\alpha}, \vec{\phi}, \vec{\theta}} \left\langle \psi(\vec{\alpha}, \vec{\phi}, \vec{\theta}) \right| \mathcal{H}(\lambda_{\mathrm{targ}}) \left| \psi(\vec{\alpha}, \vec{\phi}, \vec{\theta}) \right\rangle. \tag{3.13}$$

We anticipated these alternative ansatz constructions may converge closer to the ground state energy with fewer VQE iterations, thereby reducing the total number of measurement shots.

3.2.2 Case Study: Transverse Field Ising Model

To assess if counter-adiabatic driving operators can improve the performance of the Hamiltonian ansatz, I performed numerical VQE calculations of the prototypical n-site Transverse Field Ising Model (TFIM),

$$\mathcal{H} = J\left((1-\lambda)\sum_{i=0}^{n} X_i + \lambda \sum_{i=0}^{n-1} Z_i Z_{i+1}\right) \tag{3.14}$$

where $J = -1$, $0 < \lambda < 1$. This represents one of the simplest extensions of a classical spin system into the quantum domain, where minimal quantum behaviour is introduced by the transverse magnetic field. The TFIM provides a useful description of certain magnetic and ferro-electric systems [202]. The quantum simulation of this model in high dimensions, to some limited extent, might be relevant for real materials applications; it is nonetheless instructive for quantum algorithm design. The commutator $[\mathcal{H}_0, \mathcal{H}_1]$ of TFIM is:

$$\begin{aligned}[\mathcal{H}_0, \mathcal{H}_1] &= \sum_{i=0}^{N} X_i \sum_{j=0}^{n-1} Z_j Z_{j+1} - \sum_{j=0}^{n-1} Z_j Z_{j+1} \sum_{i=0}^{n} X_i \\ &= -2i \sum_{j=0}^{n-1} (Y_j Z_{j+1} + Z_j Y_{j+1})\end{aligned}$$

and so, the first-order approximation of the gauge potential is:

$$i\frac{d\lambda}{dt} a_1 [\mathcal{H}_0, \mathcal{H}_1] = 2\frac{d\lambda}{dt} a_1 \sum_{j=0}^{n-1} (Y_j Z_{j+1} + Z_j Y_{j+1}) \tag{3.15}$$

In my research with my collaborator, I started by comparing the conventional TFIM ansatz with the Trotterised counter-adiabatic driving ansatz. In my numerical experiments, I used single, double or triple layered ansätze. Specifically, the normal TFIM Hamiltonian ansatz is:

$$U^{\mathrm{Ham}}(\vec{\alpha}, \vec{\phi}, \vec{\varphi}) = \prod_{j}^{N=1,2,3} U(\alpha_j, \phi_j, \varphi_j), \tag{3.16}$$

whereby in each layer, the 2-site operators corresponding to interactions between even-to-odd sites and odd-to-even sites are given distinct parameters, as the two sets cannot be implemented simultaneously at a circuit level:

$$U(\alpha_j, \phi_j, \varphi_j) = \prod_{k=1}^{n} e^{i\alpha_j X_k} \prod_{k=1}^{n/2} e^{i\phi_j Z_{2k-1} Z_{2k}} \prod_{k=1}^{n/2} e^{i\varphi_j Z_{2k} Z_{2k+1}}. \tag{3.17}$$

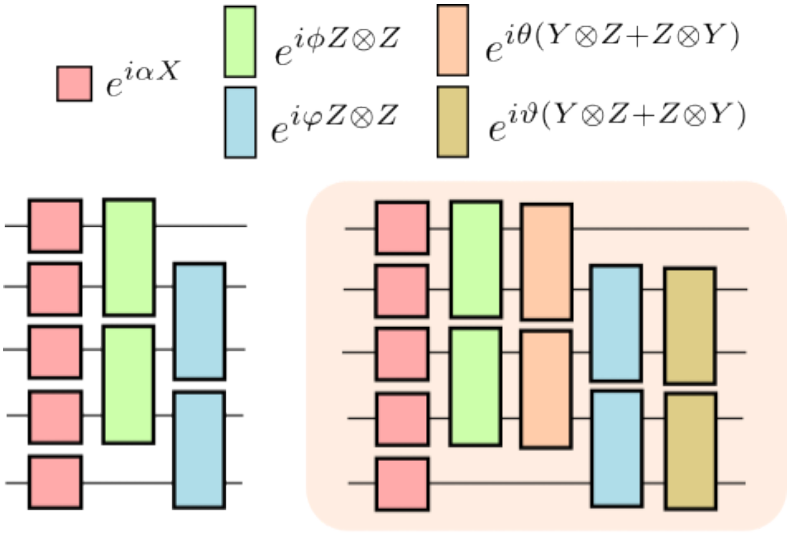


Figure 3.1: The Hamiltonian ansatz U^{Ham} versus the Trotterised counter-adiabatic driving ansatz U^{CDTrot} for the Transverse Field Ising Model. The circuit diagram on the left is U^{Ham} and on the right U^{CDTrot}. The single qubit gate $e^{i\alpha X} \equiv R_X(-2\alpha)$ and each 2-qubit gate represents a Pauli gadget. Original graphic created by me.

The 2-qubit phase rotations can be trivially implemented with Pauli gadget circuits, where each would cost 2 CNOT gates. Similarly, the Trotterised counter-adiabatic driving ansatz is

$$U^{\text{CDTrot}}(\vec{\alpha},\vec{\phi},\vec{\varphi},\vec{\theta},\vec{\vartheta}) = \prod_{j}^{N=1,2,3} U(\alpha_j,\phi_j,\varphi_j,\theta_j,\vartheta_j), \tag{3.18}$$

where each layer is

$$U(\alpha_j,\phi_j,\varphi_j,\theta_j,\vartheta_j) = \prod_{k=1}^{n} e^{i\alpha_j X_k} \prod_{k=1}^{n/2} e^{i\phi_j Z_{2k-1}Z_{2k}} \prod_{k=1}^{n/2} e^{i\theta_j(Z_{2k-1}Y_{2k}+Y_{2k-1}Z_{2k})} \prod_{k=1}^{n/2} e^{i\varphi_j Z_{2k}Z_{2k+1}} \prod_{k=1}^{n/2} e^{i\vartheta_j(Z_{2k}Y_{2k+1}+Y_{2k}Z_{2k+1})}.$$

In this case, the $e^{i\theta(ZY+YZ)}$ gates can also be implemented with single Pauli gadgets, and would contribute 2 additional CNOT gates each. Per ansatz circuit layer, the U_{T} ansatz therefore has two additional parameters and requires four additional CNOT gates. This circuit structure is illustrated in Fig. 3.1. Similar comparisons of these TFIM Hamiltonian ansatz variants have been explored for QAOA [200, 201], and the counter-adiabatic driving ansatz enhancement has also been used in adiabatic quantum computing on available hardware [203].

The implementation of counter-adiabatic driving operators as additional Pauli gadgets increases the circuit gate depth of each ansatz layer. However, in this case the operators $e^{i\theta(ZY+YZ)}$ and $e^{i\phi ZZ}$ act on the same pairs of qubits – allowing the 2-qubit unitary:

$$U_{12}(\phi,\theta) = e^{i(\phi Z_1Z_2+\theta Y_1Z_2+\theta Z_1Y_2)} \tag{3.19}$$

to be recompiled into a circuit subroutine that can be performed exactly. This general strategy of recompiling the desired unitaries into a more compact circuit is often used in near term quantum algorithms. My collaborator and I therefore derived a 2-qubit gate sequence that is able to perform $U_{12}(\phi,\theta)$ using only 2 CNOT gates. In the TFIM case, one has to first recognise the conjugation relation of the following form:

$$e^{-i(\alpha X_1+\beta X_2)}U_{12}(\phi,\theta)e^{i(\alpha X_1+\beta X_2)} = e^{i(A_{\alpha\beta}(\phi,\theta)Z_1Z_2+B_{\alpha\beta}(\phi,\theta)Y_1Y_2)}. \tag{3.20}$$

We can see this is the case by noting that:

$$e^{-i(\alpha X_1+\beta X_2)}e^{iM}e^{-i(\alpha X_1+\beta X_2)} = e^{i[e^{-i(\alpha X_1+\beta X_2)}Me^{-i(\alpha X_1+\beta X_2)}]} \tag{3.21}$$

and

$$e^{i(\alpha X_1+\beta X_2)} = e^{-i\alpha X_1}e^{i\beta X_2}. \tag{3.22}$$

Substituting $M = \phi Z_1Z_2 + \theta Y_1Z_2 + \theta Z_1Y_2$ in Eq. (3.21), we end up with the following exponent:

$$\begin{aligned}
&Y_1Y_2\left(\frac{\phi\left(\cos\left(2\left(\alpha-\beta\right)\right)-\cos\left(2\left(\alpha+\beta\right)\right)\right)}{2}-\theta\sin\left(2\left(\alpha+\beta\right)\right)\right)\\
+&Y_1Z_2\left(\frac{\phi\left(-\sin\left(2\left(\alpha-\beta\right)\right)-\sin\left(2\left(\alpha+\beta\right)\right)\right)}{2}+\theta\cos\left(2\left(\alpha+\beta\right)\right)\right)\\
+&Z_1Y_2\left(\frac{\phi\left(\sin\left(2\left(\alpha-\beta\right)\right)-\sin\left(2\left(\alpha+\beta\right)\right)\right)}{2}+\theta\cos\left(2\left(\alpha+\beta\right)\right)\right)\\
&+Z_1Z_2(4\phi\sin^2\left(\alpha\right)\sin^2\left(\beta\right)-2\phi\sin^2\left(\alpha\right)-2\phi\sin^2\left(\beta\right)+\phi\\
&-4\theta\sin\left(\alpha\right)\sin\left(\beta\right)\sin\left(\alpha+\beta\right)+2\theta\sin\left(\alpha\right)\cos\left(\alpha\right)+2\theta\sin\left(\beta\right)\cos\left(\beta\right))
\end{aligned}$$

which can be quickly simplified when we focus on the case where the coefficients of Z_1Y_2 and Y_1Z_2 are 0 (the right hand side of Eq. (3.20)). In this case, $\alpha = \beta$ and $\tan(4\alpha) = 2\theta/\phi$, and we obtain the following expression:

$$U_{12}(\phi,\theta) = e^{-i\alpha(X_1+X_2)}e^{i(A(\phi,\theta)Y_1Y_2+B(\phi,\theta)Z_1Z_2)}e^{i\alpha(X_1+X_2)}. \tag{3.23}$$

And thus:

$$A(\phi,\theta) = \left(\frac{\phi\left(1-\cos\left(4\alpha\right)\right)}{2} - \theta\sin\left(4\alpha\right)\right)$$
$$B(\phi,\theta) = \left(4\phi\sin^4\left(\alpha\right) - 4\phi\sin^2\left(\alpha\right) + \phi - 4\theta\sin^2\left(\alpha\right)\sin\left(2\alpha\right) + 4\theta\sin\left(\alpha\right)\cos\left(\alpha\right)\right)$$
$$\alpha = \arctan(2\theta/\phi)/4$$

This observation is useful to us because we can then invoke a known prescription for implementing gates of the form e^{XX+YY} using only 2 CNOT gates from Ref. [204]. This can be obtained from Eq. (3.23) with an additional layer of single qubit Y rotations

$$\begin{aligned} & e^{i(A(\phi,\theta)Z_1Z_2+B(\phi,\theta)Y_1Y_2)} \\ &= e^{-i\frac{\pi}{4}(Y_1+Y_2)}e^{i(A(\phi,\theta)X_1X_2+B(\phi,\theta)Y_1Y_2)}e^{i\frac{\pi}{4}(Y_1+Y_2)} \end{aligned} \tag{3.24}$$

Putting it altogether, the 2-qubit gate $U_{12}(\phi,\theta)$ can be implemented with single qubit rotations and 2 CNOT gates only via:

$$U_{12}(\phi,\theta) = e^{-i\alpha(X_1+X_2)}e^{-i\frac{\pi}{4}(Y_1+Y_2)}e^{i(A(\phi,\theta)X_1X_2+B(\phi,\theta)Y_1Y_2)}e^{i\frac{\pi}{4}(Y_1+Y_2)}e^{i\alpha(X_1+X_2)} \tag{3.25}$$

This machinery, illustrated in the circuit diagram of Fig. 3.2, is incorporated into the ansatz

$$U^{\text{CDRecomp}}(\vec{\alpha},\vec{\phi},\vec{\varphi},\vec{\theta},\vec{\vartheta}) = \prod_{j}^{N=1,2,3} U(\alpha_j,\phi_j,\varphi_j,\theta_j,\vartheta_j), \tag{3.26}$$

where

$$U(\alpha_j,\phi_j,\varphi_j,\theta_j,\vartheta_j) = \prod_{k=1}^{n} e^{i\alpha_j X_k} \prod_{k=1}^{n/2} U_{2k-1,2k}(\phi_j,\theta_j) \prod_{k=1}^{n/2} U_{2k,2k+1}(\varphi_j,\vartheta_j). \tag{3.27}$$

See Fig. 3.3 for an illustration of a single layer of this ansatz.

During my investigations, we also asked the question 'what if additional parameterisation of the ansatz was structured, but not physically informed?' By this we mean that instead of adding the gauge potential to each layer, an arbitrary operator is added instead. I tested the case where the $\exp(\phi ZZ + \theta(ZY + YZ))$

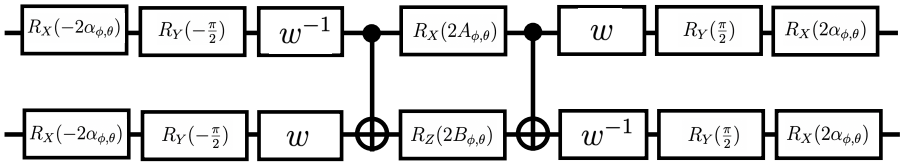


Figure 3.2: Compilation of the $U_{12}(\phi, \theta)$ unitary into a 2-qubit gate with 2 CNOTs and single qubit rotations. In this case $w = (I - iX)/\sqrt{2}$. See the main text for $\alpha_{\phi,\theta}, A_{\phi,\theta}, B_{\phi,\theta}$. Original graphic created by me.

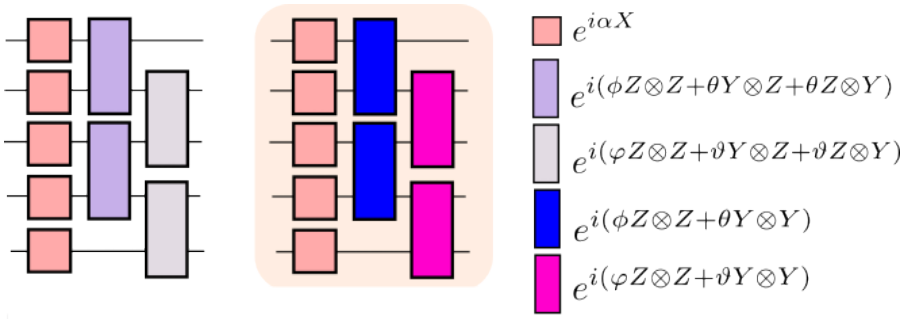


Figure 3.3: The recompiled counter-adiabatic driving ansatz U^{CDRecomp} for the Transverse Field Ising Model and the U^{ZYRecomp} ansatz where the gauge potential term is replaced with the YY operator. See the main text for the circuits implementing the 2-qubit gates. Original graphic created by me.

operator was replaced by the $\exp(\phi ZZ + \theta YY)$ operator (implemented at a circuit level using Eq. (3.24)), which has no discernible physical basis:

$$U^{\text{ZYRecomp}}(\vec{\alpha}, \vec{\phi}, \vec{\varphi}, \vec{\theta}, \vec{\vartheta}) = \prod_{j}^{N=1,2,3} U(\alpha_j, \phi_j, \varphi_j, \theta_j, \vartheta_j), \tag{3.28}$$

where

$$U(\alpha_j, \phi_j, \varphi_j, \theta_j, \vartheta_j) = \prod_{k=1}^{n} e^{i\alpha_j X_k} \prod_{k=1}^{n/2} e^{i(\phi_j Z_{2k-1}Z_{2k} + \theta_i Y_{2k-1}Y_{2k})} \prod_{k=1}^{n/2} e^{i(\varphi_j Z_{2k}Z_{2k+1} + \vartheta_j Y_{2k}Y_{2k+1})}. \tag{3.29}$$

See Fig. 3.3 for an illustration of the ansatz circuit layer.

3.2.3 Numerical Experiments

All of the experiments reported here were obtained from numerical emulations of ideal qubits written in `Julia`, using the closed-source packages of Phasecraft Limited to construct the ansatz circuits. Where I implemented custom unitary gates, specifically the $U_{12}(\phi, \theta)$ gate, I used the Yao [205] package to specify the unitary which is then fed into the layers of the ansatz circuit for simulation. I used the

NLopt [206] and Optim [207] packages to implement the L-BFGS algorithm [208] for optimisation of the ansatz parameters; their initial values were set to 0.1 for all cases. I tested the various ansätze on Ising models with $n = 5, 11, 17$ sites, across a range of λ values between 0 and 1. The simulations used either the $|0\rangle$ state, the $|+\rangle$ state (a Hadamard gate applied to each qubit) or random computational basis states.

Trotterised counter-adiabatic driving ansatz

Across the system sizes and λ values tested, the energy convergence of U^{CDTrot} was generally better than that of the conventional ansatz U^{Ham}. The ground state estimate obtained from e.g. only two ansatz layers of the former can be as good as, and occasionally better than, three layers of the latter. This was the case for when the initial state was either the ground state of $\mathcal{H}_0$ or $\mathcal{H}_1$ (the $|+\rangle$ or $|0\rangle$ state respectively). This is not unexpected – having an extra parameter in the circuit provides greater flexibility for the variational optimisation, allowing the ansatz to get closer to the ground state. Nevertheless, the additional counter-adiabatic driving operators rendered deeper circuits and more variational parameters – leading to more optimisation steps for convergence on average (the L-BFGS algorithm was implemented using NLopt). This is summarised in Fig. 3.4.

A clear, expected pattern emerged when we superimpose the results derived from the $|0\rangle$ initial state and those from the $|+\rangle$ state. The converged energy curves cross and form a 'cusp' in the region where contributions of $\mathcal{H}_0$ and $\mathcal{H}_1$ are roughly equal. It would be prudent therefore, in practice, to use an informed choice of initial state that depends on system properties like λ. This is in line with the general wisdom that exploiting the structure of $\mathcal{H}$ and heuristics to pick the right initial state can bias the ansatz towards the region of the Hilbert space which contains the desired eigenstate, and result in much better performance than using a sophisticated ansatz on a sub-optimal initial state.

Recompiled ansatz

I then implemented the $U_{12}(\phi, \theta)$ gates and used them to construct the U^{CDRecomp}. To speed up the optimisation, packages such as NLopt require inputs of the

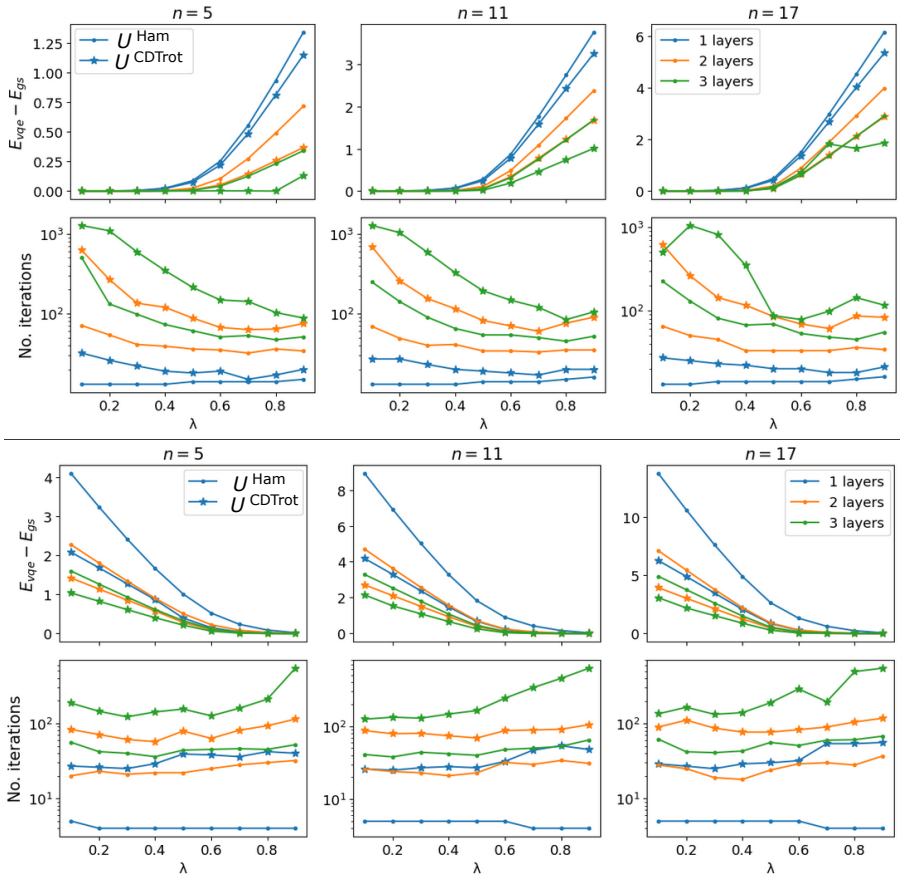


Figure 3.4: Comparing the VQE performance of the Hamiltonian ansatz with the Trotterised counter-adiabatic ansatz for the Transverse Field Ising Model. The top two rows represent where the initial state was $|+\rangle$, and the bottom two rows represent where the initial state was $|0\rangle$. The shape of the markers represent the ansatz type, and colours represent the number of ansatz layers. Original graphic created by me.

parameterised unitary as well as natural gradients of those unitaries with respect to their parameters i.e. $\frac{\partial U_{12}}{\partial \phi}$ and $\frac{\partial U_{12}}{\partial \theta}$. However, while the unitary matrix was trivial to implement, because of the fact that there are two distinct parameters, coding the natural gradients turned out to be more involved; due to time constraints, this was eventually put to the side and results were generated using a different optimiser package Optim that performed L-BFGS without needing natural gradients.

When compared against the U^{CDTrot} ansatz, the U^{CDRecomp} ansatz performed just as well, if not slightly better, in the energy convergence. The number of iterations necessary to reach energy convergence is also comparable between them. In some scenarios, using 2 layers of the U^{CDRecomp} ansatz outperforms 3 layers of the U^{Ham} ansatz – which reflected 4 fewer CNOT gates and 1 fewer ansatz

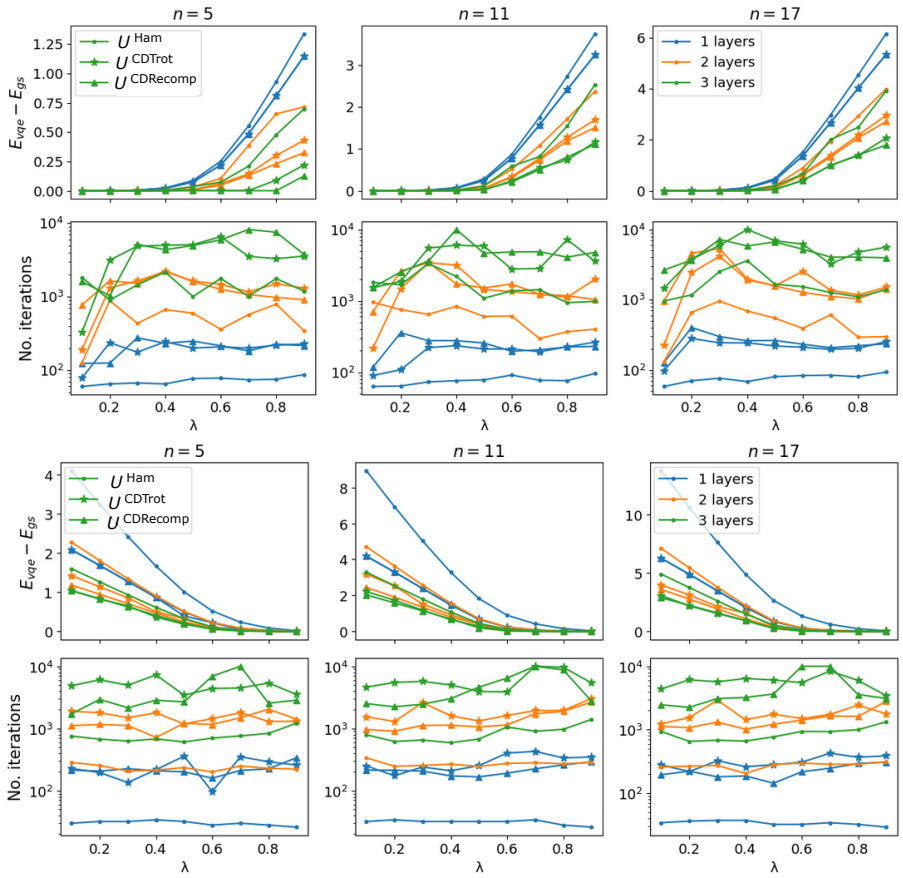


Figure 3.5: Comparing the VQE performance of the recompiled counter-adiabatic ansatz with other ansätze for the Transverse Field Ising Model. The top two rows represent where the initial state was $|+\rangle$, and the bottom two rows represent where the initial state was $|0\rangle$. The shape of the markers represent the ansatz type, and colours represent the number of ansatz layers. Original graphic created by me.

parameter in the circuit in total. Overall, it does look like the addition of the gauge potential operator is at least consistently helpful due to greater expressiveness of the ansatz, and using the U^{CDRecomp} gives the added benefit of a more frugal CNOT gate count. This is summarised in Fig. 3.5.

Is Counter-adiabatic Driving Responsible?

To verify that the ansatz is actually responsible for the performance gains, I also applied U^{Ham} and U^{CDRecomp} to four arbitrary 11-qubit initial states that are not physically motivated. As expected, the energy converged is significantly worse than if the proper initial state was used, but the key point illustrated is that using N layers of the U^{CDRecomp} was essentially always as good as, and at times better than,

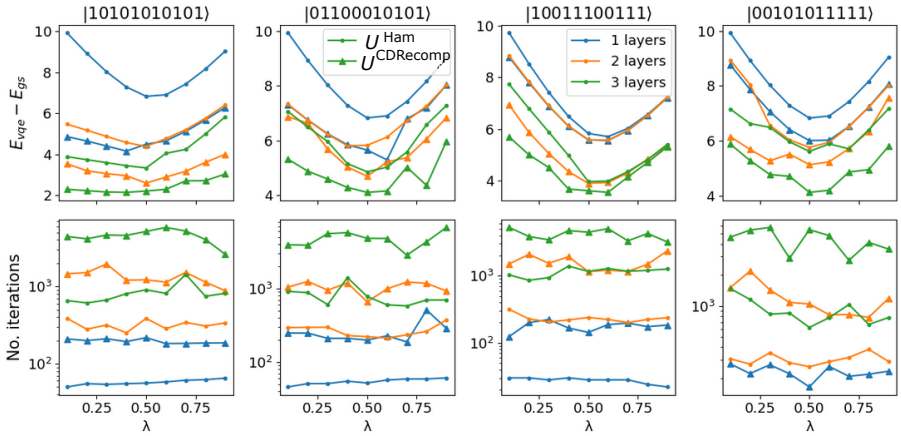


Figure 3.6: Performing VQE using the Hamiltonian ansatz or the recompiled counter-adiabatic ansatz for the 11-qubit Transverse Field Ising Model, applied to a collection of random initial states. The shape of the markers represent the ansatz type, and colours represent the number of ansatz layers. Original graphic created by me.

using $N + 1$ layers of the U^{Ham} ansatz (see Fig. 3.6).

I also further verified that the counter-adiabatic driving is responsible by replacing the approximate gauge potential with a Pauli YY operator instead (see Eq. (3.24)) in U^{ZYRecomp}. While there are scenarios where U^{ZYRecomp} behaves worse than the U^{Ham} ansatz, in general it is only a little less performant than U^{CDRecomp}. One interpretation of these results is that, while additional operators in a Hamiltonian ansatz introduces parameters that allow the VQE optimisation to reach better energies, choosing to do so via the counter-adiabatic gauge potential is a reasonably principled approach to determining which operators to add, and can on the whole be better than adding any arbitrary operator (see Fig. 3.7).

3.2.4 Beyond the TFIM

To explore if the addition of the gauge potential operator remains relevant beyond the TFIM, I performed simulations of the 17-qubit TFIM+F model – which includes an extra field along Z and represents the simplest extension of the TFIM:

$$\mathcal{H}_\theta(\lambda) = (1-\lambda)\sum_i(\cos(\theta)X_i + \sin(\theta)Z_i) + \lambda\sum_i Z_iZ_{i+1} \tag{3.30}$$

where in this case θ determines the ratio between the field along X and along Z. In this case, the first-order approximation of the gauge potential is identical to

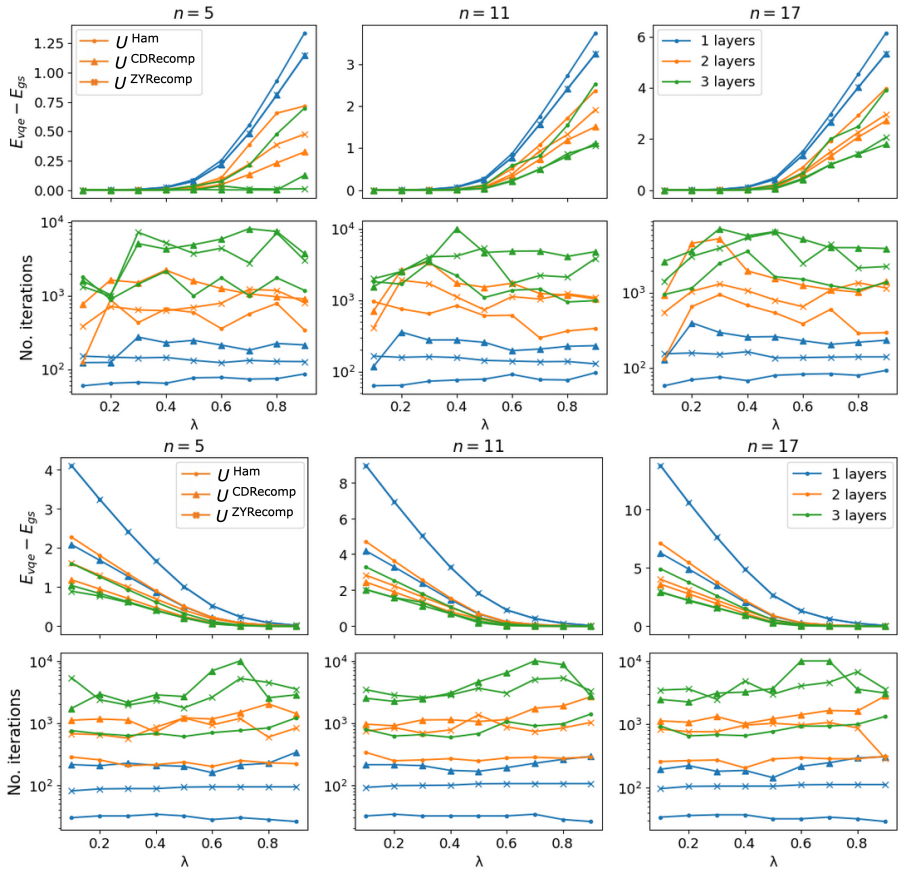


Figure 3.7: Comparing the Hamiltonian ansatz, the recompiled counter-adiabatic ansatz and an ansatz where the gauge potential operator is replaced with the YY operator for VQE. The shape of the markers represent the ansatz type, and colours represent the number of ansatz layers. Original graphic created by me.

the TFIM case. I tested the cases where $\theta \in \{\frac{\pi}{8}, \frac{\pi}{4}, \frac{3\pi}{8}\}$. The initial states tested were either $|0\rangle$ (the ground state of $\sum_i Z_i Z_{i+1}$) or $R_Y(\theta/2)\,|+\rangle$ (the ground state of $\sum_i(\cos(\theta)X_i + \sin(\theta)Z_i)$) on every qubit.

For the case where the initial state was $|0\rangle$, the U^{CDRecomp} comfortably outperformed the U^{Ham} ansatz, but was comparable with U^{ZYRecomp} beyond 1 layer. It also seems to be an adequate initial state across the entire range of λ values tested, performing well in both extremes of $\lambda \to 0$ and $\lambda \to 1$. In the cases where initial states formed of $R_Y(\theta/2)\,|+\rangle$ applied to every qubit were used, there are some regimes ($\lambda \sim 0.6$) where a single layer of U^{ZYRecomp} ansatz actually outperforms a single layer of U^{CDRecomp} ansatz, but overall these two had similar performance. This is summarised in Fig. 3.8. We can be therefore reasonably confident that

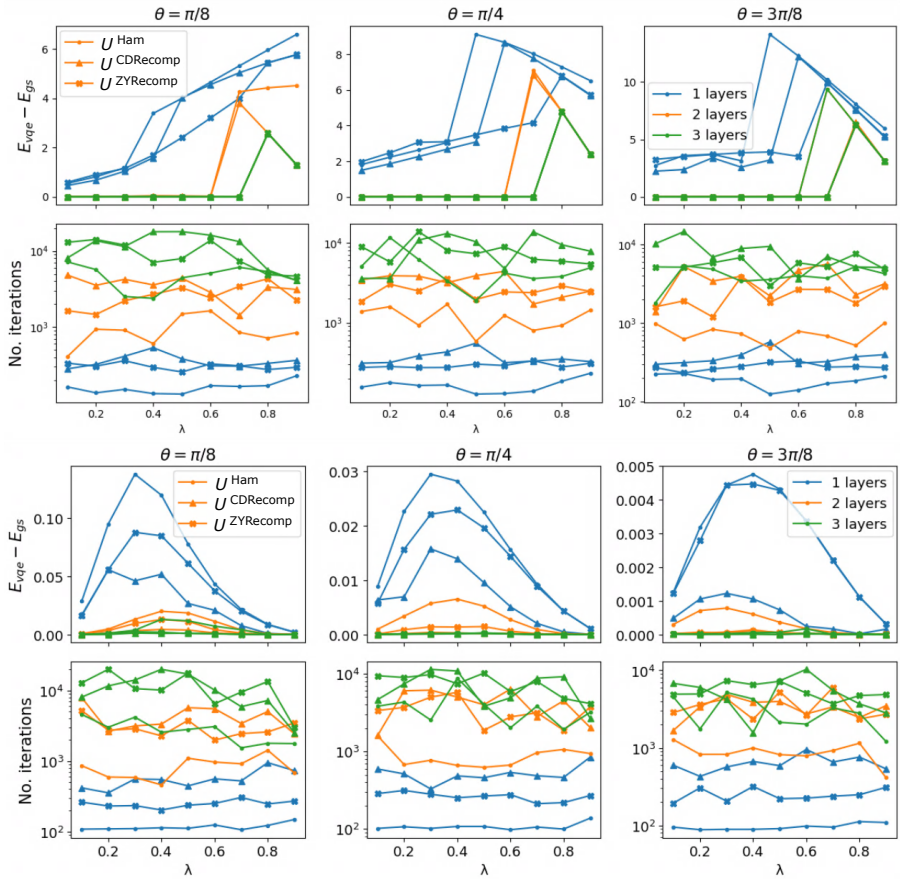


Figure 3.8: Comparing the Hamiltonian ansatz with the recompiled counter-adiabatic ansatz for the TFIM+F system. The top two rows represent where the initial state was $|+\rangle$, and the bottom two rows represent where the initial state was $|0\rangle$. The shape of the markers represent the ansatz type, and colours represent the number of ansatz layers. Original graphic created by me.

adding the first-order approximation of the gauge potential to any Hamiltonian ansatz layers can, in general, improve its energy convergence properties.

3.2.5 Discussion

The brief research described in this Section explored the synergies between adiabatic state preparation and VQE. The results suggest that adding (first order approximations of) the counter-adiabatic driving gauge potential to each ansatz layer can improve the ansatz's energy convergence performance, but this increases the number of parameters and more optimisation steps (thus the number of shots) are needed to reach convergence. Where possible recompilation techniques were used to reduce the number of CNOT gates. Additional experiments were performed

to determine if the counter-adiabatic driving operators the optimal operators to add. These results support the notion that, within the VQE paradigm, some of the most dominant factors determining how accurately an ansatz circuit can prepare ground states are (1) how much overlap the initial state has with the ground state and (2) how many parameters the ansatz has, which increases its expressiveness. That the gauge potential is the best possible operator to append to an ansatz is not clear cut – however, there is a clear physical motivation and reason behind the choice of the gauge potential.

To further explore whether ideas from adiabatic state preparation can help with the construction and optimisation of Hamiltonian ansätze, it might be worthwhile to study the use of higher-order approximations of the gauge potential. Understanding if the optimisation of a Hamiltonian ansatz can be made to reproduce the adiabatic trajectory may also provide performance gains and insights into why adiabatic gauge potential terms might help the Hamiltonian ansatz.

3.3 Symmetry-restricted Variational Ansätze for Excited States

Understanding how visible light interacts with even relatively small molecules is relevant in a great plethora of technological applications [110]. Simulations of the underlying photochemical processes often starts with calculating not only the ground electronic state but also excited electronic state energy profiles to the similar levels of accuracy. Conventional *ab initio* methods adequate for obtaining accurate excited state profiles are however computationally expensive even for these small molecule targets [209]. One might therefore conclude that it will be possible to use more limited quantum computers to demonstrate genuine quantum advantage on a small molecule with an excited-state computation task, as opposed to the ground state computation of larger systems for which conventional, classical approaches can in many cases be very accurate.

Variational Quantum Deflation (VQD) [210] is an extension of VQE designed to prepare excited states – it does so by first completing a VQE step for obtaining

the ground state, then performs a VQE calculation on a new ansatz circuit where the overlap between the new ansatz and the ground state ansatz is simultaneously minimised. In order to use the same ansatz circuit structure for both the ground and excited states, while accommodating the difference between them, the authors of Ref. [210] proposed using fixed-length UCC ansätze that use both physical and non-physical excitation operators (referred to as generalised excitations), but these will map to deep circuits. Furthermore, the prior art does not make use of symmetry of the ground and excited states.

Efforts in reducing the quantum resource cost, in particular the circuit depth, of wavefunction circuit ansätze for ground state VQE can be applicable in making excited state wavefunction ansätze more resource-frugal as well in VQD. One such innovation transfer was applying the ADAPT-VQE technique to VQD, which my collaborators and I reported in Ref. [2]. My collaborators and I suggested that using a greedy approach to 'grow' the unique set of operators needed to encode each eigenstate (as per the ADAPT-VQE method) can reduce the depth of circuits needed.

While I developed many of the ideas for ADAPT-VQD before my DPhil, I built upon them and added spin-adaptation which improved its performance for computing excited state energy profiles on NISQ devices. For the remainder of this chapter, I present this spin-restriction strategy that I created and developed with my collaborators, with results that suggest the strategy is useful for enforcing point group symmetries as well.

3.3.1 Methods

Modern quantum chemistry methods have long made use of symmetry to simplify computations. For example, in a photochemistry process, we might only be interested in excited states of a specific spin symmetry. Prior to my DPhil, my ADAPT-VQD results, and all other variational-deflation-based excited state quantum algorithms as far as we knew, would simply find the next highest energy excited state indiscriminately in a step-wise fashion without accounting for whether or not a given state has a particular spin symmetry.

Recall the quantum circuit wavefunction ansatz structure from Eq. (3.1) i.e. an easily prepared initial state $|\psi_0\rangle$ with a parameterised quantum circuit applied onto it $U(\vec{\theta})$. If one was to impose restrictions on the circuit such that properties like the spin of the initial state was preserved, it would be possible to confine the wavefunction Hilbert space explored by the ansatz during VQE optimisation to a smaller spin-restricted subspace, thereby exercising greater control over which eigenstate the ansatz converges to and, hopefully, reducing the number of iterations it takes to converge. We saw a similar type of ansatz biasing in action in the previous Section when initial states informed by the relative contributions of $\mathcal{H}_0$ and $\mathcal{H}_1$ in a bipartite Hamiltonian perform better; in the present case, we account for Fermionic spin symmetry and attempt to preserve it within the ansatz structure.

The initial reference state $|\psi_0\rangle$ defines the starting symmetry of the ansatz. Consider a molecule in the second quantised representation, and n denotes the number of spin orbitals in the basis set used to represent the system and m is the number of electrons in the molecule. For a singlet spin state, the closed-shell Hartree-Fock reference state $|\psi_{\text{HF}}\rangle$ (corresponding to the computational basis state that is the n-digit binary representation of $2^m - 1$) would be a suitable singlet initial state. The open-shell triplet reference $|\psi_{T_1}\rangle$ representing the orbtial-occupation of the lowest energy triplet state would be a suitable triplet initial state.

The subsequent parameterised ansatz circuit can be constructed to preserve the spin symmetry set out by the reference. Within the paradigm of UCC-type ansätze, we can build the ansatz out of excitation operators that respect the symmetry of the spin preference; this would include, for example, excitation operators which correspond to electrons hopping only from $\alpha \to \alpha$ or $\beta \to \beta$ spin orbitals. These ansätze would look like:

$$\left|\Psi^S(\vec{\theta})\right\rangle = \prod_k e^{\theta_k \hat{A}_k^S} |\psi_{\text{HF}}\rangle \tag{3.31}$$

$$\left|\Psi^T(\vec{\theta})\right\rangle = \prod_k e^{\theta_k \hat{A}_k^S} |\psi_{T_1}\rangle \tag{3.32}$$

where $\hat{A}_k^S \in \{E_i^a - E_a^i, E_{ij}^{ab} - E_{ab}^{ij}\}$ are *spin-adapted* single- and double-excitation operators which consist of unitary group generators ($E_i^a = a_{i\alpha}^\dagger a_{a\alpha} + a_{i\beta}^\dagger a_{a\beta}$ and

$E_{ij}^{ab} = a_{i\alpha}^\dagger a_{j\alpha}^\dagger a_{b\alpha} a_{a\alpha} + a_{i\alpha}^\dagger a_{j\beta}^\dagger a_{b\beta} a_{a\alpha} + a_{i\beta}^\dagger a_{j\alpha}^\dagger a_{b\alpha} a_{a\beta} + a_{i\beta}^\dagger a_{j\beta}^\dagger a_{b\beta} a_{a\beta}$ where i, j, a, b are indices of molecular spatial orbitals). There are two sets of double excitation operators which maintain the S_2 symmetry [211]; here I include only one for brevity. This scheme can in principle be applied to any UCC-based ansatz.

In practice, the spin adapted operators need to be transformed into a representation that can be encoded as quantum circuit, such as the mapping of Fermionic operators to Pauli strings under the Jordan-Wigner (JW) transform. The resulting sequence of Pauli strings in an ansatz (or Pauli gadgets when implemented in a circuit) may no longer respect the spin-adaptation imposed by the Fermionic operators. For example, the Pauli strings resulting from the JW transform of some of the spin-adapted double excitations $E_{ij}^{ab} - E_{ab}^{ij}$ do not commute, and therefore the spin adaptation is broken. Some of the set of generalised single and double excitations are expected to not observe spin-restrictions when encoded into a quantum circuit form, whereas the set of all generalised single and paired-double excitations are expected to do so.

Accounting for these caveats, one could construct spin-adapted UCC-based ansätze for extracting excited state eigenstates from a protocol like VQD – but these circuits will likely quite intractable for NISQ hardware. In Ref. [2], I combined the gate resource efficiency provided by the ADAPT-VQE method with physically motivated spin-adaptation, to reduce the ansatz depth and the number of measurement shots that would've been needed if ADAPT-VQE was naïvely translated to VQD.

Recall that 'canonical' ADAPT-VQE iteratively grows an ansatz circuit by selecting operators $\hat{A}_m$ which are expected to recover the most significant amount of electron correlation from a predefined pool, and sequentially appending them onto the ansatz. The operator that is selected would have the largest absolute energy gradient

$$\frac{\partial E}{\partial \theta_m} = 2\mathcal{R} \left\langle \psi(\vec{\theta}) \right| \hat{A}_m \mathcal{H} \left| \psi(\vec{\theta}) \right\rangle \tag{3.33}$$

i.e. that, which when added, influences the energy the most. The circuit parameters are optimised each time a new operator is added. When a predetermined convergence criterion is met e.g. that the next candidate operator with the largest gradient is below a threshold ϵ, no new operators are added and the ansat parameters are optimised for the last time. I extended this prescription to capture excited states by including orthogonal constraints during the operator selection and ansatz optimisation steps. In VQD, the effective Hamiltonian of the j^{th} eigenstate with the orthogonal penalisation is:

$$\mathcal{H}_j = \mathcal{H} + \sum_{i=0}^{j-1} \lambda_i \left|\Phi_i\right\rangle \left\langle\Phi_i\right| \tag{3.34}$$

where $|\Phi_i\rangle$ is the i^{th} previously optimised lower energy state and λ_i is a constant that should be larger than the energy gap between neighbouring eigenstates. To make VQD adaptive, this effective Hamiltonian is substituted into the energy gradient with respect to candidate operators (Eq. (3.33)):

$$\begin{aligned} \frac{\partial E}{\partial \theta_m} &= 2\mathcal{R} \left\langle\psi(\vec{\theta})\right| \hat{A}_m^S \mathcal{H} \left|\psi(\vec{\theta})\right\rangle \\ &\quad +2\mathcal{R} \sum_{i=0}^{j-1} \lambda_i \left\langle\psi(\vec{\theta})\right| \hat{A}_m^S \left|\Phi_i\right\rangle \left\langle\Phi_i|\psi(\vec{\theta})\right\rangle \end{aligned} \tag{3.35}$$

where $|\psi(\vec{\theta})\rangle$ is the currently sought state, and $\hat{A}_m^S \in \{E_i^a - E_a^i, E_{ij}^{ab} - E_{ab}^{ij}\}$ is a candidate spin-adapted excitation operator in the pool. The energy gradient is thus separable into two parts; the energy gradient component used in the original ADAPT-VQE, and an overlap gradient component. Both components can be measured in parallel across multiple quantum processors for each candidate operator in the pool. The pool of candidate operators in this case is chosen to be spin-adapted operators, and the initial states from which each ADAPT-VQE step starts from is either a singlet or triplet initial state (see earlier text). While the pool of generalised excitation operators can be large, there are fewer of them in the spin-adapted set, which reduces the number of operator gradients that needs to be checked. Nevertheless the number of terms, thus gradients that need to be computed, would still grow as $\mathcal{O}(n^4)$. This spin-adapted ADAPT-VQD is illustrated in Fig. 3.9.

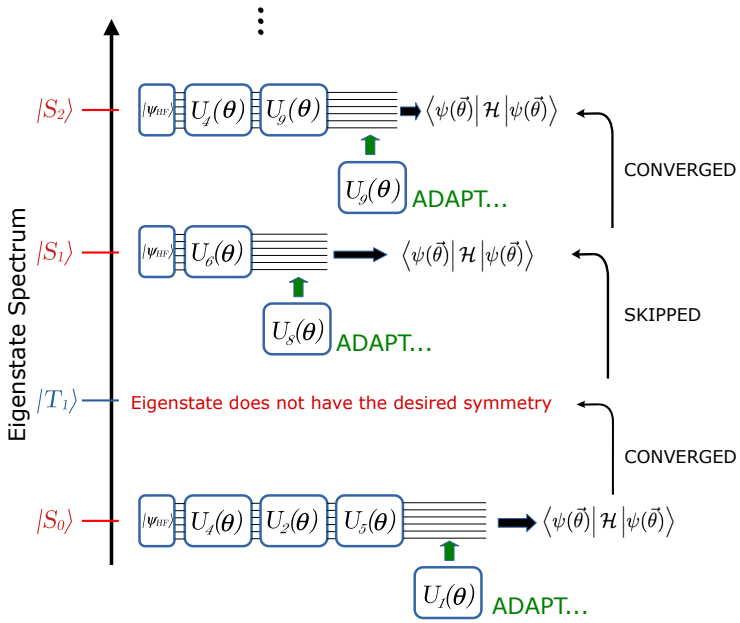
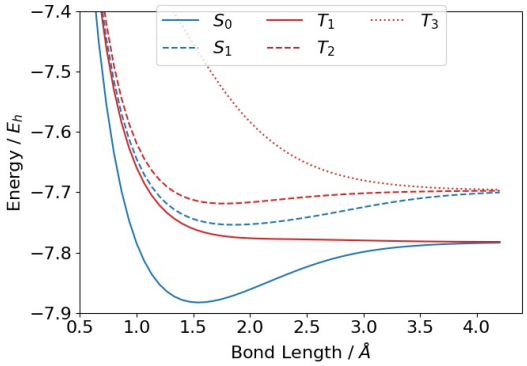


Figure 3.9: Illustration of the ADAPT-VQD approach with symmetry-restriction for computing excited states on NISQ computers. Spin-symmetry restrictions are imposed at the initial reference state level and also at the operator level – the operators that are used to grow the ansatz are spin-adapted single and double excitations. Figure adapted from [2] – I was its original creator.

3.3.2 Numerical Validation

In Ref. [2], I numerically emulated VQD for computing the lowest-lying excited state energies of a lithium hydride (LiH) molecule, using the spin-restricted k-UpCCGSD ansatz and spin-restricted ADAPT ansatz circuits to verify that the spin adaptation exercised control over VQD. In `python`, I used PySCF [212] and Libcint [213] for extracting the required one- and two-electron integrals and initialising the molecule, Open-Fermion[214] for creating the second quantised Hamiltonian and transformation of UCC excitation operators into Pauli gadgets, Pytket[215] for construction and compilation of ansatz circuits, and Qulacs [216] for GPU-accelerated circuit simulation. Additionally, scipy[217] was employed for variational minimisation of ansatz expectation energies. Although the Pytket library has the capacity to compile circuits for NISQ computers that are available today, the state of the art technology remains too noisy for demonstrating the proof-of-principle in this work. I therefore only emulated the quantum calculations on high performance conventional computers.



State	Irr. Rep.	Excitation
T_3	A_1	$\sigma \to p_z$
T_2	E_{1x}, E_{1y}	$\sigma \to p_x, p_y$
S_1	A_1	$\sigma \to \sigma^*$
T_1	A_1	$\sigma \to \sigma^*$
S_0	A_1	-

Figure 3.10: The exact electronic structure of the lowest-lying eigenstates along the Li-H bond length in the STO-3G basis. Red curves indicate triplet states and blue curves singlet states. The table on the right provides additional detail on their irreducible representation and the excitation that would generate the corresponding excited state. Figure reproduced from [2] – I was its original creator.

The minimal STO-3G basis was chosen to represent LiH. Although this is mostly a emulation memory constraint, the number of qubits available on NISQ machines will be $\mathcal{O}(10^2)$ [218], so it is conceivable that the first experimental demonstration of the methods related to this work will also be limited to small systems and basis sets. The Full Configuration Interaction (FCI) method, as implemented in PySCF[212], was used to calculate the exact eigenstate energies for Li-H bond lengths between 0.40Å and 4.20Å at 0.07Å intervals (see Fig. 3.10). The ordering of the energy states were in agreement with previous measurements from Ref. [219], although the values reflect a different basis set used.

I chose only a small number of representative geometries to benchmark the VQD methods due to the high emulation cost. For the spin-restricted k-UpCCGSD ansatz, calculations of the lowest two singlet and triplet eigenstates were simulated, at bond lengths of 0.91Å, 1.00Å, 1.80Å, 2.02Å and 2.98Å. For the spin-restricted ADAPT-VQD, we could afford to calculate eigenstates using ansätze grown at nine bond lengths between 1.00Å and 4.2Å at 0.40Å intervals. The system has 6 α- and 6 β-spin orbitals, so 12-qubit circuits were simulated for all cases. At such a small basis for such a small system, the emulated UCC-type ansätze are expected to yield energies close to the FCI. The value of $\lambda = 3.0$ from Eq. (3.34) for all VQD calculations. For ADAPT-VQD, the operator norm convergence threshold was set

to be $\epsilon = 0.01$, which was deemed adequate for small molecules in similar previous studies [95, 220]. The pool of spin-adapted, generalised single and paired-double excitations (operators that would be used to generate the UpCCGSD ansatz) were used. Operators are not exhausted and the same operator can in principle be added to an ansatz multiple times. Pytket's recompilation tools[221, 222] were used to reduce the circuit depth in every ansatz construction. L-BFGS-B was the optimiser of choice for the variational minimisation of the ansatz parameters.

To verify that the final optimised ansätze obey the applied spin restrictions, their S^2 spin expectation values were calculated in the same manner as the energy expectation calculations. I used the non-parallelity error (NPE) to characterise the consistency of each computed eigenstate energy, a metric that has been employed in the variational quantum algorithm literature for comparing performance of different ansätze [90, 93, 215]. The NPE is the difference between the maximum and minimum error relative to FCI across all computed geometries of a given eigenstate:

$$\text{NPE} = \max(E_{\text{VQE}} - E_{\text{FCI}}) - \min(E_{\text{VQE}} - E_{\text{FCI}}) \tag{3.36}$$

Spin-Restricted k-UpCCGSD-VQD

In this investigation the 1-UpCCGSD and 2-UpCCGSD ansätze were tested. The 1-UpCCGSD ansatz had 30 excitation operators, corresponding to 180 Pauli gadgets and a total 470-gate deep state preparation circuit. The simulations were not able to recover any of the excited eigenstates, likely due to the reduced expressiveness of the ansatz.

The ansatz at $k = 2$ doubled the resources at 60 excitation operators, 360 Pauli gadgets, and a circuit depth of 941 gates and CX gate depths of 580. Fig. 3.11 shows that use of the spin adaptation was successful in recovering singlet and triplet eigenstates separately. The errors of the first singlet and triplet excited states were within chemical accuracy consistently across all computed bond lengths at both S_1 and T_2, reporting NPE of 0.21 and 0.18 respectively (Table 3.1). The S^2 values were 0.00 for all singlet states and 2.00 for all triplet states, confirming successful

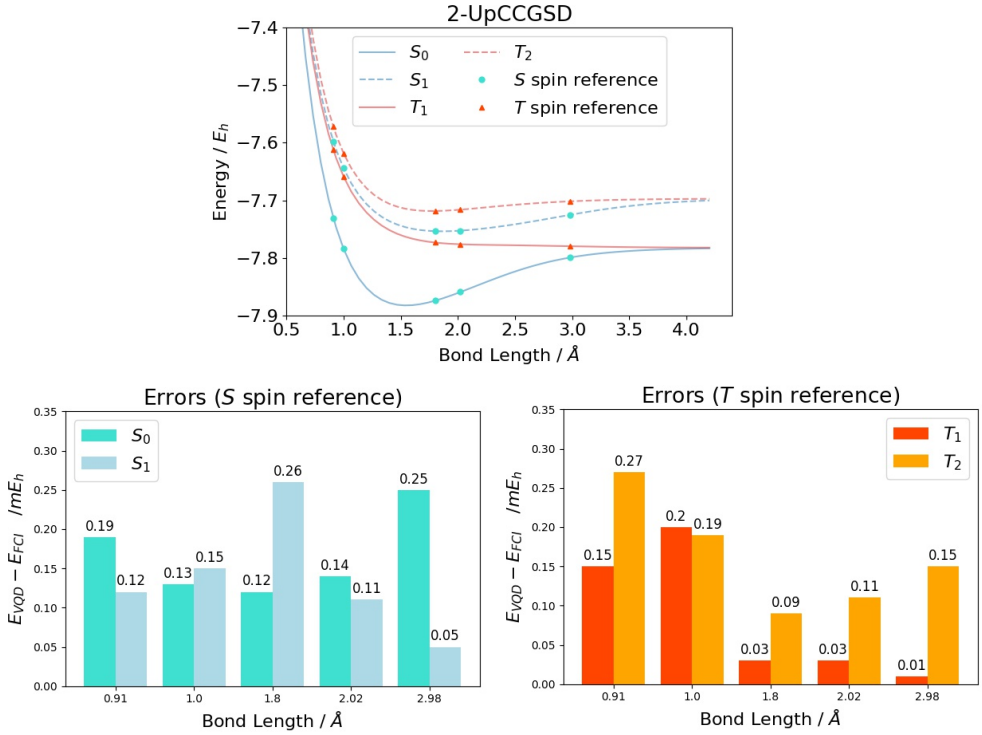


Figure 3.11: Quantum simulations of the singlet and triplet electronic excited states of LiH along its bond length using the 2-UpCCGSD ansatz. The top plot is the potential energy profile (in Hartree), where solid lines indicate FCI calculations and the markers refer to simulated 2-UpCCGSD values. The bottom plots show energy errors (in milliHartree) of singlet (left) and triplet (right) VQD energy calculations. Figure reproduced from [2] – I was its original creator.

States	0.91		1.00		1.80		2.02		2.98		NPE /mE_h
	FCI	VQD	FCI	VQD	FCI	VQD	FCI	VQD	FCI	VQD	
T_2	-7.57227	-7.57200	-7.61856	-7.61837	-7.71859	-7.71850	-7.71629	-7.71618	-7.70177	-7.70162	0.18
S_1	-7.59780	-7.59767	-7.64450	-7.64435	-7.75366	-7.75341	-7.75275	-7.75264	-7.72531	-7.72526	0.21
T_1	-7.61242	-7.61227	-7.65893	-7.65873	-7.77343	-7.77340	-7.77637	-7.77634	-7.77980	-7.77979	0.19
S_0	-7.73064	-7.73044	-7.78446	-7.78433	-7.87452	-7.87440	-7.85959	-7.85945	-7.79952	-7.79927	0.11

Table 3.1: The values of the 2-UpCCGSD-VQD energies at the computed bond lengths (in Hartree), and the NPE of recovered eigenstates (in milliHartree). Table reproduced from [2] – I was its original creator.

application of spin-restriction. On average, the number of optimisation steps for reaching convergence was 200 for 2-UpCCGSD-VQD.

Spin-Restricted ADAPT-VQD

With the singlet spin reference, the singlet ground state was first recovered as per the original ADAPT-VQE prescription. The spin references continue to be successful in isolating the relevant Hilbert space, and the next adaptively grown

ansatz converged to the first singlet excited state S_1, successfully avoiding the T_1 state. The energy was in good agreement with the exact solution, lying within 1.59 mE_h across all bond lengths in both the S_0 and S_1 states (see Fig. 3.12). The S_0 and S_1 states used, on average across calculated bond lengths, 10.9 and 14.1 excitation operators, culminating to ansatz circuit depths of 228.8 and 280.2 gates respectively. This was a 3 fold reduction over the 2-UpCCGSD circuit depth. For the S_0 ground state, the errors are similar to those from the 2-UpCCGSD-VQD calculation. For the S_1 excited state, the errors increase at shorter bond lengths, giving a NPE of 0.97 mE_h, but still within chemical accuracy. Like the 2-UpCCGSD method, the S^2 values for both states were 0.00. Note that the errors reported for the ADAPT results are specific to the operator norm convergence threshold of $\epsilon = 0.01$. We expect that as the operator norm threshold is tightened the error should reduce, but this will increase the circuit depth. The triplet spin reference was also able to adaptively recover the triplet ground state T_1. Upon convergence however, the next excited eigenstate selected by the spin restricted ADAPT-VQD was not the next highest in energy (the doubly degenerate T_2 state), but the T_3 state above, a behaviour not seen in the other ansatz which recovered the T_2 state using the same initial reference circuit.

The states recovered with the ADAPT method share the same irreducible representation A_1, whereas the doubly degenerate T_2 state which the method missed had an irreducible representation of E_1. In the presence of molecular point group symmetry, due to the active selection of excitation operators in ADAPT, the eigenstates that an initial reference state can actually express are limited to those constructed from excitations that are symmetry-allowed from the reference; an excitation is symmetry-allowed if the direct product of the irreducible representations of the molecular orbitals involved i, j, a, and b ($\Gamma_i \otimes \Gamma_j \otimes \Gamma_a \otimes \Gamma_b$) contains the irreducible representation of the point group of the desired exited state [90, 223, 224]. Consequently, addition of only symmetry-allowed single excitation operators (which can be thought of as rotations of the reference state in Hilbert space) in

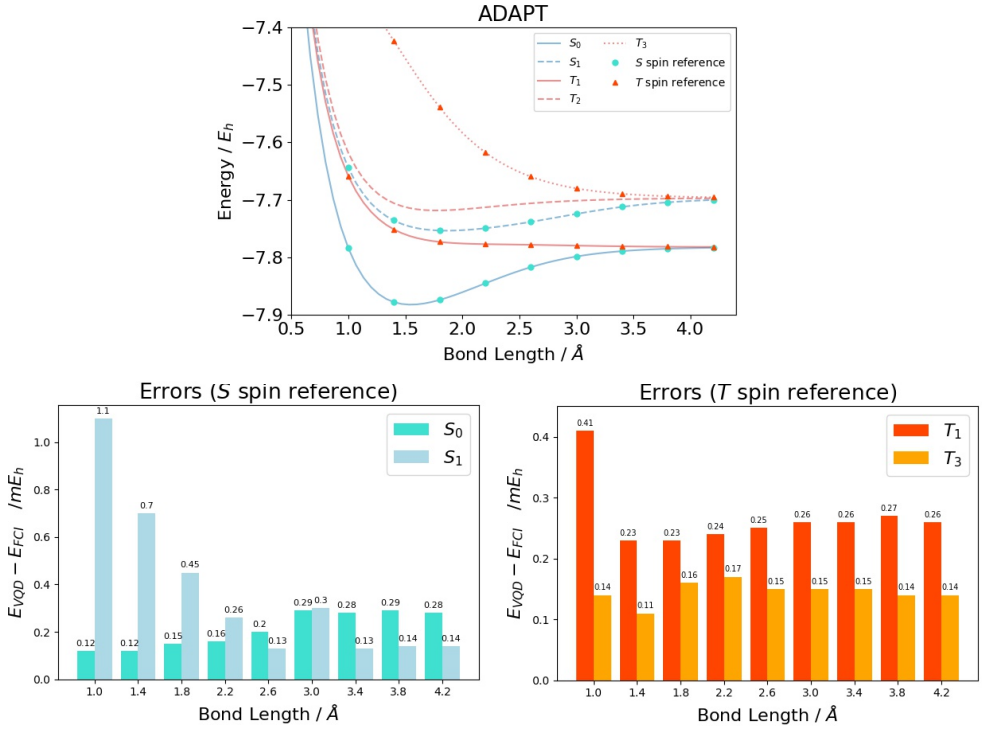


Figure 3.12: Quantum simulations of the singlet and triplet excited states of LiH at 5 bond lengths using adaptively grown ansätze, unique for each geometry and each eigenstate. The top plot shows the potential energy profile (in Hartree), where solid lines indicate FCI calculations and the markers refer to simulated values. Although the singlet spin restriction recovered both the lowest lying singlet states, the triplet spin reference only recovered the T_1 and T_3 states, missing the doubly degenerate T_2 state. The bottom plots show the errors associated with each VQD energy calculation. Energy errors (in milliHartree) of singlet (left) and triplet (right) VQD energy calculations. Figure reproduced from [2] – I was its original creator.

the presented demonstration rotates the state such that it remains in the totally symmetric manifold. Hence, the resulting ansätze were only able to recover A_1 states.

In contrast, the preceding ansatz seems to have sufficient flexibility to effectively mix occupied orbitals with virtual orbitals to obtain a different symmetry reference, which is then correlated. The greater degree of control over states of which symmetry is recovered in ADAPT-VQD is an advantageous feature over the other. Moving forward, these symmetry observations must be extended to the irreducible representations of the direct product group $\Gamma \otimes SU(2)$, which simultaneously describes point group and spin symmetry.

Apart from the convergence to the T_3 state, other performance metrics of the

States	1.00		1.40		1.80		2.20		2.60	
	FCI	VQD	FCI	VQD	FCI	VQD	FCI	VQD	FCI	VQD
T_3	-7.29798	-7.29784	-7.42335	-7.42323	-7.53882	-7.53866	-7.61737	-7.61720	-7.65974	-7.65959
S_1	-7.64450	-7.64340	-7.73624	-7.73554	-7.75366	-7.75321	-7.74960	-7.74934	-7.73847	-7.73834
T_1	-7.65893	-7.65852	-7.75180	-7.75156	-7.77343	-7.77321	-7.77729	-7.77705	-7.77845	-7.77820
S_0	-7.78446	-7.78434	-7.87845	-7.87833	-7.87452	-7.87438	-7.84568	-7.84552	-7.81740	-7.81720

States	3.00		3.40		3.80		4.20		NPE /mE_h
	FCI	VQD	FCI	VQD	FCI	VQD	FCI	VQD	
T_3	-7.68023	-7.68008	-7.68962	-7.68948	-7.69383	-7.69369	-7.69570	-7.69556	0.06
S_1	-7.72461	-7.72431	-7.71235	-7.71222	-7.70444	-7.70430	-7.70032	-7.70018	0.97
T_1	-7.77987	-7.77961	-7.78112	-7.78086	-7.78187	-7.78160	-7.78222	-7.78195	0.18
S_0	-7.79884	-7.79855	-7.78950	-7.78922	-7.78535	-7.78505	-7.78359	-7.78331	0.17

Table 3.2: The values of the ADAPT-VQD energies at the computed bond lengths (in Hartree), and the NPE of recovered eigenstates (in milliHartree). Tables reproduced from [2] – I was their original creator.

triplet calculations were comparable with the singlet calculations. Fig. 3.12 shows that the triplet energy accuracy was very even across all computed geometries. The ansatz depth for calculating eigenstates in the triplet manifold even benefited from further reductions to the circuits for adaptive singlet states, circuits of only 192.2 and 213.4 gates deep on average for T_1 and T_3 before convergence. The S^2 value was 2.00 for every triplet calculation.

One important consideration for ADAPT-VQD, which was not directly studied here, is the circuit measurement overhead for operator selection. The reductions in circuit depth for ADAPT come at the cost of more measurements at each ADAPT iteration for operator selection. Nonetheless, the significant reduction in circuit depth for encoding excited state wavefunctions with the ADAPT-VQD method would suggest that it can be more practical for ameliorating decoherence and circuit noise in the NISQ regime, whereby the measurement overhead can be mitigated using, for example, parallel gradient computation. Table 3.3 summarise the quantum circuit resources in the methods investigated with spin adaptation.

3.3.3 Discussion

Using quantum simulations on classical computers, we suggest that ADAPT-VQD will be more practical for calculating molecular excited states on near-term NISQ computers than other VQD implementations that use fixed-length ansätze such as UCCGSD and k-UpCCGSD, because it is able to prepare excited state wavefunctions on quantum computers with significantly reduced circuit depths than the other

LiH (STO-3G)	Operators	Circuit depth	CX depth
singlet 2-UpCCGSD (fixed)	60	941	580
S_0 ADAPT (average)	10.9	227.0	149.2
S_1 ADAPT (average)	14.1	340.9	227.6
T_1 ADAPT (average)	7.3	192.2	129.4
T_3 ADAPT (average)	8.7	213.4	143.0

Table 3.3: Summary of quantum circuit resource needed to compute excited states of LiH using spin restriction and different ansätze. The ansatz depth is fixed for 2-UpCCGSD, whereas for the adaptively grown ansätze the values are averages across the 5 calculated bond lengths. The ansatz convergence threshold is taken to be $\epsilon = 0.01$ for the ADAPT methods. Table reproduced from [2] – I was its original creator.

methods, while also maintaining chemical accuracy across different calculated nuclear geometries. We also showed that segregated computation of states in the singlet and triplet manifold can be achieved with spin restrictions at the excitation operator and reference circuit level, which reduces the number of states that need to have converged before reaching higher states.

However, I also confirmed that in the ADAPT-VQD, a reference circuit can only calculate eigenstates of a given irreducible representation of the molecular point group, because only excitations that are symmetry-allowed get selected. More work needs to be done on how one might recover other states efficiently, as well as applying the suggested strategies to larger molecular systems expressed in more realistic basis sets.

3.4 Concluding Thoughts

In this Chapter, I reported two different research projects that I undertook which sought to improve the performance of VQE, targeting NISQ level hardware. Both exploited properties about the underlying system that we were trying to solve, to reduce the quantum resource use and also to exert greater control over the ansatz convergence.

While both algorithm developments reported here were intended for NISQ hardware, neither directly considered the effects of circuit errors, the effects of using a finite number of measurement shot repeats (shot noise), nor were they

implemented on real quantum hardware that exist at the time of writing. For one, the quantum circuit depths and the number of shots (in both applications described) are expected to be considerable on NISQ hardware that is available – obtaining meaningful insight, with a reasonable cost of accessing hardware distributed by cloud-quantum-computing providers, seemed unlikely. More numerics testing the robustness of the algorithm in the presence of simulated noise models as well as shot noise might be useful. Ultimately, the best test of noise robustness will be to perform experiments on future NISQ hardware that might have lower noise levels.

Taking one step back, it might also be possible to demonstrate quantum advantage on NISQ computers using VQE if we critically examine how accurate the computed properties need to be to tackle a given many-body problem. For example, it has been suggested that sketching phase diagrams or obtaining thermal properties of many-body systems has lower quantum resource requirements than computing the exact eigenstates (see e.g. Ref. [225]). Where molecular excited state computations are used to determine the energy *difference* between eigenstates for example, having an accurate estimate for every eigenstate may not be necessary, provided assumptions can be made about all excited state computations sharing the same error.

Alice sighed wearily. 'I think you might do something better with the time,' she said, 'than waste it in asking riddles that have no answers.'

'If you knew Time as well as I do,' said the Hatter, 'you wouldn't talk about wasting it!'

— Lewis Carroll [226]

4

Non-Unitary Evolution on Early Fault-Tolerant Quantum Computers

Contents

4.1 Introduction 88
4.2 Methods 89
4.2.1 Non-unitary evolution via Quantum Eigenvalue Transform 89
4.2.2 Key Features and Comparison with Other Methods . . 91
4.2.3 Applications 93
4.3 Numerical Validation 99
4.3.1 Ground state preparation 101
4.3.2 Gibbs state preparation 105
4.3.3 Open Quantum System Dynamics 107
4.3.4 Hamiltonian scaling: Bounds of extremum eigenvalues . 109
4.4 Concluding Thoughts 110

In this chapter, I use text from the arXiv preprint 2303.06161 [3] verbatim – I was the first and original author of the text. The work was done in collaboration with coauthors Nathan Fitzpatrick and David Muñoz Ramo at the University of Oxford and the company Quantinuum Limited. My coauthors and I developed the methods therein as a group, based on my initial ideas. I performed all of the numerical simulations and analysis reported. It is currently under review at Physical Review Letters.

4.1 Introduction

The efficient preparation of a meaningful quantum many-body state on a qubit register will likely be a key first step in any quantum simulation. As discussed in Chapter 2, Section 2.2.5, a special type of quantum state is one representing a system that has undergone non-unitary imaginary time evolution (ITE). Performing ITE on a quantum state is often discussed in the context of preparing the lowest energy eigenstate of a many-body Hamiltonian. It is useful in other more complex simulation tasks as well, for example, modelling quantum states that are entangled to another as part of a subsystem. This has applications in predicting finite-temperature statistical mechanics as well as the non-unitary dynamics of open quantum systems.

In Ref. [3], my collaborators and I developed a quantum algorithm for preparing quantum states that have experienced ITE using the Quantum Eigenvalue Transform with Unitary block-encoding (QET-U) [148, 150] approach, a recently-introduced routine within the class of Quantum Signal Processing (QSP) methods (see Chapter 2, Section 2.2.5). It leverages stochastic measurement of only a single ancilla qubit, the ability to perform ancilla controlled time evolution, and efficient QSP phase factor evaluation. QSP approaches are generally considered to be robust and elegant means of performing sophisticated operations on quantum computers, but require both deep circuits and large qubit overheads, leading many to consider them quantum algorithms requiring mature, fault-tolerant quantum computers. The QET-U variant instead applies the QSP machinery to early fault-tolerant quantum computers that may have a limited number of error-corrected qubits. It also shares basic quantum circuit '*motifs*' with other state preparation algorithms, and I was able to manipulate these features such that the Wick-rotated Schrödinger equation could be simulated exactly (up to the polynomial approximation inherent in QSP and Trotter errors incurred during the 'encoding' of the system Hamiltonian). I validated this technique using numerical qubit emulations, targeting three key simulation tasks: Ground state preparation, thermal state preparation, and the quantum simulation of open quantum system dynamics, and found that it shows promise for preparing imaginary-time-evolved states on early fault-tolerant quantum computers.

4.2 Methods

4.2.1 Non-unitary evolution via Quantum Eigenvalue Transform

Reiterating Chapter 2 Section 2.2.5, the normalised imaginary-time or Wick-rotated Schrödinger Equation for an n-body Hamiltonian $\mathcal{H}$ is

$$|\psi(\tau)\rangle = A(\tau)e^{-\mathcal{H}\tau}\,|\psi_0\rangle \tag{4.1}$$

where $\tau = it$ is the imaginary time unit and $A(\tau) = 1/\sqrt{\langle\psi_0|\, e^{-2\mathcal{H}\tau}\,|\psi_0\rangle}$ is the normalisation constant. The main insight of the work in this Chapter, similarly employed in other works as [150], is that using only one ancilla qubit in addition to the qubits needed to represent the n-body system, the QET-U methodology (Chapter 2 Section 2.2.5) can be used to compile a polynomial transform of $\mathcal{H}$ which reproduces the non-unitary exponential function:

$$\mathcal{P}(\mathcal{H}) \approx f(\mathcal{H}) = e^{-\mathcal{H}\tau}. \tag{4.2}$$

The probability of successfully realising the desired polynomial approximation (measuring the QSP ancilla qubit in the 0 state) is given by $\langle\psi_0|\, e^{-2\mathcal{H}\tau}\,|\psi_0\rangle$ as per the normalisation constant $A(\tau)$ – which, by inspection, is analogous to the overlap between the imaginary time projected state and the initial input state $|\psi_0\rangle$.

QET-U achieves this with an alternating circuit pattern typical of QSP methods: This consists of a 'block encoding' element where $\mathcal{H}$ is embedded into a single-qubit subspace via the ancilla-qubit-controlled time evolution oracle $e^{-i\mathcal{H}}$ (which, because $\mathcal{H}$ is typically a sum of non-commuting terms, will typically be implemented in approximation via Trotterisation), and 'signal processing' orthogonal phase rotations applied to that single ancilla qubit – this circuit construction is recapitulated in Fig. 4.1. Because QET-U can only manifest the desired eigenvalue transform if the eigenvalues of $\mathcal{H}$ are between η and $\frac{\pi}{2} - \eta$ (η being a small constant used to provide numerical stability when evaluating the sequence of QSP phase rotations), the time evolution oracle in practice embeds a scaled Hamiltonian $\mathcal{H}^{scale} = \mathcal{H}t + \sigma$,

where $\sigma = \eta - \tilde{E}_0 t$ and $t = (\pi - \eta)/(\tilde{E}_{N-1} - \tilde{E}_0)$; the terms $\tilde{E}_0, \tilde{E}_{N-1}$ represent bounds on the extremum eigenvalues of $\mathcal{H}$.

Statement 1. *By compiling QSP phase factors approximating the Hamiltonian function*

$$f(\mathcal{H}^{scale}) = \frac{\exp\Big[-\tau(\mathcal{H}^{scale} - \sigma)/t\Big]}{\exp[-\tau(E_M^{scale} - \sigma)/t]}, \tag{4.3}$$

where the 'scale' superscripts represent the appropriately scaled system Hamiltonian as per Eq. (2.61) *(Chapter 2 Section 2.2.5) and its extremum eigenvalues E_M^{scale} are scaled to either the small constant η or $\frac{\pi}{2} - \eta$, QET-U can be used to apply the exact Wick-rotated Schrödinger Equation Eq.* (4.1) *up to the accuracy of the polynomial approximation $\mathcal{P} \approx f$ and the Trotter error from the encoding of $\mathcal{H}$.*

In Ref. [3], I found that Statement 1 can be achieved by effectively undoing the scaling of $\mathcal{H}^{\text{scale}}$ using the QSP sequence of phase factors. The transformation $\exp\Big[-\tau(\mathcal{H}^{\text{scale}} - \sigma)/t\Big]$ guarantees that the proportion of exponential decay experienced between eigenvalues is virtually the same *as if* the transform $e^{-\mathcal{H}\tau}$ was applied i.e. true imaginary time evolution was performed. This technique is equivalent to first scaling the input domain of a function, then accordingly scaling the function itself. Division by constant values do not affect the ratio of exponential decay between eigenvalues, and therefore do not render the non-unitary dynamics inexact.

Recall from Chapter 2 Section 2.2.5 that QET-U and related QSP approaches involves a probabilistic ancilla projection to correctly realise the state transform. Furthermore, since ITE is a 'lossy' non-unitary process, the probability of a successful ancilla measurement can be expected to decay exponentially with the overlap between the initial state and the imaginary-time-evolved state, or when τ is large. I devised an approach to maximise the probability of the ancilla measurement yielding a successful outcome: We normalised the eigenvalue transform by the exponent of the scaled ground state η (if we were performing forward imaginary time) or the maximally excited state $\frac{\pi}{2} - \eta$ (if we were performing reverse imaginary time), motivating the use of $\exp\Big[-\tau(E_M^{\text{scale}} - \sigma)/t\Big]$ in Eq. (4.3). This places the computational effort on accurately scaling the Hamiltonian by its extremum

eigenvalues and thus obtaining $\mathcal{H}^{\text{scale}}$ (I will discuss this challenge later), and reasonable time stepping to perform controlled real time evolution $e^{-i\mathcal{H}}$ for the block encoding of the Hamiltonian. Assuming these can be achieved, this tactic avoids imposing exponential attenuation on the ground/maximum energy eigenstate, and mirrors the 'eigenstate filtering' ground state preparation strategy proposed in Ref. [148] (see next Section).

4.2.2 Key Features and Comparison with Other Methods

As the operator $e^{-\mathcal{H}\tau}$, and therefore the duration of imaginary time τ, is determined only by the polynomial approximation $\mathcal{P}$ manifested by the QSP sequence of phase factors, we can in principle subject an initial quantum state to any amount of ITE, subject to the single-ancilla measurement at the very end of the QET-U circuit yielding a successful outcome. Furthermore, because the approach is performing exact ITE up to a polynomial approximation, it is useful beyond ground state preparation, and can be straightforwardly applied to preparing quantum states that have experienced non-unitary evolution. We can reason, by considering the Taylor expansion of the exponential function $e^{-\mathcal{H}\tau}$ for when τ is large, that its accurate representation will require a large degree of the polynomial $\mathcal{P}(\mathcal{H})$, which necessarily increases the length of the QSP sequence and thus circuit depth. Classical pre-computation of the phase factors for these deep QSP sequences given a target transform can be extremely efficient, especially in the case of using a symmetric set of phase factors described in Refs. [147, 148] – these methods can compute phase factors that accurately approximate polynomials of up to degree $\mathcal{O}(10^4)$. I have used this strategy to generate results in this Chapter.

This approach can be interpreted as a generalisation of the Probabilistic Imaginary Time Evolution (PITE) method [54] in Fig. 4.1, which approximates short imaginary-time steps by (similar to the present approach) implementing cosine functions of the Hamiltonian using a linear combination of forward and reverse real-time evolution, facilitated by a stochastic single-ancilla measurement (see Chapter 1 Section 2.2.5). This strategy of sequentially applying cosine of the

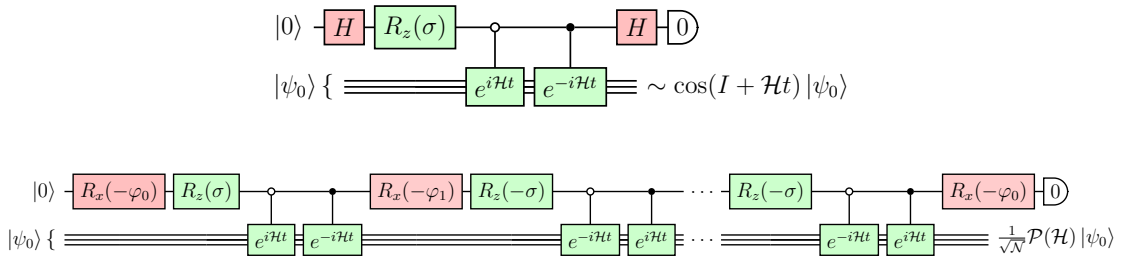


Figure 4.1: Circuit for QET-U and relation to the PITE circuit. (BOTTOM) The QET-U circuit manifests an polynomial transform $\mathcal{P}$ on the eigenvalues of matrix $\mathcal{H}$. $\mathcal{P}$ is determined by the sequence of phase factors rotated along R_x. (TOP) The PITE approach performs a $\propto \cos(I + \mathcal{H}t)$ transform to approximate imaginary-time evolution up to first order. Figure reproduced from [3] – I was its original creator.

Hamiltonian has similarly been explored by others [154, 155] for state preparation – indeed, the authors of QET-U also note the similarity between the circuit to the Iterative Phase Estimation or Hadamard test [148], which likewise applies a cosine transform of the Hamiltonian. In methods such as PITE however, one would have to sequentially apply PITE stochastic measurements in order to achieve a long imaginary time sequence, which can result in a low cumulative success probability overall. The authors of PITE initially noted this as a 'deplorable drawback' because the overall success probability can be vanishingly small even with less than only 10's of PITE steps [54], but have since proposed amplitude amplification as a meaningful strategy to counteract this at the expense of a deeper circuit [227].

I also explored the division of the total imaginary time into N evenly spaced imaginary-time fragments $e^{-\mathcal{H}\tau} = (e^{-\mathcal{H}\Delta\tau})^N$ [148, 150], and each ITE fragment $e^{-\mathcal{H}\Delta\tau}$ is realised stochastically using QET-U. This can be seen as a hybridisation with PITE, but the $\Delta\tau$ that can be performed with QET-U is much longer than time-steps that can be performed with methods like PITE.

Compared with other quantum algorithms for performing ITE, this approach also presents additional advantages; it avoids cumbersome classical computational steps such as variational optimisation necessary in the Variational Imaginary Time Evolution (VITE) approach, and Pauli tomography necessary in the Quantum Imaginary Time Evolution (QITE) approach. However, like PITE and QITE, the query complexity of simulating imaginary-time of τ with error ϵ is expected to scale

linearly with the imaginary time as $\propto \mathcal{O}(\tau \log \epsilon^{-1})$ [150], and in the case of my QET-U for ITE technique the corresponding circuits will likely remain too deep for near-term hardware when compared with near-term focused methods like VITE.

4.2.3 Applications

Ground state preparation

Given an initial state $\psi_0 = \sum_i a_i \Psi_i$ (linear combination of Hamiltonian eigenstates $H\Psi_i = E_i\Psi_i$) with non-zero overlap with the ground state Ψ_0, the ITE propagator exponentially attenuates the contribution of high energy eigenstates faster than the ground state:

$$e^{-\mathcal{H}\tau} |\psi_0\rangle = \sum_i a_i e^{-E_i\tau} |\Psi_i\rangle \tag{4.4}$$

(where $E_i \geq 0 \quad \forall i$), thus in the $\tau \to \infty$ limit the operator prepares the ground state of $\mathcal{H}$.

The original authors of the QET-U method built on their previous designs of a near-optimal ground state preparation algorithm using QSP and the Linear Combination of Unitaries (LCU) [228], to perform ground state preparation within the QET-U paradigm via an eigenstate filtering procedure [148]; using an additional lower bound for the first excited eigenstate E_1, they evaluated phase factors which approximately transformed the Hamiltonian eigenvalues by a Heaviside step function, with the step centred at the spectral gap Δ between the (scaled) ground and first excited state energies E_0^{scale} and E_1^{scale}:

$$P(\mathcal{H}) \approx f(\mathcal{H}^{\text{scale}}) = \begin{cases} 1 & \text{if } E \leq \mu \\ 0 & \text{if } E > \mu \end{cases}, \tag{4.5}$$

where $\mu = (E_0^{\text{scale}} + E_1^{\text{scale}})/2$. They claimed that this approach is near-optimal with query complexity $\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$ and also proposed a modification with query complexity $\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-2})$ [148], where γ is the overlap between the initial state and the ground state (the notation $\tilde{\mathcal{O}}(n)$ is $\mathcal{O}(n \log^k n)$). The main drawback of the eigenstate filter is, of course, the need for a bound of the first excited state energy; even an estimate of such an eigenvalue can be challenging to obtain.

In contrast, performing ITE using QET-U does not require knowledge of the first excited state; it only needs bounds of E_0 and E_{N-1} for the QET-U eigenvalue scaling. The complexity of using it for ground state preparation will however still depend on γ, as a good initial overlap reduces the duration of imaginary time propagation needed for convergence. Furthermore, the ability to perform ITE is helpful in other applications which I discuss presently.

Gibbs state and partition function

The study of finite-temperature systems on quantum computers, in particular computing partition functions and preparing Gibbs states, is considered important tasks relevant to simulation, quantum Boltzmann machine learning and quantum optimisation [54, 160]. Different techniques for performing these tasks have been proposed [54, 160, 164, 229–231]. I demonstrated that QET-U can provide a natural framework for preparing thermal states quantum states on quantum computers, and is a meaningful early fault-tolerant alternative to variational [160, 229, 230] and PITE [54] methods for thermal state quantum simulation.

The target is to prepare the Gibbs state density operator $\rho_{\text{sys}}(\beta)$ of a system with Hamiltonian $\mathcal{H}_{\text{sys}}$ at the inverse temperature $\beta = 1/k_B T$:

$$\rho_{\text{sys}}(\beta) = \frac{e^{-\mathcal{H}_{\text{sys}}\beta}}{\operatorname{Tr} e^{-\mathcal{H}_{\text{sys}}\beta}} = \frac{e^{-\mathcal{H}_{\text{sys}}\beta}}{Z(\beta)}, \tag{4.6}$$

on a register of qubits, where k_B is the Boltzmann constant and $Z(\beta)$ is the partition function. This can be achieved by considering a density matrix ρ which is a composite of the system of interest and an 'environment' subsystem. Note that this does not model the physical surroundings within which the system is embedded. But as I will show momentarily, it allows us to take the partial trace over the 'environment' subsystem Tr_{env} and obtain the reduced density matrix $\rho_{\text{sys}} = \text{Tr}_{\text{env}}\, \rho$.

One could start by generating the maximally entangled state of a bipartite $2n$-qubit register [54, 160] using the circuit U_M in the bottom panel of Figure 4.2:

$$|\psi_0\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle_{\text{sys}} |i\rangle_{\text{env}}\,, \tag{4.7}$$

where the first n-qubits encode the system of interest with Hamiltonian $\mathcal{H}_{\text{sys}}$, and the remainder encode the surrounding environment with Hamiltonian $\hat{I}_{\text{env}}$, making the total Hamiltonian $\mathcal{H}_{\text{tot}} = \mathcal{H}_{\text{sys}} \otimes \hat{I}_{\text{env}}$. We wish to evolve this maximally entangled state under the total Hamiltonian $\mathcal{H}_{\text{tot}}$ by $\tau = \beta/2$:

$$\begin{aligned} |\psi\rangle &= e^{-\mathcal{H}_{\text{tot}}\tau} |\psi_0\rangle \\ &= \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} e^{-(\mathcal{H}_{\text{sys}} \otimes \hat{I}_{\text{env}})\tau} |i\rangle_{\text{sys}} |i\rangle_{\text{env}} . \end{aligned} \tag{4.8}$$

Using a Taylor expansion of the exponential operator, we can show that $e^{-(\mathcal{H}_{\text{sys}} \otimes \hat{I}_{\text{env}})\tau} = e^{-\mathcal{H}_{\text{sys}}\tau} \otimes I_{\text{env}}$. Therefore, it is only necessary to imaginary time evolve the subsystem $\mathcal{H}_{\text{sys}}$ which I approximate using QET-U (see Figure 4.2):

$$\begin{aligned} |\psi\rangle =& \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} \frac{1}{\sqrt{\mathcal{N}}} P(\mathcal{H}_{\text{sys}}) |i\rangle_{\text{sys}} \otimes \hat{I}_{\text{env}} |i\rangle_{\text{env}} \\ \approx& \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} \frac{e^{-\mathcal{H}_{\text{sys}}\tau}}{\sqrt{\mathcal{N}}} |i\rangle_{\text{sys}} \otimes \hat{I}_{\text{env}} |i\rangle_{\text{env}} \end{aligned} \tag{4.9}$$

For now ignore the QET-U ancilla and assume its measurement yields a successful outcome. The resulting state vector of the $2n$-qubit pure state is approximately the density operator:

$$\begin{aligned} \rho =& |\psi\rangle \langle\psi| \\ =& \frac{1}{\mathcal{N} 2^n} \sum_{i,j=0}^{2^n-1} (e^{-\mathcal{H}\tau} |i\rangle \langle j| e^{-\mathcal{H}\tau})_{\text{sys}} \otimes (\hat{I} |i\rangle \langle j| \hat{I})_{\text{env}} \end{aligned} \tag{4.10}$$

(where $e^{-\mathcal{H}_{\text{sys}}\tau}$ and $\hat{I}_{\text{env}}$ are both Hermitian). At this point, we can discard the n-qubits of the environment register (equivalent to the partial trace over the environment register):

$$\begin{aligned} \text{Tr}_{\text{env}}\, \rho =& \frac{1}{\mathcal{N} 2^n} \sum_{i,j=0}^{2^n-1} (e^{-\mathcal{H}\tau} |i\rangle \langle j| e^{-\mathcal{H}\tau})_{\text{sys}} \otimes \text{Tr}(|i\rangle \langle j|)_{\text{env}} \\ =& \frac{1}{\mathcal{N} 2^n} \sum_{i,j=0}^{2^n-1} (e^{-\mathcal{H}\tau} |i\rangle \langle j| e^{-\mathcal{H}\tau})_{\text{sys}} \delta_{ji} \\ =& \frac{1}{\mathcal{N} 2^n} \sum_{i=0}^{2^n-1} e^{-\mathcal{H}_{\text{sys}}\beta} |i\rangle \langle i| , \end{aligned} \tag{4.11}$$

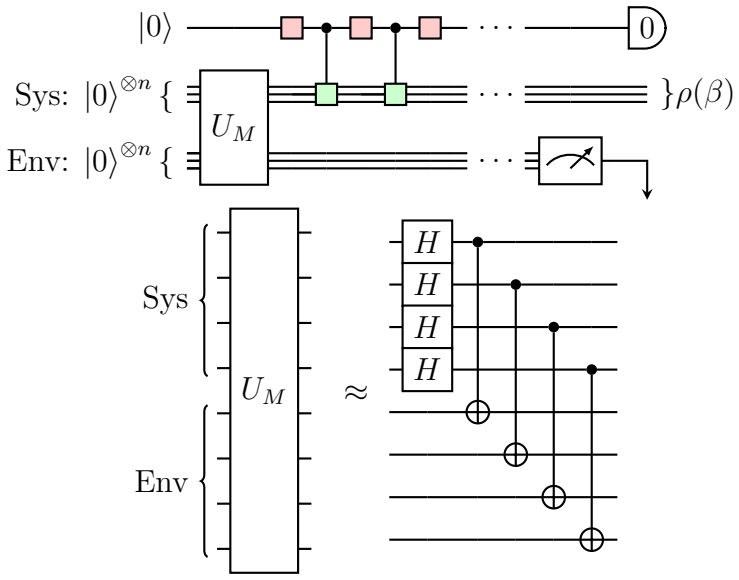


Figure 4.2: Using single-ancilla quantum eigenvalue transform for Gibbs state preparation. The initial state maximally entangles the Sys and Env registers, and QET-U is performed on the Sys register to manifest imaginary-time evolution. Tracing out the Env register yields the Gibbs state density matrix. Quantum gates in red represent the QSP phase sequence, quantum gates in green represent controlled, Trotterised time evolution of the system in question. Figure reproduced from [3] – I was its original creator.

where the final state on the n-qubit system register is now a reduced density operator that is inexpressible as a state vector, which is proportional to the Gibbs state ρ_{sys} up to the normalisation factor $\mathcal{N}$ from QET-U.

I will now show that this normalisation factor, hence the probability of measuring a successful outcome for QET-U, is proportional to the partition function. Reiterating Eq. (2.57), the state vector normalisation factor from the polynomial approximation of $e^{-\mathcal{H}_{\text{tot}}\tau}$ is $\mathcal{N}(\tau) \approx \langle\psi_0| e^{-2\mathcal{H}_{\text{tot}}\tau} |\psi_0\rangle$. Combining Eq. (4.7) and Eq. (4.8):

$$\begin{aligned}
\langle\psi_0| e^{-2\mathcal{H}_{\text{tot}}\tau} |\psi_0\rangle &= \frac{1}{2^n} \sum_{i,j=0}^{2^n-1} \langle i| e^{-2\mathcal{H}\tau} |j\rangle_{\text{sys}} \langle i| \hat{I}^2 |j\rangle_{\text{env}} \\
&= \frac{1}{2^n} \sum_{i=0}^{2^n-1} \langle i| e^{-2\mathcal{H}_{\text{sys}}\tau} |i\rangle \\
&= \frac{1}{2^n} \operatorname{Tr}\left[e^{-\mathcal{H}_{\text{sys}}\beta}\right] = \frac{Z(\beta)}{2^n}
\end{aligned} \tag{4.12}$$

shows that, similar to PITE [54], the probability of a successful ancilla measurement times the factor 2^n gives an estimate of the partition function. Given these

properties, one could for example calculate the free energy $F = -\beta^{-1} \ln Z(\beta)$, which is important in predicting chemical reaction thermal properties.

There are a number of important caveats to note here. In many practical applications, the primary interest lies in the calculation of expectation values of observables with respect to the Gibbs state, rather than the Gibbs state itself. Later in this Chapter, I will justify how if one wishes to extract expectation values of the Gibbs state, the environment register actually need not be necessary. Furthermore, because of the factor of 2^n in Eq. (4.12), the quantum resources (number of shot measurements) needed to estimate the partition function grows exponentially in n. As far as we know, other probabilistic methods for quantum computing the partition function e.g. Ref. [54] suffers from the same exponential cost.

Dynamics of Open Quantum Systems

The final application of QET-U for ITE I explored is in the non-unitary dynamics of open quantum systems. While quantum algorithms for simulating open quantum systems have been proposed (see e.g. Refs. [229, 232–234]), it has been observed recently that by interleaving real and imaginary time evolution, one can approximate the dynamics of open quantum systems on a quantum computer (see Refs. [235, 236]).

The Lindblad equation models the dynamics of a Markovian quantum system with weak entanglement to the surrounding environment:

$$\frac{d\rho}{dt} = -i[\mathcal{H}, \rho] + \sum_k \left(\hat{L}_k \rho \hat{L}_k^\dagger - \frac{1}{2}\{\hat{L}_k^\dagger \hat{L}_k, \rho\} \right) \tag{4.13}$$

where ρ is a $2^n \times 2^n$ density operator representation of the quantum system, and $\hat{L}_k$ are Lindblad operators representing the channels that couple the system to the environment; dynamics of the environment is not explicitly captured in this model. The main insight of Refs. [235, 236] is that the equation can be rewritten in a 'Schrödinger' form:

$$\frac{d\,|\rho\rangle}{dt} = \hat{\mathcal{L}}\,|\rho\rangle\,, \tag{4.14}$$

with a non-Hermitian Linbladian Hamiltonian:

$$\hat{\mathcal{L}} = -i\hat{I}_A \otimes \mathcal{H}_B + i\mathcal{H}_A^T \otimes \hat{I}_B + \sum_k \hat{L}_{k\,A}^* \otimes \hat{L}_{k\,B} - \frac{1}{2}\hat{I}_A \otimes (\hat{L}_k^\dagger \hat{L}_k)_B - \frac{1}{2}(\hat{L}_k^T \hat{L}_k^*)_A \otimes \hat{I}_B$$

admitting the integral form

$$|\rho(t)\rangle = e^{\hat{\mathcal{L}}t} |\rho(0)\rangle . \tag{4.15}$$

The superscripts $^T, ^*, ^\dagger$ correspond to the transpose, complex conjugate and conjugate transpose respectively.

The 'master Hamiltonian' $\hat{\mathcal{L}}$ is composed of operators which act on two effective subsystems (labelled with subscripts A and B), similar to the previous thermal state example where the total Hamiltonian is a tensor product of two operators acting on two subsystem Hilbert spaces. It follows that the $2^n \times 2^n$ density matrix $\rho(t)$ of the open system must be represented as a statevector $|\rho(t)\rangle$ across two equally-sized registers of n-qubits, and we can consider the time evolution of $\hat{L}$ as operators acting on either the first or second set of n-qubits.

As $\hat{\mathcal{L}}$ is typically non-normal, it cannot be embedded using the time evolution oracle directly for QET-U, because its evolution is not unitary. Fortunately, $\hat{\mathcal{L}}$ admits the following partition into Hermitian and anti-Hermitian parts:

$$\hat{\mathcal{L}} = i\mathcal{H}_1 + \mathcal{H}_2 \tag{4.16}$$

defined by:

$$\mathcal{H}_1 = \frac{1}{2i}(\hat{\mathcal{L}} - \hat{\mathcal{L}}^\dagger) \tag{4.17}$$

$$\mathcal{H}_2 = \frac{1}{2}(\hat{\mathcal{L}} + \hat{\mathcal{L}}^\dagger). \tag{4.18}$$

In the examples considered presently, I was able to straightforwardly write down $\mathcal{H}_1$ and $\mathcal{H}_2$ by simply inspecting which operators have an imaginary coefficient, a strategy expected to be scalable.

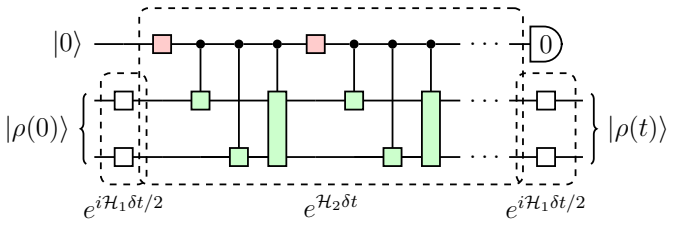


Figure 4.3: Using single-ancilla quantum eigenvalue transform for non-unitary dynamics simulation. The state vector encoded is a vectorised density matrix representing an open quantum system. The evolution is Trotterised into unitary and non-unitary parts at each time step – QET-U is used to perform the non-unitary part. Quantum gates in red represent the QSP phase sequence, quantum gates in green represent controlled, Trotterised time evolution of the system in question. Figure reproduced from [3] – I was its original creator.

The time evolution of the 'column-stacked' density operator under $\hat{\mathcal{L}}$ can thus be modelled by splitting the unitary and non-unitary evolution and time-stepping:

$$|\rho(t)\rangle = e^{i\mathcal{H}_1 t + \mathcal{H}_2 t} |\rho(0)\rangle \tag{4.19}$$

$$= \left(e^{i\mathcal{H}_1\delta t/2} e^{\mathcal{H}_2\delta t} e^{i\mathcal{H}_1\delta t/2}\right)^N |\rho(0)\rangle + \mathcal{O}(\delta t^3). \tag{4.20}$$

I performed the non-unitary part of each time step using QET-U. In this application, the desired result is a reconstruction of the density matrix of the state, and so full state tomography will need to be performed at the time steps where we are interested in the density matrix, with a measurement overhead of $O(2^{2n})$ per time step. Because this method of open quantum system dynamics simulation only preserves $\mathrm{Tr}(\rho^2)$, the amplitudes of the state vectors are actually divided by $\mathrm{Tr}(\rho)$, which can be trivially obtained once an estimate of the state vector is known.

4.3 Numerical Validation

I emulated small, noise-free quantum computers using significant classical computing resources to show that the proposed methodology can in principle be applied to the described applications. Qubit simulations were performed in `Python` using the quantum computational chemistry package InQuanto [237] and pytket [238], and benchmark calculations for Hamiltonian diagonalisation and density matrix computations were performed using SciPy [217]. The pre-computed Hamiltonian

eigenvalues were used as inputs for scaling and shifting the Hamiltonian eigenvalues in QET-U – I highlight that this is strictly for demonstration purposes; in practice the eigenvalues will not be accessible *a priori*. I used the `MATLAB` QSPPACK [147] package to evaluate the coefficients to the basis of Chebyshev functions, and thus QSP phase factors, for approximating a given target polynomial. The value of $\eta = 0.001$.

Proof of Concept

I started by performing illustrative calculations to confirm that the desired transform can be achieved. I did this by making use of the equivalence between the general form of QSP and performing arbitrary rotations on a single qubit using alternating orthogonal rotations. I will refer the reader to Chapter 2 Section 2.2.5 for more background, but provide a quick summary here; in essence, one can define a fixed *signal operator* which is an R_z rotation by a fixed angle θ, and a set of *signal processing operators* which are R_x rotations, each with a unique, pre-computed phase rotation factor $\{\varphi_i\}$. If one were to apply an alternating sequence of $R_z(\theta)$ and $\{R_x(\varphi_i)\}$ gates (a 2×2 unitary $U(\theta)$) on a single qubit initialised in state $|0\rangle$, the resulting amplitude of $\langle 0| U(\theta) |0\rangle$ (i.e. the top left element of $U(\theta)$) is a polynomial transform $\mathcal{P}(\theta)$ characterised by the phase factors $\{\varphi_i\}$. I therefore compiled sequences of symmetric phase factors φ_i to approximate the normalised exponential decay target function $\langle 0| U(\theta) |0\rangle = \mathcal{P}(\theta) \approx e^{-\theta\tau}$ at a fixed value of imaginary time τ, for $\theta \in [\eta, \frac{\pi}{2} - \eta]$. The left plot of Fig. 4.4 shows the resulting amplitude given QSP circuits approximating $8, 16, 24$-degree polynomials – it is clear that large polynomals, thus deep QSP sequences, are necessary to converge to the desired transform.

The right plot of Fig. 4.4 gives a visual representation on the transformation of eigenvalues of a Hamiltonian according to the exponential function. It demonstrates how, to manifest this same transform on the eigenvalues of a Hamiltonian in the QSP setting, the Hamiltonian must be scaled and shifted such that its entire energy spectrum (red dots) fits between the range $\eta, \eta - \frac{\pi}{2}$ within which the

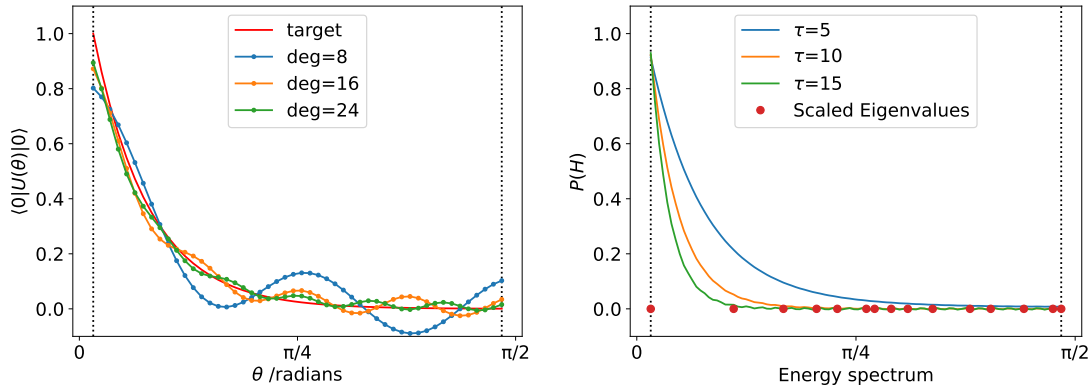


Figure 4.4: Quantum eigenvalue transform applied to a scaled Hamiltonian. (LEFT) Low-degree polynomial approximations of the target exponential decay function, applied between the range $\theta = [\eta, \frac{\pi}{2} - \eta]$. (RIGHT) Once we are able to encode eigenvalues of the Hamiltonian (red dots) on the QSP ancilla qubit, the desired polynomial transform of its eigenvalues can be applied. Figure reproduced from [3] – I was its original creator.

correct transformation is achieved (solid lines superposed in the plot, in this case corresponding to scaled exponential decay functions).

4.3.1 Ground state preparation

To validate that the QET-U imaginary time approach would work for ground state preparation, I studied a demonstrative 4-site, half-filled, linear Fermi-Hubbard chain, which is an 8-qubit problem with the Hamiltonian:

$$\mathcal{H} = -t \sum_{i,\sigma} \left(\hat{c}^{\dagger}_{i,\sigma} \hat{c}_{i+1,\sigma} + \hat{c}^{\dagger}_{i+1,\sigma} \hat{c}_{i,\sigma} \right) + U \sum_{i} \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}, \tag{4.21}$$

where I chose different ratios of on-site electron repulsion U to one-electron hopping term between adjacent sites t, representative of weak electron correlation ($U << t$) and strong electron correlation ($U >> t$) regimes.

Increasing imaginary time

I first compiled QSP phase factors for 350-, 250-, and 150-degree polynomial approximations of Eq. (4.3), for τ between 0 and 5, choosing the Hubbard Hamiltonian where $U = t$. The corresponding QET-U circuits are prepared, where every call to the time evolution oracle uses 25 Trotter steps. It is important to highlight that the only difference between circuits implementing different values of τ is

the phase factors used. Each circuit is used to prepare the ground state from a mean-field Hartree-Fock initial state, and the expected energies of each prepared post-selected state are computed.

In general, the convergence to the ground state is the same as if one were performing exact ITE (see top left panel of Fig. 4.5). However, as expected, for low-degree polynomial approximations of Eq. (4.3), the energy does not properly converge. As explained previously, this is because at long imaginary times, the sharper exponential decay requires higher order terms to represent accurately. The probability of successfully preparing the post-selected state is displayed in the top right panel of Fig. 4.5, which shows that it levels off around 10% at long imaginary times.

Comparison with Eigenstate Filter

Using the same QET-U circuit, but substituting instead the phase factors for the Heaviside function (Eq. (4.5)) as per Ref. [148], I was able to prepare a state with an energy that is 1.77mE_h from the exact ground state energy. It similarly shared a 10% probability of successful post-selection. These numerical results suggest that, for given a circuit depth budget restriction, performing ITE with QET-U can actually provide a better ground state energy convergence than performing eigenstate filtering.

The strength of the eigenstate filter approach is that it should simply converge to the ground state by increasing the polynomial degree, provided a good first excited state energy is known; in contrast, our ITE approach requires the user to make an informed choice on what is a sufficiently long imaginary time. As shown earlier, it may be counterproductive to increase the imaginary time without commensurately increasing the polynomial degree. Therefore, whereas the necessary number of measurement shots for a successful ground state preparation using the eigenstate filter will be determined solely by the success probability, an additional shots overhead might be necessary for probing what is the appropriate duration of imaginary time.

Imaginary Time 'Fragmentation'

Within the choice of algorithm parameters, the QET-U circuits required to converge the ground state for this relatively trivial system are considerable. In an effort to reduce the circuit depth, I explored a fragmentation strategy proposed in Ref. [150]. From a QSP perspective, we can also choose lower degree polynomial approximations compared with earlier demonstrations to perform each imaginary-time fragment. Unlike the above experiments though, the metric for probability of preparing a target state is now the cumulative probability of success in applying each ITE fragment; in the case of PITE, this total probability has been found to be vanishingly small even after a small number of PITE steps [54]. The circuit depth also is no longer constant but will grow linearly with the number of fragments. However, we will see momentarily that because the cost of performing each time step can be reduced by choosing a lower degree approximation of the transform, the total cost can actually be less than the previous strategy.

I divided the total imaginary-time into 4 time steps ($\Delta\tau = 1.25$), and each imaginary-time step is approximated with only a 30-degree polynomial. The top panels of Fig. 4.5 show the energy convergence behaviour and the cumulative success probability of the sequential QET-U ancilla measurements, respectively. The results suggest that fragmentation offers a similar performance in terms of convergence and probability of success. Critically, the final circuit depth of performing these three QET-U imaginary-time fragments is less than half that is necessary to perform the 350-degree, single-shot imaginary-time QET-U from Section 4.3.1. This strategy of performing fragmentation hybridises the best features of PITE and QET-U, and can be a desirable option for ground state preparation in terms of success probability, circuit depth, minimum pre-computation cost and energy convergence.

Severe dependence on the initial trial state

While these results show that the method can converge to the ground state, two other important metrics for non-unitary state preparation are the quantum circuit depth and the probability of success. Using the same Hamiltonian as before, I

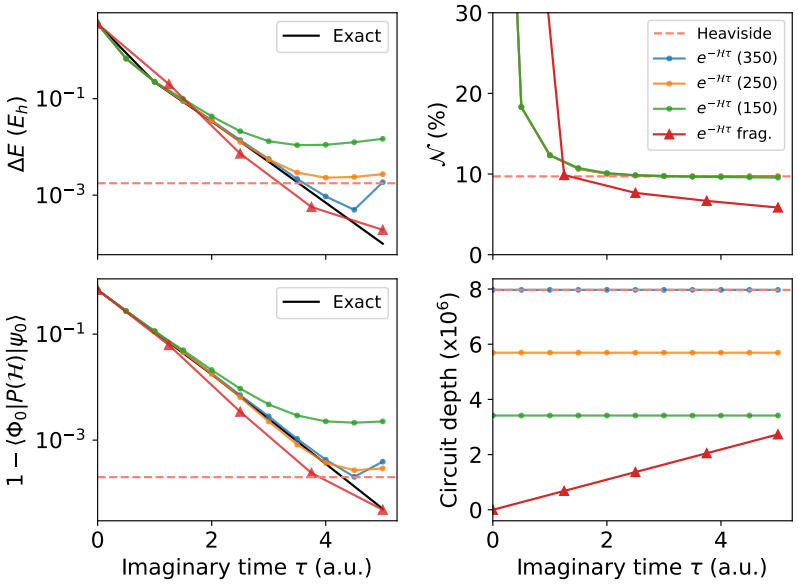


Figure 4.5: Ground state convergence properties by performing imaginary-time evolution using QET-U. Different coloured lines indicate either achieving the target imaginary time transform in a single QSP measurement, or the fragmented imaginary time evolution. (TOP LEFT) Convergence of the prepared state's energy with the ground state energy. (BOTTOM LEFT) Overlap between the ground state and the prepared state. (TOP RIGHT) The probability of successfully preparing the desired state – for the fragmented version, the probability of success is the cumulative success probability up to that point in imaginary-time. (BOTTOM RIGHT) Depth of the state preparation circuit. Figure reproduced from [3] – I was its original creator.

performed the transform of Eq. (4.3) using all possible 4-body single determinant initial reference states, and confirmed that the probability of success scales as the square of the overlap between the initial state and the projected state $\mathcal{O}(\gamma^2)$ – see leftmost panel of Fig. 4.6.

Although this is expected from Eq. (2.57), it does highlight the importance of selecting an appropriate initial state – particularly in the Fermi-Hubbard model, which is a popular choice for the study of strongly-correlated systems (in the $U >> t$ regime). In these regimes, the overlap between the single reference mean-field state and the ground state is notorious for being prohibitively small. As a result, the probability of success can vanish. Fig. 4.6 shows the exact simulations of the fragmented ITE using the transform Eq. (2.61) for values of U/t =5, 10, 20, 40, which confirm that the probability of success can become prohibitively small.

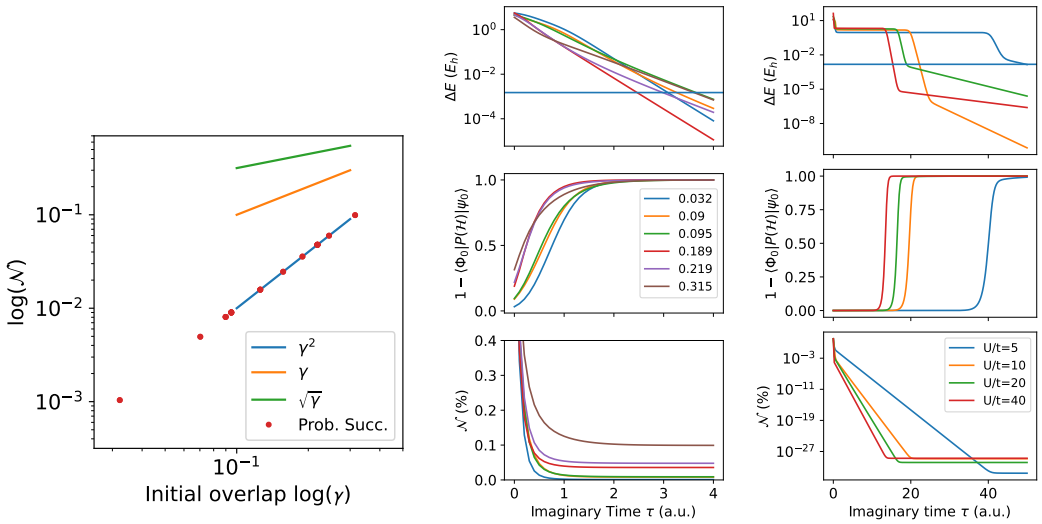


Figure 4.6: Probability of success and different initial trial states in the 4-site linear Fermi-Hubbard model. (LEFT) Imaginary time evolution of the $U = t$ Hamiltonian with initial states of different overlaps. (CENTRE) Convergence time of the $U = t$ Hamiltonian with initial states of different overlaps. (RIGHT) Imaginary time evolution of different U/t ratios. Horizontal lines represent chemical accuracy. Figure reproduced from [3] – I was its original creator.

These results echo previous conclusions on the severe dependence of the initial state for quantum state preparation, and further highlight the need for improved initial state preparation strategies suitable for strongly correlated systems. It is however possible to boost the probability of success using amplitude amplification, at the cost of deepening the circuit. How much imaginary time it takes for a state to converge, and thus the required circuit complexity, depends less on the initial trial state; I found that states with unremarkable initial overlaps can still converge to a given level of accuracy much quicker than one with a high overlap. However, the probability of success will be the limiting factor.

4.3.2 Gibbs state preparation

To validate that the technique discussed can be useful in the preparation of a thermal Gibbs state, I performed circuit emulations for the 4-site transverse-field Ising model with the Hamiltonian:

$$\mathcal{H} = -J \sum_i Z_i Z_{i+1} + g \sum_i X_i, \tag{4.22}$$

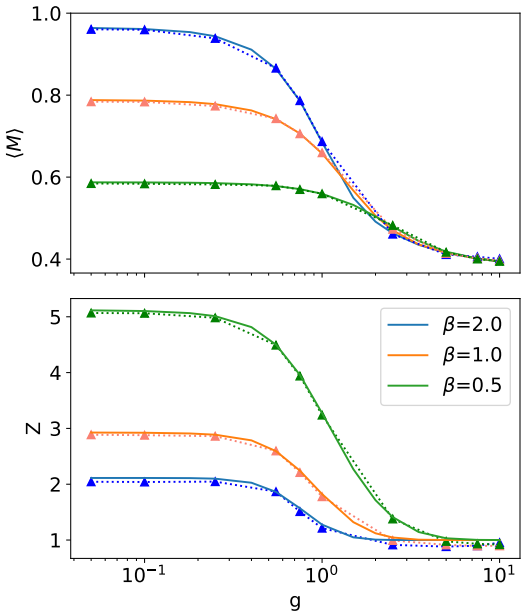


Figure 4.7: Gibbs state preparation and partition function calculation using QET-U. Solid lines represent exact values, and markers with dotted lines show the values generated using QET-U emulation. (TOP) Magnetisation expectation computed from the prepared Gibbs state. (BOTTOM) The partition function computed from multiplying the success probability of preparing the target state by 2^4. Figure reproduced from [3] – I was its original creator.

where I also chose different ratios of the nearest neighbour spin interaction J and the coupling strength of the external transverse field g, exploring the $g << 1$ to $g >> 1$ regimes. I chose inverse temperature β to be 0.5, 1.0 and 2.0. 8-qubits and 50-degree polynomial approximations were used to prepare the corresponding Gibbs states, and each call to the controlled time evolution unitary is divided into 8 steps. I measured the expected magnetisation $M = \sum_i Z_i$ of the prepared density matrix, as well as the success probability which is tied to the partition function.

I found that the value of the expected magnetisation agrees well with exact values across the regimes tested. When I computed the success probabilities of preparing the desired Gibbs states and multiplied them by 2^4, they also matched the exactly evaluated partition function.

As I mentioned, it is possible to reduce the qubit cost to n. For example, it was shown in Ref. [164] that, if we were interested in extracting the classical shadow of a Gibbs state from an n-qubit simulation instead of the Gibbs state

itself, QET-U would be equally applicable. Furthermore, as noted in Eq. (4.10), if we were performing standard sampling of the mixed Gibbs state, or measuring the expectation value of a unitary (as is the case here), the second 'Env' register is actually not necessary; this is because if we measured the second 'Env' register and found it in the binary state $|b_k\rangle$, the 'Sys' register will also be in state $|b_k\rangle$ – which can also be obtained by randomly sampling $|b_k\rangle$ states from a n-qubit register. The second register is only necessary for algorithms that require the use of a purified state, such as coherent amplitude estimation.

Despite this qubit cost saving, which is significant especially if the target class of hardware is small but fault-tolerant quantum computers, the circuit depth required to converge across the range of g values studied in this small, 4-site Ising system was still high at 6750 gates. The exponential shot measurement cost that comes with the estimation of the partition function also suggests that more work will be necessary to reap any potential quantum advantage in the quantum simulation of partition functions.

4.3.3 Open Quantum System Dynamics

Finally, I considered the basic two-level open quantum system with the Hamiltonian:

$$\mathcal{H} = -\frac{\delta}{2}\hat{\sigma}_z - \frac{\Omega}{2}\hat{\sigma}_x \tag{4.23}$$

and a single amplitude damping channel:

$$\hat{L} = \sqrt{\gamma}\hat{\sigma}_- = \sqrt{\gamma}(\hat{\sigma}_x - i\hat{\sigma}_y), \tag{4.24}$$

where δ is the detuning parameter, Ω is the Rabi frequency and γ is the damping constant or spontaneous emission rate. The density matrix of this system is a 2×2 matrix, therefore its vectorised form can be encoded on a 2^2-element state vector, corresponding to 2 qubits which are respectively labelled A and B.

If we first consider the case where $\delta = \Omega = 0$, then

$$\hat{\mathcal{L}} = \hat{L}_A^* \otimes \hat{L}_B t - \frac{1}{2}\hat{I}_A \otimes (\hat{L}^\dagger \hat{L})_B t - \frac{1}{2}(\hat{L}^T \hat{L}^*)_A \otimes \hat{I}_B t.$$

Because Eq. (4.24) is a real matrix, $\hat{L}^* = \hat{L}$ and $\hat{L}^\dagger = \hat{L}^T$, and therefore,

$$\begin{aligned}\hat{L}_A \otimes \hat{L}_B = \gamma(&\hat{\sigma}_{xA} \otimes \hat{\sigma}_{xB} - i\hat{\sigma}_{xA} \otimes \hat{\sigma}_{yB} \\ &- \hat{\sigma}_{yA} \otimes \hat{\sigma}_{yB} - i\hat{\sigma}_{yA} \otimes \hat{\sigma}_{xB})\end{aligned}$$

and

$$\begin{aligned}\hat{L}^\dagger \hat{L} &= \gamma(\hat{I} - i\hat{\sigma}_x\hat{\sigma}_y + i\hat{\sigma}_y\hat{\sigma}_x + \hat{I}) \\ &= 2\gamma(\hat{I} + \hat{\sigma}_z).\end{aligned}$$

We can therefore determine the Hermitian and anti-Hermitian components Eq. (4.17) and Eq. (4.18) of the master Hamiltonian:

$$\begin{aligned}\mathcal{H}_1 &= -\gamma(\sigma_{xA} \otimes \sigma_{yB} + \sigma_{yA} \otimes \sigma_{xB}) \\ \mathcal{H}_2 &= \gamma(\sigma_{xA} \otimes \sigma_{xB} - \sigma_{xA} \otimes \sigma_{xB} - \sigma_{zA} - \sigma_{zB} - 2\hat{I}).\end{aligned}$$

In this case, the ITE of $\mathcal{H}_2$ must reflect the exact non-unitary evolution, and the imaginary-time is reversed. I therefore used QET-U to perform Eq. (4.3) which undoes the scaling, with $\tau > 0$ and normalised to the maximum eigenvalue. I used only a 6th degree polynomial to approximate the transform, and two Trotter steps for each call of the time evolution oracle in the QSP sequence. I therefore performed time evolution using a circuit implementation of Eq. (4.20), for a total time of 10 a.u. with $\delta t = 1.0$ a.u..

In this scenario, the populations of the excited state and ground state are given by the square roots of the probabilities of measuring the $|00\rangle$ and the $|11\rangle$ states respectively. These are plotted in Fig. 4.8. The populations agree well with the exact solution. Reintroducing the Hamiltonian dynamics, which involves only two additional terms to $\mathcal{H}_1$:

$$\mathcal{H}_1 = -\gamma(\sigma_{xA} \otimes \sigma_{yB} + \sigma_{yA} \otimes \sigma_{xB})$$

produces evolution that agrees well with the exact propagation.

Because the non-unitary dynamics has to be Trotterised, more shots will be required for long time evolution as the cumulative success probability decays. In

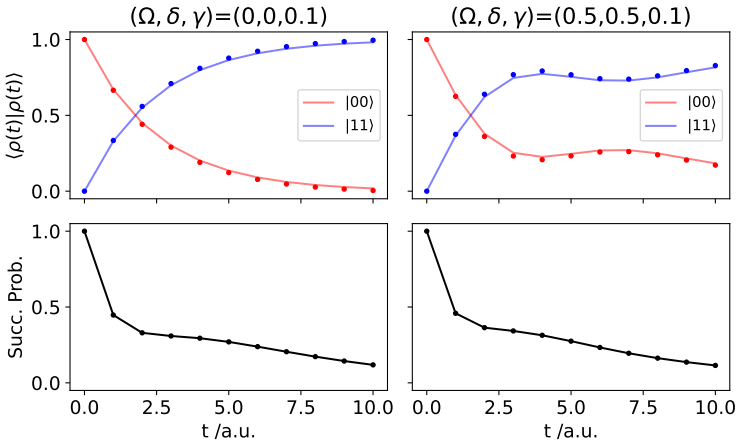


Figure 4.8: Propagating open quantum system dynamics using QET-U. (TOP) The populations of the ground and excited states are encoded as the qubit amplitudes, and in practice will be obtained by state tomography. (BOTTOM) The success probability is cumulative, and for the time scales I explored remained above 10%. Figure reproduced from [3] – I was its original creator.

this simple test case as shown in Fig. 4.8, the cumulative success probability remains above 10% even at long times. At each time step, the circuit depth is 904 gates; predicting the total 10 a.u. dynamics described herein requires a gate depth of $O(10^4)$.

4.3.4 Hamiltonian scaling: Bounds of extremum eigenvalues

The QET-U approach used here requires pre-computation of the extremum eigenvalues, which are necessary for eigenvalue scaling such that they span $\eta, \frac{\pi}{2} - \eta$. For these demonstrations, I pre-computed the exact values E_0 and E_{N-1} for eigenvalue shifting and scaling in the reported QET-U emulation. In general, this is of course an extremely challenging task and estimating good bounds is dependant on the applicability of this algorithm. I will conclude by considering some strategies to address this step.

The simplest approach for calculating the approximate eigenvalue bounds of the operator when it is supplied as a linear combination of terms $\mathcal{H} = \sum_k a_k \mathcal{P}_k$ (which is the case for second quantised chemistry and physics Hamiltonians) is given by

$$E_{\pm} = \pm \sum_k |a_k| \cdot \|\mathcal{P}_k\|. \tag{4.25}$$

Where the $\|\mathcal{P}_k\|$ spectral norm of operator $\mathcal{P}_k$. Other methods to find the lower bound of the ground state $\tilde{E}_0$ can be provided by computing the energy of the initial reference state $\tilde{E}_0 = \langle\psi_0| \mathcal{H} |\psi_0\rangle$. The accuracy of this bound naturally depends on the overlap between the initial state and the ground state; this holds its own unique challenges in the strong correlation regime as discussed. An upper bound of the maximum eigenvalue $\tilde{E}_{N-1}$ can in turn be estimated using Gershgorin's Circle Theorem [239], whereby discs bounding eigenvalues on the complex plane are constructed using diagonal elements of $\mathcal{H}$ (the origins of said discs) and sum of off diagonal elements on a corresponding row (the radius of said discs). The fact that eigenvalues of $\mathcal{H}$ are real and symmetric can be exploited to simplify upper bound determination. Furthermore, like the LCU Hamiltonian input model, it should be possible to divide the coefficients of Pauli terms by the norm of the Hamiltonian, and perform controlled time evolution of this Hamiltonian instead.

4.4 Concluding Thoughts

In Ref. [3], I showed that QET-U can in principle be leveraged to perform non-unitary evolution, with potential applications on early fault-tolerant quantum computers as it only requires a single additional ancillary qubit. By using recent advances in efficient computation of phase factors, I was able to generate high-degree polynomial approximations of the ITE propagator. While normalisation of the ITE propagator gave a boost to the probability of preparing the desired imaginary-time-evolved state in the systems tested, as the number of qubits increase or as systems become more strongly correlated, the probability of success can become prohibitively small, potentially offsetting any quantum speedup in the ability to perform many-body non-unitary evolution.

Given that the method seems most suited as a means of preparing a state that has undergone non-unitary evolution or a thermal state, further work should seek to manage this probability of success so that it does not compound the measurement cost when subsequent steps in quantum simulation are accounted for e.g. quantum measurement, which may further increase the number of shots. As mentioned, a

natural extension to counteract this is to incorporate amplitude amplification, at the expense of additional quantum gates [227]. Hybridising this approach with other 'cheap' state preparation algorithms or initial states to 'warm-start' the ITE would also improve the probability of success by reducing the amount of imaginary-time necessary for the task of e.g. ground state preparation.

We did not specifically explore how the errors in the Hamiltonian rescaling or how different Trotterisation schemes of the Hamiltonian during its encoding might affect the performance of ITE. A more concrete understanding of how these errors might influence the quality of the prepared state will be valuable.

The utility of the described approach will ultimately depend on the non-unitary simulation task that the user wishes to perform. For example, if one already has access to an accurate estimate of the ground state or the composition of eigenstates in the initial state, well established methods like iterative phase estimation might in fact be more useful than the methods explored here for ground state preparation. Applications to open quantum system dynamics investigated here seem promising for large scale Lindbladian dynamics simulations, though the diminishing cumulative probability of success needs to be resolved and trying to obtain the density matrix of the state instead of some meaningful observable may prove prohibitive due to the requirement for state tomography. Nonetheless, QET-U for ITE can trivially be translated to first-quantized e.g. real-space grid quantum simulations for ground state preparation (see next Chapter), and indeed even to molecular geometry optimisation as recently proposed in Ref. [240]. Extension to preparing state vectors from the time-integration of non-unitary differential equations is clear.

Boy: "Do not try and bend the spoon—that's impossible. Instead, only try to realise the truth."

Neo: "What truth?"

Boy: "There is no spoon."

— The Matrix (1999)

5

Real-Space Grid Physics Simulations on Fault-Tolerant Quantum Computers

Contents

5.1 Introduction 114
5.2 Methods 119
5.2.1 Split-Operator Hamiltonian Simulation under the Coulomb potential 119
5.2.2 Energy Observable 125
5.2.3 Attenuation and Scattering 126
5.2.4 State Preparation 129
5.2.5 Augmented Split-Operator 136
5.3 Results 138
5.3.1 Trotter Evolution and the Relationship between Spatial and Temporal Resolution 140
5.3.2 Quantum Dynamics Experiments 144
5.3.3 Initial State Preparation 149
5.3.4 Antisymmetrisation of the initial state 152
5.3.5 Augmented Split-Operator 154
5.3.6 Informal Resource Estimate 158
5.3.7 Quantum Hardware and Architecture 165
5.3.8 Additional Analysis and Results 167
5.4 Concluding Thoughts 169

In this chapter, I use text from the Science Advances paper "Grid-based methods for chemistry simulations on a quantum computer" [4] verbatim – I was the first

and original author of the text. The work was done in collaboration with coauthors Richard Meister, Tyson Jones, David Tew and Simon Benjamin at the University of Oxford. Simon Benjamin and I developed all of the methodology in this work and performed the majority of the numerical simulations.

5.1 Introduction

As discussed in previous Chapters, it is generally believed that if the technical challenges of scaling qubits can be overcome, fault-tolerant and error-corrected quantum computers can one day outperform classical computers in the simulation of quantum many-body physics. While the quantum physics community waits for these technological breakthroughs to be realised, there is naturally a debate on how one might best utilise future fault-tolerant quantum computers for many-body physics simulations.

I was interested in exploring how a *first-quantised, real-space grid* encoding of a many-body wavefunction on a quantum computer (See Chapter 2, Section 2.2.3) might be advantageous for chemical and material simulations in the fault-tolerant quantum era. In contrast to the compact, conventional second-quantised representation of (Fermionic) many-body states, which uses the occupation of spin orbitals that implicitly respect particle spin (anti-)symmetry, the first-quantised grid approach instead represents wavefunctions in real space on a discretised grid of points, and particle symmetry is explicitly encoded. At first glance, this representation is appealingly intuitive; many-body quantum physics concern dynamical processes that take place in physical, three-dimensional space, and naturally we might choose to simulate such processes within a finite-sized 'simulation world' or 'box' divided up into a number of volumetric pixels (or *voxels*). Representing a quantum many-body system in such a way has a cost that grows exponentially in the number of particles, and has therefore been dismissed in most classical numerical methods. In contrast, it can be achieved very efficiently on a digital quantum computer, even if the way we might extract properties from this simulation has to conform to the same quantum measurement constraints of the modelled system. Some researchers even

claim this representation will in fact offer the optimal quantum resource scaling for modelling large, complex and interesting molecules and materials [51, 53, 241] compared to second quantisation. Some have further suggested that first-quantised simulation can efficiently encode both nuclear and electronic degrees of freedom on an equal footing, potentially addressing the gap in simulating non-Born-Oppenheimer processes in modern chemical physics [47].

Real-space grid methods have been used with classical computers since at least the 80's [102–104, 242, 243]. However, the resource cost grows exponentially with the number of pixels and thus in practice only simplified models with wavefunctions that are, so to speak, 'heavily pixelated' can be stored and processed within conventional Random Access Memory (RAM). Even with quantum computers, first-quantised methods will require many more qubits and deeper circuits for meaningful realisations compared to second-quantisation, making them impractical on the noise-burdened quantum computers of today – hence the expectation that first-quantisation approaches for many-body physics simulation will only be useful in the fault-tolerant regime.

Because of these limitations, most prior studies of these first-quantised, grid-based methods have therefore focused on theoretical 'pen and paper' analysis of the quantum resource costs. Building on the works of Zalka [46] and Wiesner [45] from the mid 1990's, it was observed just over a decade ago by Kassal *et al.* [47] that it will be advantageous to store high-dimensional spatial grid representations of molecular wavefunctions on digital quantum computers using qubits that scale only linearly with the degrees of freedom. They also showed that, by exploiting the remarkably efficient quantum Fourier transform (QFT), a quantum version of the SO method (the SO-QFT) will be able to simulate the dynamics of a chemical Hamiltonian in polynomial time. These authors challenged the *status quo* of limiting the split-operator method to propagating nuclei on electronic potential energy surfaces.

Further work has been done since to explore the quantum propagation of one-dimensional Gaussian wavepackets [48], quantum harmonic oscillator [50], non-adiabatic dynamics across coupled, pre-computed potential energy surfaces

[52], imaginary-time evolution for molecular geometry optimisation [59] as well as ground state, Gibbs state and partition function determination [54]. Work geared specifically towards understanding the scaling of the SO-QFT method at mid-to-large implementations include e.g. (1) that of Jones *et al.*, which looked at the resource necessary for a fault-tolerant implementation [49], (2) the extensive investigation into first quantized basis-set simulations for chemistry of Su *et al.* that also explored the block-encoding the SO-QFT Hamiltonian for simulation using qubitization and interaction picture techniques [53], and (3) analysis by Childs *et al.* which provided concrete upper bounds for the gate complexity of SO-QFT to arbitrary order [55]. Other attempts at direct quantum simulations in real-space grids include using finite difference stencils to approximate the kinetic operator in second quantization [244], and in first quantization with propagation based on the truncated Taylor series algorithm [51], as well as a Cartesian component-separated approach [56].

In Ref. [4], my collaborators and I instead took a different approach; we deployed very substantial classical computing resources to perform exact emulations of small but noise-free quantum computers. These emulated computers then simulate representative case studies that reflect potential target applications for achieving quantum advantage in the first-quantised grid paradigm. Some key applications include:

Application I: Computing the electronic structure, especially the ground state, of large many-body systems that are out-of-reach of conventional quantum chemistry. This underpins any efforts in using computers to design and explore new chemicals and materials. One key question might be how the number of qubits and circuit operations scale compared with the second quantised representation for eigenvalue determination. In this work, I explored a number of established quantum algorithmic techniques applied to extracting energy properties from grid-based quantum simulations.

Application II: Simulation of molecular dynamics in the presence of external fields. Ultimately, efforts in this direction encompasses topics such as photochemistry and laser excitation. Applications that surpass the capabilities of classical computing

include comprehensively modelling the electron dynamics of photosynthesis, which would involve the modelling of highly complex molecules e.g. the Fenna-Matthews-Olson complex [245]. The coherent quantum control of small molecules using laser driven dynamics has also been considered one of the 'holy grails' of chemical science [246]. The exploratory work here involves a suddenly-applied external field, which results in a dipole oscillation and ionisation of a single bound electron.

Application III: Simulating the dynamics of particle scattering events. In general, electron-molecule scattering is relevant in spectroscopy, astrochemistry, atmospheric chemistry, as well as manufacturing processes [247]. While part of the computational challenge is scanning through possible initial energies of the incoming electron, predicting what happens upon collision and scattering is a highly quantum dynamical process difficult to model classically. Many such events involve reaction intermediates with fleeting lifetimes that are hard to observe, and occur under conditions experimentally challenging to access. The exploratory work here involves an incident electron scattering from a bound electron and potentially ionising it.

The cost of emulation limited us to very modest-sized quantum computers (my collaborator use at most 36 'perfect' qubits), and so we never directly simulated the 3 applications listed above on meaningful molecules. Nonetheless, we were able to explore a number of informative scenarios within this restriction: 2D and 3D simulations of one- and two-electron systems. By examining these emulations, my collaborators and I were able to produce approximate costs and performance measures.

My collaborators and I performed wavepacket time-propagation with the split-operator quantum Fourier transform (SO-QFT) Hamiltonian simulation approach, based on the Lie-Suzuki-Trotter product formula (see Chapter 2). Specifically, we modelled subatomic particles (electrons) interacting directly via the Coulomb potential; a prerequisite for ultimately treating electrons and nuclei on a fully quantum basis. Compared with other Hamiltonian simulation methods (see e.g. linear combination of unitaries based on the truncated Taylor series algorithm [51] or qubitization [53, 55]), SO-QFT can require the fewest qubits to implement time evolution [47]. Other attempts at direct quantum simulations in real-space grids

include using finite difference stencils to approximate the kinetic operator in second quantization [244], as well as a Cartesian component-separated approach [56].

Beyond the inherent value in emulating previously-proposed quantum algorithms for the first time in the real-space context, and thus determining performance measures that hitherto could only be estimated, my collaborators and I made a number of contributions:

- To perform scattering and ionisation modelling within the finite 'simulation box' of the grid-based method, my collaborator and I co-created a non-unitary wavepacket attenuation approach, inspired by complex absorbing potentials (CAPs) from classical simulation. The method uses measurement of a single entangled ancilla qubit to attenuate particles that are ejected from the finite simulation environment, preventing them from returning at the periodic boundary and interfering with the simulation. Sampling the outcome of the ancilla measurement also doubled as a means of tracking escape probabilities of wavepackets, and has the potential for quantum computing reaction rates.

- I found that the use of the unmitigated Coulomb singularity posed a challenge in the spatial resolution, which my collaborator and I addressed by co-creating the augmented split operator (ASO) approach. An additional small quantum circuit was used to correct the Trotter error incurred at every SO-QFT time evolution step when low spatial and temporal resolution is used.

- I assessed prior methods for state preparation and made additional contributions which addressed the challenges therein:

  - My collaborator and I investigated using the single-ancilla iterative phase estimation (IPE) measurement to project out excited states;

  - I investigated an adaption of the single-ancilla probabilistic imaginary time evolution (PITE) method [54] for approximating small imaginary time evolution steps;

  - My collaborator and I built upon existing work [60] to co-create a method that explicitly generated the correct particle (anti)symmetry for first-quantised simulations.

- Finally, I performed informal estimations of the quantum resource costs (time and hardware scale) for modelling interesting molecules representative of the aforementioned applications. My collaborator and I also indicated the hardware layout of a suitable quantum computer for first-quantised grid-based simulations.

5.2 Methods

5.2.1 Split-Operator Hamiltonian Simulation under the Coulomb potential

The split operator approach performs many-body Hamiltonian simulation by splitting a Hamiltonian $\mathcal{H}$ into its kinetic $\mathcal{H}_{\text{kin}}$ and potential interaction $\mathcal{H}_{\text{int}}$ parts, then invokes the Lie-Trotter-Suzuki approximation and applies evolution of each part separately in 'bursts' across a short time step δt:

$$e^{-i\mathcal{H}\delta t} = e^{-i(\mathcal{H}_{\text{kin}}+\mathcal{H}_{\text{int}})\delta t} = e^{-i\mathcal{H}_{\text{kin}}\delta t}e^{-i\mathcal{H}_{\text{int}}\delta t} + O(\delta t^2) = U_{\text{SO}}(\delta t) + O(\delta t^2). \quad (5.1)$$

The above equation shows a first-order splitting; second-order ($O(\delta t^3)$) Trotterisation split-operators have well-documented numerical advantages [100, 105, 248, 249], and higher order Lie-Trotter-Suzuki sequences can also be relevant [250, 251]. The U_{SO} then effects an approximate time step evolution δt onto an state:

$$|\psi(t+\delta t)\rangle = U_{\text{SO}}(\delta t)\ |\psi(t)\rangle$$

In Chapter 2, Section 2.2.3, I mentioned that $\mathcal{H}_{\text{kin}}$ and $\mathcal{H}_{\text{int}}$ are diagonal or approximately-diagonal in the momentum-space and the real-space grid-based representation respectively – which we can efficiently switch between using the quantum Fourier transform U_{QFT}. Thus, the two parts of the split-operator

quantum Fourier transform (SO-QFT) cycle can be implemented each in their preferred, diagonal basis;

$$U_{\mathrm{SO}}(\delta t) = e^{-iD_{\mathrm{int}}\delta t}\left(U_{\mathrm{QFT}}^{\dagger}\; e^{-iD_{\mathrm{kin}}\delta t}\; U_{\mathrm{QFT}}\right). \tag{5.2}$$

Here D_{kin} and D_{int} are real, diagonal matricies. The classical variant SO-FT algorithm which uses the discrete Fourier transform has had decades of demonstrated success in nuclear wavepacket propagation [120] on electronic potential energy surfaces, but as far as my collaborators and I know was never used with Coulomb potentials.

The most involved part of SO-QFT is the time evolution under the Coulomb interactions between particles. Consider first $\mathcal{H}_U$, which represents the energy of single particles in 'classical' fields. In this case, we are primarily interested in an attractive Coulomb potential representing atomic nuclei (although I will discuss a variation including a static electric field presently);

$$\mathcal{H}_{\mathrm{U}} = \sum_{p=1}^{P} \frac{Q_p}{|\mathbf{r}|_p} = Q\sum_{p=1}^{P} \frac{1}{\sqrt{x_p^2 + y_p^2 + z_p^2}}. \tag{5.3}$$

For simplicity, I took the nuclear charge Q outside of the sum since all Q_p are the same in the atomic and molecular systems I looked at; however there would be no difficulty in retaining distinct values. As before, the time evolution operator can be written as

$$U_{\mathrm{U}}(\delta t) = \prod_{p=1}^{P} \exp\left(-i\frac{Q\delta t}{\sqrt{x^2 + y^2 + z^2}}\right).$$

The operations are also independent between the sub-registers corresponding to different particles, but not independent between registers corresponding to different spatial dimensions assigned to a given particle.

In the real space representation when the corresponding dynamics is implemented, the computational basis states of each sub-register will thus correspond to locally-peaked 1D single-particle basis wavefunctions,

$$\phi_n(x) = P_{n_r}^{n}(x) \quad \leftrightarrow \quad |n\rangle$$

with each such 'pixel function' having its primary peak at $x_n = nL/2^{n_r} = n\delta r$. Similar expressions describe the roles of the y sub-register states $|m\rangle$, peaked at y_m, and the z sub-register states $|l\rangle$ peaked at z_l.

One could then proceed to apply a period of time evolution under $\mathcal{H}_U$ using an approximation that *would* be exact in the limit of infinite spatial resolution. In that case, one could apply a series of phase changes to our quantum registers to properly describe any time step δt

$$\begin{aligned} |n\rangle\,|m\rangle\,|l\rangle \quad &\Rightarrow \quad \exp\left(\frac{-iQ\,\delta t}{\sqrt{x_n^2 + y_m^2 + z_l^2}}\right)|n\rangle\,|m\rangle\,|l\rangle \\ &= \quad \exp\left(\frac{-iQ\,\delta t}{\delta r\sqrt{n^2 + m^2 + l^2}}\right)|n\rangle\,|m\rangle\,|l\rangle\,. \end{aligned} \tag{5.4}$$

The task of implementing the latter is only modestly more complex than the former – this time, one would need to compute on an entire register of all three spatial sub-registers and the phases required are inverse-square-root instead of merely square. Efficient evaluation of functions such as the inverse-square-root on quantum computers is an active area of development. In particular, Kassal *et al.* [47], Jones *et al.* [49] and Häner *et al.* [252] proposed procedures for reversible computation of the inverse-square-root using fixed-point arithmetic and Newton-Raphson iteration. Polynomial function interpolation using quantum read-only memory (QROM) type circuits is another approach [253]. Recent work from Poirier *et al.* [56, 254] combined with other formulations of the Coulomb potential may be another option for computing the inverse-square-root. I will refer the reader to Refs. [49, 252] for resource estimates, but here it suffices to note that the number of gates required can potentially scale as $\mathcal{O}(n_r^2)$.

This would be the exactly correct process if the spatial basis functions were Dirac delta functions. These states, in fact, have finite spatial extent, but we can expect that for any sufficiently short time δt and given a smoothly varying potential, an adequate spatial resolution δr^{-1} will result in dynamics that converge to the exact behaviour. However, the Coulomb potential is singular at $\mathbf{r} = \mathbf{0}$, and consequently the behaviour in this region needs to be considered carefully. I will show later a

method that my collaborator and I developed that makes it possible to augment the basic SO-QFT cycle of kinetic and potential evolution with an additional third phase, whose purpose is to stabilise behaviour at the core singularity.

Then, consider the part of the Hamiltonian corresponding to particle-particle interactions $\mathcal{H}_{\mathrm{V}}$. This is straightforward to handle using the same approximation as described above, i.e. by updating registers *as if* the basis functions they represent are Dirac delta functions. Then the operator,

$$U_{\mathrm{V}}(\delta t) = \exp\left(i\delta t \sum_{p,q=1;p\neq q}^{P} \frac{q_{p,q}}{|\mathbf{r}_p - \mathbf{r}_q|} \right) \tag{5.5}$$

will be approximated as having diagonal form (and will therefore commute with the $U_U(\delta t)$ operator). Consider two particle sub-registers, each composed of three spatial registers, which represent a given pair of particles $p = 1$ and $p = 2$. Then basis states will be updated as

$$|n_1\, m_1\, l_1\rangle\, |n_2\, m_2\, l_2\rangle \;\Rightarrow\; e^{-i\Theta}\, |n_1\, m_1\, l_1\rangle\, |n_2\, m_2\, l_2\rangle$$

where

$$\begin{aligned} \Theta &= \frac{Q\,\delta t}{\sqrt{(x_{n_1} - x_{n_2})^2 + (y_{m_1} - y_{m_2})^2 + (z_{l_1} - z_{l_2})^2}} \\ &= \frac{Q\,\delta t}{\delta r \sqrt{(n_1 - n_2)^2 + (m_1 - m_2)^2 + (l_1 - l_2)^2}} \end{aligned} \tag{5.6}$$

One means of computing this operator is to pair every particle with a partner and perform the following process (which can occur simultaneously over each pair of registers), before repeating with a different pairing, and so on until all pairings are considered. For each pair, we can perform the computation

$$\begin{aligned} &|n_1, m_1, l_1\rangle\, |n_2, m_2, l_2\rangle \;\Rightarrow \\ &\quad |n_1 - n_2, m_1 - m_2, l_1 - l_2\rangle\, |n_2, m_2, l_2\rangle \end{aligned} \tag{5.7}$$

which is a straightforward instance of quantum arithmetic on pairs of sub-registers. One suitable method uses the QFT which of course our machine will already have been optimised for [255]. It is also frugal in terms of the qubit count – only a single

additional qubit is needed to fully represent each difference $n_1 - n_2$. With the registers in this form, the same procedure employed for the single-particle phases can be used. Once the desired phases are applied, the reverse of the computation Section 5.2.1 is performed and the whole process is repeated with another set of particle pairings. For a system of P particles, the total number of particle pairings is $\frac{1}{2}P(P-1)$. However, there is a natural parallelism so that the time cost should scale only linearly with P. Moreover, there are methods to trade greater hardware resources to reduce the time cost; Jones *et al.* introduced a parallelised scheme for speeding up the computation of the particle pairings using only $\mathcal{O}(\log P)$, or even $\mathcal{O}(1)$, circuit depths [49]. Regardless of the implementation method used, the elements of the diagonal matrix D_{int} from Eq. (5.2) are thus simply sums of phases.

In the approach just described, the relative coordinates of each particle pairing can be found. There was no need to compute the complete transformation by mapping the second particle's register to $|n_1 + n_2, m_1 + m_2, l_1 + l_2\rangle$. However, this would have been possible. A motivation for completing this full transformation is that general, non-diagonal operations could be performed if we wished; this observation relates to the augmented split-operator discussed later on in the text.

Finally, time evolution under purely the kinetic part of a P-particle Hamiltonian $\mathcal{H}_{\text{kin}}$ i.e.

$$U_{\text{kin}}(\delta t) = e^{-i\mathcal{H}_{\text{kin}}\delta t}$$

can be expressed as a product of operators acting independently on each particle and in each dimension:

$$U_{\text{kin}}(\delta t) = \prod_{p=1}^{P} \prod_{q\in\{x,y,z\}} \exp\left(i\frac{\hbar^2\delta t}{2m_p}\frac{\partial^2}{\partial q^2}\right)$$

because these components commute – in other words, $\mathcal{H}_{\text{kin}}$ is separable and exactly local (diagonal) in a k-space grid. For one particle, each x-basis state, and thus computational basis state, represents

$$\phi_k(x) = L^{-\frac{1}{2}} \exp\left(\frac{i2\pi kx}{L}\right) \quad \leftrightarrow \quad |k\rangle\,.$$

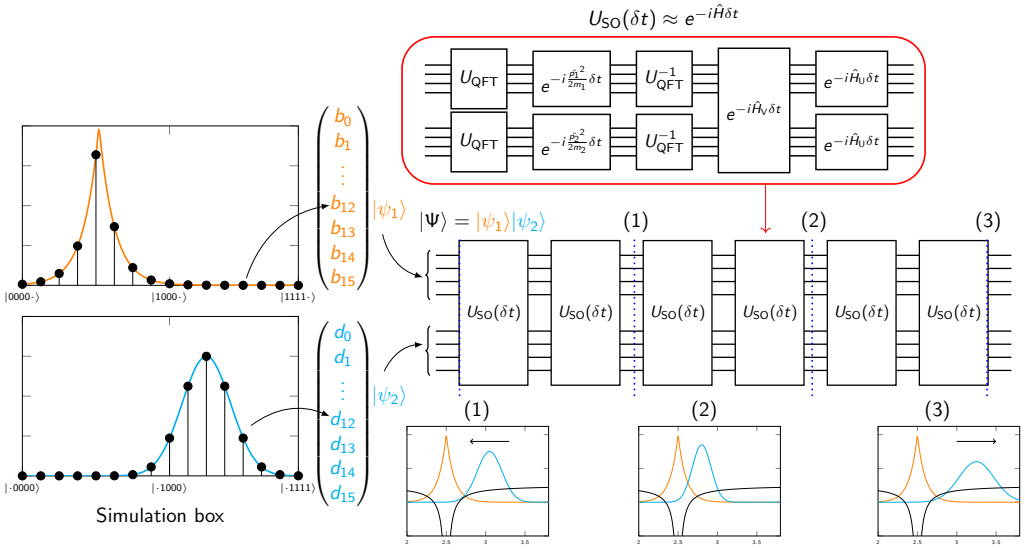


Figure 5.1: Quantum simulation using the SO-QFT approach in the real-space representation. In this cartoon, 2 electrons in a simulation box are represented: One is modelled as a Gaussian wavepacket, and the other as the bound ground state electron of the hydrogen atom. The two interact via a repulsive Coulomb potential, and the atomic nucleus is modelled as a classical attractive Coulomb potential. The particles are digitised into regularly spaced pixels, encoded into the state of two identically-sized qubit sub-registers. The starting state is a tensor of these two states (a Hartree product), and upon antisymmetrisation of the two particles (which would generate a Slater determinant in quantum chemistry terminology), the initial state can then be propagated in time by repeated application of the SO-QFT short-time evolution operator $U_{\mathrm{SO}}(\delta t)$ (depicted by the top circuit). A possible sequence of the subsequent dynamics, where the unbound electron scatters with the bound electron, is plotted at the bottom. Figure reproduced from [4] – I was the original creator.

The proper way to account for the action of $U_{\mathrm{kin}}(\delta t)$ on the molecule is to introduce phases onto our computational basis states according to

$$|k\rangle \;\Rightarrow\; e^{iC\delta t\, k^2}\,|k\rangle \tag{5.8}$$

with constant $C = 2\hbar^2\pi^2/(L^2 m_p)$. Thus the task is simply to apply phase operations to the quantum computer, independently for each particle sub-register (and in parallel if the hardware has that capability). Each element of the matrix D_{kin} appearing in Eq. (5.2) is just the sum of $C(k_x^2+k_y^2+k_z^2)$ over all particles. It is trivial to implement the required phases as, for example, a sequence of single- and two-qubit phase gates (see Refs. [49, 52]). The number of gates required scales as $\mathcal{O}(n_r^2)$.

5.2.2 Energy Observable

I will now review standard energy measurement approaches that my collaborators and I have applied to grid-based quantum simulations.

Energy Expectation

In the grid-based approach it is possible to split the Hamiltonian and calculate the kinetic energy expectation in k-space, followed by the potential energy expectation in real-space

$$\langle E \rangle = \langle \psi | \, U_{\mathrm{QFT}}^{\dagger} \mathcal{H}_{\mathrm{kin}} U_{\mathrm{QFT}} \, | \psi \rangle + \langle \psi | \, \mathcal{H}_{\mathrm{int}} \, | \psi \rangle \tag{5.9}$$

where here $|\psi\rangle$ is in the spatial representation.

Iterative Phase Estimation

The energy of a state can also be extracted by tracking the global phase it acquires during time propagation. This is given by the phase of the autocorrelation function, which is a discrete-time series of the inner product between a propagated state and the initial state. Suppose we time propagate an eigenstate Ψ_n of $\mathcal{H}_{\mathrm{tot}}$ with energy E_n. The corresponding autocorrelation signal is

$$\langle \Psi_n(t) | \Psi_n(t=0) \rangle = e^{-iE_n t}, \tag{5.10}$$

where the absolute value $| \langle \Psi(t=0) | \Psi(t) \rangle |$ should be close to unity and I used it as one measure of a simulation's veracity.

On a quantum computer, this otherwise-not-observable global phase is efficiently extracted using phase estimation. Its utility in the context of SO-QFT and first quantised simulation is well recognised (see e.g. [47]). In the Iterative Phase Estimation (IPE) approach, even a single ancilla [50, 256] is sufficient to learn this phase; a resource cost saving that will be welcome in the early fault-tolerant regime. The method is summarised in the left panel of Fig. 5.3. The approach is to conditionally apply N SO-QFT steps $U^N(\delta t)$ controlled by an ancillary qubit in the $|+\rangle$ state. At the N^{th} step, measure the ancillary qubit in the $|+\rangle$ basis,

at which point the state is discarded. It can be see that for an eigenstate, the global phase information is encoded in the relative phase between the $|0\rangle$ and $|1\rangle$ state of the ancillary qubit

$$\begin{aligned} |+\rangle\,|\Psi_n(0)\rangle &\Rightarrow \frac{1}{\sqrt{2}}\Big(\,|0\rangle\,|\Psi_n(0)\rangle + |1\rangle\,|\Psi_n(t)\rangle\,\Big) \\ &= \frac{1}{\sqrt{2}}\Big(\,|0\rangle + e^{-iE_n t}\,|1\rangle\,\Big)\,|\Psi_n(0)\rangle\,. \end{aligned} \tag{5.11}$$

The probability of finding the ancillary qubit in state $|+\rangle$ fluctuates as the phase. I used this to extract a periodic time signal $a(t)$ where the frequency is proportional to the energy of the simulated wavefunction

$$a(t) = \cos^2\left(\frac{E_n}{2}t\right). \tag{5.12}$$

Because the number of qubits I could classically emulate is limited, using the single-ancilla IPE for our demonstrations herein is a natural choice. In later sections I report the exact evolution of $a(t)$ plotted at regular time points; this is straightforward since classically-emulated quantum processors were used. On a real device, because the single-ancilla projection probability is statistical in nature, the time propagation and measurement will have to be repeated multiple times. If more ancilla qubits are available, this extends to the standard, canonical Fourier phase estimatio (see Chapter 2).

5.2.3 Attenuation and Scattering

The use of the k-space mapping in SO-QFT imposes a periodic boundary condition in real space. For non-periodic systems, this introduces artificial interactions and state interference through the boundary, and one of the requirements on the simulation box is that it is sufficiently large that these have negligible impact on the dynamics. Given that, for fixed spatial resolution, the width can be increased with only a logarithmic cost both in the number of additional qubits and the 'wall clock' execution time this may not pose a problem. However, there are simulation tasks where simply enlarging the box may not be an ideal solution. For example, a

study of a scattering process where the quantum states of scattered particles have contributions that travel at speed towards the edges of the simulation box.

One alternative is to attenuate wavefunction amplitude near the boundary. My collaborator and I co-developed a quantum circuit analogue of complex absorbing potentials (CAPs), a well-established technique from classical simulation [257]. I now present this method and show that it can also be used to track the probabilities of wavepackets being found at specific pixels, important to the determination of reaction rates in molecular processes. Attenuation of wavefunction contribution on a qubit register would naturally be a non-unitary process involving measurement and implies some probabilistic element; however the cost associated with any need for repetition may be preferable to the increase in the number of qubits required in the former approach. My collaborator and I tested these techniques in the ionisation of a single bound electron by a strong applied electric field and two-particle scattering scenarios, which I present in Section 5.3.2.

The interaction potential part of the Hamiltonian $\mathcal{H}_{\text{int}}$ is modified to contain both a real and (negative) imaginary part

$$\mathcal{H}_{\text{int}} = \mathcal{H}_{\text{U}} + \mathcal{H}_{\text{V}} - i\mathcal{V} \tag{5.13}$$

The time evolution operator is thus

$$|\psi(\mathbf{r}, t)\rangle = e^{-i\mathcal{H}t} e^{-\mathcal{V}t} |\psi(\mathbf{r}, 0)\rangle \tag{5.14}$$

The absorbing potential $\mathcal{V} = V(\mathbf{r})$ should be balanced between sharp, abrupt attenuation which reduces the size of the attenuation region while suppressing transmission, and smooth, gradual attenuation which suppresses reflection. For brevity, I will only state that it is possible to construct complex absorbing potentials which are completely reflection and transmission free [258], and refer the reader to the rich body of classical quantum dynamics literature on complex potentials [106, 259–263].

On a quantum computer, the objective is to achieve the following non-unitary transformation in (though not limited to) the real space representation at every SO-QFT step

$$\sum a_n |n\rangle \Rightarrow \sum a_n e^{-V(x_n)\delta t} |n\rangle \,, \tag{5.15}$$

followed of course by the necessary renormalization. Here x_n is the location of the peak of the pixel function $P^n_{n_r}(x)$ as per Eq. (2.27). The time resolution δt is merely a constant, corresponding to the temporal granularity of the simulation. The analogous mapping should occur for spatial dimensions y and z, and for each particle in the system – ideally this should be done in a fashion that perfectly preserves any relevant particle exchange symmetry.

To achieve this, one could proceed as follows for each x_n that lies within the attenuating region (i.e. where $V(x_n) \neq 0$). At the end of each spatial part of the SO-QFT cycle, prepare an ancilla qubit in state $|0\rangle$ and perform a multi-controlled rotation $e^{i\sigma_x\theta}$ on that ancilla. The rotation occurs if and only if the sub-register is exactly in state $|n\rangle$ which corresponds to the pixel function peaked at x_n. The ancilla is then measured. If the outcome $|1\rangle$ is obtained, we say that the particle in question has been measured to be that point, and we may either halt the simulation or proceed with the remaining particles – both scenarios are examined in the numerical emulations presented. However, if outcome $|0\rangle$ is obtained then the simulation simply proceeds, except that the amplitude of state $|n\rangle$ has been attenuated from a_n to $a_n \cos(\theta)$ (and the state has been renormalised). By repeating this process for all x_n within the attenuating region, selecting the proper $\theta = \arccos\left(e^{-V(x_n)\delta t}\right)$ in each case, we implement Eq. (5.15) (see centre panel of Fig. 5.3). In parallel the same process can be applied for other sub-registers and other particles.

If the attenuation region is narrow, corresponding to only a few values of x_n, then the above protocol can be relatively efficient. Remarkably frugal constructions for multi-controlled Paulis do exist [264], and of course a controlled Pauli can implement the general θ rotation described above with the use of additional single-qubit rotations. However if the attenuation region is a main portion of the simulation

box then the method would be exponentially inefficient as it requires action for each element of the spatial superposition. Fortunately however, one can simply set a uniform attenuation strength V over a range of x_n values and implement it in one step using an appropriate subset of register qubits in the control process. The simplest example occurs when the width of the attenuating region is $2^{-m}L$ for some integer m (L being the width of the simulation box). In this case, only the most significant m qubits from the sub-register control the rotation. This was the approach taken in the numerical emulations presented herein.

By repeating the entire simulation multiple times and keeping track of the probability that the ancilla qubit yields $|1\rangle$, the time-dependent probability that a particle has reached a given pixel in the simulation box is also being recorded. One could in principle obtain this information by generating the probability distribution of particles at different times through direct sampling of particle sub-registers; indeed, plots that would be obtained this way are generated. However, on a real quantum computer, selecting specific pixels and entangling them with a single qubit that would then be measured is a much more cost efficient statistical approach. One might envisage the determination of reaction rates by selecting key pixels, much like the selection of a reaction coordinate in traditional reaction dynamics, and using an equivalent approach to track the probability that a particle wavepacket crosses said pixels. Later on, I show a proof-of-concept of this approach by directly measuring particle scattering and ionisation probabilities, in combination with CAPs.

5.2.4 State Preparation

Preparing an appropriate starting state is a crucial and non-trivial component of quantum simulation. To be consistent with the performance of the subsequent time-evolution, we should ideally perform state preparation in polynomial time. Furthermore, since we are working in the first-quantised picture, the correct particle (anti-)symmetry therein must be explicitly enforced.

The canonical quantum algorithmic approach is to start by initialising an easily prepared reference state, and then drive it towards e.g. the ground state. In this

Section, I will introduce the state preparation methods my collaborators and I assessed in Ref. [4]; a quantum algorithm for preparing antisymmetric states, state preparation with IPE described in Section 5.2.2, and the probabilistic imaginary-time evolution (PITE) [54] for preparing ground states.

Particle Antisymmetrisation

For many-body systems in which the majority of particles are indistinguishable, the challenge of preparing an initial state with proper exchange symmetry is non-trivial. Here, I assume the preparation of Fermionic antisymmetric states, but the methods to prepare Bosonic symmetric states are near-identical. The choice of method to accomplish antisymmetrisation determines the options for actually loading the initial single-particle basis states, as I now explain.

Specifically, my collaborator and I built on the work of Ref. [60], wherein the authors propose performing an antisymmetrisation on a product of single particle states using an adaption of the methods in Ref. [71] (i.e. permuting a product state by co-sorting a state of the form $|\text{prm}\rangle$). However, they did not explicitly specify how the resulting SWAP flags are to be un-computed (moving from Eqn.(4) to Eqn.(5) in Ref. [60] appears non-trivial).

My collaborator and I introduced two simple solutions to this issue. We rely on the following assumption (which may itself be non-trivial); we must specify a Hamiltonian $\mathcal{H}_{\text{synth}}$ with the property that

$$\mathcal{H}_{\text{synth}} |\psi_i\rangle = E_i |\psi_i\rangle , \qquad E_{i+1} > E_i. \tag{5.16}$$

This Hamiltonian is synthetic in the sense that it need not correspond to any physically legitimate scenario. The condition implies that the states $|\psi_i\rangle$ must be orthogonal; moreover we desire well-separated energies.

Assume that there is some set of P single-particle basis states ψ_i, each of which we know how to prepare on a register. One means of obtaining $\mathcal{H}_{\text{synth}}$ would be to start from a description of this system, and then introduce modifications to conveniently localise Hamiltonian terms and to break any degeneracies in the

single-particle solutions. In this way, the available basis states $|\psi_i\rangle$ will be close to the canonical choice of basis states that might be made in, e.g., chemistry modelling with a conventional computer. While this is an interesting topic to consider further, I will not do so presently but simply assume that $\mathcal{H}_{\text{synth}}$ can be found.

Given that $\mathcal{H}_{\text{synth}}$ is synthetic, it can of course be scaled and shifted arbitrarily. We can align the energies conveniently with respect to the binary states that can be represented by a register of t bits. For example we can set $E_0 = 0$ and $E_{P-1} = 2^t - 1$ where t is sufficiently large that the smallest energy gap $\min(E_{i+1} - E_i)$ is at least unity.

We can proceed in at least two distinct ways, and I will describe the first in detail. It requires a modest number of additional qubits: For each of the P particle registers, introduce an 'tag' ancilla register of t qubits.

1. Prepare the particle registers as $|\psi_0\rangle\,|\psi_1\rangle\,...\,|\psi_{P-1}\rangle$.

2. Set the i^{th} tag register to the the integer closest to E_i, which we write as E'_i.

3. Permute the entire object (the particle registers and their tag registers) into antisymmetric form by any means that can permute an initially-ordered list of integers; for example, the inverse sorting network method of Ref. [72]. We simply apply this method to the tags while 'dragging along' the particle registers 'for the ride'. The result is the state

$$\frac{1}{\sqrt{P!}} \sum_{i}^{P!} \text{perm}_i\Big((|E'_0\rangle\,|\psi_0\rangle)...(\Big|E'_{P-1}\Big\rangle\,|\psi_{P-1}\rangle)\Big) \tag{5.17}$$

   with the notation (|tag⟩ |particle state⟩).

4. Erase the tag registers through phase estimation: apply a QFT to every tag register, and then apply operations of the form controlled-$\exp(i\mathcal{H}_{\text{synth}}C\pi)$ from each tag register qubit to its main register, where C is an appropriate power of two.

5. Measure out the tag register qubits in the x basis; with high probability, all should be in the $|+\rangle$ state and this outcome is our success criterion.

See Fig. 5.2 for a diagrammatic representation of this approach. If all energies E_i can be exactly represented with t bits, then $E'_i = E_i$ for all i and the method will succeed with certainty. Moreover, even in the in the event that the energies cannot be perfectly represented in this binary form (as will be the case, for example, if two energies differ by an irrational number) then the only consequence is that success criterion in Step 5 will have a reduced probability of occurring. Given that it does occur, the final state remains ideal. For a system of P particles, assuming we can construct $\mathcal{H}_{\text{synth}}$ such that each of the single particle states have a unique, integer valued eigenvalue, then the number qubits in each tag register $t \sim \mathcal{O}(\log P)$. Implementation of the phase estimation tag erasure is massively parallel and can happen independently for each particle – we need to implement QFT and inverse QFT subroutines to each tag register with known gate complexity $\mathcal{O}(2t \log t)$, and accounting for the number of controlled time evolution steps the total complexity is $\mathcal{O}(t + 2t \log t)$. As $\mathcal{H}_{\text{synth}}$ is a Hamiltonian of our choosing, it is conceivable the simulation of $e^{-i\mathcal{H}_{\text{synth}}\pi}$ will not suffer from critical scaling issues.

While not directly tested, I outline a secondary approach for preparing an antisymmetric state that also exploits $\mathcal{H}_{\text{synth}}$. Instead of fully representing all the tag registers, we can consider what happens if we compute and un-compute the energies of the states $|\psi\rangle$ 'on the fly'. We assume that we have available a sorting network formed of swap operations and 'controller qubits' that determine whether each swap operation will occur. Moreover we take it that these controller qubits are already properly primed in the appropriate superposition to create an antisymmetric permutation (as per the method of Berry *et al.* [72]). As we push our initially canonically-ordered state 'backwards' through this network, as per Berry, we must erase each controller qubit as we go: wherever two inputs $|A\rangle\,|B\rangle$ have indeed been swapped to $|B\rangle\,|A\rangle$, we must now flip the controller qubit from $|1\rangle$ to $|0\rangle$.

Ref. [72] describes this erasure for the case of integer labels, noting that if $|A\rangle$ and $|B\rangle$ are any two superpositions of integers, but all integers in $|A\rangle$ are smaller than all those in $|B\rangle$, then the reordering implied by the SWAP can be used to erase the controller. We must instead erase the controller qubit with an analogous

condition based on the energies of our grid-based states. This can be achieved by introducing a set of t ancilla qubits in the state $|+\rangle^{\otimes t}$ and performing a phase estimation process, where the two registers involved in the swap operation are subject to controlled-$e^{i\mathcal{H}_{\text{synth}}C\pi}$ and controlled-$e^{-i\mathcal{H}_{\text{synth}}C\pi}$, respectively. Here, as earlier, C denotes the appropriate power of two for each ancilla qubit. As a result, the ancilla now encodes the QFT of a number that is either positive or negative; by performing an inverse-QFT, that sign is made manifest as the most-significant bit, and can be used to unset the SWAP controller qubit before the entire computation is reversed.

This latter approach has the merit that not all tag registers need to coexist simultaneously; the process can occur sequentially using only t rather than $P\,t$ qubits (albeit incurring a time cost). However the method is more computationally costly, requiring more phase-estimation processes: Multiple uses to compute and un-compute at *each* SWAP in the sorting network. An interesting distinction is that there will be more validating measurements, one confirming each controller qubit is in state $|0\rangle$ after its SWAP process is complete.

Refining the initial state: Ground state preparation

The remarks above concern the initialisation of our grid-based representation into a state that we understand and can describe analytically. Naturally, we may then wish to drive this state into a more accurate representation of e.g. the ground state of the real system prior to exploring its dynamics, assuming non-zero overlap between the initial reference and the target state. While many quantum algorithms have been proposed for this, my collaborator and I demonstrated two techniques which exploit the non-unitary nature of single-ancilla measurements.

- *State Editing* Use of ancilla measurements in phase estimation is a well-established technique for ground state preparation [72, 180, 265].

  My collaborator and I considered the single-ancilla variant which uses IPE described in Section 5.2.2, where we regard the ancilla measurement as a mid-point rather than an end-goal. We emulated an approach which uses this measurement to remove a known state from superposition to 2D hydrogen.

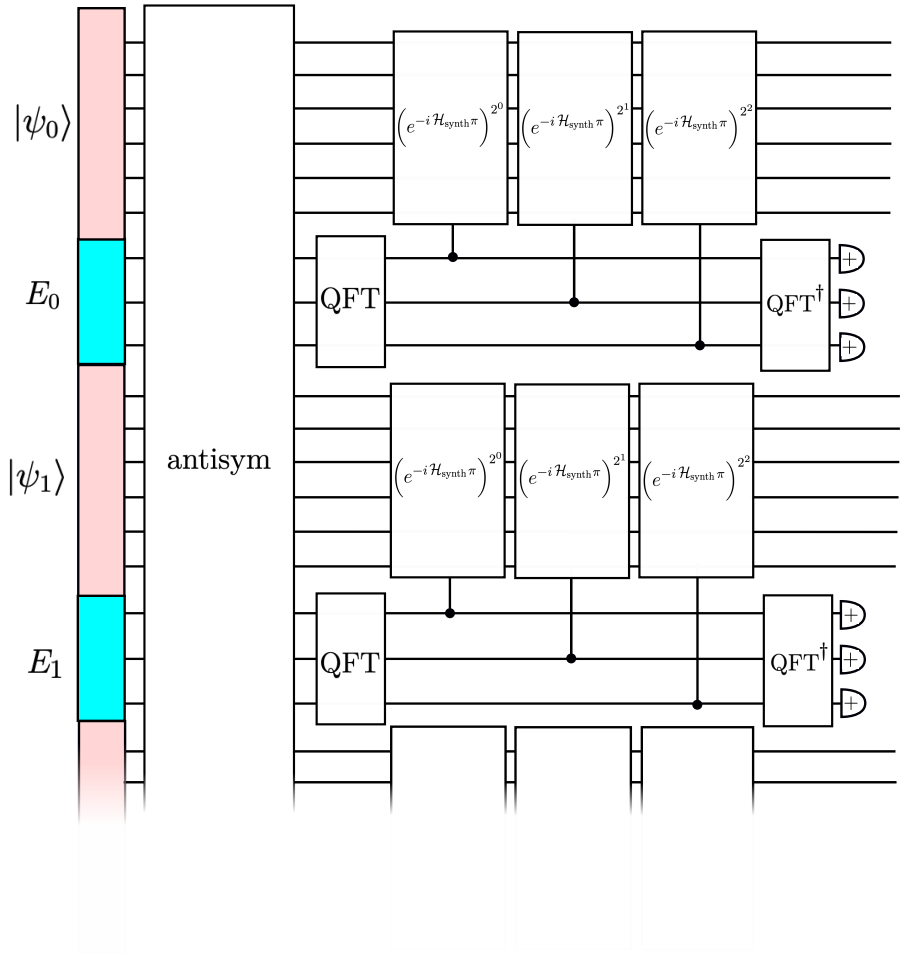


Figure 5.2: Un-computing tag ancilla registers used in the preparation of an antisymmetric state in first quantisation. $|\psi_i\rangle$ represent single particle initial states, and we assume we can construct a fictitious Hamiltonian $\mathcal{H}_{\text{synth}}$ such that $\mathcal{H}_{\text{synth}}|\psi_i\rangle = E_i|\psi_i\rangle$. In this case, the integer eigenvalues of the fictitious Hamiltonian E_i are encoded exactly in binary in the tag registers – we can then use phase estimation to un-compute these registers with a 100% success probability.

Suppose we can prepare an initial state which contains a superposition of eigenstates including our desired state. Assume that we know the energy E_κ of the state we wish to remove, which can be revealed using the method of Section 5.2.2 and classical Fourier analysis. We prepare our ancilla in the $|+\rangle$ state and then perform conditional evolution previously described. To remove state(s) Ψ_κ with eigenvalue E_κ, we require measurement at T_κ, such that $e^{-iE_\kappa T_\kappa} = -1$, thus

$$T_\kappa = N_\kappa \delta t = \frac{\pi}{E_\kappa} \tag{5.18}$$

Now the undesired component has zero probability of yielding a $|+\rangle$ from the

ancilla measurement, so post-selecting on that outcome will entirely eliminate its contribution from the register.

- *Probabilistic Imaginary-Time Evolution (PITE)* The second state preparation method I explored was to perform imaginary-time evolution [54, 266]. In general, it does not require *a priori* knowledge of Hamiltonian eigenvalues; however it does suffer from an exponentially vanishing success probability, as the authors of Refs. [54, 266] noted.

  Suppose again that we have an initial state which contains a superposition of eigenstates $|\Psi\rangle = \sum_n c_n |\Psi_n\rangle$. Imaginary-time evolution applies the non-unitary evolution operator $e^{-\mathcal{H}\tau}$ to the initial state

$$|\Psi(\tau)\rangle = e^{-\mathcal{H}\tau} |\Psi\rangle = \sum_n c_n e^{-E_n\tau} |\Psi_n\rangle \tag{5.19}$$

  and therefore approaches the lowest energy state in the initial superposition exponentially fast at the long τ limit. The PITE approximates, to 1$^{\text{st}}$ order, the Hermitian small imaginary-time evolution operator $m_0 e^{-\mathcal{H}\delta\tau}$ (where $0 < m_0 < 1$ and $m_0 \neq 1/\sqrt{2}$) applied to the state on the main register, by controlling both forward and reverse real-time evolution operators $U(\delta t)$ and $U^\dagger(\delta t)$ with a single ancillary qubit prepared in superposition. Post-selecting on the ancilla measurement outcome of $|0\rangle$ applies the desired evolution on the register. In emulation, I prepared the same $|+\rangle$ ancilla state as per the IPE, and started from the circuit in Figure 5.3, noting that it achieves the exact same non-unitary operation proposed in Ref. [54] but with fewer gates.

  Like state-editing by phase estimation, this approach is also probabilistic and preparing an initial state with large overlap with the ground state is key to maximising the rate of success. In this case however, the probability that every sequential ancilla measurement yields the desired outcome decays exponentially with the number of measurements. The reader can refer to Ref. [54], which introduced the technique and applied it to wavepackets in parabolic potentials, for details of this drawback and possible ways to address

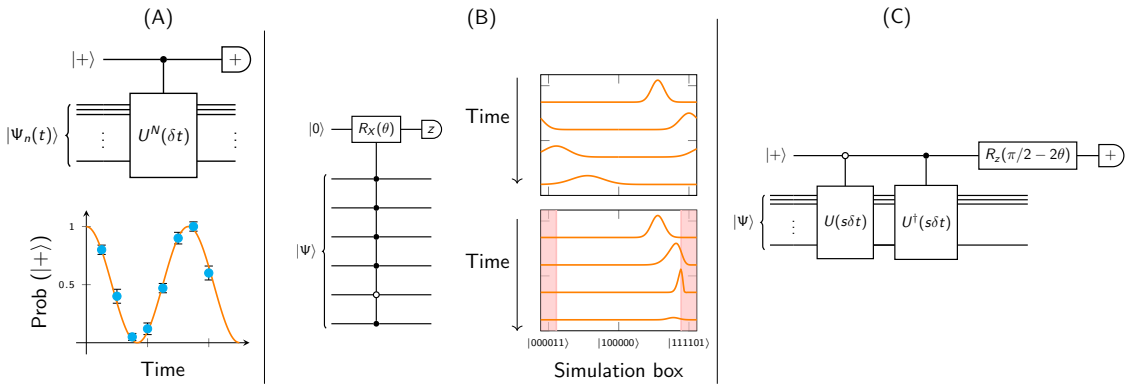


Figure 5.3: Three quantum circuit techniques used for real-space quantum simulation. **(A)** The single-ancilla IPE method. The global phase is encoded in the probability of measuring an entangled ancilla qubit, which controls the application of N split-operator cycles, in the $|+\rangle$ state. To obtain this probability one must repeat the propagation and measurement at each point in time where one wishes to sample the signal. **(B)** Attenuation of a wavepacket on a qubit register. The top right panel depicts the dispersion of a Gaussian wavepacket across the periodic boundary. Addition of complex absorbing region in pink (bottom right panel) attenuates the scattered wavepacket by reducing its norm. In the illustrated case, the process is not quite perfect: There is some reflection caused by the attenuation being too severe. The left panel is a circuit that performs probabilistic wavepacket attenuation at a select pixel. The pixel $|111101\rangle$ is selected, corresponding to the attenuating region on the right hand side of the simulation box in the figure. **(C)** Preparing ground states using the PITE circuit. The filled circles indicate 'control by $|1\rangle$' and the open circles 'control by $|0\rangle$'. Post-selecting on the $|+\rangle$ outcome yields a state with an imaginary time evolution step applied (to 1^{st} order). Given choice of an appropriate parameter m_0, the time rescaling factor $s = m_0/\sqrt{1-m_0^2}$ and the rotation $\theta = \kappa \arccos\left[m_0 + \sqrt{1-m_0^2}/\sqrt{2}\right]$, where $\kappa = \text{sgn}(m_0 - 1/\sqrt{2})$.

it. I highlight that, while the current method may be impractical on a real quantum computer, imaginary-time evolution approaches in general (see Chapter 4) can still be useful for state driving.

5.2.5 Augmented Split-Operator

I now describe an approach that addresses the simulation of states that may have large amplitudes at the Coulomb singularity, instead of using a 'brute force' increase in spatial and temporal resolution. The approach, called augmented split-operator (ASO), involves introducing an additional step in the SO-QFT cycle.

Following the implementation of the exponential operators in a SO-QFT time step, and while still in the real-space representation, we apply a *core stabilization*

process. This is intended to 'fix' the most severe deviations of the applied unitary from the ideal unitary $e^{-i\mathcal{H}_{\rm tot}\delta t}$. We emphasise that the fix is at the operator level, and is agnostic to the specific state that is being modelled. However, the motivation is to make the dynamics of core-peaked states more accurate. We aim to reduce the overall time cost of a simulation with a given accuracy, without increasing the qubit count.

We describe the process for a simulation of a single particle and comment on the generalisation to N particle systems further below. The following is performed on a classical computer, as a one-off preparation for our simulation. We make a fixed choice of δt, and we must repeat the analysis for each value of δt we wish to use.

1. Compute the ideal iterate $U_{\rm ideal} = e^{-i\mathcal{H}_{\rm tot}\delta t}$ and the actual iterate defined by Eq. (5.2) i.e.
$$U_{\rm SO}(\delta t) = e^{-iD_{\rm int}\delta t}\left(U_{\rm QFT}^{\dagger}\; e^{-iD_{\rm kin}\delta t}\; U_{\rm QFT}\right).$$
These objects are computed in the full basis of pixel functions for the particle (and are therefore far from diagonal).

2. Examine $U_{\rm repair}$ which is defined by
$$U_{\rm ideal} = U_{\rm repair}U_{\rm SO}(\delta t)$$
and we will find that $U_{\rm repair}$ is close to the identity except for elements corresponding to pixel functions that are near to the Coulomb singularity. Therefore we select a set of Q such pixels, and derive a small $Q \times Q$ unitary $U_{\rm core}$ that closely matches $U_{\rm repair}$ in that subspace.

3. Define augmentation $U_{\rm aug}$ as the identity operator except for the Q-state subspace, where it corresponds to $U_{\rm core}$.

4. Find a circuit $C_{\rm aug}$ that can implement $U_{\rm aug}$ to a good approximation. This circuit is not merely a set of phase operations since $U_{\rm aug}$ is not diagonal in the Q-state subspace; $C_{\rm aug}$ will cause a flow between the amplitudes of those states.

This classical analysis is non-trivial but tractable since the unitaries involved are 'only' of size $2^{n_r D}$ for a D-dimensional problem where n_r is the number of qubits per sub-register. Moreover there are helpful symmetries in the analysis, and notwithstanding the description in Step 1 we need not compute U_{ideal} and U_{SO} entirely but only in the subspace where we are likely to identify the Q core pixels. With these steps completed, we are in a position to perform simulations with the augmented SO iterate:

$$\begin{aligned} U_{\text{ASO}}(\delta t) &\equiv U_{\text{aug}}\, U_{\text{SO}}(\delta t) \\ &= U_{\text{aug}}\, e^{-iD_{\text{int}}\delta t} \left(U_{\text{QFT}}^{\dagger}\; e^{-iD_{\text{kin}}\delta t}\; U_{\text{QFT}} \right). \end{aligned} \tag{5.20}$$

The generalisation of the ASO method to the P-particle case is straightforward except for a caveat. The main time cost of the SO-QFT method lies in the implementation of the $e^{-i\mathcal{H}_{\text{int}}\delta t}$ part, because of the $P(P-1)/2$ particle pairings involved. The ASO method can be efficiently implemented by applying an U_{aug} augmentation after each such pairing, subsequent to implementing the $e^{-iC\delta t/|\mathbf{r}|}$ and while the registers are still representing relative coordinates $x_i - x_j$, $y_i - y_j$ etc. Because U_{aug} is not diagonal, ensuring the proper evolution requires implementing the full orthogonal transformation to relative and central coordinates:

$$\{x_i, x_j\} \Rightarrow \{x_i - x_j, x_i + x_j\} \text{ rather than } \{x_i - x_j, x_i\}$$

and similarly for y and z. Fortunately the additional arithmetic step that this implies, is trivial.

5.3 Results

Before I detail the results of this investigation, I will layout the hardware and software configurations used. All of my collaborators' and my numerical results were ultimately obtained through the Quantum Exact Simulation Toolkit (QuEST) [267], though interfaced through the `Python` pyQuEST [268] and `Mathematica` QuESTlink [179] packages. This enabled multithreaded and Graphical Processing Unit (GPU)

simulation of the modestly sized systems, such as a 25-qubit 2D scattering simulation, directly within QuESTlink.

Results for the larger scale systems considered were generated with pyQuEST, which enabled relatively low-level access to QuEST simulation primitives through a high-level `Python` interface. Importantly, pyQuEST inherits the capacity to run distributed tasks, where the numerical representation of a quantum state is partitioned between compute nodes cooperating over a network. This allowed both the representation of states too large to fit into the memory of any single compute node, and their concurrent simulation – each multicore node is further able to parallelise its local simulation tasks through multithreading. I used distributed pyQuEST for some of my emulations; the largest simulation performed in relation to this Chapter was performed by my collaborator at 36 qubits, employing up to 32 compute nodes of the Oxford Advanced Research Computing (ARC) facility [269]. Each node contains 48 CPU cores, and took roughly 52 hours ($\approx 50\,000$ core hours) to process its 64 GiB partition of the full 1 TiB quantum state-vector.

Throughout this Chapter, I frequently use the 2D hydrogenic system described by the Hamiltonian

$$\mathcal{H}_{\text{tot}} = -\frac{\hbar^2}{2m_e}\nabla_{\mathbf{r}} - \frac{e^2}{4\pi\epsilon_0|\mathbf{R}-\mathbf{r}|}, \tag{5.21}$$

where the atomic nucleus is modelled as a classical discretised Coulomb potential, clamped with the origin between two pixels. Analytic solutions to Eq. (5.21) have been reported in Refs. [270, 271]. I used an equation from the former,

$$\begin{aligned}\Psi_{n,m}(r,\theta) =& \sqrt{\frac{q_0^3(n-|m|)!}{\pi(n+|m|)!}} \times (2q_0 r)^{|m|} \times e^{-q_0 r} \\ & \times L_{n-|m|}^{2|m|}(2q_0 r) \times e^{im\theta},\end{aligned} \tag{5.22}$$

where $q_0 = \frac{1}{n+1/2}$ and L are the generalised Laguerre polynomials. The quantum numbers $n = 0, 1, 2, \ldots$, and there are $(2n+1)$ values of m. The energy eigenvalues are

$$E_n = -\frac{1}{2\left(n+\frac{1}{2}\right)^2}. \tag{5.23}$$

Two of the eigenstates used in this work are shown in Fig. 5.4. In a two-particle, 36 qubit simulation, my collaborator also use states related to the well-known 3D hydrogenic eigenstates

$$\begin{aligned}\Psi_{n,l,m}(r,\theta,\phi) =& \sqrt{\left(\frac{2Z}{n}\right)^3 \frac{(n-l-1)!}{2n(n+l)!}} \times \left(\frac{2Zr}{n}\right)^l \\ & \times e^{-Zr/n} \times L_{n-l-1}^{2l+1}\left(\frac{2Zr}{n}\right) \\ & \times Y_l^m(\theta,\phi),\end{aligned} \tag{5.24}$$

where $Y_l^m(\theta,\phi)$ is a spherical harmonic and Z is the central nuclear charge. Finally, my collaborators and I also use Gaussian wavepackets of the form:

$$\Psi(x) = e^{\mathrm{Im}(\gamma)} \left(\frac{2\,\mathrm{Re}(\alpha)}{\pi}\right)^{1/4} e^{-\alpha(x-x_c)^2 + ip_c(x-x_c)+i\gamma} \tag{5.25}$$

where x_c, p_c, α and γ are continuous parameters. Simulation results are reported in Hartree atomic units, where the reduced Planck constant, electron mass, elementary charge and Bohr radius are treated to be unity $\hbar = m_e = e = a_0 = 1$.

5.3.1 Trotter Evolution and the Relationship between Spatial and Temporal Resolution

The model's spatial resolution δr^{-1} and the temporal resolution of the dynamics δt^{-1} are crucial in determining accuracy of the grid-based simulation; the qubit count per spatial dimension, n_r, is logarithmically related to the former (see Eq. (2.29) in the Methods Section of this Chapter). To explore these requirements independent of how the states are prepared, I started with propagation of eigenstates of the 2D hydrogen with different choices of δr^{-1} and δt^{-1} in exact simulation.

I loaded the discretised analytic ground and excited states $\Psi_{0,0}$ and $\Psi_{1,1}$ into emulated qubit registers with an increasing n_r, then performed time-evolution experiments using the 1$^{\mathrm{st}}$-order SO-QFT. The eigenstates were propagated for a total duration of 1.5 atomic time units at different time step resolutions. Fig. 5.4 summarises these results.

In the ideal case, the eigenstates would remain static up to a global phase. Thus the final absolute value of the autocorrelation for each propagation sequence,

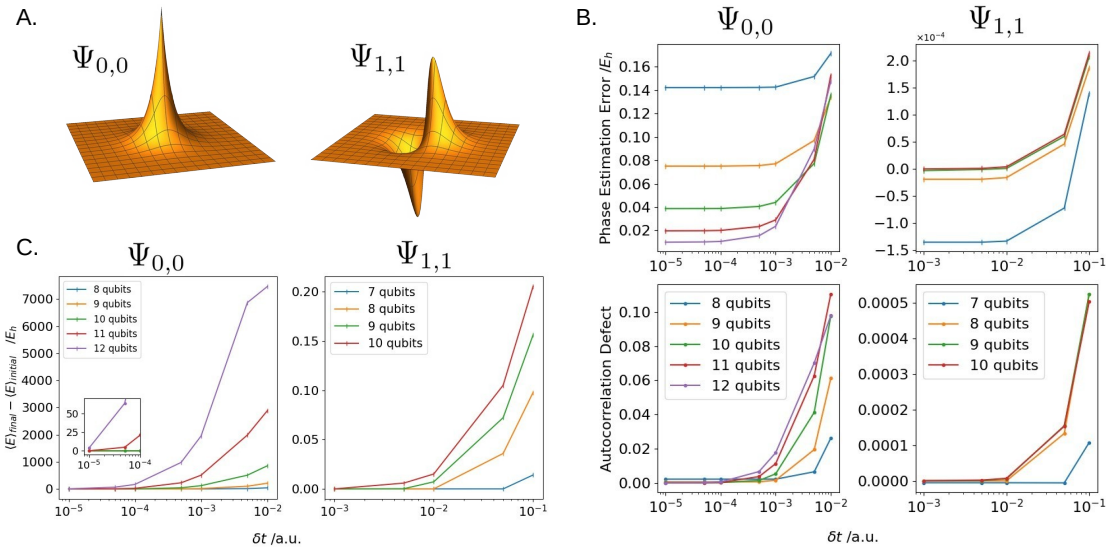


Figure 5.4: Relationship between spatial and temporal resolution in a SO-QFT simulation. **(A)** The ground $\Psi_{0,0}$ state (left) and a first excited $\Psi_{1,1}$ state (right) of two-dimensional hydrogen. The plots here do not reflect the choice of simulation box size and are not to scale. **(B)** Top panels represent difference between the energy from phase estimation and the analytic energy of 2D hydrogen. Bottom panels capture deviation of the simulation fidelity at the end of the propagation. In this series of experiments we initialised the ground state $\Psi_{0,0}$ centred in a simulation box with $L = 10$ a.u., such that the origin of the Coulomb singularity lies half way between two central grid points. Each sub-register has a budget of $8 \leq n_r \leq 12$ qubits to store the wavefunction, corresponding to spatial resolutions of $0.039 \geq \delta r \geq 0.002$ a.u.. I also initialised the $\Psi_{1,1}$ excited state in a simulation box with sides of length 40 a.u., with budgets of $7 \leq n_r \leq 10$ qubits per sub-register and corresponding resolutions of $0.313 \geq \delta r \geq 0.039$ a.u.. These two states, represented at different spatial resolutions, are all time propagated using the first-order SO for 1.5 atomic time units. I used time steps between $0.00001 \leq \delta t \leq 0.01$ (150000 to 150 SO steps) for the $\Psi_{0,0}$ state, and between $0.001 \leq \delta t \leq 0.1$ (1500 to 15 SO steps) for the $\Psi_{1,1}$ state. **(C)** The difference between the final and initial energy expectation value, measured by direct sampling of the state, of the ground state (left panels) and the first excited state (right panels), propagated at different spatial and time resolutions. The left inset plot zooms in on the energy error at high temporal resolutions for the ground state. Figure reproduced from [4] – I was the original creator.

specifically the deviation from unity, is a suitable fidelity metric for the accuracy of the simulation. The lower pair of plots in Fig. 5.4(B) displays this metric for $\Psi_{0,0}$ (left) and $\Psi_{1,1}$ (right). When a higher spatial resolution is used, correspondingly finer time steps are needed to achieve a high fidelity. This is true of both cases but the relationship is more dramatic for $\Psi_{0,0}$, given that its amplitude is peaked at the origin of the Coulomb singularity.

The upper panels in Fig. 5.4(B) plots the deviation of the energy extracted

from an emulated IPE experiment from the exact analytic result. One observes qualitatively the same behaviour as for the autocorrelation; there is a strong divergence when the δt is insufficient, and that depends on both the δr and the nature of the modelled state. For the ground $\Psi_{0,0}$ state, the error in the most accurate energy prediction (attained with the smallest time step propagation) halves when we increment n_r. While deviation of the energy inferred from phase estimation versus the exact value falls exponentially with the number of qubits n_r, chemical accuracy is not yet reached with $n_r = 12$.

In additional simulations, I found that by quartering the size of the simulation box and increasing n_r by 1 (effectively an additional 4 qubits), then propagating an initial state at $\delta t = 10^{-5}$ a.u. for 4 a.u., an error of 0.652 mE_h from the exact energy was achieved. In this case 400,000 SO-QFT cycles were used to cover only ≈ 0.1 femtoseconds of physical process. While these time requirements for accurate simulation of core-peaked states like $\Psi_{0,0}$ may seem daunting, it is worth reiterating that the approach taken here is not optimised, and only the 1$^{\text{st}}$-order Trotter sequence. Moreover, as shown in Section 5.3.5, our Augmented Split-Operator approach can obtain accurate phase estimation with far fewer qubits and lower time resolution than the 'brute force' method reported here. It is also noteworthy that accurate modelling of the $\Psi_{1,1}$ state is remarkably tolerant of low resolutions.

Cautionary Tale – a 'Bad' Energy Measurement

I also examined a sampling-based method of estimating the system's energy; I am referring to measuring the energy expectation of a state as given by Eq. (5.9) i.e. $\langle E \rangle = \langle \mathcal{H}_{\text{kin}} \rangle + \langle \mathcal{H}_{\text{int}} \rangle$, and supposing that this would be obtained from our quantum computer as follows: Generate the desired state at time t, measure in the k-space representation, and repeat this many times in order to estimate the first term; then apply the same process but measuring in the real-space representation to estimate the second term.

This method, like any quantum algorithm of which the output is a quantum state, can become inefficient because of the sampling cost, and is therefore already

unattractive compared to phase estimation, but more problematically it is very sensitive to the resolution parameters. As shown in Fig. 5.4, the error between the initial and final expected energy grows exponentially with the time resolution δt; for the ground state $\Psi_{0,0}$, the energy difference was more than 7000 E_h accumulated over less than 40 attoseconds of simulated time in the worst offending case. In the $\Psi_{1,1}$ simulation, the non-conservation of expected energy is also apparent when the temporal resolution does not match the spatial resolution, but is more contained relative to the ground state (in the worst case, it goes up to 0.20 E_h). It is evident that the core-peaked nature of the $\Psi_{0,0}$ state is a key issue.

Upon further investigation, I found that it is the kinetic energy term $\langle \mathcal{H}_{\text{kin}} \rangle$ that exhibits this diverging behaviour. This is because the imbalance between the extreme potential and extreme kinetic energy near the Coulomb origin, inevitable in our discrete grid representation, can allow a small amount of the amplitude to diffuse towards high frequency states in the plane wave representation. The extent of this diffusion may be small relative to the initial state so that the fidelity of the state remains high (see Fig. 5.4) and thus phase estimation methods can perform well. However, simply sampling $|k\rangle$ and squaring it to estimate $\langle \mathcal{H}_{\text{kin}} \rangle$ gives direct weight to this error that actually worsens as we improve the fineness of our resolution. Reducing the separation of spatial pixels δr can reduce the leakage of amplitude (increasing the state's fidelity) but the maximum momentum, thus kinetic energy, that the basis can represent goes as $1/(\delta r)^2$. Amplitude leakage declines less rapidly than the rate at which the energy of the leaked-states increases – thus the problem worsens. One must use extremely high time resolution to ameliorate the effect.

A higher order Trotter formula would also presumably help, in the sense that a more modest time resolution could control the leakage. Nevertheless, I anticipate that this approach to estimating energy will always be inferior to phase estimation because of the sensitivity to inevitable imperfections in the model. While it is expected to converge to the correct expected energy once spatial and temporal resolution are sufficiently high, I found it to be profoundly inaccurate at resolutions where the IPE can already provide reliable, well-converged results.

Placement of Coulomb singularity

I want to briefly report additional emulations that I performed which suggested that placement of the origin of a classically-treated nucleus between grid points does not significantly affect the dynamics (summarised in Table 5.1). When the $\Psi_{1,1}$ state is time-propagated with the Coulomb potential centred at different fractions between two neighbouring grid points, the state visibly oscillates where grid points lie asymmetrically around the Coulomb origin. However, the magnitude of the loss in fidelity is limited. Most importantly, the energy from IPE is minimally influenced by the skew. Discrepancies in the placement of grid points relative to the system can, if needed, be alleviated with a brute force increase in the spatial resolution and correspondingly the time resolution, but its influence on observables related to state fidelity can be negligible.

	0.5	0.4	0.3	0.2	0.1
Max loss	7×10^{-7}	0.00085	0.00021	0.0042	0.0022
ϵ ($10^{-5}E_h$)	-1.84	-1.11	0.29	-1.66	5.98

Table 5.1: Placement of the Coulomb nuclear potential between grid points and its influence on simulation fidelity and phase estimation. Table reproduced from [4] – I was the original creator.

5.3.2 Quantum Dynamics Experiments

I describe proof-of-concept, real-space grid simulations relevant to two of the scenarios introduced in the beginning of this Chapter: ionisation by strong external field, and electron-electron scattering. The corresponding data are shown in Fig. 5.5. In both these studies, my collaborator and I employed the complex absorbing potential approach detailed in Section 5.2.3.

Scenario I: Electric field ionisation

In panel (A) of Fig. 5.5 I show the performance of a 19 qubit emulated quantum computer, modelling a single 2D particle ($9 + 9$ qubits represent the state, one

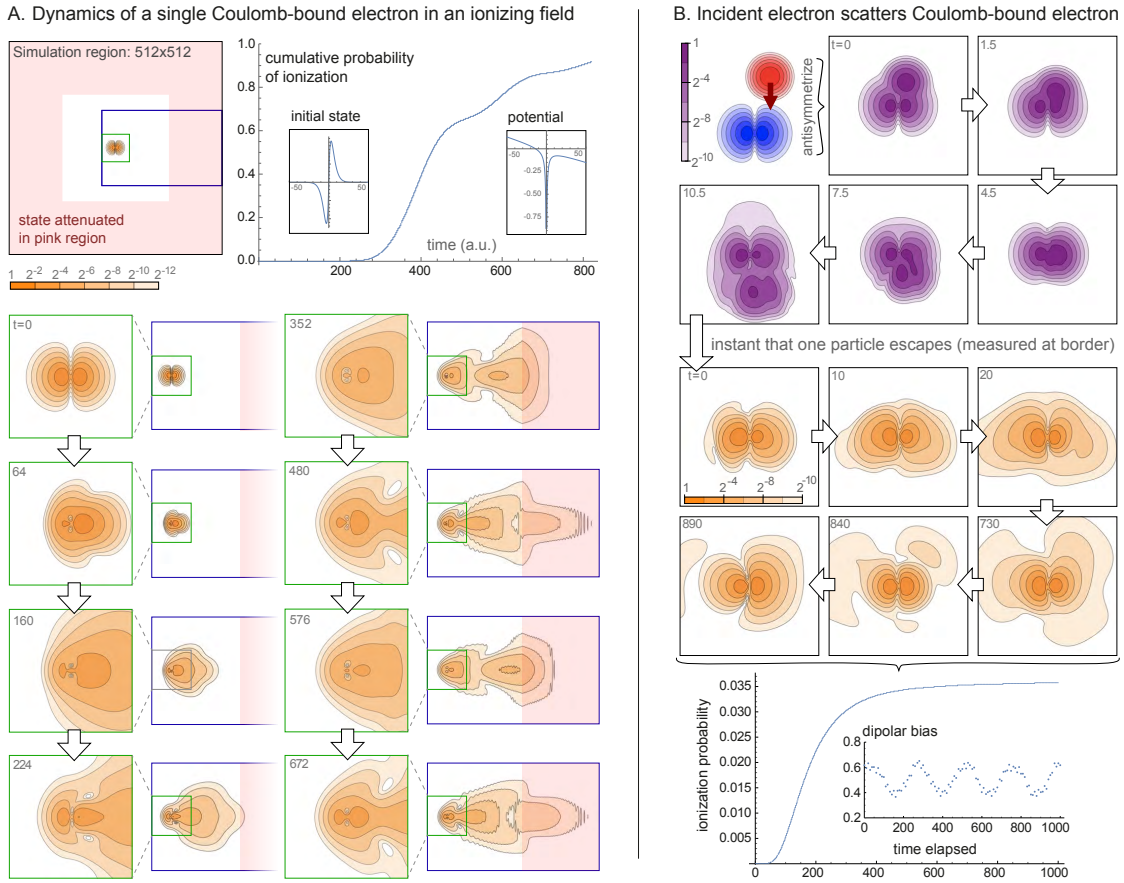


Figure 5.5: Real-space simulation of single and double-electron dynamics. **(A)** Figures showing a 2D simulation of a single electron ionised by a strong electric field, executed on a $1 + 2 \times 9$-qubit quantum computer. The initial state is a low-lying bound state of 2D hydrogen ($\Psi_{1,1} + \Psi_{1,-1}$). The panel in the upper left shows the initial state as a contour plot; it occupies a small region in the centre of the entire simulation box (black boundary). The regions marked with green and blue borders are the zones into which we zoom in the lower panels, which are a series of contour plots with the time index shown in the corner of each plot. The pink region corresponds to attenuation (a complex potential). The total probability of the particle having been found in the outer, pink region is shown by the upper right plot. **(B)** 2D simulation of two-particle scattering using a 25-qubit emulated quantum computer. The initial state is the antisymmetrised version of the following: an electron (blue) in a low-lying bound state of 2D Hydrogen while a second electron (red) is in a Gaussian state and moving in the negative y-direction (downward in the plots). The simulation reveals the scattering event (purple plots) until we deem one particle to have exited. A second phase of simulation shows the evolution of the state of the remaining electron (orange) which has a probability of about 3.5% of ultimately ionising (escaping the nuclear potential). Interestingly, in the event that the remaining electron does not ionise, it nevertheless remains in an state where it oscillates back and forth; the inset plot shows the quantity $\langle x_L \rangle >$ as defined in the text. Figure reproduced from [4] – created by Simon Benjamin and me.

qubit is used for the weak measurements). The modelled system at $t = 0$ is a state within the first excited manifold of 2D hydrogen, specifically

$$\frac{1}{\sqrt{2}}(\Psi_{1,1} + \Psi_{1,-1}).$$

An additional E_x term in the Hamiltonian corresponds to a strong, static electric field applied in the horizontal direction; combination of the Coulomb potential and the electric field is shown the Figure inset. Because of the electric component, the initial state is no longer an eigenstate, and the simulation can determine whether the electron will indeed be removed from the nucleus.

The initial state occupies only a small central region in the simulation box. The contour plots in the main part of the figure show the evolution, both in the centre part of the simulation box (the box bordered in green) and a larger region encompassing the centre and the region to the right (the box bordered in blue). The pink region, constituting the outermost 50% of the simulation box in both the x and y directions, is the region that is monitored by weak measurements.

The figure also shows the cumulative probability that the particle will indeed have 'escaped' i.e. that it will have been measured to be in the pink region. The curve ultimately approaches unity, indicating that the particle will eventually ionise with certainty. Interestingly, we observe oscillations in this curve, which we can account for by examining the contour plots of particle density shown below.

Note that the contour plots show the particle's probability density post-selected on it not having yet escaped; therefore the normalisation is unity in each case. Focusing on the green boxes, those zoomed in close to the nucleus, we observe that the part of the state that remains close to the nuclear core actually oscillates in a dipole-like fashion. Note that the sequence of green panels, labelled 352 to 672, exhibit this: panel 352 is similar to 576 while 480 is similar to 672. Moreover, examining the corresponding blue regions we observe waves of density propagating away from the nucleus, synchronised with the dipole oscillation; whenever the oscillation favours the 'down stream' right side, there is an enhanced probability that the particle will escape; in due course this leads to a fluctuation in the probability

of observing the particle in the pink, attenuating region. It is interesting to note that this fluctuation probability is reminiscent of the bond-breaking of sodium iodide observed with femto-second pulsed lasers [272], an experiment recognised by the Nobel Prize in Chemistry in 1999.

On a real quantum processor, the contour plots of Fig. 5.5 can be obtained by repeated sampling, simply by measuring the state of the particle's register at a given time t. Obtaining these outputs through 'brute force' sampling would obviously represent a multiplicative cost depending on the accuracy with which we require the plots. The plot of the cumulative probability can also be obtained by repeatedly executing the simulation; however, it only requires measurement on a single ancilla, and directly produces particle location data that can be used to determine, for example, rates in a chemical reaction dynamics simulation. We argue therefore that this approach is more useful for studying real chemistry problems on early fault-tolerant quantum computers than direct state sampling.

Scenario II: Two-Particle Scattering

Panel (B) of Fig. 5.5 shows the results of a two-particle scattering simulation. There are two simulation phases. In the initial phase, we have the both interacting particles present: an electron initially in a bound state of 2D hydrogen corresponding to $(\Psi_{1,1} + \Psi_{1,-1})/\sqrt{2}$, and an incident electron in a Gaussian state (but with the total state properly antisymmetrised). This simulation runs until one of the particles is measured to have 'escaped' our simulation box. We then proceed to a second phase of simulation where we study the dynamics of the surviving particle. We find that it has a small ($\sim 3.5\%$) probability of ionising due to the perturbation of the prior 'impact' with the incident particle.

In the first phase we use a $25 = 1 + 4 \times 6$ qubit emulated quantum processor, where the x and y coordinates of each particle are represented with $2^6 = 64$ states. As in the case of the electric field ionisation described, we monitor with weak measurement a set of spatial pixels near the boundary of the simulation box. In the initial phase of this two-particle simulation the width of that region is 25% rather

than the 50% width used in the electric field case. The contour plot panels in the figure show the central, non-attenuated region. The state is antisymmetrised so that the probability density plots do not distinguish one particle from the other, but informally we can say that the incident particle interacts with the bound particle before passing on, away from the nucleus. In the case shown in the figure, we deem that a particle has been detected in the attenuation region exactly at the point when the cumulative probability of detection reaches 50% (this is an arbitrary choice; in a real quantum processor the user would of course be unable to specify this). This event occurs shortly after the last of the panels in the upper part of the figure.

The simulation to that point would not teach us much about the nature of the scattering event. We could in principle measure, e.g, any deflection in the trajectory of the incident particle, and we can confirm (from longer simulations) that there is near-100% probability of a particle exiting the simulation region; the incident particle does not become bound. However, it is more interesting to study the subsequent behaviour of the remaining particle. Because this particle is represented by a register in the quantum processor that has not 'collapsed', we can simply continue to simulate its evolution; the component U_V in our SO-QFT cycle, corresponding to particle-particle interaction, will no longer be applied. We can now choose to vary other parameters such as δt, anticipating that further dynamics are on a slower timescale since the high energy particle has exited the simulation. Moreover, we can reallocate some of the qubits that were previously used to model the now-exited particle, re-purposing them to model a larger simulation box. Given that the outer regions of such a box have zero amplitude associated with them at the moment the first particle exits, there is no difficulty in simply introducing those qubits. (In the case that we employ two's complement, so that the coordinate $(0, 0)$ is in the centre of the simulation box, we simply append the new qubits to the high-order end of each sub-register and perform a CNOT gate on each new qubit controlled by the prior highest-order qubit.) This was indeed performed in the simulation shown in the right panel of Fig. 5.5, and the second phase of the simulation uses $1 + 2 \times 8$ qubits to provide a much larger attenuation region.

In this second phase we observe that the remaining particle has been perturbed by the passage of the incident particle: its distribution at $t = 0$ (now measuring time relative to the exit of the other particle) is noticeably more irregular than the simple symmetric initial form. The probability distribution is lopsided, favouring the left side. Over the remaining period of simulation, the particle exhibits mild dipolar oscillation between a left-favoured and a right-favoured distribution. Defining $P_{\text{left}}(t)$ as the probability that the particle would be found left-of-centre, the oscillation shown in the inset to the time-series plot in the right panel of Fig. 5.5 is observed. Moreover, as the particle oscillates it sheds probability – i.e. there is a finite probability that the particle will escape the simulation region. In contrast to the electric field simulation, this probability is shed symmetrically (both left and right) and it converges to a small cumulative probability of about 3.5% (see main plot).

5.3.3 Initial State Preparation

State Editing

Earlier plots (Fig. 5.4) have presented the results of IPE for energy estimation. Here I demonstrate preparation of an initial wavepacket state on a set of qubits using IPE (Section 5.2.4). Fig. 5.6 shows the results of a simulation where the initial state (shown lower left and inset (i)) is a superposition of two eigenstates of 2D hydrogen:

$$\frac{1}{\sqrt{2}}\left(\Psi_{1,1} + \Psi_{2,2}\right).$$

The state's amplitude is not symmetric about the nucleus, and when SO-QFT cycles are applied a rotation of the state can be observed, due to the different rates at which the two superposed eigenstates acquire phase. At the time marked (ii), the total phase acquired by more tightly bound state $\Psi_{1,1}$ (having a more negative energy) has reached π; and thus, if this state were the sole one present, a control ancilla initially in state $|+\rangle$ would now certainly be in state $|-\rangle$. The ancilla was indeed measured at this point, but I post-selected on seeing the outcome $|+\rangle$. The probability of the desired outcome depends on the probability associated with the target $\Psi_{2,2}$ within the initial superposition (which was 0.5) and the probability

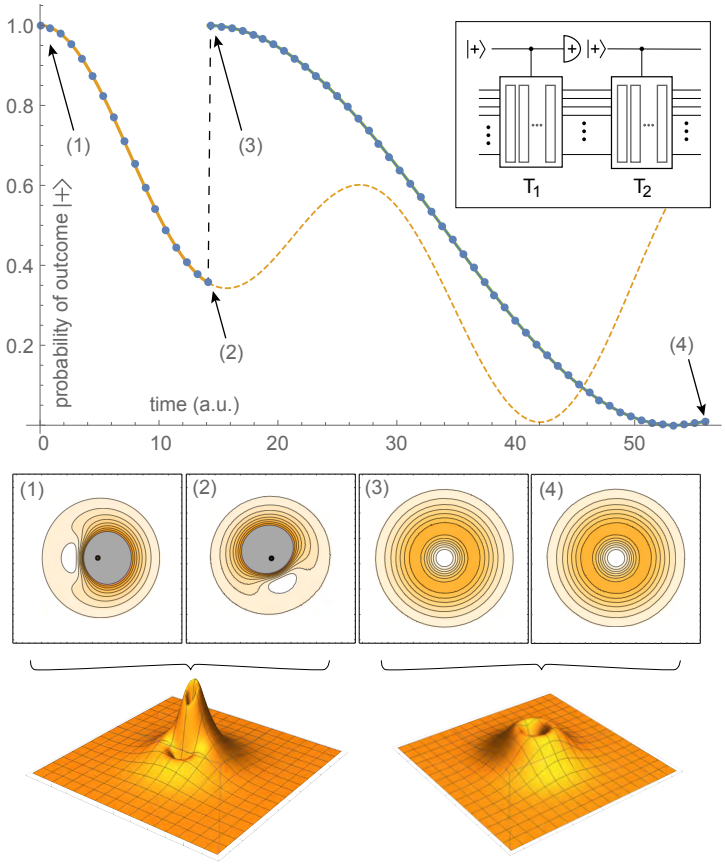


Figure 5.6: Demonstrating state editing by ancilla phase measurement. I emulated a $1 + 2 \times 8$-qubit quantum computer. The initial state was a superposition of non-degenerate eigenstates of 2D hydrogen. The representation used 8 qubits for each of the two dimensions, with a total simulation box width of $L = 56$ a.u. The state evolves, conditional on a controlling ancilla, for time T_1 chosen such that $T_1 E_1 = \pi$; in this period the conditional evolution is a rotation of the state due to the accumulated phase difference between states (lower left plots). At T_1 the controlling ancilla is measured in the $|+\rangle$ as shown in the upper right inset. This projects the state into the $\Psi_{2,2}$ state. Figure reproduced from Ref. [4] – created by Simon Benjamin and me.

that this state, had it been prepared alone, would yield a $|+\rangle$ outcome at this point. The latter is 0.713 in the present case. In a scenario where the states to be distinguished are closer in energy, it may be optimal to simulate for several complete cycles of the undesired state before measuring.

In my emulation, I assumed that the desired $|+\rangle$ state is obtained and time propagation continued. The new evolution of the ancilla state (green curve in the figure) is exactly that of the pure $\Psi_{2,2}$ state. The contour plots of the simulated state (insets (iii) and (iv)) confirm that a pure state had been prepared.

Fidelity with respect to $\Psi_{2,2}$ was essentially identical to an initial state prepared directly in that state.

This is a demonstration of the practically-useful capability to take an initial state that is not fully understood, and remove from it the components corresponding to states with energies that *are* understood. More generally, one could use Fourier analysis to identify the components in the plot of Prob($|+\rangle$) for the full state, and then use the post-selection method to stochastically isolate given components.

Probabilistic Imaginary-Time Evolution

I compared the state editing approach with an alternative that instead approximates non-unitary imaginary time evolution (Section 5.3.3). As before, I modelled an attractive nucleus centred in a square simulation box. Instead of starting from an explicitly defined superposition of eigenstates as in the previous example, I initialised a Gaussian wavepacket centred about the origin of the Coulomb potential. The initial Gaussian wavepacket can be expanded in the eigenstate basis, which I assume to have a large component corresponding to the ground state of the Hamiltonian.

I then propagated the state using PITE for 300,000 steps, using $m_0 = 0.9$ and $\delta t = 2 \times 10^{-5}$ (note the actual imaginary-time $\delta\tau$ rescales δt). Fig. 5.7 shows the state evolving under the approximate imaginary-time evolution. The main plot shows the overlap of the state with analytic eigenstates. Only eigenstates peaked at the origin (the equivalent of s-states in 3D) have large overlap with the state throughout the propagation, with the $n = 1$ excited state contributing more to the initial Gaussian wavepacket than the ground state. In the long time limit, the ground state overlap approaches unity, whereas its overlap with higher energy states decays. The overlap signal does not go exactly to 1, and nor do the contributions of higher energy states go exactly to 0; this is because the state prepared is the ground state of the pixelated model Hamiltonian, which nonetheless has a high overlap with the true analytic ground state $\Psi_{0,0}$ digitised to the same spatial resolution. This disparity should vanish with higher spatial resolution.

The probability distribution of the state, plotted at the bottom, also shows that the broad initial wavepacket contracting to a sharp state peaked at the origin. Visually it would appear that at a short evolved time, the state already resembles the ground state with a singular central peak. However, superimposing the sampled distributions (top right panel), the subsequent long time evolution appears necessary to render increased sharpness.

I further assessed the state prepared from PITE by subjecting it to real-time SO-QFT propagation for 4 a.u., and estimate its energy via IPE. The fidelity does not drop below 4×10^{-8}, and the estimated phase agrees with the converged energy at this spatial resolution reported in the previous section, further confirming that the PITE converges to the ground state of the model at this resolution.

This scenario however demonstrates clearly the main drawback of PITE: At every measurement, there is a substantial probability of measuring the undesired outcome; this is then a failure of the procedure. In the case described here that probability is about 0.33; this means the cumulative success probability falls below 10^{-4} after only about 23 measurement steps. The demonstration here, with 300,000 steps, would therefore have (essentially) zero success probability on a real quantum computer. The method may be useful in 'quantum-inspired' classical algorithms given its attractive feature of not requiring *a priori* knowledge of the states. Moreover the authors of Ref. [54] suggest that amplitude amplification methods might address the issue of vanishing success probability.

5.3.4 Antisymmetrisation of the initial state

In Section 5.2.4, I described an approach my collaborator and I created for the preparation of an initial state in a first-quantised grid-based simulator with proper antisymmetrisation of the electrons. We explored it in emulation on inputs of up to $P = 5$ particles.

I used 30 emulated qubits: $t = 3$ tag qubits for each particle, and a further 3 qubits to represent each corresponding ψ state. My collaborator and I considered several choices of spectrum E_i and proceeded in each case as follows. We prepared

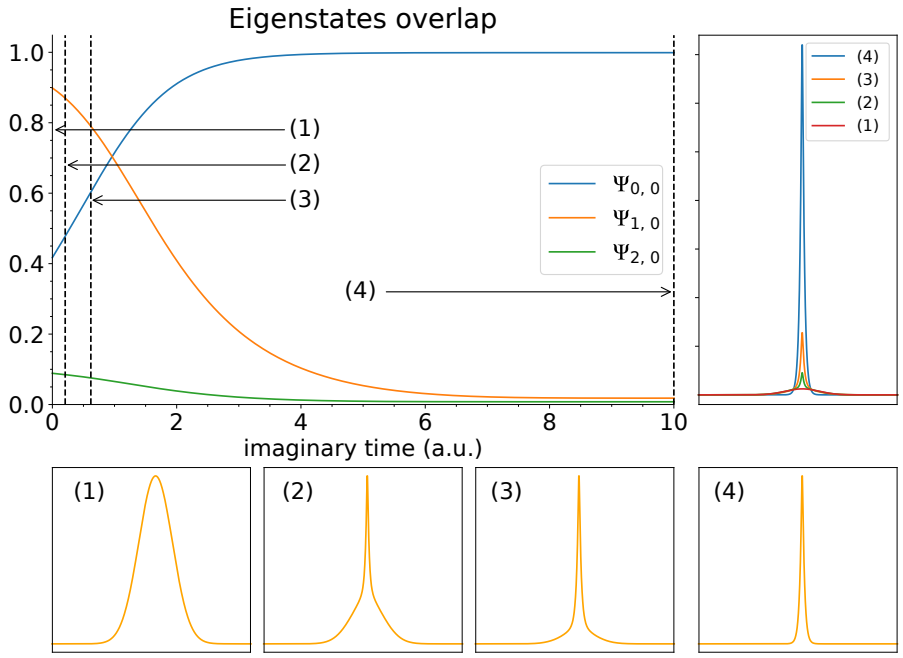


Figure 5.7: Ground state preparation of the 2D hydrogen atom using the PITE technique. The method was emulated on a $1 + 2 \times 10$-qubit quantum computer. Post-select the successful outcome at all times. The main plot shows the overlap of the propagated state with analytic eigenstates. Bottom panels are scaled cross-sections of the electron probability density sampled at labelled points during the imaginary time evolution. Top right panel shows the same probability densities plotted on the same scale. Figure reproduced from Ref. [4] – I was the original creator.

the state given in Eq. (5.17) directly in the emulator's RAM: It is readily verifiable by inspection that the first three steps indeed lead to this state regardless of the choice of E_i, whereas the additional qubits required to permute via the sorting network would have greatly increased the emulation cost. We then performed the final steps explicitly and noted the performance.

For $P = 5$ particles, with the ideal case that all energies correspond to integers, my collaborator and I confirmed the expected success probability of unity. For cases where the all energies differed from an integer, we found success probabilities of 0.990, 0.960, and 0.850 for deviations of 0.025, 0.05 and 0.1 respectively, showing an anticipated quadratic behaviour. However, when only one of the energies differed from an integer by some specified discrepancy, then we found that the success probability was essentially constant regardless of P. Thus the performance for a large number of particles will depend on how many of the energies differ from an

integer, and to what extent (with respect to 2^{-t}). Note that in each emulated scenario, we confirmed that given the 'all-$|+\rangle$' success criterion, the resulting state is ideal (up to a meaningless global phase).

5.3.5 Augmented Split-Operator

My collaborator and I assessed the ASO method by using it to simulate the dynamics of states peaked at the singularity, i.e. the most challenging cases.

Fig. 5.8 shows numerical results from our study of the ASO method. We considered the ground state of 2D hydrogen $\Psi_{0,0}$ which is peaked (with discontinuous gradient) at the origin of the classical Coulomb potential. A relatively modest resolution corresponding to $n_r = 6$ qubits per register (i.e. a 64×64 grid of spatial pixels) was used to represent the state , and the simulation box was optimal (see upper right graphic in the Figure).

Step 2 in Section 5.2.5 states that, having decided that our core patch will involve a subspace of Q pixels, we 'derive a small $Q \times Q$ unitary U_{core} that closely matches U_{repair} in that subspace'. The operator U_{repair} is simply the matrix that maps from the SO cycle as actually applied, to the ideal time increment operator. We therefore begin by calculating these matrices, each an $2^{2n_r} \times 2^{2n_r}$ object, explicitly using `Mathematica`. Having thus obtained U_{repair} we can proceed.

In simulations reported in Fig. 5.8 we considered two cases: core patches of size $Q = 2 \times 2$ and $Q = 4 \times 4$ pixels. For each case, we wrote down the $Q \times Q$ matrix M_{core}, which is composed of the elements of U_{repair} that lie in the Q-pixel subspace. M_{core} will not be unitary and therefore could not be implemented deterministically on our emulated quantum computer; we needed a unitary matrix U_{core} that is close to M_{core}. This was obtained by performing a standard singular value decomposition of M_{core} and using the components to construct U_{core}:

$$M_{\text{core}} = u\,\Sigma\,v^\dagger \quad \text{then} \quad U_{\text{core}} \equiv u\,v^\dagger.$$

Note that Σ matrix would be the identity if M_{core} were unitary; it is not, but by simply omitting Σ we generated our unitary approximation. The final step 4 of

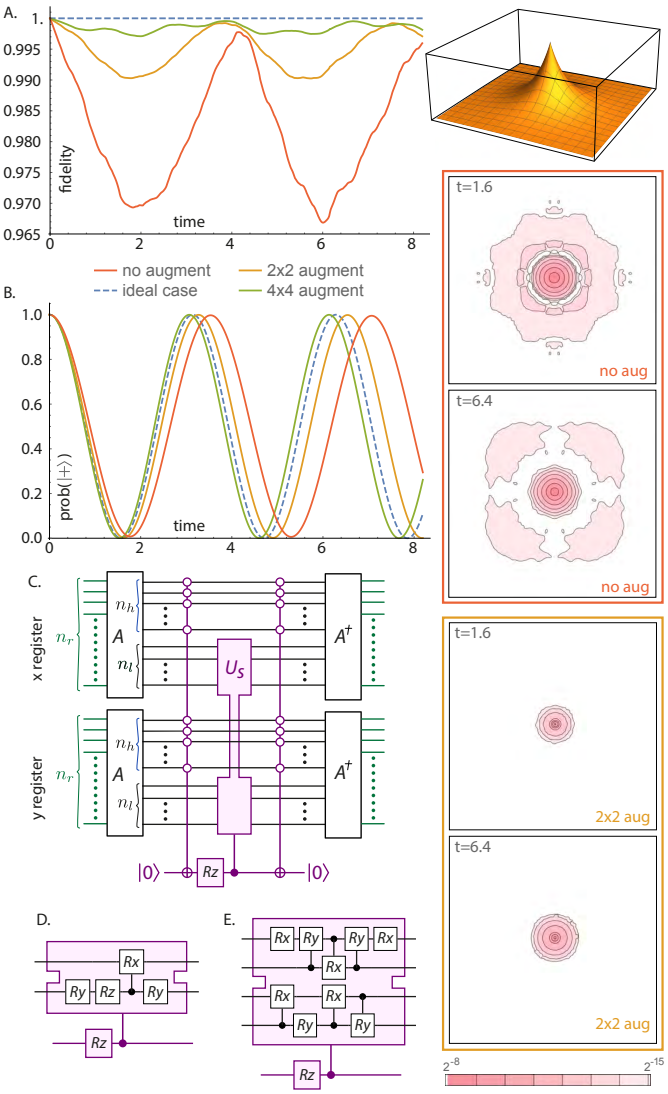


Figure 5.8: The performance of the Augmented Split Operator (ASO) technique. The emulated quantum computer has $13 = 1+6\times6$ qubits. **(A)** The 3D inset depicts eigenstate $\Psi_{0,0}$ of 2D hydrogen, within its simulation box. The graph shows the autocorrelation of the state at time t with respect to the initial state, which ideally remains at unity (blue dashed line). The red, orange and green lines are respectively the cases of no augmentation, a 2×2 augmentation and a 4×4 augmentation. **(B)** The result of phase estimation for the same three cases, with the ideal again shown with a blue dashed line. The contour plots on the right show the absolute difference in probability density, with respect to the initial state, for the case of no augmentation (red) and the 2×2 augmentation (orange). **(C)** Shows generic circuit employed, while **(D)** and **(E)** specify the particular circuits used for our 2×2 and 4×4 augmentations, respectively. The A operators simply increment the indexing of the spatial 'pixels' so that the lower right pixel is (0,0). In the case of the 2×2 augmentation, the increment used is $G = 1$ so that A maps the indices of the pixels of interest, i.e. $(-1,-1),(-1,0),(0,-1),(0,0)$, to $(0,0),(0,1),(1,0),(1,1)$ respectively. These indices are now exactly those states for which the most-important $n_r - 1$ qubits are zero, facilitating the application of the small augmentation circuit to *only* the four target states. Figure reproduced from Ref. [4] – created by Simon Benjamin and me.

finding a circuit to implement the stabilisation was performed using the circuit synthesis tool described in Ref. [155]; for the sizes used here, a trivial task.

We performed a series of standard and ASO simulations, in all cases fixing the time resolution at $\delta t = 0.004$, and monitor the absolute value of the autocorrelation. Graph (A) in Fig. 5.8 shows the result for three cases: The simple SO-QFT protocol (red), and protocols with a small (orange) and a medium (green) scale core stabilization augmentation. The small augmentation involves a circuit that modifies only the amplitudes associated with the 2×2 spatial pixels that are closest to the singularity (note that we align the spatial pixel lattice such that the singularity is mid-way between the four central pixels). The medium augmentation involves a larger set of 4×4 spatial pixels. There is a dramatic improvement in the autocorrelation, by an order of magnitude between the red and green plots.

The plots in graph (B) of Fig. 5.8 show the result of the IPE method. Ideally, the autocorrelation plot would match the dashed blue curve, which corresponds to phase acquisition according to the analytically derived energy. We can see that the red, orange and green lines are again progressively closer to the ideal. I should emphasise that the ASO method is agnostic to the state simulated; therefore, while it would be trivial to 'cheat' and apply an exactly compensating phase to obtain the blue dashed curve, in fact the circuits employed are derived without foreknowledge of the specific simulation task (i.e. $\Psi_{0,0}$ in no way enters the derivation of U_{core}).

Finally, a third lens on the simulation fidelity is provided by evaluating the difference between the $t = 0$ probability density (over the grid of 64^2 spatial pixels) and the density at a later time. Ideally, of course, this difference would be zero. In the contour plots on the right of Fig. 5.8, I contrast the case with no augmentation with the small augmentation scenario. While there is still a discrepancy, it is far more localised and stable (the un-augmented simulation involves wider, more dramatic fluctuations in the probability distribution).

The circuits that implemented both the small and the medium augmentations are shown in the bottom left of Fig. 5.8. They were derived through the process

explained in Section 5.2.5. Note the preparatory step of applying a unitary A that simply adds an integer to all states in the spatial representation, i.e.

$$A\,|n\rangle = |n + G\rangle$$

where the addition is understood to be modulo 2^{n_r}. The main figure shows the implementation of this for a general 'size' of the augmentation patch n_l, where in the numerical studies $n_l = 1$ for the small case and $n_l = 2$ for the medium case. The figure caption defines the shift for the case of the 2×2 augmentation, where $G = 1$ was used. The A operators can be avoided entirely if we simply define the origin of our spatial coordinates to a bounding corner of our core stabilisation region. The implementation of the spatial part of the SO-QFT cycle would need to be corrected for such shift, but that may prove to have lower total cost. Regardless, I included the A operators in the figure so as to make the it directly consistent with the other circuits and expressions in the present paper.

The multi-controlled NOT operation appearing in the circuits of Fig. 5.8 can be compactly realised by circuits reported in e.g. Ref. [264] that involve $4n - 6$ T-gates (single-qubit phase gates) and a comparable number of CNOT's, together with an ancilla that is measured during the process.

It is anticipated that the time cost of moving from the simple SO-QFT approach to the ASO method should be far less than the cost of the 'brute force' increase to the temporal resolution needed to assure proper behaviour of core-peaked states. For this demonstration, the performance of the $n_r = 6$ qubit registers was able to approximately match that of the $n_r = 11$ qubit registers under a 'brute force' approach (Fig. 5.4, upper left panel) and moreover this was achieved with a time step δt an order of magnitude greater, so facilitating shorter runtimes. The ASO method is therefore highly relevant when one operates with the strict Coulomb interaction, as in all the numerical studies in the present Chapter. It is distinct from, and compatible with, the use of higher order Trotter sequences, but whether it would also be useful in other real-space Hamiltonian simulation approaches remains to be seen.

5.3.6 Informal Resource Estimate

Based on the numerical results detailed above, I now present an estimate of the required resources for two molecules which are relevant examples for Applications II and III detailed at the start of this Chapter:

- A modest molecule well-worthy of a full-dimensional quantum dynamics study in the context of field-molecule interactions would be **ammonia ($\mathbf{NH_3}$)**, investigated in the context of selective hydrogen atom removal [273]. If quantum modelling of its dynamics under laser control were to reveal new synthesis options, the consequences could be profound since ammonia use lies at the heart of modern agriculture.

- For electron scattering processes, an example currently beyond the reach of full-dimensional quantum dynamics simulation is **hexafluoro ethane ($\mathbf{C_2F_6}$)**, a representative example of fluorocarbons [274] relevant in the chemistry of the ozone layer and in plasma etching.

These resource estimate are extrapolations and speculative - and does not constitute a rigorous scaling analysis. Asymptotic analyses have been made in the last few years for real-space grid first-quantised Hamiltonian simulation [49, 51, 53] and specifically the SO-QFT method [55].

Space Cost

The lowest possible qubit count (for a given simulation accuracy) results from selecting the smallest adequate 'simulation box' length L, and setting the spatial resolution δr (Eq. (2.29)) to be just sufficient to capture the most curved elements of the wavefunction. The number of qubits per particle, per spatial dimension, is then $n_r = \log_2(L/\delta r)$. The quantity L specifies the region of space outside of which our multi-particle wavefunction only has negligible amplitude 'clipped' by the boundary. Moreover, I have presented methods that allow us to study processes that would, over the simulation time, go beyond the simulation box due to scattering or ionisation. It is reasonable to assume, as in Ref. [54] for example, that the

volume L^3 goes linearly with the number of particles P; a molecule with twice the number of particles requires roughly twice the volume.

Meanwhile, the severity of the wavefunction's curvature should scale directly with $Z_{\max}$, the maximum nuclear charge of any of the nuclei in the system. Thus, increasing the molecule's size without increasing $Z_{\max}$ should not require any adjustment in δr, but the simulation box may have to be larger to accommodate more particles. Importantly, the accuracy with which we model the particle is essentially unaffected by this change.

In view of the above observations, we can expect that the number of logical qubits necessary scales, for suitable constants C_i, according to

$$\begin{aligned} n_r = \log_2\left(\frac{L}{\delta r}\right) &\approx \log_2\left(\frac{C_1 P^{\frac{1}{3}}}{C_2 Z_{\max}^{-1}}\right) \\ &= C_3 + \log_2(Z_{\max}) + \frac{1}{3}\log_2(P). \end{aligned} \tag{5.26}$$

The total qubit count will be $3Pn_r$ for P particles in 3D. Note that this simple expression for n_r does not account for the potentially helpful fact that atomic radii have a highly non-linear (and sub-linear) dependence on the number of electrons [275]. Accounting for this would lower the outcome. Since n_r really depends on more properties of the problem besides $Z_{\max}$ and P, the constant C_3 will not be an absolute value but rather fluctuate with other system properties such as the geometry and electron configuration.

I will now provide values for C_3 for each example molecule. First, note that in the numerical modelling reported, remarkably accurate and stable simulations can generally be achieved with as few as 6 qubits per x, y or z sub-register (thus, 18 qubits per 3D particle). Methods such as IPE, requiring only one additional logical qubit, can then obtain eigenenergies with accuracy up to 6 decimal places. Using the ASO method, even electrons with substantial weight at the nuclear core can be modelled with only a modest increase in resolution. I therefore assume that $n_r = 7$ can suffice for systems with $P = Z = 1$. For larger values of P and Z, more pixels per dimension are necessary to compensate for differently sized

simulation boxes and higher curvature of core states. I discuss momentarily how to estimate these required changes.

One may wish to use a quantum simulator to explore molecular systems that have not been experimentally charatcterised. Therefore, I will adopt a first-principles argument based only on the constituent atoms and their presumed locations. We can collocate the origin of the highest-occupied hydrogen-like electron wavefunction of each atom at the coordinates of each nucleus, and calculate its radial charge density according to e.g. Ref. [276], which also semi-empirically accounts for nucleus shielding. Summing up the contributions from all atoms gives an approximation to the total charge density, assuming the electrons of different atoms were not interacting with each other. We then find the maximum of this total density on the surface of the simulation box, which informs us about how strongly the electrons will interact with its boundary. For molecules of interest here, the atomic locations in their equilibrium geometry are presented in Ref. [277].

To get an idea of what an acceptable charge density at the box surface could be, I utilised the simulation results of a 3D helium atom without electron-electron interaction performed by my collaborator (see next Section). From the electron configuration and box size of that simulation, I derived its maximum surface electron density ρ_0 using the approximate method described above. Consequently, the simulation box for any other molecule might be considered sufficiently large whenever the method above yields a maximum surface electron density that does not exceed ρ_0. Taking a relatively optimistic line in the resource estimation, I will refer to this rule; but note that because it is only a rough approximation to the extent of the electron cloud, and specific scenarios might necessitate an increase. Interesting dynamics with substantial numbers of moving particles may require significantly larger boxes, so as to avoid significant amplitude collision with the simulation boundary (or to allow an adequate attenuation regions as described in the main paper).

Besides varying simulation box sizes, we should also account for changes in the required spatial resolution. For $Z > 1$, the features of hydrogenic wave functions

shrink by a factor of exactly Z^{-1}, which is usually also a good approximation for low-lying core states. To accurately resolve these electronic states, we must therefore increase the number of grid points per unit length by a factor of $Z_{\max}$.

For ammonia (NH_3), the more modest of the two molecules identified, an optimistic resource audit can proceed as follows: The highest charge in the molecule is $Z_{\max} = 7$. Therefore, the number of grid points must increase to 7 times that of our estimate for $Z = 1$. At the same time, the above mentioned method to determine the box size yields a side length of ≈ 1.1 times the length required for a single hydrogen atom, giving a total factor of $\approx 1.1 \times 7 = 7.7 < 8 = 2^3$, which means we must increase n_r by 3 from its reference value of 7, giving $n_r = 10$ qubits per particle and dimension. Substituting these values in Eq. (5.26) and rounding where appropriate leads to a value of $C_3^{\mathrm{NH_3}} \approx 6$. For 14 particles (10 electrons and 4 nuclei) in 3 dimensions, we arrive at a total of $3 \times 10 \times 14 = 420$ qubits. One might round this to 450, recognising that multiple ancillas may be required even in a very frugal implementation.

For C_2F_6, the highest-charge nuclei are those of fluorine, giving us $Z_{\max} = 9$, thus increasing the required spatial resolution 9-fold. Surprisingly, as the electron densities of C and F drop off quite rapidly with increasing distance from the atom, the above described method suggests that the required side length of the simulation box could be as small as 0.85 times that of a single hydrogen atom. This gives a total factor of $\approx 0.85 \times 9 = 7.65 < 8 = 2^3$, again meaning that 3 qubits must be added to the reference value, giving a total of $n_r = 10$. Using Eq. (5.26) again, I arrive at $C_3^{\mathrm{C_2F_6}} \approx 5$. The total number of required qubits is $3 \times 10 \times 74 = 2220$, which might be round up to 2250 to allow for a frugal ancilla overhead.

From these two examples, a cautious estimate of C_3 is on the order of ~ 10, but note that it might be lower for advantageous circumstances as in the examples given, or may also increase in unfavourable situations.

Time Costs

It is difficult to provide meaningful estimate of total 'wall clock' time, even to within an order of magnitude. I will nonetheless use some broad assumptions to provide generally indicative numbers for the time cost.

The timescale of the dynamics depends heavily on the energy scales across which the events probed occur; the models herein (electron scattering, ionisation) correspond to high energy 'fast events' that occur on a sub-femtosecond timescale, whereas resulting dynamics (e.g. bond breaking, molecular fragmentation, energy transfer) are 'slow events' that happen at $O(1)$ even up to $O(10)$ picosecond timescales. This can be postulated to be independent of system size.

The resolution by which one would have to discretise this total simulation length by δt, and thus the number of SO-QFT steps, depends only on the δr sufficient to describe the most tightly bound electrons of the element with highest nuclear charge, which also does not change with system size. Moreover, the methods explored here, and indeed proposals of e.g. Ref. [228], indicate that energy estimation and state preparation could also be achieved using a number of time steps that does not scale significantly with system size.

Many of the simulations presented used $\mathcal{O}(10^3)$ SO-QFT cycles to generate dynamics with good precision. This count would diminish with the adoption of more sophisticated Trotter sequences than the lowest-order pattern my collaborators and I employed; one might expect an order of magnitude reduction is possible here. However, one might then ask whether it is possible to encapsulate a 'slow event' like the subsequent dynamics of a molecule immediately after an interesting 'fast event'. Simplistically, this might seem to imply the need for a corresponding multiplicative factor in the number of SO-QFT time steps required for the complete simulation. Fortunately however, as explored in the two-particle scattering simulation, it might be natural to use coarser time steps according to the speed of the physics that is unfolding. Accounting for these various considerations, let's assume that $\mathcal{O}(10^5)$ SO-QFT cycles should suffice to reveal longer time quantum molecular dynamics.

The depth of quantum gates, thus time, used to complete one SO-QFT time step *does* scale with system size. There is a modest scaling due to the computation associated with introducing proper phases (equivalent to binary addition/subtraction, squaring, and inversion) which scales with the number of qubits in each register – although not with the number of qubits in total. For example, the QFT implementation used herein scales as $\mathcal{O}(n_r^2)$. Moreover, other operations also require n_r^2 gates, such as the application of the $\propto k^2$ phases for the kinetic operator. A full SO-QFT cycle may therefore involve operations equivalent to several QFTs, since they constitute an approach to the addition and squaring operations employed. If we therefore take $10n_r^2$ as a plausible suggestion for the gate depth of each SO-QFT sub-step in a parallelised implementation, then for the C_2F_6 molecule, the total gate depth is roughly $260 \times 10 \times 10^2$ which one can conservatively call $\mathcal{O}(10^6)$ gates required for each SO-QFT time step (and $260/2 = 130$ such processes occurring in parallel throughout that cycle). One could therefore assume that the algorithmic gate depth for understanding a 'fast event' such as electron ionisation is hundreds of millions, and evolving such a system in the slower picosecond timescale is about a hundred billion.

Since the n_r scales only logarithmically with simulation box size L, total wall time is expected to be poly-logarithmic in e.g. the number of particles $\mathcal{O}(\log^k P)$. A more significant factor comes from the need to compute the electron-electron interaction for all possible pairings of electrons. The task can be divided into a series of stages within each of which all electrons are paired-off, and we can assume that computations relevant to electron-electron dynamics occur in parallel (see [49]). The number of such stages is $P - 1$ for P electrons, so that the total time for a complete SO-QFT time evolution is of the order $O(P \log^k P)$. Considerations such as the time required to move from one pairing to another should also be borne in mind, but one finds that this should scale well even in devices with restricted connectivity.

As the methods here will almost certainly require a fault-tolerant quantum computer, the gate-count metric relevant for approximating the time complexity of implementing a SO-QFT step may be the T-gate count. This is a measure of how

many steps in the algorithm correspond to the costly non-Clifford operations that cannot be directly performed in stabiliser codes on fault-tolerant quantum computers, and efforts to minimise the count for standard subroutines are ongoing [278]. For example, it was observed that the number of T-type gates needed for multi-controlled rotations can be reduced to a remarkably frugal level [264]; such rotations are key in e.g. both the attenuation and ASO techniques. More generally, the trade-off between time and qubit count is a research topic in its own right and recent papers have shown how dramatically this balance can be adjusted [279, 280].

Given the algorithmic gate depth for each SO-QFT step therefore, the translation to execution time will vary dramatically depending on the native physical gate error rate (typically assumed to be 10^{-3} or 10^{-4}), the speed of a stabiliser cycle and the option to trade computation speed for higher qubit overheads [280]. For surface code based implementations with solid state platforms, a credible stabiliser speed [281] is 1μs (with faster speeds being conceivable). Assuming the state preparation procedure only requires a polynomial-scaling overhead, and is thus not the dominant cost, I therefore give an extremely speculative estimate of the total clock time on the order of minutes for 'fast', small molecule simulations, while the challenging long-timescale process would require on the order of a day.

I will also note that generating certain interesting plots (equivalent to those shown in this paper) will require many repeated executions of the simulation. Fortunately such repetitions can be perfectly parallelised over independent quantum processors, which need not have quantum interlinks or even be co-located.

The same full-dimensional real-space grid simulation of the reaction on classical machines of course is not possible, simply because the sheer number of pixels in the simulation is beyond memory available on most high-performance computer clusters. An equivalent classical procedure of selecting reduced reaction coordinates where the dynamics is expected to be relevant, computing the corresponding electronic potential energy surfaces, and subsequent dynamics propagation, can require months of effort; finalising the electronic structure itself can often be the main bottleneck. Thus, this work suggests that quantum Hamiltonian simulation

approach presented can be more efficient than an equivalent classical method, based on the fact that a more complete picture is used and the effort of computing potential energy surfaces can be circumvented.

5.3.7 Quantum Hardware and Architecture

The preceding Section obtained back-of-the-envelope estimates of logical qubit counts ranging from several hundred to several thousand. Presently, the most well-understood error correction codes can require many hundreds of physical qubits per logical qubit, assuming relatively deep algorithms and physical error rates comparable to today's best QC prototypes. Even if one makes the very optimistic assumption that some form of error mitigation can suffice in place of full fault tolerance, at least for small molecular simulations, the more powerful forms of mitigation can still require a multiplicative increase in the number of physical qubits [282, 283]. Thus the number of physical qubits required for the modelling considered in the preceding section could easily reach the high thousands or millions. This raises the question of what kind of quantum computing architecture would be needed.

In particular, my collaborator and I were motivated to explore whether some form of network architecture might be compatible with the SO-QFT (or ASO) approach. Such a network might involve quantum computers interlinked within a building, analogously to a conventional HPC facility, and relevant methods of linking processors have been experimentally realised. Alternatively, for suitably compact platforms the network might correspond to linked QC processors on a single chip, analogous to today's multicore CPUs – which has been explored for quantum computing [284]. In either case, it is realistic to assume that in a network of processor nodes the inter-node operations are slower than the intra-node ones.

Fortunately, the SO-QFT method is quite compatible with a network paradigm; there is a natural partitioning of the problem into nodes that each contain the registers (or sub-registers) associated with a given simulated particle. While data

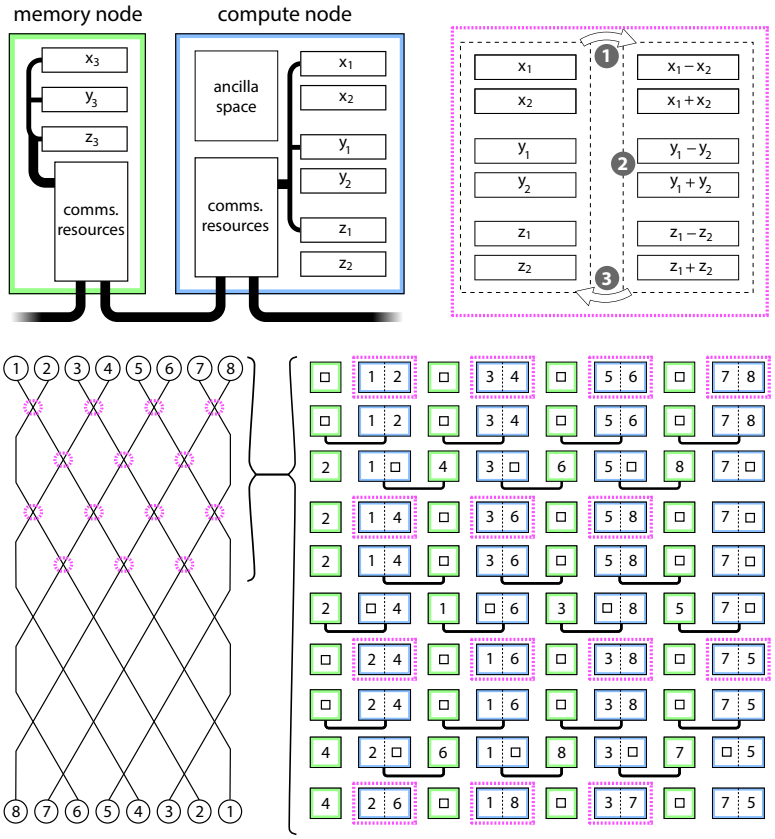


Figure 5.9: A possible 'multicore' quantum computer architecture for grid-based simulation. Based on distributing the P-particle state over $P/2$ memory nodes ($3n_r$ qubits) and $P/2$ compute nodes ($6n_r$ qubits) which are interlinked either on-chip [284] or at the macro-scale. The upper-right panel indicates that three steps occur in a compute node. In step 1, quantum addition and subtraction are used to move to the relative and total coordinates (and step 3 will reverse this). In the middle step 2, the relative coordinate is used to apply phase shifts required by the SO-QFT cycles, and optionally we apply an augmentation step from ASO. The lower panels show how parallelised phases alternatively process and permute the data. The numbers within the circles and squares are the labels of particles; they correspond to the subscripts of the x, y and z symbols in the upper-right panel. Figure reproduced from Ref. [4] – created by Simon Benjamin and me.

exchange between nodes is obviously required, it need not be a dominant component of the overall resource costing even if the physical links are slow.

In Fig. 5.9, I show one possible partitioning – it is not the most granular, since one could assign individual particle sub-registers to cores, but it does strike a good balance between the intra- and inter-core operations. Note that the required connectivity between cores is merely linear and nearest-neighbour. We supposed that there are two forms of core: the compute nodes are responsible for all the

processing that is associated with the SO-QFT method, while the simpler memory nodes only store registers transiently.

Each node includes communication resources which facilitate transfer of quantum information between nodes, for example through the use of teleportation enabled by high-quality shared Bell pairs. Such processes can occur independently of the main computation and simultaneously with it, and need not involve a large number of qubits; see for example the analysis in Refs. [284, 285].

In the scheme illustrated in Fig. 5.9, the transfer of a register from one node to another involves writing into an 'empty' register where all qubits are in state $|0\rangle$. If the individual qubits are in fact encoded logical states formed of many physical qubits, then this would introduce a multiplicative factor in the Bell pair count, but it would remain linear in the register size. Ideally the communications hardware would be capable of generating the required Bell pairs within the time that the compute node requires for a full implementation of the particle-particle interaction component (and any augmentation as per ASO). Given that this computation will require a gate depth of at least n_r^2, there is scope for a factor of $10 - 100$ in the relative speeds of the intra-core computations and the inter-core Bell pair preparation before the latter would become a bottleneck.

5.3.8 Additional Analysis and Results

For completeness, I want to briefly mention some key results my collaborators contributed in Ref. [4].

The SO-QFT emulation required the repeated simulation of circuits within which a series of contiguous gates could be more efficiently effected by a single invocation of a bespoke function. By leveraging the fact that the Coulomb potential operators involved are diagonal in the real representation, it is possible to numerically simulate them cheaper than even a *single* general unitary gate, in an *embarrassingly parallel* manner. Tyson Jones implemented this optimisation to accelerate the simulations reported in this Chapter, which has since been integrated into the QuEST emulators. The interested reader may view these documented facilities here. This insight has

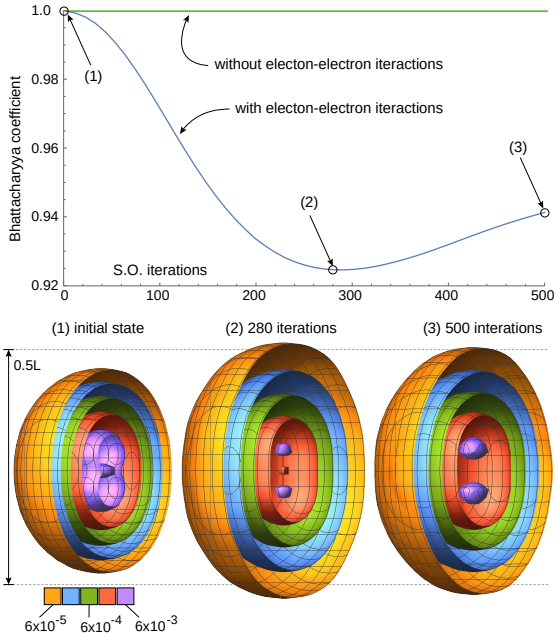


Figure 5.10: Simulation of a Helium atom in real space. The top plot is the Bhattacharyya coefficient (with and without enabled electron-electron interactions), and the bottom is the real-space electron density distributions during the time evolution of the Helium atom simulation. Coloured shells are electron probability density isosurfaces, within a simulation box with $L = 25$ a.u. Distributions are shown for the initial state, at the time where it is maximally spread out, and at the end of the simulation. The 500 SO cycles correspond to propagation of 25 a.u. (≈ 0.6 fs). Figure reproduced from [4] – created by Richard Meister.

allowed my collaborators and me to perform the simulations without incurring impractical time costs. However, it is important to stress that the use of these high level functions, which did not involve specifying a circuit-level description, did not compromise the exact nature of the emulations – it is simply an expediency allowing us to arrive at the exact behaviour with less classical resources than would be required if explicit gate-level operations were used throughout. In the real device, such circuit-level prescriptions must of course be used.

Richard Meister built on top the exploratory simulations I did and extended the low-dimensional results to perform the largest grid-based simulation reported in Ref. [4]. He modelled the dynamics of the helium atom: Two electrons interacting via a repulsive Coulomb interaction, both bound by a central attractive Coulomb potential, in three spatial dimensions (see Fig. 5.10). These results were used in the informal resource estimates I performed.

5.4 Concluding Thoughts

In this work, I explored the quantum simulation of many-body physics expressed in real 3D space, using the SO-QFT approach performed on exactly-emulated qubits, and tested a variety of resource-frugal techniques that facilitate augmentation and monitoring of the simulation. My collaborators and I tested known as well as new quantum circuit techniques, covering all key aspects of simulation: State preparation, Hamiltonian simulation, and the extraction of physical observables. From these emulated experiments, we characterised the resources and suggested possible quantum computer architectures necessary to realise 'digital experiments' of quantum molecular dynamics [241] on small, error-corrected quantum computers. While certain restrictions and caveats were identified, generally the grid-based method performed very well; the results were consistent with the view that first-quantised many-body simulation paradigms can be significant from the early fault-tolerant quantum computing era onward.

While combining SO-QFT with other Hamiltonian simulation methods may lead to hybrid quantum-classical approaches for real-space simulations that can be useful in the near term, the high qubit count necessary for encoding the first-quantised grid representation and generally deep circuits will likely prevent its applicability in the NISQ era. Nevertheless, small executions of these methods might offer synergies with real-space electronic structure approaches such as density functional theory (DFT) in the early fault-tolerant regime; one can imagine using small exact calculations enabled by real-space quantum simulations to improve DFT functionals, or using particle densities provided by an initial DFT precomputation to inform the single particle functions that are loaded into particle registers.

SO-QFT itself leads to robust methods for measuring observables such as phase estimation. It is also flexibile enough for us to check certain properties quite cheaply at any time, analogous to stabilisers of a code, and their cost of evaluation can be low even if preparing the state was not – in the PITE and attenuation cases, frequent measurement can be used to modify the evolution of the system

in a non-unitary manner. In light of this, it is interesting to ask whether the grid method can be inherently robust to errors. It may be possible to craft a version of the algorithm where the majority of harmful errors will cause the state to fail a validation check, while less damaging errors are mitigated by a non-unitary component in the dynamics.

A natural next step is to explore the use of multi-resolution grids as employed in e.g. MADNESS [243, 286], the highly successful classical computing program for real-space grid simulations. Given that the present paper has revealed remarkable variation in resolution requirements – even across different states of a single system, one might be well-motivated to incorporate such ideas to reduce the computational resources required. Moreover, particles within many-body systems can often be considered as localized, presenting the opportunity for more resource efficient representations based on that locality. However, while multi-resolution grids might reduce the qubit count considerably, it would likely be at the cost of more sophisticated time propagators and basis transformations. Thus, it is important to establish whether the methods in e.g. [243, 286] can indeed be translated successfully to the quantum context.

It is obvious that the SO-QFT simulation techniques presented here can be generalised to modelling systems beyond quantum chemistry. The solution to any Cauchy type initial value problem

$$\frac{\partial}{\partial t}\Psi(x,t) = \hat{\mathcal{D}}(t)\Psi(x,t) \tag{5.27}$$

with a time-dependent differential operator that may be separated into operators that are respectively diagonal in position- and momentum-space

$$\hat{\mathcal{D}}(t) = \hat{\mathcal{D}}_1(t) + \hat{\mathcal{D}}_2(t) \tag{5.28}$$

can be approximated with the SO-QFT approach [287]. Many problems of academic and commercial interest can be modelled with Cauchy type partial differential equations. For example, Dirac and Klein-Gordon equations, which are of Cauchy form, reconcile quantum mechanics with special relativity and may be used for

modelling of high energy particles [135]. A very different application is financial engineering. Quantum advantage is often promised for the modelling of how the prices of assets, such as options and derivatives, evolve over time [288, 289]. This is key to executing purchases and sales that maximise the eventual payoff from trading such assets. These assets often have complex underlying dependence on random variables, which have in practice been modelled using computationally expensive stochastic Monte Carlo methods. In the same manner as pixelating real-space wavefunctions and storing them in the computational basis states of a quantum computer, probability distributions corresponding to asset prices can be discretised and loaded into quantum registers. A very relevant model for time-propagation of these distributions is the Black-Scholes-Merton equation [290, 291], which is a Cauchy type partial differential equation. Beyond these use cases, it remains to be seen how non-unitary operations achieved through ancilla measurements can extend the applications of the SO-QFT model.

To confront a person with his shadow is to show him his own light. Once one has experienced ... what it is like to stand judgingly between the opposites, one begins to understand what is meant by the self. Anyone who perceives his shadow and his light simultaneously sees himself from two sides and thus gets in the middle.

— By C. G. Jung [292]

6

Shadow Spectroscopy: Measurement Protocol for Any Quantum Technology

Contents

6.1 Introduction 174
6.2 Method 175
6.2.1 Quatum Measurement of Many Signals 175
6.2.2 Classical Post-Processing 178
6.2.3 Summary of Shadow Spectroscopy 182
6.3 Key Features of Shadow Spectroscopy 182
6.3.1 Robustness to shot noise 183
6.3.2 Robustness to algorithmic or gate errors 183
6.3.3 Compatibility with Quantum Chemistry 184
6.3.4 Comparison with other quantum algorithms 188
6.4 Experimental Validation: NISQ Scenario 189
6.4.1 Proof-of-Concept 190
6.4.2 Hardware Demonstration 193
6.5 Experimental Validation: Fault Tolerant Scenario . . . 199
6.5.1 Utility of Local Pauli Operators 199
6.5.2 Methylene Example 200
6.6 Additional Analysis and Experiments 203
6.7 Concluding Thoughts 204

In this chapter, I use text from the arXiv preprint 2212.11036 [5] verbatim – I was the first and original author of the text. The work was done in collaboration with coauthors Richard Meister, Matthew Goh and Bálint Koczor at the University of

Oxford. My coauthors and I developed the methods therein as a group, based on initial ideas from Bálint Koczor. I performed a significant portion of the numerical simulations, all of the quantum hardware demonstrations, and developed quantum-chemistry specific ideas. It is currently under review at Physical Review X.

6.1 Introduction

The quantum resource cost e.g. the number of circuit measurement shots or the number of qubits needed to extract useful information from a quantum simulation can diminish, even cancel out, advantages or speedups offered by quantum computers. In Ref. [5], my collaborators and I created a novel method for the accurate estimation of energy gaps in any many-body Hamiltonian – information that is of key importance for applications in chemistry, materials science and beyond. It uses comparatively very few measurement shots to a quantum computer and requires no additional ancilla qubits beyond those used for encoding a system and performing Hamiltonian simulation. It can in principle work on any quantum simulator platform (digital or analogue), as well as use any time-evolution method of choice. The initial state only needs to have sufficient support for the set of eigenstates which we want to know the energy gaps for, instead of a large overlap with one target eigenstate as is the case for methods like phase estimation.

The key innovation is to harness powerful classical shadow measurement techniques (as prefaced in Section 2.2.6) which enables the efficient estimation of many observables of a quantum state using orders of magnitude fewer shots than e.g. VQE. Specifically, our strategy is to analyse how observables of a quantum state, as reconstructed from classical shadows, change as the state evolves in time. Due to the fundamental laws of quantum mechanics, oscillations of all observable properties in a quantum system contain exactly the same fundamental harmonic frequencies that precisely correspond to energy differences in the Hamiltonian. These can then be extracted to form a spectrum of energy gaps via classical post-processing techniques.

I tested *shadow spectroscopy* on a broad set of applications, including the use of real state-of-the-art quantum hardware. My results and those of my coauthors

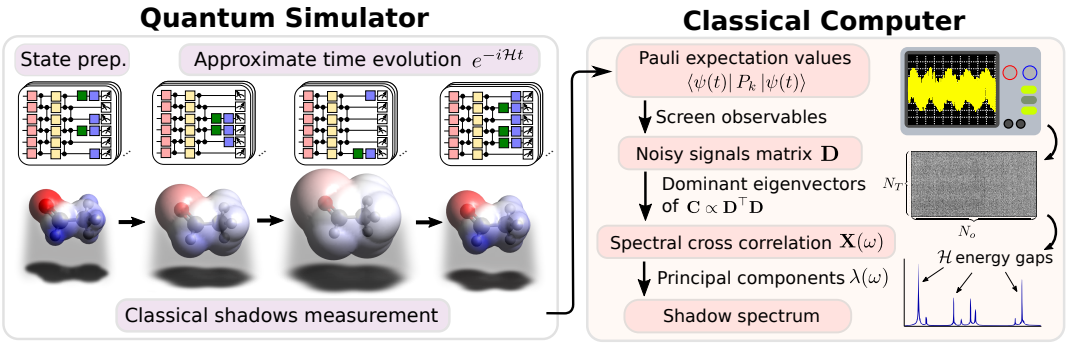


Figure 6.1: Flowchart of algorithmic shadow spectroscopy. A plethora of well-established quantum simulation algorithms can be used to prepare time-evolved quantum states $|\psi(t_n)\rangle$, from which a series of classical shadows are extracted using randomised measurements —in this cartoon a mean-field many-electron wavefunction of the acetaldehyde molecule is propagated. The shadows are stored classically as a list of binary numbers and measurement basis indexes. These are post-processed using conventional computing resource by first generating a large noisy observables signal matrix $\mathbf{D}$ from the classical shadows, then a square matrix $\mathbf{C}$, and finally the shadow spectrum representing energy differences of the system; these steps are elaborated upon in the text. Figure reproduced from Ref. [5] – I was its original creator.

suggest that it can provide high quality estimations of energy gaps in quantum systems while being particularly robust to the limitations of quantum measurements. It was also very general and easy to use in practice, and stands to be a promising option for quantum computing excitation energies, applicable in the near-term, early fault-tolerant and fully fault-tolerant eras of quantum computing.

6.2 Method

6.2.1 Quatum Measurement of Many Signals

The time-evolution of any quantum state under Hamiltonian dynamics introduces periodic oscillations precisely determined by differences of eigenvalues in the Hamiltonian operator. The following statement summarises the well-known property that expected values of observables oscillate with these frequencies under a closed Hamiltonian evolution [293, 294].

Statement 2. *Given any d-dimensional Hamiltonian $\mathcal{H}$ with eigenvectors $|\psi_k\rangle$, one can define any quantum state as $|\psi\rangle = \sum_{k=1}^{d} c_k\,|\psi_k\rangle$. Measuring the expected value*

of any observable O in the time-evolved quantum state $|\psi(t)\rangle := e^{-it\mathcal{H}}|\psi\rangle$ yields the signal $S(t)$ as

$$S(t) := \langle\psi(t)|O|\psi(t)\rangle = \sum_{k,l=1}^{d} I_{kl}\, e^{-it(E_l - E_k)}. \tag{6.1}$$

The above Fourier components have frequencies that correspond to differences of eigenvalues $E_l - E_k$ of the Hamiltonian. The intensity of the corresponding peak in the Fourier spectrum is $I_{kl} = c_k^* c_l \langle\psi_k|O|\psi_l\rangle$.

Proof. All d-dimensional Hamiltonians admit the spectral decomposition $\mathcal{H} = \sum_{k=1}^{d} E_k|\psi_k\rangle\langle\psi_k|$ with eigenvectors $|\psi_k\rangle$ spanning the full space. Thus, the evolution operator acts as

$$|\psi(t)\rangle = e^{-it\mathcal{H}}|\psi\rangle = \sum_{k=1}^{q} c_k e^{-itE_k}|\psi_k\rangle,$$

and the time-dependent $\langle\psi(t)|O|\psi(t)\rangle$ expected value is then obtained via a direct calculation. The generalisation of this statement to density matrices as well as to infinite-dimensional but bounded Hamiltonians follows trivially. □

An important consequence of this is that any time-dependent expected values contain the same harmonic components $e^{-it(E_l - E_k)}$ *regardless* of the observable measured. These frequencies precisely correspond to energy gaps $(E_l - E_k)$ in a Hamiltonian. In principle, this can be determined by naïvely choosing one observable O and measuring its time-dependent expectation value $S(t)$, then Fourier transforming the signal.

In practice, multiple factors determine the intensity of the signal $S(t)$. First, it depends on the amplitudes of the statevector c_k; if the initial state is, e.g., a random state or a uniform superposition over all eigenstates, the Fourier spectrum may contain an exponentially large number of overlapping peaks and prohibit us from learning anything from the spectrum. If we can instead initialise the quantum state such that only a handful of lower-energy eigenstates contribute, i.e., $|c_k| \gg |c_l|$ where all $k \in \mathcal{Q}$ and $l \notin \mathcal{Q}$, then we can isolate $\mathcal{O}(|\mathcal{Q}|^2)$ dominant peaks. In this work, for example, my collaborators and I employed several known strategies for preparing initial states such that $|c_1| \approx 1$ and $|c_k| \gg |c_l|$ for all $2 \leq k \leq q$ and $l > q$.

This means only q dominant peaks of non-zero frequency are expected, from which we can learn the lowest q excitation energies as $E_2 - E_1$, $E_3 - E_1$ etc.

Second, the signal intensity depends on the matrix elements between eigenstates as $\langle\psi_k|O|\psi_l\rangle$ which can be arbitrarily small for any chosen observable O. One can identify problem specific observables that are expected to yield dominant signals, such as excitation operators in quantum chemistry as I discuss in Section 6.3.3. However, even if we can guarantee that a certain observable yields a non-zero matrix element, on a real quantum simulator the signal intensity might be so low that a large number of circuit repeats and measurements are still required to obtain a sufficiently accurate estimate of the expectation value i.e. suppression of the background shot noise.

For this reason, my collaborators and I chose to use classical shadows, a measurement method designed specifically to reconstruct a large set of observable expectation values $\{\langle\psi| O_i |\psi\rangle\}$ using few measurements to the quantum state $|\psi\rangle$. Specifically, we propose performing classical shadows measurements at every time step during a Hamiltonian evolution sequence on a quantum computer, and then classically reconstructing time-dependent expectation values $S_i(t)$ for many local observables O_i simultaneously, in the form of Pauli strings i.e. $\{O_i\} = \{P_i\}$. We then search for observables that yield intense signals, and extract their dominant harmonic frequencies to reveal the energy gaps in the Hamiltonian.

The complexity of estimating time-dependent signals this way is as follows:

Statement 3. *Using the classical shadows technique, determining N_o independent signals (each with N_T time increments) corresponding to the time-dependent expected values of at most q-local Pauli strings requires a number of circuit repetitions (shots) as*

$$N_s \in \mathcal{O}[N_T 3^q \epsilon^{-2} \log(N_o)]. \tag{6.2}$$

Proof. Following Theorem 1 of Ref. [173], N_o Pauli strings P_i of locality q, can be estimated to precision parameters ϵ, δ via the number of batches $K = 2\log(2N_o/\delta)$ where the number of samples in the individual batches is $N_{batch} = \frac{34}{\epsilon^2} \max_i \|P_i\|^2_{shadow}$.

This results in an overall number of samples $N_s = N_{batch}K$ while the norm is given in Lemma 3 in Ref. [173] as $\|P_i\|^2_{shadow} = 3^q$. As such, at each time step, we use this procedure once to estimate N_o Pauli strings of locality q which has a sample complexity upper bounded by

$$\frac{68}{\epsilon^2} 3^q \log(2N_o/\delta). \tag{6.3}$$

The procedure is repeated for each time step and is thus used overall N_T times from which the total sample complexity follows. □

The above statement guarantees that we can obtain a large number of signals with logarithmic efficiency. However, each signal is still burdened with shot noise $S_i(t) + \mathcal{E}_i(t)$ such that the variance of the random fluctuations $|\mathcal{E}_i(t)| \leq \epsilon$ is globally bounded by the precision parameter ϵ inherited from classical shadows. One might be inclined to average the signals as $\sum_i S_i(t)$; but because each signal is phase shifted by an amount that is specific to each observable as

$$S_i(t) = \text{const} + \sum_{\substack{k,l=1 \\ k<l}}^{q} |I_{ikl}| \cos[t(E_l - E_k) + \phi_{ikl}] \tag{6.4}$$

(where $I_{ikl} = |I_{ikl}|e^{-i\phi_{ikl}} = c_k^* c_l \langle k|P_i|l\rangle$), and the individual phase shifts are not known in advance, averaging over phase-shifted cosines may even cancel out the signal. The problem then becomes that of finding the most intense common frequency components across a collection of noisy signals that share common harmonics.

6.2.2 Classical Post-Processing

A myriad of classical signal processing techniques can be useful for efficiently extracting these common harmonics to generate a spectrum of energy gaps. I now present and justify the sequence of classical processing steps that my collaborators and I designed and used to obtain shadow spectra.

Standardise and Isolate Signals

We first standardise the signal of the Pauli string P_k using the sample mean μ_k and the empirical standard deviation σ_k as

$$f_k(n) = \frac{\langle\psi(t_n)|P_k|\psi(t_n)\rangle - \mu_k}{\sigma_k}. \tag{6.5}$$

The above statement also includes the discrete time step index $n \leq N_T$ and the time variable $t_n = n\Delta t$. To reduce the size of our dataset and speed up subsequent processing steps, we could reconstruct observables that are expected to yield strong signals, as well as filter out signals $f_k(n)$ that are not sufficiently different from statistical noise – a simple autocorrelation test, such as the Ljung-Box test [295, 296], can work very well in practice. This step requires linear computational time $\mathcal{O}(N_o)$ in the worst case, but can be fully parallelised because assessing whether a signal is statistically significant can be done independently of other signals. These signals are then stacked together into a rectangular data matrix with N_T columns and N_o rows i.e. $\mathbf{D} \in \mathbb{R}^{N_o \times N_T}$ such that $[\mathbf{D}]_{kn} = f_k(n)$ are row vectors, and the column dimension can be assumed to be much larger than the row dimension $N_o \gg N_T$.

Further Dimensionality Reduction

At this point, one could employ generalised coherence techniques in classical signal processing [297–299] to extract common frequencies from the data matrix $\mathbf{D}$ directly. But since $\mathbf{D}$ might still contain a large number of observables $N_o \gg N_T$, doing so will likely be too prohibitive. We thus choose to then perform an additional dimensionality reduction step that is analogous to well-established subspace methods in signal processing, such as the MUSIC algorithm [300]. In these methods, one estimates a square signal autocorrelation matrix $\mathbf{C}$, then constructs a signal subspace that is spanned by the eigenvectors of the matrix $\mathbf{C}$ with the largest c eigenvalues and a noise subspace spanned by its complement, from either of which signals are extracted.

In the method we reported, we first obtain a formal equivalent of this signal autocorrelation matrix $\mathbf{C} \in \mathbb{R}^{N_T \times N_T}$ via the matrix-matrix product $\mathbf{D}^T\mathbf{D}$ as

$$[\mathbf{C}]_{mn} = \frac{1}{N_o}\sum_{k=1}^{N_o}[\mathbf{D}]_{mk}[\mathbf{D}]_{kn} = \frac{1}{N_o}\sum_{k=1}^{N_o} f_k(m)f_k(n), \tag{6.6}$$

which can be performed in linear time as $\mathcal{O}(N_o N_T^2)$. We then need to find the c dominant eigenvectors $v_1, v_2, \dots v_c$ of $\mathbf{C}$ to form the signal subspace – we can understand why these eigenvalues are important by considering a time signal $v(n) \in \mathbb{R}^{N_T}$, for which we calculate the expression with respect to the square matrix $\mathbf{C}$ as

$$\begin{aligned} N_o\, v^T\mathbf{C}v &= \sum_{m,n=1}^{N_T} v(m)C_{mn}v(n) \\ &= \sum_{m,n=1}^{N_T}\sum_{k=1}^{N_o} f_k(m)f_k(n)v(m)v(n) \\ &= \sum_{k=1}^{N_o}[\sum_{m=1}^{N_T} f_k(m)v(m)][\sum_{n=1}^{N_T} f_k(n)v(n)] \end{aligned}$$

I now use the notation $\langle f_k, v\rangle := \sum_{n=1}^{N_T} f_k(n)v_n$ for scalar products between the k-th observable's signal f_k and the probe signal v – this gives

$$v^T\mathbf{C}v = \frac{1}{N_o}\sum_{k=1}^{N_o}|\langle f, v\rangle_k|^2 \tag{6.7}$$

as the average squared overlap between the individual signals and the probe signal $v(n)$. It can be seen, therefore, that the dominant eigenvector v_1 of the real, symmetric matrix $\mathbf{C}$ must have the largest average squared overlap with the individual signals as

$$v_1^T\mathbf{C}v_1 = \max_{\|v\|_2=1} v^T\mathbf{C}v = \max_{\|v\|_2=1}\frac{1}{N_o}\sum_{k=1}^{N_o}|\langle f, v\rangle_k|^2.$$

The complexity of an exact diagonalization is $\mathcal{O}(N_T^3)$ which is independent on the number of observables – it only depends on the number of timesteps N_T. Overall, the complexity of dimensionality reduction is thus $\mathcal{O}(N_T^3 + N_T^2 N_o)$.

Spectrum Extraction

We can finally obtain the shadow spectrum by essentially Fourier transforming the dominant eigenvectors $v_1, v_2, \ldots v_c$ that form the signal subspace. My collaborators found that, in practice, extracting the spectral density from the eigenvectors gives a better signal-to-noise ratio than direct Fourier transform.

The cross correlation $X_{kl}(m)$ between the individual signals $v_k(n)$ and $v_l(n)$ is a signal with temporal index m which quantifies the overlap between the signals $v_k(n)$ and $v_l(n)$ after a time lag of m:

$$X_{kl}(m) = \sum_{n=1}^{N_T-m-1} v_k(n+m)v_l(n).$$

We calculate the spectral cross correlation matrix $\mathbf{X}(\omega_n)$ between the set of eigenvectors $v_1, v_2, \ldots v_c$ as

$$[\mathbf{X}(\omega_n)]_{kl} = \mathcal{F}[X_{kl}](\omega_n), \tag{6.8}$$

where we use the discrete Fourier Transform, or the Fast Fourier Transform, of the cross correlation signals. If the vectors $v_1, v_2, \ldots v_c$ were infinitely long, then elements of the spectral cross correlation matrix are simply products of the individual Fourier transforms as

$$[\mathbf{X}(\omega_n)]_{kl} \propto [\mathcal{F}[v_k](\omega_n)]^* \times \mathcal{F}[v_l](\omega_n) \tag{6.9}$$

and the matrix $\mathbf{X}(\omega_n)$ would be rank one with only 1 non-zero singular value $\lambda(\omega_n) = \sum_k |\mathcal{F}[v_k](\omega_n)|^2$ which is just the squared sum of the individual spectra. In practice, $v_1, v_2, \ldots v_c$ are noisy and finite, and so the matrix $\mathbf{X}(\omega_n)$ is not simply rank 1. By calculating the dominant singular values $\lambda(\omega_n)$ of the $\mathbf{X}(\omega_n)$ at each frequency, we can estimate the spectral density of the ideal noise free component, which is our shadow spectrum.

The computational complexity of this step is as follows. Calculating the cross correlations between the signals takes $\mathcal{O}(c^2 N_T)$ time while calculating the matrices $\mathbf{X}(\omega_n)$ through Fast Fourier Transforms takes $\mathcal{O}(c^2 N_T \log N_T)$. Finally, obtaining the dominant singular values scales, in the worst case, as $\mathcal{O}(c^3 N_T)$. Therefore, the worst-case scaling of this procedure is $\mathcal{O}(c^3 N_T)$.

6.2.3 Summary of Shadow Spectroscopy

I now summarise the steps of shadow spectroscopy and its computational complexity:

1. Measure N_o time-dependent signals with classical shadows on a quantum computer, with sample compelxity $\mathcal{O}(\log N_o)$.

2. Optionally reject signals that are not statistically different from noise via e.g. a Ljung-Box test – this scales in the worst case as $\mathcal{O}(N_o)$ but can be trivially parallelised.

3. Stack the signals as row vectors into a matrix $\mathbf{D}$.

4. Calculate the square matrix $\mathbf{C} \propto \mathbf{D}^T\mathbf{D}$ – this is the time-determining step with a complexity $\mathcal{O}(N_o)$.

5. Find the dominant eigenvectors $v_1, v_2, \ldots v_c$ of $\mathbf{C}$ – we need to choose a cutoff, e.g, $c = 4$ in Fig. 6.2, based on the eigenvalue distribution of $\mathbf{C}$.

6. Finally, we effectively Fourier transform these dominant eigenvectors; we obtain the spectral (Fourier) intensity at frequency ω by calculating the spectral cross correlation matrix $\mathbf{X}(\omega)$ from the dominant eigenvectors $v_1, v_2, \ldots v_c$

7. The spectrum $\lambda(\omega)$ is then obtained as the principal component of each matrix $\mathbf{X}(\omega)$.

6.3 Key Features of Shadow Spectroscopy

My collaborators and I found that, as a result of combining classical shadow measurements and time-signal processing as described above, shadow spectroscopy has a number of features that make it particularly favourable for mitigating limitations of quantum computers, and specific applications in many-body physics.

6.3.1 Robustness to shot noise

For example, my collaborators built analytical models to understand how statistical shot noise due to finite circuit repetition affects the time signal intensity and the dimensionality reduction, and thus the resulting shadow spectrum. Their models suggest that the shot noise is essentially suppressed as the number of observables increases: Specifically, they found that the signal-to-noise ratio (SNR) of a shadow spectrum peak i.e. the gap between the spectral baseline and the peak, relative to the standard deviation of the baseline, is

$$\mathrm{SNR} \propto N_T N_s \overline{I^2} \sqrt{N_o}, \tag{6.10}$$

where $\overline{I^2} := \frac{1}{N_o} \sum_{i=1}^{N_o} I_i^2$ is the average intensity in the group of N_o observables. They also found that the dimensionality reduction step is not expected to significantly influence this result.

Quantum protocols in which the required quantum resources scale inversely proportionally with the precision achievable are often refereed to as admitting Heisenberg scaling. The precision of energy difference estimation in shadow spectroscopy is expected to exhibit Heisenberg scaling in the sense that the evolution time T required for an energy gap Δ to be resolved is $T \sim 1/\Delta$; it follows that quantum resources (both circuit depth and the number of shots required – assuming at least one shot taken per time step) scales similarly as $\sim 1/\Delta$. This is in contrast with the standard shot-noise scaling of energy estimation in VQE, where the precision scales with the inverse square root of the number of shots.

6.3.2 Robustness to algorithmic or gate errors

It is also crucial to consider how the ideal signal, as the time-dependent expected value of an observable $S(t_n) := \langle\psi(t_n)|P_k|\psi(t_n)\rangle$, is altered when the quantum states $|\psi(t_n)\rangle$ are not time-evolved exactly: this can be because they are prepared with a noisy quantum circuit and/or the time evolution unitary is approximated. My collaborators presented a proof which states that, under a broad class of typical

gate-error models, the noisy spectrum $\mathcal{F}[S_k^{noisy}]$ decomposes into a sum of the ideal spectrum $\mathcal{F}[\eta S_k]$ and an additive $\mathcal{F}[(1-\eta)W_k]$ artefact.

The additive artefacts $\mathcal{F}[(1-\eta)W_k]$ have magnitudes bounded by how close the noise model is to global depolarising (white) noise. They are specific to each Pauli string P_k and, assuming they are uncorrelated, their effect is suppressed by a factor N_o in the shadow spectrum. In addition, these bounds on the magnitude of $\mathcal{F}[(1-\eta)W_k]$ decrease with an increasing number of qubits. On the other hand, the value of η in the additive factor is either a constant or a time-dependent exponential envelope, depending on how the time evolution is approximated – and this is expected to actually offset the peak position from the exact energy difference.

Critically, this result guarantees that the spectral peak positions are not shifted by reasonable levels of incoherent noise or indeed a broad class of noise models, The noise may slightly shrink the intensity of the ideal peaks, but does not cause a broadening of their lineshape. My experimental demonstration of shadow spectroscopy measurements on current noisy quantum hardware in Section 6.4 would also suggest that shadow spectroscopy is equally robust against errors that go beyond the noise models considered.

6.3.3 Compatibility with Quantum Chemistry

Finally, I want to highlight how shadow spectroscopy is particularly amiable to answering questions about the excited state profile of Fermionic many-body systems, specifically those in quantum chemistry.

Fermionic operators

Consider the example of a transition between the electronic ground state and the first excited state of a molecule. The peak intensity with respect to an observable O (Statement 2) is calculated via an expression formally resembling the conventional transition dipole moment as

$$I_{g\to e} \propto \langle \psi_g | \, O \, | \psi_e \rangle \, . \tag{6.11}$$

I will expand both the exact ground and excited states as a Configuration Interaction (CI) expansion from quantum chemistry, i.e., a sum of the HF Slater determinant and post-HF corrections:

$$|\psi_s\rangle = c_0^{(s)} |\mathrm{HF}\rangle + \sum_{pq} c_{pq}^{(s)} a_p^\dagger a_q |\mathrm{HF}\rangle + \sum_{pqrs} c_{pqrs}^{(s)} a_p^\dagger a_q^\dagger a_r a_s |\mathrm{HF}\rangle + \ldots,$$

where $|\mathrm{HF}\rangle$ is the mean-field Hartree-Fock state [1] and $a_p^\dagger a_q, a_p^\dagger a_q^\dagger a_r a_s$ are single- or double-excitation Fermionic operators respectively. Thus Eq. (6.11) can be expanded as

$$\begin{aligned}\langle\psi_g| O |\psi_e\rangle =& c_0^{(g)*} c_0^{(e)} \langle\mathrm{HF}| O |\mathrm{HF}\rangle \\ &+ c_0^{(g)*} \sum_{pq} c_{pq}^{(e)} \langle\mathrm{HF}| O a_p^\dagger a_q |\mathrm{HF}\rangle \\ &+ c_0^{(e)} \sum_{pq} c_{pq}^{(g)*} \langle\mathrm{HF}| a_q^\dagger a_p O |\mathrm{HF}\rangle + \ldots\end{aligned} \tag{6.12}$$

I only go up to the single-excitation CI terms, but it is clear that if O cancels out the Fermionic excitation operators, we get $\langle\mathrm{HF}|\,\mathrm{HF}\rangle = 1$ and the corresponding CI coefficients would contribute to $\langle\psi_g| O |\psi_e\rangle$.

If we choose O as the Hermitian operator that corresponds to the single excitation $O = a_m^\dagger a_n + a_n^\dagger a_m$ with m acting on occupied and n acting on virtual orbitals, the leading terms in Eq. (6.12) can be evaluated analytically as

$$\begin{aligned}\langle\mathrm{HF}| O |\mathrm{HF}\rangle &= \langle\mathrm{HF}| a_m^\dagger a_n + a_n^\dagger a_m |\mathrm{HF}\rangle = 0, \\ \langle\mathrm{HF}| O a_p^\dagger a_q |\mathrm{HF}\rangle &= \langle\mathrm{HF}| a_m^\dagger a_n a_p^\dagger a_q |\mathrm{HF}\rangle + 0 = \delta_{np}\delta_{mq}, \\ \langle\mathrm{HF}| a_q^\dagger a_p O |\mathrm{HF}\rangle &= 0 + \langle\mathrm{HF}| a_q^\dagger a_p a_n^\dagger a_m |\mathrm{HF}\rangle = \delta_{np}\delta_{mq}.\end{aligned} \tag{6.13}$$

Thus, in leading order, the intensity of the observable $\langle\psi_g| a_m^\dagger a_n + a_n^\dagger a_m |\psi_e\rangle$ is determined by the corresponding single-excitation coefficients $c_0^{(g)*} c_{mn}^{(e)}$ and $c_0^{(e)} c_{mn}^{(g)*}$, as all other terms shown explicitly in Eq. (6.12) give null contributions. Of course,

[1] Typically the HF solution has the dominant weight as $|c_0^{(g)}| > |c_{pq}^{(g)}|$ (even if the HF Slater determinant is not a good approximation) and thus in Eq. (6.11) the first term in the expansion is the dominant term.

higher order CI terms in Eq. (6.12) can also contribute, though they might have lower contributions.

If we instead choose O to be a higher-order excitation Hermitian Fermionic operator, the intensity of the corresponding signal will be determined in leading order by its higher-order CI coefficients in Eq. (6.12). For example, by choosing $O = a_p^\dagger a_q^\dagger a_r a_s + a_s^\dagger a_r^\dagger a_q a_p$, the intensity in leading order is determined by the corresponding expansion coefficient $c_{pqrs}^{(s)}$.

Thus, if we are able to simultaneously estimate expected values of all q-local Fermionic excitation operators, we would get a large number of signals but the intense signals will come from Fermionic operators that have dominant expansion coefficients in Eq. (6.12). While it is possible to use machinery like Fermionic shadows [175, 176] to obtain expectation values for these q-local Fermionic excitation operators, I now show how, in fact, using the original near-term friendly classical shadow measurement approach to reconstruct q-local Pauli observables (even for $q \sim 4$) is sufficient to reveal single electron transitions between meaningful eigenstates.

Jordan-Wigner Encoding

Under the JW-transform, excitation operators map to non-local Pauli strings e.g. the single excitation $a_p^\dagger a_q \mapsto i A_p^\dagger A_q \prod_{k=p+1}^{q-1} Z_k$, where $A = (X+iY)/2$ is the qubit lowering operator. In spite of this, we find that in practice the signal intensities of observables like 2-local Pauli strings e.g. X_pX_q, X_pY_q, or indeed any variant where we additionally append one or more single-qubit Z operators, yield dominant signals. The reason is that the intensity of these local observables is determined directly by the single-excitation coefficients in the Configuration Interaction expansion of the time-evolved states, which is the same as if the observable $a_p^\dagger a_q$ was measured instead.

Statement 4. *Choosing the Pauli bitflip observable $O = X_mX_n$, the signal intensity $\langle\psi_g| X_mX_n |\psi_e\rangle$ is determined in leading order by the single-excitation expansion coefficients $c_0^{(g)*}c_{mn}^{(e)}$ and $c_0^{(e)}c_{mn}^{(g)*}$ of the state from Eq.* (6.13)*. Similarly, if we choose any other operator $O(m,n)$ by replacing any of the two Pauli X operators in O with*

Pauli Y operators and we additionally append Pauli Z operators on any qubit other than m or n then the intensity is still determined by the same expansion coefficient.

Proof. In general, the excitation operators in the JW encoding are mapped to

$$a_p^\dagger \mapsto \alpha_p^\dagger := Z_{:p} A_p^\dagger, \qquad a_p \mapsto \alpha_p := Z_{:p} A_p$$
$$a_p^\dagger a_q \mapsto i A_p^\dagger Z_{q:p} A_q$$

where $Z_{:p} := \prod_{k=1}^{p-1} Z_k$ and $A = (X + iY)/2$ is the qubit lowering operator and $Z_{q:p} = \prod_{k=p+1}^{q-1} Z_k$. Furthermore, the HF state becomes a computational basis state $|\mathrm{HF}\rangle \mapsto |11\ldots1100\ldots\rangle \equiv |b\rangle$ where b is simply a binary number.

Thus for any Pauli string $O = \bigotimes_{k=1}^{N} P_{j_k}$ the expected value in Eq. (6.13) splits up as a product of single-qubit expected values as

$$\langle b| O(m,n) \alpha_p^\dagger \alpha_q |b\rangle = \langle b| O(m,n) A_p^\dagger Z_{q:p} A_q |b\rangle$$
$$= \langle b_p| P_{j_p} A^\dagger |b_p\rangle \langle b_q| P_{j_q} A |b_q\rangle \prod_{k\neq p, k\neq q} \langle b_k| P_{j_k} Q_{pq}(k) |b_k\rangle ,$$

where p_p is the p^{th} bit of the binary number b and $Q_{pk}(k) = Z$ if $q < k < p$ and the identity operator otherwise. We can evaluate the product of coefficients as

$$|\langle b_p| P_{j_p} A^\dagger |b_p\rangle \langle b_q| P_{j_q} A |b_q\rangle| = \delta_{np}\delta_{mq}/4,$$

assuming that $m < n$ without loss of generality and we only show the absolute value of the expression for ease of notation. The above expression evaluates (in absolute value) to 1 only when the two Pauli X (or Y) operators in the observable $O(m,n)$ act on the same two qubits where A and $A^\dagger$ act, i.e., on qubits p and q. Furthermore, $\prod_{k\neq p,k\neq q} \langle b_k| P_{j_k} Q_{pq}(k) |b_k\rangle \in \{\pm 1\}$ is only a sign factor that is determined by the bitstring b and by the number of Pauli Z operators in P and in $Z_{p:q}$. Thus we conclude that the expected values from Eq. (6.13) in the JW encoding evaluate to

$$\langle b| O(m,n) |b\rangle = 0, \tag{6.14}$$
$$|\langle b| O(m,n) \alpha_p^\dagger \alpha_q |b\rangle| = \delta_{np}\delta_{mq}/4$$
$$|\langle b| \alpha_q^\dagger \alpha_p O(m,n) |b\rangle| = \delta_{np}\delta_{mq}/4.$$

The signal intensity $\langle\psi_g| O(m,n) |\psi_e\rangle$ of the observable $O(m,n)$ is determined in leading order by the single-excitation expansion coefficients $c_0^{(g)*}c_{mn}^{(e)}$ and $c_0^{(e)}c_{mn}^{(g)*}$ of the state from Eq. (6.13). □

The crucial observation we make in the above statement is that even though Fermionic excitation operators $a_p^\dagger a_q \mapsto iA_p^\dagger Z_{q:p} A_q$ are mapped to non-local Pauli strings in the JW encoding, we can still resolve signals of the corresponding operators just by measuring 2-local Pauli strings e.g. X_pX_q or any variant where we additionally append Pauli Z operators.

This insight can be used to inform the user which specific Pauli string observable to reconstruct. Furthermore, one could use e.g. classical MP2 simulations to estimate which excitations will be dominant, and thus fetch the corresponding families of Pauli strings. If we know which operators to measure, they can be ordered into a small number of commuting groups, and then measure them simultaneously using well-established techniques e.g. Ref. [194]. Of course we need not perform any *a priori* prediction of expected operator contributions, because classical shadows allow us to estimate the signal of every q-local Pauli string very efficiently (as the data matrix $\mathbf{D}$) and the statistical autocorrelation test can isolate meaningful signals for us.

6.3.4 Comparison with other quantum algorithms

In shadow spectroscopy, the intensity of the signals is expressed as $I_{kl} = c_k^* c_l \langle\psi_k|O|\psi_l\rangle$. As such, we aim to apply operator pools that give non-negligable overlaps $\langle\psi_k|O|\psi_l\rangle$ – this is loosely analogous with subspace expansion (see Chapter 2) where similar operators are required. However, the difference in efficiency between the present approach and subspace expansion is crucial; in subspace expansion one estimates the overlaps one-by-one resulting in a linear complexity $\mathcal{O}(N_o)$ whereas the effective use of classical shadows reduces the quantum resources to only logarithmic $\mathcal{O}(\log N_o)$. Nevertheless, literature on subspace expansion will be relevant in identifying problem specific observables O that give significant overlap considered in Section 6.3.3, e.g., excitation operators in quantum chemistry or certain local observables in spin problems. Furthermore, a significant advantage of the present approach is that

explicit knowledge of which operators give intense signals is not necessary as we can efficiently reconstruct all q-local Pauli strings and identify the useful signals via statistical autocorrelation tests. Compared with Krylov subspace approaches, our approach likewise features time evolution of an initial state as well as classical signal post-processing and matrix diagonalisation. However, it does not have the burden of evaluating Krylov matrix elements, nor does it require an initial state with good ground state overlap.

Phase estimation techniques (see Chapter 2) estimate a single signal that requires non-trivial quantum resources, such as controlled time-evolution and state-preparation oracles, and then perform classical post-processing whose computational cost is cheap. In contrast here we only use logarithmic quantum resources to cheaply estimate a large number of signals that we then classically analyse in a non-trivial post-processing step.

6.4 Experimental Validation: NISQ Scenario

To experimentally verify these theoretical results, my collaborators and I first chose to demonstrate the measurement of a shadow spectrum on NISQ-level hardware with a 1D Heisenberg ring system with random local magnetic fields and nearest-neighbour interactions

$$\mathcal{H} = -J \sum_{j=1}^{N-1} \vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + \sum_{j}^{N} h_j \sigma_j^z, \tag{6.15}$$

($\vec{\sigma}_j = [\sigma_j^x, \sigma_j^y, \sigma_j^z]$ is a vector of Pauli x, y and z matrices on the jth qubit, J is a coupling constant, $h_j \in [-h, h]$ is sampled from a uniform distribution for disordered strength h, and periodic boundary conditions are imposed). Heisenberg spin systems are relevant in studies of many-body localization and self-thermalization [301, 302], but their phase transitions are not well understood due to the challenge of classically emulating these systems for larger sizes; even if, compared with e.g. Fermionic systems, it is more straightforward to represent on quantum computers. It is thus a promising candidate for demonstrating quantum advantage even on early NISQ-level hardware.

6.4.1 Proof-of-Concept

My collaborator and I started with performing exact time evolution on spin chains and computing their shadow spectrum numerically, using QuESTlink[179] and pyQuEST[268]. This allowed us to test whether or not the idea made sense, independent of how a particular initial state might be prepared, or the influence of algorithmic and gate errors. We always started from an initial state that is assumed to have a fixed overlap with the first three eigenstates of the n-qubit Hamiltonian as $|\psi(0)\rangle \propto (1, \frac{1}{10}, \frac{1}{10}, 0, 0 \dots)$, which was evolved in $\mathcal{O}(10^2)$ time steps. We then estimated l-local Pauli observables at each timestep using classical shadows formed from between $10 - 100$ snapshots, and applied the post-processing steps described above.

The results from my 14-qubit simulation, with $N_T = 200$ time evolution steps, nicely illustrates the classical post-processing steps for extracting shadow spectra (Fig. 6.2). In the experiment shown, every at-most 3-local Pauli string ($N_o = 10689$) was estimated from classical shadows of $N_s = 100$ snapshots at each time step – forming the data matrix $\mathbf{D}$. Converting that into $\mathbf{C} \propto \mathbf{D}^T\mathbf{D}$, it can be seen that time-periodic correlations are clearly visible. With the subsequent post processing steps, peaks in the spectrum of the dominant eigenvectors of $\mathbf{C}$ reveal energy differences in the Hamiltonian. The rightmost panel of Fig. 6.2 also illustrates the theoretical predictions of the signal-to-noise ratio; when the number of snapshots was reduced to $N_s = 10$ and the time steps increased proportionally to $N_T = 2000$, the signal to noise ratio remain unchanged because the product $N_T N_s$ stayed the same.

From the same set of numerical signals, I randomly selected 100 groups of N_o observables out of the $\tilde{N}_o = 10689$ at most 3-local Pauli strings ($N_o < \tilde{N}_o$), to assess how the SNR relates to the number of observables used. The expected value of the average signal intensity for groups of random N_o signals is $\mathbb{E}[\overline{I^2}] = \frac{1}{\tilde{N}_o} \sum_{i=1}^{\tilde{N}_o} I_i^2$ i.e. the average over the full set, which is a constant and is independent of the number of selected terms. The average of the SNR, shown in Fig. 6.3(a), confirms that it indeed increases with $\propto \sqrt{N_o}$ as my collaborator predicted in Eq. (6.10).

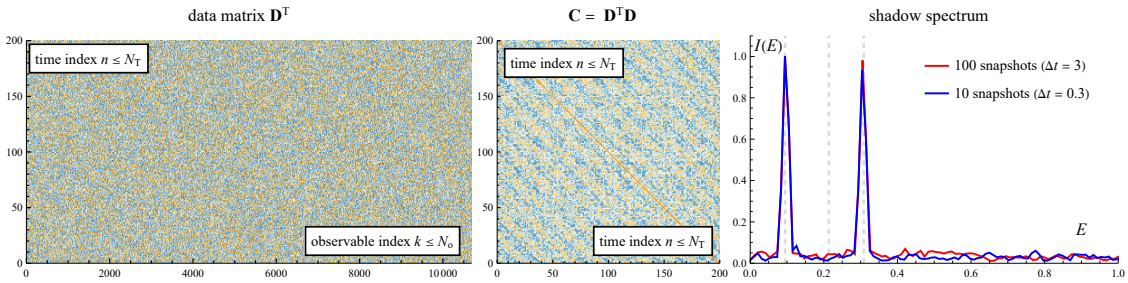


Figure 6.2: Classical post-processing step in shadow spectroscopy. (Left) The discrete time signal of all 3-local Pauli observables for a 14-qubit spin ring, reconstructed from classical shadows at each of the N_T time steps. (Centre) Converting $\mathbf{D}$ into $\mathbf{C} \propto \mathbf{D}^T\mathbf{D}$ to reveal harmonics – note that no autocorrelation test was used to pre-screen low intensity signals. (Right) The shadow spectrum is obtained by effectively Fourier transforming the dominant eigenvectors of $\mathbf{C}$ (see details in main text). The dashed lines are exact energy gaps as $E_0 - E_1$ and $E_0 - E_2$. A third peak of $E_1 - E_2$ does not appear as it is suppressed due to the initial state's small support for the two eigenstates. Figure reproduced from Ref. [5] – created by Bálint Koczor and me.

Such a random selection of observables can naturally be considered as a worst-case scenario when no information about the observables is assumed – we, of course, made well-informed decisions regarding which signals to process based on e.g. the aforementioned Ljung-Box test. I therefore next tested the SNR's dependence on how many of the statistically most significant and intense signals are used to generate a shadow spectrum. Performing a Ljung-Box test on each signal isolated about 2000/10689 signals that were distinguishable from pure shot noise i.e. signals whose Ljung-Box test returned statistical p values satisfying $p < 0.01$. These were then ranked, and even by just picking $\sim 10^2$ of the highest-rank signals we can get a significant boost to the SNR, as shown in Fig. 6.3(b). In this example, the SNR improves up to an optimal point at $N_o = 407$, beyond that point the SNR starts to decrease as the intensity of the signals with higher p values decreases. However, we note that, where using more observables allow for the resolution of additional lower-intensity peaks, one may not mind a decreased average SNR.

I then reduced the system size one qubit at a time, and used the same strategy (exact time evolution of initial states that have fixed overlaps with the three lowest-energy eigenstates, and reconstruction of up-to 3-local Pauli observable signals from shadows) to explore the scalability with respect to qubits. For each system

size, I repeated the previous analysis used to retrieve the SNR as a function of significant signals used (i.e. the curve of Fig. 6.3 (b)), and computed the average SNR and its standard deviation. I found that the average intensity among local observables actually grows favourably with the number of qubits. While this suggests a reassuring scalability of the approach, we did always start with initial states that have a constant overlap with the first 3 eigenstates – which would, in the worse case, be exponentially difficult to guarantee for general Hamiltonians given.

We did additional experiments to verify that up-to 3-local Pauli strings are well-suited for shadow spectroscopy. When I used only up-to 1-local Pauli observables for each simulation, the number of signals that have p values less than 0.01 only grew as $\mathcal{O}(N^{0.4})$; this confirmed that not all 1-local Pauli operators yield meaningfully intense signals, as expected from the entangled nature of the spin ring's low-lying eigenstates. In contrast, measuring all up-to 3-local Pauli strings yields a significant growth in the number of statistically significant signals as $\mathcal{O}(N^{2.5})$ – see Fig. 6.3. This matched up with the SNR data; most of the statistically significant signals are, in fact, 3-local Pauli observables. In all subsequent experiments in this work we therefore kept reconstructing only up-to 3-local Pauli observables. Prior results also suggest that local operators indeed yield dominant signal intensities, e.g., in translation invariant, gapped Hamiltonians of spin systems the intensity $|\langle\psi_k|O_i|\psi_l\rangle|$ is an exponentially decreasing function of the locality of O_i [303].

Finally, from an exact evolution sequence of fixed $N_T = 1000$, we estimated the absolute time required to perform the entire classical signal processing step for an increasing number of observables reconstructed. I found that the computation time of the sequence is asymptotically linear in the number of observables – see Fig. 6.3(e) – which is in good agreement with my collaborator's prediction of a polynomially-scaling classical post-processing. Deviation from linearity is expected at low N_o because calculating the eigenvectors of the constant-sized matrix $\mathbf{C}$ dominates the time. The absolute time can be quite reasonable even for a very large number of observables, i.e., less than an hour computed on a desktop PC even for $N_o \propto 10^8$.

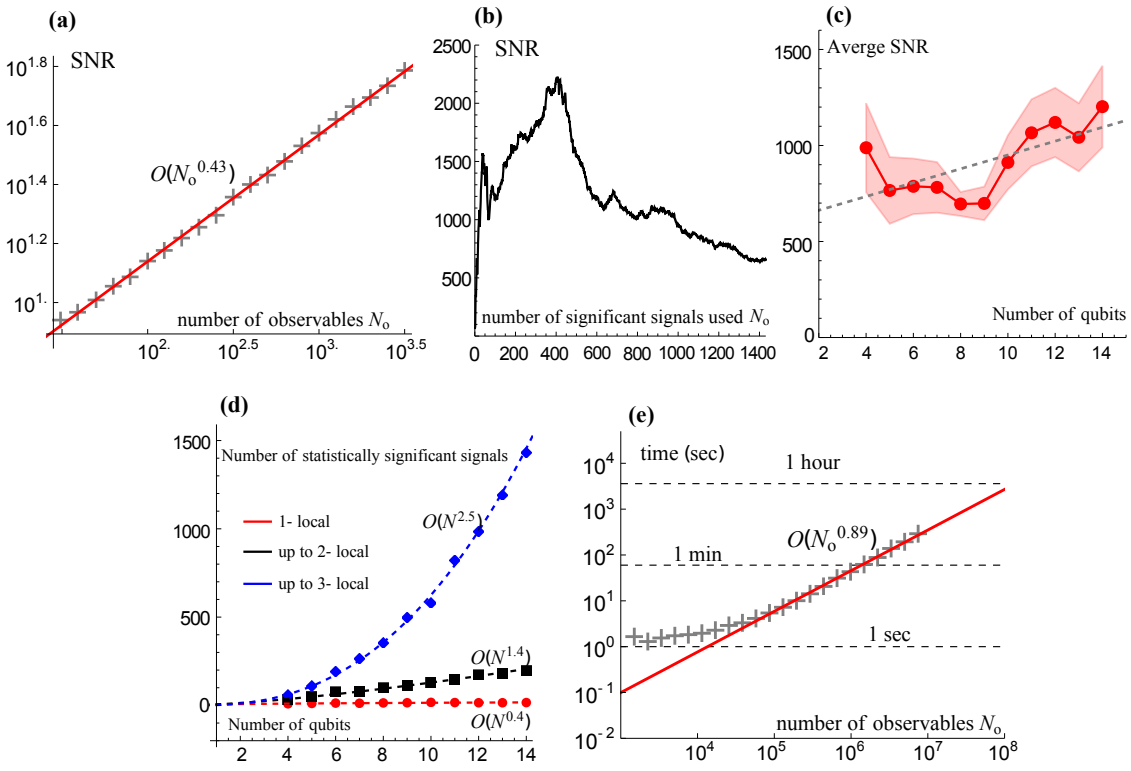


Figure 6.3: Scaling properties of shadow spectroscopy. (a) The reconstructed Pauli observable signals from a 14-qubit spin ring simulation were randomly selected into groups of N_o to form shadow spectra. The SNR of these spectra improved with the number of signals used as $\propto \sqrt{N_o}$. (b) When the signals from the 14-qubit simulation were instead ranked in terms of their statistical p values according to the Ljung-Box test, the SNR of the shadow spectrum was boosted significantly even when using only ~ 100 of the most statistically significant signals. The SNR reaches a maximum value then decays, as more signals with lower intensity are used. (c) The average SNR of shadow spectra (red dots) when using statistically significant signals with respect to an increasing qubit count – the final 14-qubit data point is the average of the SNR values in panel (b). The red shading represents the standard deviation. (d) The number of signals with a Ljung-Box test statistical significance $p < 0.01$ using all 1-local Pauli strings (red) or using up to all 3-local Pauli strings (blue). (e) The computation time of the signal post-processing step is asymptotically linear in the number of observables. Computed on a desktop PC at a fixed $N_T = 1000$ and increasing N_o. All plots reproduced from Ref. [5] – created by Bálint Koczor and me.

6.4.2 Hardware Demonstration

Motivated by these results, we decided to move towards demonstrating shadow spectroscopy on existing quantum hardware. Because NISQ devices generally have limited circuit depth and gate connectivity, we had to approximate time-evolved states of spin systems by an ansatz $\left|\psi(\vec{\theta})\right\rangle = U(\vec{\theta})\,|\psi(t=0)\rangle$, where $\vec{\theta}$ is a vector

of real classical parameters and $U(\vec{\theta})$ is a 5-layer hardware-efficient ansatz circuit with a fixed circuit depth – see Fig. 6.4 (a) for one such layer. My collaborator first implemented variational quantum simulation [304, 305] numerically to find $U(\vec{\theta})$ that approximates every time step; a quantum computer would have been used to estimate a small matrix A_{ij} and a small vector C_i is at each time step, and the resulting linear system of equations would then be solved classically to obtain an updated set of parameters $\vec{\theta}$ that best approximate the time-evolved state – an emulated noiseless quantum computer was used instead. The initial state at $t = 0$ was obtained using a variational imaginary-time evolution approach [156, 306, 307] which started from random ansatz parameters and was intentionally terminated before it optimised to the ground state; this resulted in an initial state that overlaps mostly with the ground state, with some support for additional low-lying excited states. The variational time evolution was performed for 1000 time steps and, once the parameterisation for each time-evolved state was found, up-to 3-local Pauli operators were then reconstructed from 3 batches of N_s=50 shadows snapshots for each time-evolved state. The numerical results my collaborator obtained revealed a clear peak corresponding to the energy gap between the ground and first excited state – see Fig. 6.5.

From these initial results, I then took the above classically pre-computed time-evolution circuits representing the 1000 time-evolved states, and prepared them on the 7-qubit superconducting quantum computer IBMQ Jakarta (Falcon r5.11H processor) and measured their classical shadows. This was done through the IBM Quantum Platform [308], where the parameterised ansatz and randomised bases measurement circuits were generated in Qiskit [309]. At each time step, 3 batches of 50 snapshots classical shadow measurements were taken by appending random basis rotations to 150 copies of an ansatz, and requesting only a single shot measurement for each circuit via the Qiskit Runtime `Sampler` primitive. Settings of all `Sampler` primitives were taken as defaults - specifically the resilience level argument was set to 1 (minimal mitigation costs) and optimisation level set to 3 (even heavier optimisation). Because of constraints in Qiskit's job queuing system, I grouped all

the shadow measurements at one time step into a standalone Qiskit `Job` i.e. 150 `Sampler` primitives each. Every 100-`Job` were then bundled into a `Session`, resulting in 10 separate submitted `Sessions` in total. To avoid device drift, I structured the `Jobs` such that 1 of every 10 time-steps were grouped into one Session (i.e. `Jobs` indexed 0, 10, 20,..., 90 and 1, 11, 21..., 91 etc. would correspond to the first two Sessions) – see Fig. 6.4(c). From this shadows data we determined expected values of all 3-local Pauli strings and performed a Ljung-Box test, retaining only the highest scoring 50 signals. The correlation matrix and its eigenvectors were then computed and only the 5 eigenvectors with the highest eigenvalues were retained. The shadow spectrum was then computed from these 5 eigenvectors.

The data generation took 1 week in total to complete (total compute time of the IBM numerical simulation was < 0.5 hour). The experimental shadow spectrum computed from the data I collected was extremely noisy compared with the noiseless simulation – see Fig. 6.5. We were certain that gate errors corrupted the time-evolution, resulting in some pronounced artificial peaks not seen in the noiseless simulation, including a prominent peak at $E = 0.0$ and a pair of peaks at $E \sim 1.6$. Nevertheless, the predicted spectral peak corresponding to a transition between the ground state to first excited state is present and has the strongest intensity. Furthermore, while the SNR might not be ideal, the accuracy of the spectral peak matched that of numerical simulation and the exact value of the energy gap. Interestingly, the imperfection in the time evolution also revealed the energy gap between the ground and the second excited state with good accuracy, even if it was a much weaker signal. The figure also shows numerical simulations conducted on the IBM Quantum Platform that use the noise model of the IBM Jakarta chip provided by the IBM. This didn't deviate significantly from the ideal noiseless simulation – this could a result of out-of-date and inaccurate noise models for the quantum chip.

To capture the periodic boundary condition of the Heisenberg ring, the ansatz circuit layers each contain a two-qubit gate that entangle the first and sixth (final) qubit. I posited that, because the quantum computer used is limited to nearest-neighbour qubit connectivity, to implement this gate in practice it had to be

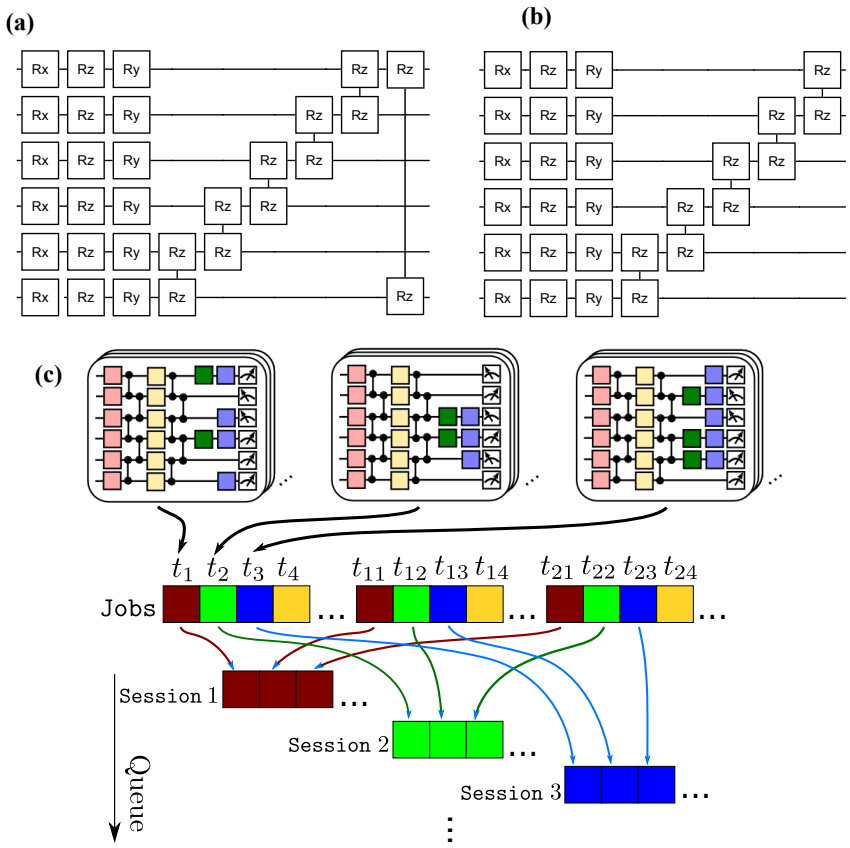


Figure 6.4: Hardware efficient ansatz circuit layers used to approximate spin system evolution, and their measurement on real quantum hardware. Each gate in the ansatz is uniquely parameterised. (a) Single ansatz layer for the spin ring simulation. (b) Single ansatz layer for the spin chain simulation. (c) I was only able to perform simulations on a single quantum device, which meant that every shadow snapshot was measured in series. Therefore, to wash out any potential device drift in the time-signal of the observables, I structured the job submission such that the shadow data for time steps t_1, t_{11}, t_{21}... were first measured, forming a sparse signal with time resolution of 10 instead of 1. This was followed by collecting the data for t_2, t_{12}, t_{22}... until each of the 100 time steps has been measured. Circuit diagrams reproduced from Ref. [5] generated by Matthew Goh. Unpublished schematic on job scheduling in the bottom panel created by me.

decomposed into multiple native CNOT gates between qubit-pairs of neighbouring index. As each CNOT gate has a given error rate, the state-preparation circuit must have accumulated significant noise before the shadow measurement.

We therefore decided to repeat the same numerical simulations and hardware experiments with a 6-site 1D Heisenberg chain instead, without the periodic boundary condition in the Hamiltonian and the corresponding problematic two-qubit gates in the hardware efficient ansatz – see Fig. 6.4(b) for the circuit. Almost all the exact same steps were taken as the last experiment, except that at each

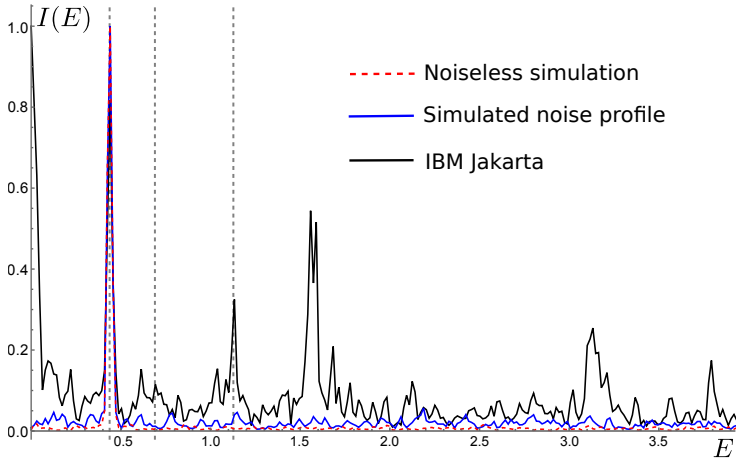


Figure 6.5: Experimental shadow spectrum of a 6-qubit Heisenberg spin ring obtained using the IBM Jakarta QPU compared with simulation. The spectra were normalised with respect to the dominant non-trivial signal. The grey vertical dashed lines represent the energy gaps between the first three eigenstates of the Hamiltonian. The noise profile spectrum (blue line) was obtained with numerical simulations on the IBM Quantum platform which uses the noise profile of the target IBM Jakarta computer. Unpublished figure created by Bálint Koczor and me.

time step 220 classical shadow measurements were taken instead. However, for reconstructing the shadow spectrum, we still only used 3×50 randomly selected shots (split into 3 batches of 50 snapshots at each timestep) for reconstruction of 3-local Pauli observables.

The data collection took about 9 days to complete. As predicted, because of the lower error rates in preparing time-evolved states, the experimental shadow spectrum I obtained in Fig. 6.6 resolves the main energy gap with much greater clarity. Perhaps surprisingly, it actually achieves the same precision in predicting the energy gap as the pure state simulator without using any error mitigation techniques. The second, faint peak seen in simulation corresponding to the transition between the ground to third excited state is not resolved in the experimental spectrum; the initial state's support for $|\psi_3\rangle$ is already very low, and this peak is likely suppressed by algorithmic errors from the approximate, variational time evolution, and here it is likely below the level of shot noise. I also found that the numerical simulation with the IBM-provided noise profile actually gave a weaker spectrum than the actual experimental result. This could again be down to a not up-to-date noise profile. My collaborators performed additional numerical experiments to

show that it would have been resolved if more shots are available (e.g. by taking 4000 timesteps vs 1000 timesteps), even under gate noise. These results clearly confirm our theoretical guarantees that reasonable levels of experimental noise do not affect the position of the peak.

My collaborator and I also analysed the data obtained from all 220 snapshots per time step collected. These snapshot were again split into 3 batches and all up-to 3-local Pauli strings were estimated using classical shadows. The observable expected values shown in the right panel of Fig. 6.6 were estimated the following way. First, an autocorrelation test (Ljung-Box test [295, 296]) was performed on the time-dependent Pauli expected values from noisy simulation and experiment, and the worst-scoring 100 observables were selected, i.e., the ones that are most well described by random fluctuations around the mean. This way we can select observables whose expected values are (nearly) the same at each timestep. We then estimated the time-average of these constant signals. Since 1000 timesteps were recorded with 200 shots at each timestep, overall 2×10^4 shots were available to predict a time-constant expected value, thus the level of shot noise is suppressed to a relatively low level (average standard deviation is 0.01, hence no error bars are shown).

My collaborators showed in a series of analyses that, under a depolarising noise model at a circuit error rate of ξ, every Pauli string's expected value is anticipated to shrink by the same factor $\eta_{exp} \approx e^{-\xi}$, resulting in the global relation between experimental and ideal expected values as $\langle P \rangle_{exp} = \eta_{exp} \langle P \rangle_{id}$. In the right panel of Fig. 6.6, we see that the noisy observable expected values indeed shrank compared to the ideal ones, but not uniformly i.e. some observables are shrunk more than others. This is consistent with the fact that the global depolarising noise model is only a crude approximation to gate errors [310]. Based on this data, I estimated the observables on average to have shrunk by a factor $\eta_{exp} \approx 0.54$ in the experiment and by $\eta_{sim} \approx 0.55$ in simulation, but with a significant variability across the different observables as per the standard deviations $\sigma_{exp} = 0.16$ and $\sigma_{sim} = 0.19$ respectively. The observed average shrink factor is comparable to a circuit error rate $\xi \approx -\ln \eta_{exp} \approx 0.62$.

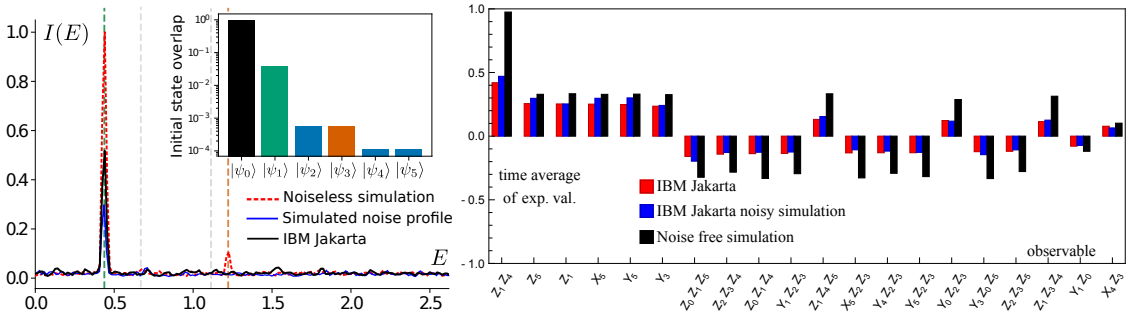


Figure 6.6: Experimental shadow spectrum of a 6-qubit disordered Heisenberg chain obtained using the 7-qubit IBM Jakarta QPU. (Left) The noisy experimental spectrum correctly resolves exact energy gaps (dashed vertical lines) between the ground state and first excited state, albeit with a lower peak intensity – the spectra are not normalised. The error in the gap estimate in the experiment is identical to that in ideal simulation as 2.8×10^{-3}. The peak corresponding to $E_3 - E_0$, which is already predicted to have a low intensity due to the small support ($< 10^{-3}$) the initial state has for $|\psi_3\rangle$, is not visible in the experimental spectrum due to the approximate time-evolution algorithm used. Inset: Overlaps $|\langle\psi_n|\psi(\theta)\rangle|^2$ between the initial state $|\psi(\theta)\rangle$ and eigenstates $|\psi_n\rangle$. (Right) The time average of the most significant Pauli observables, comparing the ideal simulation with the experimental and simulated noise results. Figure reproduced from Ref. [5] – created by Bálint Koczor, Matthew Goh and me.

6.5 Experimental Validation: Fault Tolerant Scenario

I then explored the utility of shadow spectroscopy in the fully fault-tolerant quantum computing regime, where algorithmic errors in time evolution can be suppressed exponentially via sophisticated post-Trotter time-evolution techniques. I did numerical experiments to investigate the discovery of a molecule's low-lying electronic transitions using shadow spectroscopy.

6.5.1 Utility of Local Pauli Operators

First, my collaborator and I tested whether up-to 3-local Pauli operators can actually give meaningful observable time signals for Fermionic systems. We performed a simple estimation by computing the matrix elements $\langle\psi_e| O_i |\psi_g\rangle$ for O_i in the set of up-to 3-local Pauli strings and $|\psi_g\rangle, |\psi_e\rangle$ are the ground and first triplet excited states of the LiH molecule (12-qubits) computed in the minimal basis using pyscf [212]. Fig. 6.7 summarises the 75 Pauli strings O_i that give the largest $\langle\psi_g| O_i |\psi_e\rangle$ values. It can be seen that the most intense observables are the

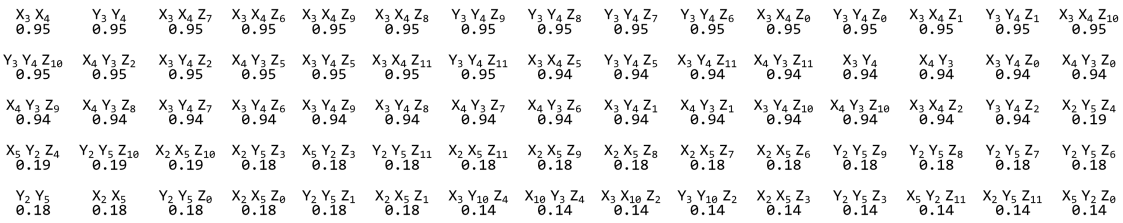


Figure 6.7: List of the at-most 3-local Pauli strings which give most intense signals for a 12-qubit LiH Hamiltonian. Figure reproduced from Ref. [5] – I was its original creator.

ones that correspond to the HOMO-LUMO triplet transition of the basis states $|111100\dots\rangle \to |111010\dots\rangle$ which include, e.g, X_3X_4, X_3Y_4, or any Pauli string obtained by adding a Pauli Z operator on any other qubit, e.g., $Y_3Y_4Z_7$. The next most intense ones also generate HOMO-LUMO transitions to the triplet excited state $|111100\dots\rangle \to |110101\dots\rangle$ including e.g. X_2Y_5 or any other string that is related to this by adding a Pauli Z operator to any other qubit. We then performed an exact 1000-step time evolution sequence of the HF initial state $|111100\dots\rangle$, and collected 3 batches of 50 snapshots at each time step. Without using any chemical information, we simply reconstructed all up-to 3-local Pauli observable signals, and sorted them according to their statistical p-values in from a Ljung-Box test. We obtained a virtually identical list of Pauli strings as Fig. 6.7. This is in agreement with our expectations I described in Section 6.3.3.

6.5.2 Methylene Example

I then explored a specific quantum chemistry example: the discovery of the energy spectrum of a bent methylene molecule CH_2 in the minimal STO-3G basis (14 qubits). Because of the challenges in isolating this molecular biradical experimentally, methylene was one of the first instances where classical computational quantum chemistry calculations predicted the properties of a molecule before it was then verified by photochemistry and spectroscopy experimentally [311–314]. It exhibits the following atypical features: First, its ground state (at equilibrium geometry) has an open-shelled triplet electron configuration instead of the usual closed-shell singlet. Second, the singlet and triplet states cross at a certain geometry, called a

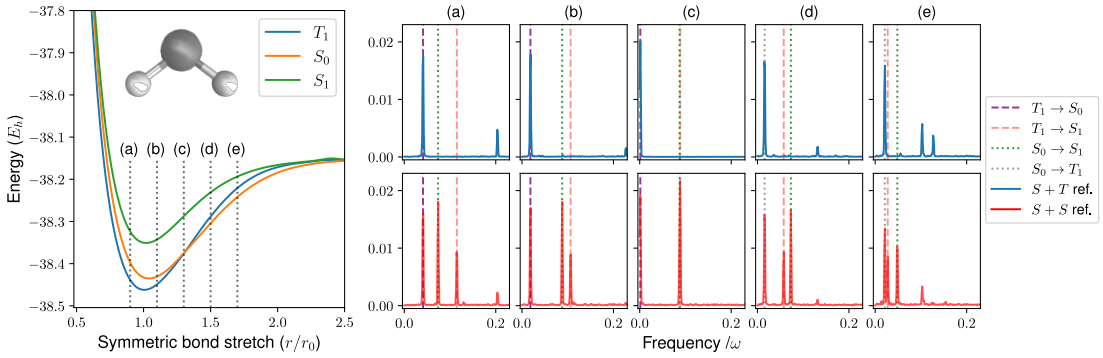


Figure 6.8: Discovering the electronic spectrum of methylene with (classically emulated) shadow spectroscopy. I fixed the bond angle to 101.89° and choose 5 symmetric C-H bond geometries between 0.5 and 3 times the equilibrium bond length of $r_0 = 1.1089$ Å [318]. The top row of spectra were generated using an initial state which was a superposition of the lowest energy singlet and triplet determinants Eq. (6.16), and the bottom row of spectra a superposition of the lowest and first excited singlet determinants Eq. (6.17). The dotted lines are the exact excitation energies from the singlet ground state, purple dashed lines correspond to those from the triplet ground state. From classical shadows of 50 snapshots, all up to 3-local Pauli strings were determined with $N_o = 10689$, out of which 2000 signals passed a statistical test (they contain a deterministic signal with p values in the range $10^{-20} \leq p \leq 0.06$). Note that a number of high energy peaks are observed too – assignment of these transitions can be achieved using techniques detailed in the main text. Figure reproduced from Ref. [5] – I was its original creator.

conical intersection in the excited-state chemistry literature, and the system becomes degenerate. I probed these features [315–317] by estimating shadow spectra at 5 evenly-spaced geometries around the conical intersection, see Fig. 6.8 (a)-(e).

Recall that in the initial state we need only have a reasonable support for eigenstates indexed k and l to observe the energy gap $E_k - E_l$: In Fig. 6.8 (top) I used an initial state as a superposition of the two basis states (Slater determinants) that correspond to the lowest-lying singlet and triplet states as

$$|\psi(0)\rangle = \frac{\sqrt{3}}{2}|11111111000000\rangle + \frac{1}{2}|11111110100000\rangle\,, \tag{6.16}$$

which can be prepared straightforwardly [2]. I time evolved these states exactly in $N_T = 500$ evenly-spaced time intervals with $\Delta t = 10$, and at each time increment

[2] using single-qubit Pauli X gates followed by a single application of a subspace-preserving 2-qubit gate, such as the fSIM gate [319] (see e.g. Ref. [320] for related strategies for preparing similar initial states).

reconstructed expected values of all 3-local Pauli strings from 50 snapshots of classical shadows.

Fig. 6.8 (top row) shows the resulting shadow spectra for each geometry. As expected, there is a dominant peak in each spectrum and their positions agree with the exact singlet-triplet gap (energy difference between the S_0 and T_1 eigenstates) to within error ± 0.4 mE$_h$ (0.25 kcal/mol), as predicted from the Heisenberg scaling $\propto \frac{2\pi}{N_T \Delta t}$ – as singlet-triplet gaps are typically on the order of few kcal/mol, the low error of shadow spectroscopy is very promising. While physical optical spectroscopy only measures a single observable (e.g. electric dipole moment) where 'spin-forbidden' transitions can prohibit resolution of the $S_0 \leftrightarrow T_1$ transition, shadow spectroscopy allows us to probe many different observables, some of which indeed lead to intense peaks as long as the initial state has support for them. At geometry (c), despite the near degenerate singlet-triplet intersection, the near-zero expected frequency is still observed, confirming that indeed shadow spectroscopy can be particularly accurate when used in combination with sophisticated post-Trotter Hamiltonian simulation, even in the strong correlation regime where multiple eigenstates are often near degenerate.

The flexibility to the initial state used is also demonstrated here: I initialised the time evolution in another superposition of basis states that approximate the two lowest-energy singlet states S_0 and S_1 as

$$|\psi(0)\rangle = \frac{1}{\sqrt{2}} |11111111000000\rangle + \frac{1}{\sqrt{2}} |11111110010000\rangle \,, \tag{6.17}$$

In Fig. 6.8 (bottom spectra), the same peak corresponding to the $T_1 \leftrightarrow S_0$ transition is present, suggesting the Slater determinants in Eq. (6.17) have sufficient fidelity with respect to the exact T_1 and S_0 eigenstates. Additional high energy peaks were also revealed; their positions match the exact transition energies of $S_0 \leftrightarrow S_1$ and $T_1 \leftrightarrow S_1$ states within error ± 0.3 mE$_h$ (0.18 kcal/mol).

It is important to note that the low-lying excited state spectrum of a more complex molecule will likely be more challenging to decipher. While in such complex scenarios phase estimation protocols are equally challenging (e.g. due to initial

state overlap with multiple eigenstates), the present approach may be helpful in reducing complexity, as different combinations of input states and Pauli strings measured only give rise to specific peaks. A reasonable strategy might be, as we have demonstrated in Fig. 6.8, to trial a few initial Slater determinants, linear combinations of them, different choices of Pauli operator observables, and multiple geometries. The difference between these scenarios should let the user assign peaks to energy transitions by elimination and obtain an excited state landscape.

As I have eluded to, another strategy is to make use of additional system knowledge; one could envisage a hybrid quantum-classical learning cycle that uses cheap e.g. density functional or perturbation theory calculations to provide initial state guesses, and from the shadow spectroscopy results update the initial reference or geometry until a confident peak assignment is found. One might also consider suppressing high frequency satellite peaks by reducing support for high energy states such as performing a partial convergence as demonstrated in Section 6.4. If one instead desires e.g. ionisation energies or electron affinities, the present approach can also in principle be straightforwardly applied by initialising in a superposition of Slater determinants with different electron-numbers (different Hamming-weights). Further explorations in any of these directions may prove fruitful and offer a new pathway to quantum computing energy gaps in molecules without ever preparing eigenstates in a quantum computer and evaluating static electronic energies.

6.6 Additional Analysis and Experiments

Before I conclude this Chapter, I want to highlight key results obtained by my collaborators which further strengthened the case for shadow spectroscopy. As mentioned briefly in Section 6.3.1, my collaborators used detailed analytic models to understand how the shot noise would propagate in the classical post-processing if we naïvely computed an average of the Fourier transforms of the signals, versus the sequence of steps that we did use – they found that our approach produced shadow spectra with greater SNR. In addition, as mentioned in Section 6.3.2, my collaborators conducted a very comprehensive study of how a broad range of

quantum gate noise models would affect the different stages of quantum hardware investigated, leading to the conclusion that reasonable gate imperfection does not influence the energy accuracy nor the lineshape of spectral peaks. This was first verified by modelling depolarising noise in numerical simulations of the 6-qubit spin chain simulations; when exact time evolution was used and depolarised noise introduced, the weaker secondary peak observed in the numerical simulation (suppressed in the experimental spectrum) actually has the same intensity as the primary $|\psi_0\rangle \rightarrow |\psi_2\rangle$ transition. This was then also confirmed in the shadow spectrum measurement of a 3×2 Hubbard model on an emulated early fault-tolerant quantum computer, where the time evolution was approximated with a first-order Trotter expansion. My collaborator showed that reasonable occurrences and levels of isotropic depolarising, anisotropic depolarising and damping gate noise may slightly reduce the intensity of the spectrum, but does not influence the accuracy. They also showed that, by obtaining shadow spectra from time evolution sequences with progressively higher time resolution, an accurate estimate of the energy gap can be extrapolated to the zero Trotter-error limit. In response to reviewer feedback regarding how well the method would scale, my collaborators also performed a 100-qubit tensor-network spin-problem emulation of shadow spectroscopy, finding that the benefits of shadow spectroscopy still hold.

6.7 Concluding Thoughts

As one of my final projects, I co-developed a general quantum algorithm that combines established techniques in signal processing and cutting-edge methods in quantum metrology to extract energy differences between the quantised levels of, in principle, any quantum many-body system. Shadow spectroscopy can be viewed as a versatile quantum-digitisation of spectroscopy; a quantum simulator only needs to perform time evolution of an initial state that has low-lying eigenstate support, and random basis measurements at each time step allows us to reconstruct signals that share eigenstate information in their harmonics.

Using classical shadows to efficiently estimate q-local Pauli observables N_o turns out to offer a number of benefits beyond the number of shots scaling only as $\mathcal{O}(3^q \log(N_o))$. Firstly, reconstructing (only) up-to 3-local Pauli observables happens to work well in practice, particularly for Fermionic systems: Of course, one caveat of classical shadows in general is that even the number of 3-local Pauli operators will grow exponentially with the system size, but as we discussed we need not reconstruct all of them. We also found that the SNR of shadow spectra can increase in the worst case as $\propto \sqrt{N_o}$ and, by using straightforward statistical correlation tests that scales as $\mathcal{O}(N_o)$, we can efficiently isolate and use only $\sim 10^2$ of the most intense signals to boost the SNR even further. Our numerical results also showed that the number of significant signals is very scalable in the number of qubits. Furthermore, the actual classical resource needed for the post-processing these signals can be quite small in practice – I showed that analysis of $\sim 10^8$ Pauli observable time signals can take $\mathcal{O}(1)$ hours, even on laptops.

Secondly, the approach appears to be fundamentally quite robust to errors in the Hamiltonian simulation. Theoretical and experimental evidence indicated that reasonable gate errors do not degrade the energy accuracy of the spectral peak. Further numerical evidence also suggested that errors from estimating the time evolution can either shift the position of the peak (as a result of Trotter error) or can suppress the intensity of some peaks (as a result of using variational approximations to time evolved states) – my collaborators found that the former can be mitigated by performing finer time evolution sequences and extrapolating the zero-Trotter error peak position, and the latter can be mitigated by increasing the number of shots.

Thirdly, shadow spectroscopy also appears to be robust to shot noise: For example, in every experiment we were using many orders of magnitude fewer shots compared with typical NISQ measurement schemes, where figures well beyond 10^6 shots for a single eigenvalue estimation are very common. In both noisy hardware experiments and numerical simulations, a few dozen circuit repetitions per time step were sufficient to give meaningful results – of course, finding time

evolved states with variational techniques on real hardware will incur additional quantum resources in added shots.

As discussed at the end of Section 6.5, if the goal is to estimate high-quality energy gap values, a key next step would be to establish a more concrete procedure controlling which peaks are revealed, using both an informed choice of initial reference state and observables (a similar sentiment to methods I presented in Chapter 3). Preparing an initial state that has specific support for eigenstates whose energy differences one aims to extract is generally a challenging task, which can be a bottleneck for even advanced, fault-tolerant phase-estimation protocols [321]. This challenge can, indeed, be attributed to general, exponential hardness results for finding ground or excited states [322]. A sensible starting point would be to compute the shadow spectra of a group of well-designed initial states, and assign eigenstate pairs to shared peaks by inspection. Furthermore, if one wishes to compute the absolute values of the energy profiles of excited states, as opposed to just the values of the energy gaps, an additional routine will be necessary to estimate the ground state energy.

A related problem which can limit the effectiveness of shadow spectroscopy is spectral gap closure; the smaller the energy gap, the deeper the necessary circuits. For example, in systems with 'strong correlation', such as the Fermi-Hubbard model where $U/t \ll 1$ or at chemical bond dissociation limits in molecules, low-lying eigenstates may get exponentially close to each other, forming a dense band with degenerate or near-degenerate states. However, resolving individual eigenvalues using phase estimation protocols likewise becomes equally difficult in these cases. Other adversarial examples which may give exponentially vanishing signal intensities include e.g. random quantum states, due to the exponential concentration of eigenvalues of traceless Hamiltonians around $E = 0$. As such, the primary application of the present approach is accurate estimation of excitation energies in gapped quantum systems where the gap is expected to close no faster than polynomially. Most practically relevant problems represent gapped systems, and for example my numerical results in Section 6.3.3 demonstrated that we could

still resolve the energy gap regime near a conical intersection, where states become degenerate. The question of whether a broad class of quantum systems is gapped in the thermodynamic limit is, in general, an undecidable problem [323].

The other main determining factor of shadow spectroscopy is the accuracy of Hamiltonian simulation. Unitary-evolution simulations are almost always only approximate, introducing algorithmic errors that shift the positions of the peaks or dampens the intensity. Understanding how artefacts in the noisy spectrum relate to the approximations used in the Hamiltonian simulation would be helpful.

One might wish instead to predict real, experimental spectra – the primary aim of shadow spectroscopy is not necessarily to faithfully simulate conventional spectroscopy, but to rather estimate differences of eigenvalues making as frugal use of quantum resources as possible. Consequently, under the current prescription it can resolve peaks corresponding to spin-forbidden transitions, which will not be observed in optical spectroscopy. To get real spectra, time evolution of a physical initial state will be necessary, and can be prepared using e.g. techniques I introduced in Chapter 4. However, getting a shadow spectrum's peak heights and lineshapes to match experimental spectra might prove challenging – further work is needed to understand how shadows in a time series can be used to reconstruct the dipole moment or time autocorrelation functions that are actually used in computational techniques for real spectra prediction.

So many worlds, so much to do,
so little done, such things to be

— Alfred, Lord Tennyson

7 Conclusion

This thesis has focused on exploring the potential use of quantum computers for many-body simulation applications. I have explored the use of NISQ computers for variational quantum eigensolvers, with an emphasis on ansatz construction for ground and excited state energies; I found that initial state and operator symmetries, as well as incorporating techniques from adiabatic state preparation especially when Hamiltonian ansätze are used, can increase the efficiency of variational algorithms, even if the number of gates and the number of qubits necessary for simulating a large, interesting molecule or spin system remains out-of-reach. I investigated the use of a single-ancilla QET-U structure to perform non-unitary time evolution. The method was able to prepare a state that has experienced ITE using fewer ancilla measurements compared with state-of-the-art methods. The method would be particularly suited for preparing a state with much ground state support on early fault-tolerant quantum computers, however its use in other non-unitary evolution applications remains to be seen. I also explored the prospects of using fully fault-tolerant quantum computers for simulations of molecules in real space. This approach scales much better than if a classical computer was used, and offers significant insights into quantum dynamics processes of many-body systems. Finally, I co-developed a general technique for extracting spectral gaps from a Hamiltonian, relevant to different levels of quantum technology. Simulation and

experimental results suggest that it is particularly robust to noise, and would be scalable to larger systems.

Because virtually all of the quantum algorithms here have been tested on state vector qubit emulators, the many-body systems I was able to simulate were small in size; once these algorithms can be implemented on larger quantum computers, they have the potential to provide practical and industrial value through the simulation of more complex molecules and materials. It is unclear when will quantum computers powerful enough to perform meaningful simulations using the above quantum algorithms be built. However, as algorithms such as those presented in this thesis continue to improve in their quantum resource use, and as quantum computing technology mature, I remain hopeful that the crossing point of when genuine quantum utility can be delivered will happen sooner, rather than later.

As explained in Chapter 1, a combination of experimental data and synthetic chemical data produced with classical and quantum simulations, perhaps via the methodologies I have presented, can become part of a machine learning and prediction cycle which augments chemistry experiments and discovery. Quantum simulations can provide accurate synthetic data sets for machine-learned emulators or a digital 'atlas' for navigating the chemical space, thereby accelerating chemical and material discovery. The potential impact on the field is immense, setting the stage for a new era of accelerated and efficient scientific and industrial exploration.

References

[1] P.A.M. Dirac and R.H. Dalitz. *The Collected Works of P. A. M. Dirac: Volume 1: 1924-1948*. The Collected Works of P.A.M. Dirac. Cambridge University Press, 1995. URL: https://books.google.hn/books?id=GBEdMdNWFJIC.

[2] Hans Hon Sang Chan et al. "Molecular excited state calculations with adaptive wavefunctions on a quantum eigensolver emulation: reducing circuit depth and separating spin states". In: *Phys. Chem. Chem. Phys.* 23 (46 2021), pp. 26438–26450. URL: http://dx.doi.org/10.1039/D1CP02227J.

[3] Hans Hon Sang Chan, David Muñoz Ramo, and Nathan Fitzpatrick. *Simulating non-unitary dynamics using quantum signal processing with unitary block encoding*. 2023. arXiv: 2303.06161 [quant-ph].

[4] Hans Hon Sang Chan et al. "Grid-based methods for chemistry simulations on a quantum computer". In: *Science Advances* 9.9 (2023), eabo7484. URL: https://www.science.org/doi/abs/10.1126/sciadv.abo7484.

[5] Hans Hon Sang Chan et al. *Algorithmic Shadow Spectroscopy*. 2023. arXiv: 2212.11036 [quant-ph].

[6] Chaomei Chen. "On the Shoulders of Giants". In: *Mapping Scientific Frontiers: The Quest for Knowledge Visualization*. London: Springer London, 2003, pp. 135–166. URL: https://doi.org/10.1007/978-1-4471-0051-5_5.

[7] Arnab Das and Bikas K. Chakrabarti. "Colloquium: Quantum annealing and analog quantum computation". In: *Rev. Mod. Phys.* 80 (3 Sept. 2008), pp. 1061–1081. URL: https://link.aps.org/doi/10.1103/RevModPhys.80.1061.

[8] Andrew J Daley. "Twenty-five years of analogue quantum simulation". In: *Nature Reviews Physics* 5.12 (2023), pp. 702–703.

[9] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010. arXiv: arXiv:1011.1669v3. URL: http://www.tandfonline.com/doi/abs/10.1080/00107514.2011.587535.

[10] Alexander M. Dalzell et al. *Quantum algorithms: A survey of applications and end-to-end complexities*. 2023. arXiv: 2310.03011 [quant-ph].

[11] Ashley Montanaro. "Quantum algorithms: an overview". In: *npj Quantum Information* 2.1 (2016), pp. 1–8.

[12] Zhenyu Cai et al. "Quantum error mitigation". In: *Rev. Mod. Phys.* 95 (4 Dec. 2023), p. 045005. URL: https://link.aps.org/doi/10.1103/RevModPhys.95.045005.

[13] Kishor Bharti et al. "Noisy intermediate-scale quantum algorithms". In: *Rev. Mod. Phys.* 94 (1 Feb. 2022), p. 015004. URL: https://link.aps.org/doi/10.1103/RevModPhys.94.015004.

[14] Simon J Devitt, William J Munro, and Kae Nemoto. "Quantum error correction for beginners". In: *Reports on Progress in Physics* 76.7 (2013), p. 076001.

[15] Yu Tong. "Designing algorithms for estimating ground state properties on early fault-tolerant quantum computers". In: *Quantum Views* 6 (July 2022), p. 65. URL: https://doi.org/10.22331/qv-2022-07-22-65.

[16] Chris N. Self, Marcello Benedetti, and David Amaro. *Protecting Expressive Circuits with a Quantum Error Detection Code*. 2022. URL: https://arxiv.org/abs/2211.06703.

[17] *New Scientist magazine article dated 04/12/2023 'IBM's "Condor" quantum computer has more than 1000 qubits'*. URL: https://www.newscientist.com/article/2405789-ibms-condor-quantum-computer-has-more-than-1000-qubits/.

[18] *New Scientist magazine article dated 24/10/2023 'Record-breaking quantum computer has more than 1000 qubits'*. URL: https://www.newscientist.com/article/2399246-record-breaking-quantum-computer-has-more-than-1000-qubits/.

[19] Google Quantum AI. "Suppressing quantum errors by scaling a surface code logical qubit". In: *Nature* 614.7949 (2023), pp. 676–681. URL: https://www.nature.com/articles/s41586-022-05434-1.

[20] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays". In: *Nature* 626.7997 (2024), pp. 58–65. URL: https://www.nature.com/articles/s41586-023-06927-3.

[21] Earl Campbell. "A series of fast-paced advances in Quantum Error Correction". In: *Nature Reviews Physics* (2024), pp. 1–2.

[22] Amit Jamadagni, Andreas M. Läuchli, and Cornelius Hempel. *Benchmarking quantum computer simulation software packages*. 2024. arXiv: 2401.09076 [quant-ph].

[23] David J. Tannor. *Introduction to Quantum Mechanics: A Time-Dependent Perspective*. University Science Books, 2007.

[24] A. Szabo and N.S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Books on Chemistry. Dover Publications, 1996. URL: https://books.google.co.uk/books?id=6mV9gYzEkgIC.

[25] I. M. Georgescu, S. Ashhab, and Franco Nori. "Quantum simulation". In: *Rev. Mod. Phys.* 86 (1 Mar. 2014), pp. 153–185. URL: https://link.aps.org/doi/10.1103/RevModPhys.86.153.

[26] Benedikt Fauseweh. "Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges". In: *Nature Communications* 15.1 (2024), p. 2123. URL: https://www.nature.com/articles/s41467-024-46402-9.

[27] Andrew J Daley et al. "Practical quantum advantage in quantum simulation". In: *Nature* 607.7920 (2022), pp. 667–676. URL: https://www.nature.com/articles/s41586-022-04940-6.

[28] Jonathan Olson et al. "Quantum Information and Computation for Chemistry". In: (2016). arXiv: 1706.05413. URL: http://arxiv.org/abs/1706.05413.

[29] Yudong Cao et al. "Quantum chemistry in the age of quantum computing". In: *Chemical reviews* 119.19 (2019), pp. 10856–10915.

[30] Sam McArdle et al. "Quantum computational chemistry". In: *Reviews of Modern Physics* 92.1 (2020), p. 015003. URL: https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.92.015003.

[31] Bela Bauer et al. "Quantum Algorithms for Quantum Chemistry and Quantum Materials Science". In: *Chemical Reviews* 120.22 (Oct. 2020), pp. 12685–12717. arXiv: 2001.03685. URL: https://pubs.acs.org/doi/full/10.1021/acs.chemrev.9b00829%20https://pubs.acs.org/doi/10.1021/acs.chemrev.9b00829.

[32] Hari P. Paudel et al. "Quantum Computing and Simulations for Energy Applications: Review and Perspective". In: *ACS Engineering Au* 2.3 (2022), pp. 151–196. eprint: https://doi.org/10.1021/acsengineeringau.1c00033. URL: https://doi.org/10.1021/acsengineeringau.1c00033.

[33] Dominic W Berry, Graeme Ahokas, and Richard Cleve. "Quantum algorithms for Hamiltonian simulation". In: *Mathematics of Quantum Computation and Quantum Technology*. Chapman and Hall/CRC, 2007, pp. 107–130.

[34] Alexander Miessen et al. "Quantum algorithms for quantum dynamics". In: *Nature Computational Science* 3.1 (2023), pp. 25–37.

[35] A Yu Kitaev. "Quantum computations: algorithms and error correction". In: *Russian Mathematical Surveys* 52.6 (Dec. 1997), p. 1191. URL: https://dx.doi.org/10.1070/RM1997v052n06ABEH002155.

[36] Bálint Koczor, John J. L. Morton, and Simon C. Benjamin. "Probabilistic Interpolation of Quantum Rotation Angles". In: *Phys. Rev. Lett.* 132 (13 Mar. 2024), p. 130602. URL: https://link.aps.org/doi/10.1103/PhysRevLett.132.130602.

[37] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Dover Publications, 1989, p. 1097.

[38] Peter J Knowles, Martin Schütz, and Hans-Joachim Werner. "Ab Initio Methods for Electron Correlation in Molecules". In: *Modern Methods and Algorithms of Quantum Chemistry Proceedings, Second Edition* (2000), p. 84. URL: http://www.fz-juelich.de/nic-series/Volume3/Volume3.html%5Cnfile:///Users/peterk/Library/Application%20Support/Papers2/Articles/2000/Knowles/Modern%20Methods%20and%20Algorithms%20of%20Quantum%20Chemistry%5CnProceedings%20Second%20Edition%202000%20Knowles.pdf%5Cnpapers2://publi.

[39] Joseph J W McDouall. "Computational Quantum Chemistry". In: *Computational Quantum Chemistry: Molecular Structure and Properties in Silico*. Royal Society of Chemistry, 2013. Chap. 1, pp. 1–62.

[40] T. Helgaker, P. Jorgensen, and J. Olsen. *Molecular Electronic-Structure Theory*. Wiley, 2014. URL: https://books.google.co.uk/books?id=lNVLBAAAQBAJ.

[41] Andrew Tranter et al. "A Comparison of the Bravyi-Kitaev and Jordan-Wigner Transformations for the Quantum Simulation of Quantum Chemistry". In: *Journal of Chemical Theory and Computation* 14.11 (2018), pp. 5617–5630. arXiv: 1812.02233.

[42] Kanav Setia et al. "Reducing qubit requirements for quantum simulation using molecular point group symmetries". In: (2019). arXiv: 1910.14644. URL: http://arxiv.org/abs/1910.14644.

[43] Ryan Babbush et al. "Low-Depth Quantum Simulation of Materials". In: *Physical Review X* 8.1 (2018), p. 11044. URL: https://doi.org/10.1103/PhysRevX.8.011044.

[44] Alexis Ralli et al. "Unitary partitioning and the contextual subspace variational quantum eigensolver". In: *Phys. Rev. Res.* 5 (1 Feb. 2023), p. 013095. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.5.013095.

[45] Stephen Wiesner. *Simulations of Many-Body Quantum Systems by a Quantum Computer*. 1996. arXiv: quant-ph/9603028 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/9603028.

[46] Christof Zalka. "Simulating quantum systems on a quantum computer". In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454.1969 (1998), pp. 313–322. eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1998.0162. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1998.0162.

[47] Ivan Kassal et al. "Polynomial-time quantum algorithm for the simulation of chemical dynamics". In: *Proceedings of the National Academy of Sciences* (2008), pp. 18681–18686. eprint: https://www.pnas.org/content/early/2008/11/24/0808245105.full.pdf. URL: https://www.pnas.org/content/early/2008/11/24/0808245105.

[48] Giuliano Benenti and Giuliano Strini. "Quantum simulation of the single-particle Schrödinger equation". In: *American Journal of Physics* 76.7 (2008), pp. 657–662. eprint: https://doi.org/10.1119/1.2894532. URL: https://doi.org/10.1119/1.2894532.

[49] N Cody Jones et al. "Faster quantum chemistry simulation on fault-tolerant quantum computers". In: *New Journal of Physics* 14.11 (Nov. 2012), p. 115023. URL: https://doi.org/10.1088/1367-2630/14/11/115023.

[50] Rolando D. Somma. *Quantum simulations of one dimensional quantum systems*. 2016. arXiv: 1503.06319 [quant-ph]. URL: https://arxiv.org/abs/1503.06319.

[51] Ian D Kivlichan et al. "Bounding the costs of quantum simulation of many-body physics in real space". In: *Journal of Physics A: Mathematical and Theoretical* 50.30 (June 2017), p. 305301. URL: https://doi.org/10.1088/1751-8121/aa77b8.

[52] Pauline J. Ollitrault et al. "Hardware Efficient Quantum Algorithms for Vibrational Structure Calculations". In: (2020). arXiv: 2003.12578. URL: http://arxiv.org/abs/2003.12578.

[53] Yuan Su et al. "Fault-Tolerant Quantum Simulations of Chemistry in First Quantization". In: *PRX Quantum* 2 (4 Nov. 2021), p. 040332. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040332.

[54] Taichi Kosugi et al. "Imaginary-time evolution using forward and backward real-time evolution with a single ancilla: First-quantized eigensolver algorithm for quantum chemistry". In: *Phys. Rev. Research* 4 (3 Aug. 2022), p. 033121. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.033121.

[55] Andrew M. Childs et al. *Quantum simulation of real-space dynamics*. 2022. URL: https://arxiv.org/abs/2203.17006.

[56] Bill Poirier and Jonathan Jerke. "Full-dimensional Schrödinger wavefunction calculations using tensors and quantum computers: the Cartesian component-separated approach". In: *Phys. Chem. Chem. Phys.* 24 (7 2022), pp. 4437–4454. URL: http://dx.doi.org/10.1039/D1CP02036F.

[57] Pauline J Ollitrault et al. *Quantum algorithms for grid-based variational time evolution*. 2022. URL: https://arxiv.org/abs/2203.02521.

[58] Taichi Kosugi, Hirofumi Nishi, and Yuichiro Matsushita. *Exhaustive search for optimal molecular geometries using imaginary-time evolution on a quantum computer*. 2022. URL: https://arxiv.org/abs/2210.09883.

[59] Hirotoshi Hirai et al. "Molecular Structure Optimization Based on Electrons–Nuclei Quantum Dynamics Computation". In: *ACS Omega* 7.23 (2022), pp. 19784–19793. eprint: https://doi.org/10.1021/acsomega.2c01546. URL: https://doi.org/10.1021/acsomega.2c01546.

[60] Nicholas J. Ward, Ivan Kassal, and Alán Aspuru-Guzik. "Preparation of many-body states for quantum simulation". In: *The Journal of Chemical Physics* 130.19 (2009), p. 194105. eprint: https://doi.org/10.1063/1.3115177. URL: https://doi.org/10.1063/1.3115177.

[61] Lov Grover and Terry Rudolph. *Creating superpositions that correspond to efficiently integrable probability distributions*. 2002. URL: https://arxiv.org/abs/quant-ph/0208112.

[62] Phillip Kaye and Michele Mosca. "Quantum networks for generating arbitrary quantum states". In: *International Conference on Quantum Information*. Optical Society of America. 2001, PB28.

[63] Mikko Möttönen et al. "Transformation of Quantum States Using Uniformly Controlled Rotations". In: *Quantum Info. Comput.* 5.6 (Sept. 2005), pp. 467–473.

[64] Gabriel Marin-Sanchez, Javier Gonzalez-Conde, and Mikel Sanz. *Quantum algorithms for approximate function loading*. 2021. arXiv: 2111.07933 [quant-ph]. URL: https://arxiv.org/abs/2111.07933.

[65] Alexei Kitaev and William A. Webb. *Wavefunction preparation and resampling using a quantum computer*. 2009. arXiv: 0801.0342 [quant-ph]. URL: https://arxiv.org/abs/0801.0342.

[66] A. Holmes and A. Y. Matsuura. "Efficient Quantum Circuits for Accurate State Preparation of Smooth, Differentiable Functions". In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct. 2020, pp. 169–179. URL: https://doi.ieeecomputersociety.org/10.1109/QCE49297.2020.00030.

[67] Almudena Carrera Vazquez and Stefan Woerner. "Efficient State Preparation for Quantum Amplitude Estimation". In: *Phys. Rev. Applied* 15 (3 Mar. 2021), p. 034027. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.15.034027.

[68] Arthur G Rattew et al. "The efficient preparation of normal distributions in quantum registers". In: *Quantum* 5 (2021), p. 609.

[69] Arthur G. Rattew and Bálint Koczor. *Preparing Arbitrary Continuous Functions in Quantum Registers With Logarithmic Complexity*. 2022. URL: https://arxiv.org/abs/2205.00519.

[70] Taichi Kosugi et al. *Qubit encoding for a mixture of localized functions*. 2024. arXiv: 2404.18529 [quant-ph].

[71] Daniel S. Abrams and Seth Lloyd. "Simulation of Many-Body Fermi Systems on a Universal Quantum Computer". In: *Phys. Rev. Lett.* 79 (13 Sept. 1997), pp. 2586–2589. URL: https://link.aps.org/doi/10.1103/PhysRevLett.79.2586.

[72] Dominic W Berry et al. "Improved techniques for preparing eigenstates of fermionic Hamiltonians". In: *npj Quantum Information* 4.1 (2018), pp. 1–7.

[73] Seunghoon Lee et al. *Is there evidence for exponential quantum advantage in quantum chemistry?* 2022. URL: https://arxiv.org/abs/2208.02199.

[74] Sahil Gulania and James Daniel Whitfield. "Young frames for quantum chemistry". In: (Apr. 2019). arXiv: 1904.10469. URL: http://arxiv.org/abs/1904.10469.

[75] Daniel Marti-Dafcik, Hugh G. A. Burton, and David P. Tew. *Spin coupling is all you need: Encoding strong electron correlation on quantum computers*. 2024. arXiv: 2404.18878 [quant-ph].

[76] M. Cerezo et al. *Variational Quantum Algorithms*. Dec. 2020. arXiv: 2012.09265 [quant-ph]. URL: http://arxiv.org/abs/2012.09265.

[77] Xiao Yuan et al. "Theory of variational quantum simulation". In: *Quantum* 3 (2019), p. 191. arXiv: 1812.08767.

[78] Alberto Peruzzo et al. "A variational eigenvalue solver on a photonic quantum processor". In: *Nature Communications* 5.1 (Sept. 2014), p. 4213. arXiv: 1304.3061. URL: http://www.nature.com/articles/ncomms5213.

[79] Abhinav Kandala et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets". In: *Nature* 549.7671 (Sept. 2017), pp. 242–246. arXiv: 1704.05018. URL: http://dx.doi.org/10.1038/nature23879%20http://www.nature.com/articles/nature23879.

[80] J. I. Colless et al. "Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm". In: *Physical Review X* 8.1 (Feb. 2018), p. 011021. URL: https://doi.org/10.1103/PhysRevX.8.011021%20https://link.aps.org/doi/10.1103/PhysRevX.8.011021.

[81] Cornelius Hempel et al. "Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator". In: *Physical Review X* 8.3 (July 2018), p. 031022. arXiv: 1803.10238. URL: https://doi.org/10.1103/PhysRevX.8.031022%20https://link.aps.org/doi/10.1103/PhysRevX.8.031022.

[82] P. J. J. O'Malley et al. "Scalable Quantum Simulation of Molecular Energies". In: *Physical Review X* 6.3 (July 2016), p. 031007. arXiv: arXiv:1512.06860v2. URL: https://link.aps.org/doi/10.1103/PhysRevX.6.031007.

[83] C. Kokail et al. "Self-verifying variational quantum simulation of lattice models". In: *Nature* 569.7756 (2019), pp. 355–360. URL: http://dx.doi.org/10.1038/s41586-019-1177-4.

[84] M. Ganzhorn et al. "Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer". In: *Phys. Rev. Appl.* 11 (4 Apr. 2019), p. 044092. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.11.044092.

[85] Google AI Quantum. "Hartree-Fock on a superconducting qubit quantum computer". In: *Science* 369.6507 (2020), pp. 1084–1089. arXiv: 2004.04174.

[86] Jarrod R McClean et al. "Barren plateaus in quantum neural network training landscapes". In: *Nature Communications* 9.1 (2018), pp. 1–6.

[87] Eric R Anschuetz and Bobak T Kiani. "Beyond Barren Plateaus: Quantum Variational Algorithms Are Swamped With Traps". In: *arXiv preprint arXiv:2205.05786* (2022).

[88] Dave Wecker, Matthew B. Hastings, and Matthias Troyer. "Progress towards practical quantum variational algorithms". In: *Physical Review A* 92.4 (Oct. 2015). URL: http://dx.doi.org/10.1103/PhysRevA.92.042303.

[89] Jan Michael Reiner et al. "Finding the ground state of the Hubbard model by variational methods on a quantum computer with gate errors". In: *Quantum Science and Technology* 4.3 (2019), pp. 1–8. arXiv: 1811.04476.

[90] Joonho Lee et al. "Generalized Unitary Coupled Cluster Wave functions for Quantum Computation". In: *Journal of Chemical Theory and Computation* 15.1 (Jan. 2019), pp. 311–324. arXiv: 1810.02327. URL: https://pubs.acs.org/doi/10.1021/acs.jctc.8b01004.

[91] Francesco A. Evangelista, Garnet Kin Lic Chan, and Gustavo E. Scuseria. "Exact parameterization of fermionic wave functions via unitary coupled cluster theory". In: *The Journal of Chemical Physics* 151.24 (Dec. 2019), p. 244112. arXiv: 1910.10130. URL: http://arxiv.org/abs/1910.10130%20http://aip.scitation.org/doi/10.1063/1.5133059.

[92] Rodney J. Bartlett and Monika Musiał. "Coupled-cluster theory in quantum chemistry". In: *Reviews of Modern Physics* 79.1 (Feb. 2007), pp. 291–352. URL: https://link.aps.org/doi/10.1103/RevModPhys.79.291.

[93] Jonathan Romero et al. "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz". In: *Quantum Science and Technology* 4.1 (Jan. 2017), pp. 1–22. arXiv: 1701.02691. URL: http://arxiv.org/abs/1701.02691.

[94] Panagiotis Kl Barkoutsos et al. "Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions". In: *Physical Review A* 98.2 (2018), pp. 1–13. arXiv: 1805.04340.

[95] Harper R. Grimsley et al. "An adaptive variational algorithm for exact molecular simulations on a quantum computer". In: *Nature Communications* 10.1 (2019), pp. 1–11. arXiv: 1812.11173. URL: http://dx.doi.org/10.1038/s41467-019-10988-2.

[96] Jules Tilly et al. "The Variational Quantum Eigensolver: A review of methods and best practices". In: *Physics Reports* 986 (2022). The Variational Quantum Eigensolver: a review of methods and best practices, pp. 1–128. URL: https://www.sciencedirect.com/science/article/pii/S0370157322003118.

[97] Oscar Higgott, Daochen Wang, and Stephen Brierley. "Variational quantum computation of excited states". In: *Quantum* 3 (2019), p. 156.

[98] Yohei Ibe et al. *Calculating transition amplitudes by variational quantum eigensolvers*. Feb. 2020. arXiv: 2002.11724. URL: http://arxiv.org/abs/2002.11724.

[99] Tyson Jones et al. "Variational quantum algorithms for discovering Hamiltonian spectra". In: *Physical Review A* 99.6 (June 2019), p. 062304. arXiv: 1806.05707. URL: http://arxiv.org/abs/1806.05707%20http://dx.doi.org/10.1103/PhysRevA.99.062304%20https://link.aps.org/doi/10.1103/PhysRevA.99.062304.

[100] Ronnie Kosloff. "Time-dependent quantum-mechanical methods for molecular dynamics". In: *The Journal of Physical Chemistry* 92.8 (1988), pp. 2087–2100. URL: https://doi.org/10.1021/j100319a003.

[101] R Kosloff. "Propagation Methods for Quantum Molecular Dynamics". In: *Annual Review of Physical Chemistry* 45.1 (1994), pp. 145–178. URL: https://doi.org/10.1146/annurev.pc.45.100194.001045.

[102] C. Leforestier et al. "A comparison of different propagation schemes for the time dependent Schrödinger equation". In: *Journal of Computational Physics* 94.1 (1991), pp. 59–80.

[103] John C. Light and Tucker Carrington Jr. "Discrete-Variable Representations and their Utilization". In: *Advances in Chemical Physics*. John Wiley and Sons, Ltd, 2000, pp. 263–310. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470141731.ch4.

[104] Barry I. Schneider, Lee A. Collins, and S. X. Hu. "Parallel solver for the time-dependent linear and nonlinear Schrödinger equation". In: *Phys. Rev. E* 73 (3 Mar. 2006), p. 036708. URL: https://link.aps.org/doi/10.1103/PhysRevE.73.036708.

[105] D.J. Tannor. *Introduction to Quantum Mechanics*. University Science Books, 2007. URL: https://books.google.co.uk/books?id=WbQOzQEACAAJ.

[106] Fabien Gatti et al. "Introduction to Numerical Methods". In: *Applications of Quantum Dynamics in Chemistry*. Cham: Springer International Publishing, 2017, pp. 201–285. URL: https://doi.org/10.1007/978-3-319-53923-2_8.

[107] M Beck. "The multiconfiguration time-dependent Hartree (MCTDH) method: a highly efficient algorithm for propagating wavepackets". In: *Physics Reports* 324.1 (Jan. 2000), pp. 1–105. URL: https://linkinghub.elsevier.com/retrieve/pii/S0370157399000472.

[108] Hans-Dieter Meyer. *Introduction To MCTDH*. July 2011. 2013.

[109] Axel U.J. Lode et al. "Colloquium: Multiconfigurational time-dependent Hartree approaches for indistinguishable particles". In: *Reviews of Modern Physics* 92.1 (2020), pp. 1–25. arXiv: arXiv:1908.03578v2.

[110] Roland Lindh and Leticia González. *Quantum Chemistry and Dynamics of Excited States: Methods and Applications*. John Wiley and Sons, Ltd, Nov. 2020. URL: https://onlinelibrary.wiley.com/doi/book/10.1002/9781119417774.

[111] Haobin Wang. "Multilayer Multiconfiguration Time-Dependent Hartree Theory". In: *The Journal of Physical Chemistry A* 119.29 (2015). PMID: 26020459, pp. 7951–7965. eprint: https://doi.org/10.1021/acs.jpca.5b03256. URL: https://doi.org/10.1021/acs.jpca.5b03256.

[112] G. A. Worth et al. "Using the MCTDH wavepacket propagation method to describe multimode non-adiabatic dynamics". In: *International Reviews in Physical Chemistry* 27.3 (2008), pp. 569–606. eprint: https://doi.org/10.1080/01442350802137656. URL: https://doi.org/10.1080/01442350802137656.

[113] Joseph E. Lawrence and David E. Manolopoulos. "Path integral methods for reaction rates in complex systems". In: *Faraday Discuss.* 221 (0 2020), pp. 9–29. URL: http://dx.doi.org/10.1039/C9FD00084D.

[114] Le Yu et al. "Trajectory-based nonadiabatic molecular dynamics without calculating nonadiabatic coupling in the avoided crossing case: trans – cis photoisomerization in azobenzene". In: *Phys. Chem. Chem. Phys.* 16 (47 2014), pp. 25883–25895. URL: http://dx.doi.org/10.1039/C4CP03498H.

[115] My Hang V. Huynh and Thomas J. Meyer. "Proton-Coupled Electron Transfer". In: *Chemical Reviews* 107.11 (2007). PMID: 17999556, pp. 5004–5064. eprint: https://doi.org/10.1021/cr0500030. URL: https://doi.org/10.1021/cr0500030.

[116] Mauro Nisoli et al. "Attosecond Electron Dynamics in Molecules". In: *Chemical Reviews* 117.16 (2017). PMID: 28488433, pp. 10760–10825. eprint: https://doi.org/10.1021/acs.chemrev.6b00453. URL: https://doi.org/10.1021/acs.chemrev.6b00453.

[117] Shaunak Mukherjee et al. "Hot Electrons Do the Impossible: Plasmon-Induced Dissociation of H2 on Au". In: *Nano Letters* 13.1 (2013). PMID: 23194158, pp. 240–247. eprint: https://doi.org/10.1021/nl303940z. URL: https://doi.org/10.1021/nl303940z.

[118] Alicia B. Magann et al. "Digital quantum simulation of molecular dynamics and control". In: (Feb. 2020), pp. 1–20. arXiv: 2002.12497. URL: http://arxiv.org/abs/2002.12497.

[119] Seth Lloyd. "Universal Quantum Simulators". In: *Science* 273 (1996), pp. 1073–1078. URL: https://www.jstor.org/stable/2899535.

[120] Jr A Fleck, JR Morris, and MD Feit. "Time-dependent propagation of high energy laser beams through the atmosphere". In: *Applied physics* 10.2 (1976), pp. 129–160.

[121] M.D Feit, J.A Fleck, and A Steiger. "Solution of the Schrödinger equation by a spectral method". In: *Journal of Computational Physics* 47.3 (1982), pp. 412–433. URL: https://www.sciencedirect.com/science/article/pii/0021999182900912.

[122] M. D. Feit and J. A. Fleck. "Solution of the Schrödinger equation by a spectral method II: Vibrational energy levels of triatomic molecules". In: *The Journal of Chemical Physics* 78.1 (1983), pp. 301–308. URL: https://doi.org/10.1063/1.444501.

[123] R. Sadeghi and Rex T. Skodje. "Spectral quantization of high energy transition state resonances in the H+H2 reaction". In: *The Journal of Chemical Physics* 99.7 (1993), pp. 5126–5140. URL: https://doi.org/10.1063/1.466014.

[124] Dong H. Zhang and John Z. H. Zhang. "Quantum calculations of reaction probabilities for HO + CO→ H + CO2 and bound states of HOCO". In: *The Journal of Chemical Physics* 103.15 (1995), pp. 6512–6519. URL: https://doi.org/10.1063/1.470377.

[125] V. N. Serov et al. "Split operator method for the nonadiabatic (J=0) bound states and (A←X) absorption spectrum of NO2". In: *The Journal of Chemical Physics* 115.14 (2001), pp. 6450–6458. eprint: https://doi.org/10.1063/1.1396854. URL: https://doi.org/10.1063/1.1396854.

[126] Xin Chen and Victor S. Batista. "Matching-pursuit/split-operator-Fourier-transform simulations of excited-state nonadiabatic quantum dynamics in pyrazine". In: *The Journal of Chemical Physics* 125.12 (2006), p. 124313. eprint: https://doi.org/10.1063/1.2356477. URL: https://doi.org/10.1063/1.2356477.

[127] André D. Bandrauk and Hai Shen. "Exponential split operator methods for solving coupled time-dependent Schrödinger equations". In: *The Journal of Chemical Physics* 99.2 (1993), pp. 1185–1193. eprint: https://doi.org/10.1063/1.465362. URL: https://doi.org/10.1063/1.465362.

[128] Samuel M. Greene and Victor S. Batista. "Tensor-Train Split-Operator Fourier Transform (TT-SOFT) Method: Multidimensional Nonadiabatic Quantum Dynamics". In: *Journal of Chemical Theory and Computation* 13.9 (2017). PMID: 28763215, pp. 4034–4042. eprint: https://doi.org/10.1021/acs.jctc.7b00608. URL: https://doi.org/10.1021/acs.jctc.7b00608.

[129] Basile F.E. Curchod and Todd J. Martínez. "Ab Initio Nonadiabatic Quantum Molecular Dynamics". In: *Chemical Reviews* 118.7 (2018), pp. 3305–3336.

[130] Yinghua Wu and Victor S. Batista. "Matching-pursuit for simulations of quantum processes". In: *The Journal of Chemical Physics* 118.15 (2003), pp. 6720–6724. eprint: https://doi.org/10.1063/1.1560636. URL: https://doi.org/10.1063/1.1560636.

[131] Yinghua Wu and Victor S. Batista. "Quantum tunneling dynamics in multidimensional systems: A matching-pursuit description". In: *The Journal of Chemical Physics* 121.4 (2004), pp. 1676–1680. eprint: https://doi.org/10.1063/1.1766298. URL: https://doi.org/10.1063/1.1766298.

[132] Vyacheslav A Trofimov and NV Peskov. "Comparison of finite-difference schemes for the Gross-Pitaevskii equation". In: *Mathematical Modelling and Analysis* 14.1 (2009), pp. 109–126.

[133] C.M. Dion, A. Hashemloo, and G. Rahali. "Program for quantum wave-packet dynamics with time-dependent potentials". In: *Computer Physics Communications* 185.1 (2014), pp. 407–414. URL: https://www.sciencedirect.com/science/article/pii/S0010465513003135.

[134] Brandon Buonacorsi, Benjamin Shaw, and Jonathan Baugh. "Simulated coherent electron shuttling in silicon quantum dots". In: *Phys. Rev. B* 102 (12 Sept. 2020), p. 125406. URL: https://link.aps.org/doi/10.1103/PhysRevB.102.125406.

[135] Matthias Ruf, Heiko Bauke, and Christoph H. Keitel. "A real space split operator method for the Klein–Gordon equation". In: *Journal of Computational Physics* 228.24 (2009), pp. 9092–9106. URL: https://www.sciencedirect.com/science/article/pii/S0021999109004975.

[136] Paul L. DeVries. "Application of the Split Operator Fourier Transform method to the solution of the nonlinear Schrödinger equation". In: *AIP Conference Proceedings* 160.1 (1987), pp. 269–271. eprint: https://aip.scitation.org/doi/pdf/10.1063/1.36847. URL: https://aip.scitation.org/doi/abs/10.1063/1.36847.

[137] J. A. C. Weideman and B. M. Herbst. "Split-Step Methods for the Solution of the Nonlinear Schrodinger Equation". In: *SIAM Journal on Numerical Analysis* 23.3 (1986), pp. 485–507. URL: http://www.jstor.org/stable/2157521.

[138] ERWAN Faou, LUDWIG Gauckler, and CHRISTIAN Lubich. "PLANE WAVE STABILITY OF THE SPLIT-STEP FOURIER METHOD FOR THE NONLINEAR SCHRÖDINGER EQUATION". In: *Forum of Mathematics, Sigma* 2 (2014), e5.

[139] R. Kosloff. "The Fourier Method". In: *Numerical Grid Methods and Their Application to Schrödinger's Equation*. Ed. by Charles Cerjan. Dordrecht: Springer Netherlands, 1993, pp. 175–194. URL: https://doi.org/10.1007/978-94-015-8240-7_11.

[140] Dominic W. Berry et al. "Simulating Hamiltonian Dynamics with a Truncated Taylor Series". In: *Physical Review Letters* 114.9 (Mar. 2015), p. 090502. arXiv: 1412.4687. URL: https://link.aps.org/doi/10.1103/PhysRevLett.114.090502.

[141] Richard Meister, Simon C. Benjamin, and Earl T. Campbell. "Tailoring Term Truncations for Electronic Structure Calculations Using a Linear Combination of Unitaries". In: (2020), pp. 1–13. arXiv: 2007.11624. URL: http://arxiv.org/abs/2007.11624.

[142] Mária Kieferová, Artur Scherer, and Dominic W. Berry. "Simulating the dynamics of time-dependent Hamiltonians with a truncated Dyson series". In: *Phys. Rev. A* 99 (4 Apr. 2019), p. 042314. URL: https://link.aps.org/doi/10.1103/PhysRevA.99.042314.

[143] Guang Hao Low and Isaac L. Chuang. "Hamiltonian Simulation by Qubitization". In: *Quantum* 3 (July 2019), p. 163. URL: https://doi.org/10.22331/q-2019-07-12-163.

[144] Ryan Babbush et al. "Quantum simulation of chemistry with sublinear scaling in basis size". In: *npj Quantum Information* 5.1 (2019), pp. 1–7.

[145] John M. Martyn et al. "Grand Unification of Quantum Algorithms". In: *PRX Quantum* 2 (4 Dec. 2021), p. 040203. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.040203.

[146] Guang Hao Low and Isaac L. Chuang. "Optimal Hamiltonian Simulation by Quantum Signal Processing". In: *Phys. Rev. Lett.* 118 (1 Jan. 2017), p. 010501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.010501.

[147] Yulong Dong et al. "Efficient phase-factor evaluation in quantum signal processing". In: *Phys. Rev. A* 103 (4 Apr. 2021), p. 042419. URL: https://link.aps.org/doi/10.1103/PhysRevA.103.042419.

[148] Yulong Dong, Lin Lin, and Yu Tong. *Ground state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices*. 2022. arXiv: 2204.05955 [quant-ph].

[149] Sam McArdle, András Gilyén, and Mario Berta. *Quantum state preparation without coherent arithmetic*. 2022. URL: https://arxiv.org/abs/2210.14892.

[150] Thais L Silva et al. *Fragmented imaginary-time evolution for early-stage quantum signal processors*. 2023.

[151] Thais de Lima Silva, Lucas Borges, and Leandro Aolita. *Fourier-based quantum signal processing*. 2022. URL: https://arxiv.org/abs/2206.02826.

[152] Kianna Wan, Mario Berta, and Earl T. Campbell. "Randomized Quantum Algorithm for Statistical Phase Estimation". In: *Physical Review Letters* 129.3 (July 2022). URL: https://doi.org/10.1103%2Fphysrevlett.129.030503.

[153] Taichi Kosugi et al. "Imaginary-time evolution using forward and backward real-time evolution with a single ancilla: First-quantized eigensolver algorithm for quantum chemistry". In: *Phys. Rev. Research* 4 (3 Aug. 2022), p. 033121. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.033121.

[154] Kenneth Choi et al. "Rodeo Algorithm for Quantum Computing". In: *Phys. Rev. Lett.* 127 (4 July 2021), p. 040505. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.040505.

[155] Richard Meister and Simon C. Benjamin. *Resource-frugal Hamiltonian eigenstate preparation via repeated quantum phase estimation measurements*. 2022. URL: https://arxiv.org/abs/2212.00846.

[156] Sam McArdle et al. "Variational ansatz-based quantum simulation of imaginary time evolution". In: *npj Quantum Information* 5.1 (June 2019), p. 75. URL: https://doi.org/10.1038/s41534-019-0187-2.

[157] Nathan Fitzpatrick, Harriet Apel, and David Muñoz Ramo. *Evaluating low-depth quantum algorithms for time evolution on fermion-boson systems*. 2021. URL: https://arxiv.org/abs/2106.03985.

[158] Marcello Benedetti, Mattia Fiorentini, and Michael Lubasch. "Hardware-efficient variational quantum algorithms for time evolution". In: *Phys. Rev. Res.* 3 (3 July 2021), p. 033083. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.033083.

[159] David Amaro et al. "A case study of variational quantum algorithms for a job shop scheduling problem". In: *EPJ Quantum Technology* 9.1 (Feb. 2022), p. 5. URL: https://doi.org/10.1140/epjqt/s40507-022-00123-4.

[160] Keisuke Matsumoto et al. "Calculation of Gibbs partition function with imaginary time evolution on near-term quantum computers". In: *Japanese Journal of Applied Physics* 61.4 (Mar. 2022), p. 042002. URL: https://doi.org/10.35848/1347-4065/ac5152.

[161] Igor O. Sokolov et al. *Orders of magnitude reduction in the computational overhead for quantum many-body problems on quantum computers via an exact transcorrelated method*. 2022. URL: https://arxiv.org/abs/2201.03049.

[162] Conor Mc Keever and Michael Lubasch. *Classically optimized Hamiltonian simulation*. 2022. URL: https://arxiv.org/abs/2205.11427.

[163] Mario Motta et al. "Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution". In: *Nature Physics* 16.2 (2020), pp. 205–210.

[164] Luuk Coopmans, Yuta Kikuchi, and Marcello Benedetti. "Predicting Gibbs-State Expectation Values with Pure Thermal Shadows". In: *PRX Quantum* 4 (1 Jan. 2023), p. 010305. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.010305.

[165] Tzu-Ching Yen, Vladyslav Verteletskyi, and Artur F. Izmaylov. "Measuring all compatible operators in one series of a single-qubit measurements using unitary transformations". In: (2019). arXiv: 1907.09386. URL: http://arxiv.org/abs/1907.09386.

[166] Pauline J. Ollitrault et al. "Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor". In: *Phys. Rev. Res.* 2 (4 Oct. 2020), p. 043140. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.2.043140.

[167] Ayush Asthana et al. *Equation-of-motion variational quantum eigensolver method for computing molecular excitation energies, ionization potentials, and electron affinities*. 2022. URL: https://arxiv.org/abs/2206.10502.

[168] Tyson Jones et al. "Variational quantum algorithms for discovering Hamiltonian spectra". In: *Physical Review A* 99.6 (2019), p. 062304.

[169] Ken M Nakanishi, Kosuke Mitarai, and Keisuke Fujii. "Subspace-search variational quantum eigensolver for excited states". In: *Physical Review Research* 1.3 (2019), p. 033062. URL: https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.1.033062.

[170] Jarrod R. McClean et al. “Hybrid Quantum-Classical Hierarchy for Mitigation of Decoherence and Determination of Excited States”. In: *Phys. Rev. A* 95.4 (Apr. 2017), p. 042308.

[171] William Kirby, Mario Motta, and Antonio Mezzacapo. *Exact and efficient Lanczos method on a quantum computer*. 2022. URL: https://arxiv.org/abs/2208.00567.

[172] Katherine Klymko et al. “Real-Time Evolution for Ultracompact Hamiltonian Eigenstates on Quantum Hardware”. In: *PRX Quantum* 3 (2 May 2022), p. 020323. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.020323.

[173] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Predicting Many Properties of a Quantum System from Very Few Measurements”. In: *Nature Physics* 16.10 (Oct. 2020), pp. 1050–1057.

[174] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Efficient Estimation of Pauli Observables by Derandomization”. In: *Physical Review Letters* 127.3 (July 2021), p. 030503.

[175] Kianna Wan et al. *Matchgate Shadows for Fermionic Quantum Simulation*. July 2022. arXiv: 2207.13723.

[176] Andrew Zhao, Nicholas C. Rubin, and Akimasa Miyake. “Fermionic Partial Tomography via Classical Shadows”. In: *Physical Review Letters* 127.11 (Sept. 2021), p. 110504.

[177] Ahmed A. Akhtar, Hong-Ye Hu, and Yi-Zhuang You. *Scalable and Flexible Classical Shadow Tomography with Tensor Networks*. Sept. 2022. arXiv: 2209.02093.

[178] Christian Bertoni et al. *Shallow Shadows: Expectation Estimation Using Low-Depth Random Clifford Circuits*. Sept. 2022. arXiv: 2209.12924.

[179] Tyson Jones and Simon Benjamin. “QuESTlink—Mathematica embiggened by a hardware-optimised quantum emulator”. In: *Quantum Sci. Techn.* 5.3 (2020), p. 034012.

[180] A. Yu. Kitaev. *Quantum measurements and the Abelian Stabilizer Problem*. 1995. arXiv: quant-ph/9511026 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/9511026.

[181] Norm M. Tubman et al. “Postponing the orthogonality catastrophe: efficient state preparation for electronic structure simulations on quantum devices”. In: *arXiv* (Sept. 2018). arXiv: 1809.05523. URL: http://arxiv.org/abs/1809.05523.

[182] Kenji Sugisaki et al. “Open shell electronic state calculations on quantum computers: A quantum circuit for the preparation of configuration state functions based on Serber construction”. In: *Chemical Physics Letters: X* 1.December 2018 (2019), p. 100002. URL: https://doi.org/10.1016/j.cpletx.2018.100002.

[183] Guoming Wang et al. “Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision”. In: *Quantum* 7 (Nov. 2023), p. 1167. URL: http://dx.doi.org/10.22331/q-2023-11-06-1167.

[184] Lin Lin and Yu Tong. "Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers". In: *PRX Quantum* 3.1 (Feb. 2022). URL: https://doi.org/10.1103%2Fprxquantum.3.010318.

[185] Zhiyan Ding and Lin Lin. "Simultaneous estimation of multiple eigenvalues with short-depth quantum circuit on early fault-tolerant quantum computers". In: *Quantum* 7 (Oct. 2023), p. 1136. URL: https://doi.org/10.22331/q-2023-10-11-1136.

[186] Nick S. Blunt et al. "Statistical Phase Estimation and Error Mitigation on a Superconducting Quantum Processor". In: *PRX Quantum* 4 (4 Dec. 2023), p. 040341. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.040341.

[187] Laura Clinton et al. "Phase Estimation of Local Hamiltonians on NISQ Hardware". In: *arXiv preprint arXiv:2110.13584* (2021).

[188] Nicholas P. Bauman et al. "Toward Quantum Computing for High-Energy Excited States in Molecular Systems: Quantum Phase Estimations of Core-Level States". In: *Journal of Chemical Theory and Computation* 17.1 (2021). PMID: 33332965, pp. 201–210. eprint: https://doi.org/10.1021/acs.jctc.0c00909. URL: https://doi.org/10.1021/acs.jctc.0c00909.

[189] Kenji Sugisaki et al. "Bayesian phase difference estimation: a general quantum algorithm for the direct calculation of energy gaps". In: *Phys. Chem. Chem. Phys.* 23.36 (2021), pp. 20152–20162.

[190] Kenji Sugisaki et al. "Quantum Algorithm for Full Configuration Interaction Calculations without Controlled Time Evolutions". In: *J. Phys. Chem. Lett.* 12.45 (2021), pp. 11085–11089.

[191] Antonio E Russo et al. "Evaluating energy differences on a quantum computer with robust phase estimation". In: *Phys. Rev. Lett.* 126.21 (2021), p. 210501.

[192] Ilia Zintchenko and Nathan Wiebe. "Randomized gap and amplitude estimation". In: *Phys. Rev. A* 93 (6 June 2016), p. 062306. URL: https://link.aps.org/doi/10.1103/PhysRevA.93.062306.

[193] T. Roosevelt. *Theodore Roosevelt: An Autobiography ...* Macmillan, 1913. URL: https://books.google.co.uk/books?id=S9lBAAAAIAAJ.

[194] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F. Izmaylov. "Measurement optimization in the variational quantum eigensolver using a minimum clique cover". In: *The Journal of Chemical Physics* 152.12 (2020), p. 124114. eprint: https://doi.org/10.1063/1.5141458. URL: https://doi.org/10.1063/1.5141458.

[195] Alexis Ralli et al. "Implementation of measurement reduction for the variational quantum eigensolver". In: *Phys. Rev. Res.* 3 (3 Aug. 2021), p. 033195. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.033195.

[196] Max Born and Vladimir Fock. "Beweis des adiabatensatzes". In: *Zeitschrift für Physik* 51.3 (1928), pp. 165–180.

[197] Alán Aspuru-Guzik et al. "Simulated Quantum Computation of Molecular Energies". In: *Science* 309.5741 (2005), pp. 1704–1707. eprint: https://www.science.org/doi/pdf/10.1126/science.1113479. URL: https://www.science.org/doi/abs/10.1126/science.1113479.

[198] Roeland Wiersema et al. “Exploring Entanglement and Optimization within the Hamiltonian Variational Ansatz”. In: *PRX Quantum* 1.2 (Dec. 2020). URL: http://dx.doi.org/10.1103/PRXQuantum.1.020319.

[199] Pieter W. Claeys et al. “Floquet-Engineering Counterdiabatic Protocols in Quantum Many-Body Systems”. In: *Phys. Rev. Lett.* 123 (9 Aug. 2019), p. 090602. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.090602.

[200] P. Chandarana et al. “Digitized-counterdiabatic quantum approximate optimization algorithm”. In: *Phys. Rev. Res.* 4 (1 Feb. 2022), p. 013141. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.013141.

[201] Jonathan Wurtz and Peter J. Love. “Counterdiabaticity and the quantum approximate optimization algorithm”. In: *Quantum* 6 (Jan. 2022), p. 635. URL: https://doi.org/10.22331/q-2022-01-27-635.

[202] R B Stinchcombe. “Ising model in a transverse field. I. Basic theory”. In: *Journal of Physics C: Solid State Physics* 6.15 (Aug. 1973), p. 2459. URL: https://dx.doi.org/10.1088/0022-3719/6/15/009.

[203] Narendra N. Hegade et al. “Shortcuts to Adiabaticity in Digitized Adiabatic Quantum Computing”. In: *Physical Review Applied* 15.2 (Feb. 2021). URL: http://dx.doi.org/10.1103/PhysRevApplied.15.024038.

[204] G. Vidal and C. M. Dawson. “Universal quantum circuit for two-qubit transformations with three controlled-NOT gates”. In: *Phys. Rev. A* 69 (1 Jan. 2004), p. 010301. URL: https://link.aps.org/doi/10.1103/PhysRevA.69.010301.

[205] Xiu-Zhe Luo et al. “Yao.jl: Extensible, Efficient Framework for Quantum Algorithm Design”. In: *Quantum* 4 (Oct. 2020), p. 341. URL: http://dx.doi.org/10.22331/q-2020-10-11-341.

[206] Steven G. Johnson. *The NLopt nonlinear-optimization package*. https://github.com/stevengj/nlopt. 2007.

[207] Patrick Kofod Mogensen and Asbjørn Nilsen Riseth. “Optim: A mathematical optimization package for Julia”. In: *Journal of Open Source Software* 3.24 (2018), p. 615.

[208] Dong C. Liu and Jorge Nocedal. “On the limited memory BFGS method for large scale optimization”. In: *Mathematical Programming* 45 (1989), pp. 503–528.

[209] Janus J. Eriksen et al. “The Ground State Electronic Energy of Benzene”. In: *Journal of Physical Chemistry Letters* 11.20 (2020), pp. 8922–8929. arXiv: 2008.02678.

[210] Oscar Higgott, Daochen Wang, and Stephen Brierley. “Variational Quantum Computation of Excited States”. In: *Quantum* 3 (July 2019), p. 156. arXiv: 1805.08138. URL: https://quantum-journal.org/papers/q-2019-07-01-156/.

[211] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen. *Molecular Electronic-Structure Theory*. 2000.

[212] Qiming Sun et al. "Recent developments in the PySCF program package". In: *The Journal of Chemical Physics* 153.2 (July 2020), p. 024109. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0006074/16722275/024109_1_online.pdf. URL: https://doi.org/10.1063/5.0006074.

[213] Qiming Sun. "Libcint: An efficient general integral library for Gaussian basis functions". In: *Journal of Computational Chemistry* 36.22 (2015), pp. 1664–1671. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.23981. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.23981.

[214] Jarrod R. McClean et al. *OpenFermion: The Electronic Structure Package for Quantum Computers*. 2017. URL: https://arxiv.org/abs/1710.07629.

[215] Igor O. Sokolov et al. "Quantum orbital-optimized unitary coupled cluster methods in the strongly correlated regime: Can quantum algorithms outperform their classical equivalents?" In: *The Journal of Chemical Physics* 152.12 (2020), p. 124107. arXiv: 1911.10864. URL: https://doi.org/10.1063/1.5141835.

[216] Yasunari Suzuki et al. *Qulacs: a fast and versatile quantum circuit simulator for research purpose*. 2020. arXiv: 2011.13524 [quant-ph]. URL: https://arxiv.org/abs/2011.13524.

[217] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272.

[218] Frank Arute et al. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779 (2019), pp. 505–510. arXiv: 1911.00577. URL: http://dx.doi.org/10.1038/s41586-019-1666-5.

[219] Kate Docken and Juergen Hinze. "LiH Potential Curves and Wavefunctions". In: *Journal of Chemical Physics* 57.11 (1972), p. 4928.

[220] Daniel Claudino et al. "Benchmarking Adaptive Variational Quantum Eigensolvers". In: *Frontiers in Chemistry* 8 (Dec. 2020), pp. 1–12. arXiv: 2011.01279. URL: https://www.frontiersin.org/articles/10.3389/fchem.2020.606863/full.

[221] Alexander Cowtan et al. "Phase Gadget Synthesis for Shallow Circuits". In: *Electronic Proceedings in Theoretical Computer Science* 318 (Apr. 2020), pp. 214–229. arXiv: 1906.01734. URL: http://arxiv.org/abs/1906.01734%20http://dx.doi.org/10.4204/EPTCS.318.13%20http://eptcs.web.cse.unsw.edu.au/paper.cgi?QPL2019.13.

[222] Alexander Cowtan, Will Simmons, and Ross Duncan. *A Generic Compilation Strategy for the Unitary Coupled Cluster Ansatz*. 2020. arXiv: 2007.10515 [quant-ph]. URL: https://arxiv.org/abs/2007.10515.

[223] Jakob S. Kottmann, Abhinav Anand, and Alán Aspuru-Guzik. "A feasible approach for automatically differentiable unitary coupled-cluster on quantum computers". In: *Chem. Sci.* 12 (2021), pp. 3497–3508. URL: http://dx.doi.org/10.1039/D0SC06627C.

[224] J. Stanton et al. "A direct product decomposition approach for symmetry exploitation in many-body methods. I. Energy calculations". In: *Journal of Chemical Physics* 94 (1991), pp. 4334–4345.

[225] Jan Lukas Bosse, Raul Santos, and Ashley Montanaro. *Sketching phase diagrams using low-depth variational quantum algorithms*. 2023. arXiv: 2301.09369 [quant-ph].

[226] L. Carroll. *Alice's Adventures in Wonderland*. Lee and Shepard, 1869. URL: https://books.google.co.uk/books?id=btIQAAAAYAAJ.

[227] Hirofumi Nishi et al. *Acceleration of probabilistic imaginary-time evolution method combined with quantum amplitude amplification*. 2022. URL: https://arxiv.org/abs/2212.13816.

[228] Lin Lin and Yu Tong. "Near-optimal ground state preparation". In: *Quantum* 4 (Dec. 2020), p. 372. URL: https://doi.org/10.22331/q-2020-12-14-372.

[229] Xiao Yuan et al. "Theory of variational quantum simulation". In: *Quantum* 3 (Oct. 2019), p. 191. URL: https://doi.org/10.22331/q-2019-10-07-191.

[230] Youle Wang, Guangxi Li, and Xin Wang. "Variational Quantum Gibbs State Preparation with a Truncated Taylor Series". In: *Physical Review Applied* 16.5 (Nov. 2021). URL: https://doi.org/10.1103%2Fphysrevapplied.16.054035.

[231] Connor Powers et al. "Exploring finite temperature properties of materials with quantum computers". In: *Scientific Reports* 13.1 (2023), p. 1986.

[232] Yuchen Wang et al. *Simulation of open quantum system dynamics based on the generalized quantum master equation on quantum computing devices*. 2022. URL: https://arxiv.org/abs/2209.04956.

[233] Hefeng Wang, S. Ashhab, and Franco Nori. "Quantum algorithm for simulating the dynamics of an open quantum system". In: *Phys. Rev. A* 83 (6 June 2011), p. 062317. URL: https://link.aps.org/doi/10.1103/PhysRevA.83.062317.

[234] Nishchay Suri et al. *Two-Unitary Decomposition Algorithm and Open Quantum System Simulation*. 2022. arXiv: 2207.10007 [quant-ph].

[235] Hirsh Kamakari et al. "Digital Quantum Simulation of Open Quantum Systems Using Quantum Imaginary–Time Evolution". In: *PRX Quantum* 3 (1 Feb. 2022), p. 010320. URL: https://link.aps.org/doi/10.1103/PRXQuantum.3.010320.

[236] Anthony W. Schlimgen et al. "Quantum simulation of the Lindblad equation using a unitary decomposition of operators". In: *Phys. Rev. Research* 4 (2 June 2022), p. 023216. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.023216.

[237] Andrew Tranter et al. *InQuanto: Quantum Computational Chemistry*. 2022. URL: https://www.quantinuum.com/products/inquanto.

[238] Seyon Sivarajah et al. "t|ket⟩: a retargetable compiler for NISQ devices". In: *Quantum Science and Technology* 6.1 (Nov. 2020), p. 014003. URL: https://dx.doi.org/10.1088/2058-9565/ab8e92.

[239] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Third. The Johns Hopkins University Press, 1996.

[240] Taichi Kosugi, Hirofumi Nishi, and Yuichiro Matsushita. *Exhaustive search for optimal molecular geometries using imaginary-time evolution on a quantum computer*. 2022. URL: https://arxiv.org/abs/2210.09883.

[241] Jarrod R. McClean et al. "What the foundations of quantum computer science teach us about chemistry". In: *The Journal of Chemical Physics* 155.15 (2021), p. 150901. eprint: https://doi.org/10.1063/5.0060367. URL: https://doi.org/10.1063/5.0060367.

[242] Charles Cerjan. *Numerical grid methods and their application to Schrödinger's equation*. Vol. 412. Springer Science & Business Media, 2013. URL: https://doi.org/10.1007/978-94-015-8240-7_11.

[243] Robert J Harrison et al. "MADNESS: A multiresolution, adaptive numerical environment for scientific simulation". In: *SIAM Journal on Scientific Computing* 38.5 (2016), S123–S142.

[244] Ryan Babbush et al. "Low-Depth Quantum Simulation of Materials". In: *Phys. Rev. X* 8 (1 Mar. 2018), p. 011044. URL: https://link.aps.org/doi/10.1103/PhysRevX.8.011044.

[245] RE Fenna and BW Matthews. "Chlorophyll arrangement in a bacteriochlorophyll protein from Chlorobium limicola". In: *Nature* 258.5536 (1975), pp. 573–577.

[246] Bern Kohler et al. "Controlling the Future of Matter". In: *Accounts of Chemical Research* 28.3 (1995), pp. 133–140. URL: https://doi.org/10.1021/ar00051a006.

[247] M. RÜHLE and M. WILKENS. "CHAPTER 11 - TRANSMISSION ELECTRON MICROSCOPY". In: *Physical Metallurgy (Fourth Edition)*. Ed. by Robert W. CAHN and Peter HAASEN†. Fourth Edition. Oxford: North-Holland, 1996, pp. 1033–1113. URL: https://www.sciencedirect.com/science/article/pii/B9780444898753500168.

[248] Seonghoon Choi and Jiří Vaníček. "Efficient geometric integrators for nonadiabatic quantum dynamics. I. The adiabatic representation". In: *The Journal of Chemical Physics* 150.20 (2019), p. 204112. eprint: https://doi.org/10.1063/1.5092611. URL: https://doi.org/10.1063/1.5092611.

[249] Julien Roulet, Seonghoon Choi, and Jiří Vaníček. "Efficient geometric integrators for nonadiabatic quantum dynamics. II. The diabatic representation". In: *The Journal of Chemical Physics* 150.20 (2019), p. 204113. eprint: https://doi.org/10.1063/1.5094046. URL: https://doi.org/10.1063/1.5094046.

[250] Masuo Suzuki. "General theory of fractal path integrals with applications to many-body theories and statistical physics". In: *Journal of Mathematical Physics* 32.2 (1991), pp. 400–407. eprint: https://doi.org/10.1063/1.529425. URL: https://doi.org/10.1063/1.529425.

[251] Andrew M. Childs et al. "Theory of Trotter Error with Commutator Scaling". In: *Phys. Rev. X* 11 (1 Feb. 2021), p. 011020. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.011020.

[252] Thomas Häner, Martin Roetteler, and Krysta M. Svore. *Optimizing Quantum Circuits for Arithmetic*. 2018. arXiv: 1805.12445 [quant-ph]. URL: https://arxiv.org/abs/1805.12445.

[253] Yuval R. Sanders et al. "Compilation of Fault-Tolerant Quantum Heuristics for Combinatorial Optimization". In: *PRX Quantum* 1 (2 Nov. 2020), p. 020312. URL: https://link.aps.org/doi/10.1103/PRXQuantum.1.020312.

[254] Bill Poirier. *Efficient Evaluation of Exponential and Gaussian Functions on a Quantum Computer*. 2021. arXiv: 2110.05653 [quant-ph]. URL: https://arxiv.org/abs/2110.05653.

[255] Thomas G. Draper. *Addition on a Quantum Computer*. 2000. arXiv: 0008033 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/0008033.

[256] Miroslav Dob ší ček et al. "Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark". In: *Phys. Rev. A* 76 (3 Sept. 2007), p. 030306. URL: https://link.aps.org/doi/10.1103/PhysRevA.76.030306.

[257] R. Kosloff and D. Kosloff. "Absorbing boundaries for wave propagation problems". In: *Journal of Computational Physics* 63.2 (1986), pp. 363–376. URL: https://www.sciencedirect.com/science/article/pii/0021999186901993.

[258] U V Riss and H -D Meyer. "Reflection-free complex absorbing potentials". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 28.8 (Apr. 1995), pp. 1475–1493. URL: https://doi.org/10.1088/0953-4075/28/8/012.

[259] Georges Jolicard and Elizabeth J. Austin. "Optical potential stabilisation method for predicting resonance levels". In: *Chemical Physics Letters* 121.1 (1985), pp. 106–110. URL: https://www.sciencedirect.com/science/article/pii/0009261485871645.

[260] Daniel Neuhauser and Michael Baer. "The time-dependent Schrödinger equation: Application of absorbing boundary conditions". In: *The Journal of Chemical Physics* 90 (Apr. 1989), pp. 4351–4355.

[261] U V Riss and H -D Meyer. "Calculation of resonance energies and widths using the complex absorbing potential method". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 26.23 (Dec. 1993), pp. 4503–4535. URL: https://doi.org/10.1088/0953-4075/26/23/021.

[262] G. G. Balint-Kurti and Á. Vibók. "Complex Absorbing Potentials in Time Dependent Quantum Dynamics". In: *Numerical Grid Methods and Their Application to Schrödinger's Equation*. Ed. by Charles Cerjan. Dordrecht: Springer Netherlands, 1993, pp. 195–205. URL: https://doi.org/10.1007/978-94-015-8240-7_11.

[263] S. Scheit et al. "On the unphysical impact of complex absorbing potentials on the Hamiltonian and its remedy". In: *The Journal of Chemical Physics* 124.3 (2006), p. 034102. eprint: https://doi.org/10.1063/1.2158991. URL: https://doi.org/10.1063/1.2158991.

[264] Craig Gidney and N. Cody Jones. *A CCCZ gate performed with 6 T gates*. 2021. arXiv: 2106.11513 [quant-ph]. URL: https://arxiv.org/abs/2106.11513.

[265] David Poulin and Pawel Wocjan. "Preparing Ground States of Quantum Many-Body Systems on a Quantum Computer". In: *Phys. Rev. Lett.* 102 (13 Apr. 2009), p. 130503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.102.130503.

[266] Tong Liu, Jin-Guo Liu, and Heng Fan. "Probabilistic nonunitary gate in imaginary time evolution". In: *Quantum Information Processing* 20.6 (2021), pp. 1–21.

[267] Tyson Jones et al. "QuEST and High Performance Simulation of Quantum Computers". In: *Sci. Reports* 9 (2019), p. 10736.

[268] Richard Meister. *pyQuEST - A Python interface for the Quantum Exact Simulation Toolkit*. July 2022. URL: https://github.com/rrmeister/pyQuEST.

[269] Andrew Richards. *University of Oxford Advanced Research Computing*. 2015. URL: https://zenodo.org/record/22558.

[270] D. G. W. Parfitt and M. E. Portnoi. "The two-dimensional hydrogen atom revisited". In: *Journal of Mathematical Physics* 43.10 (2002), pp. 4681–4691. eprint: https://aip.scitation.org/doi/pdf/10.1063/1.1503868. URL: https://aip.scitation.org/doi/abs/10.1063/1.1503868.

[271] X. L. Yang et al. "Analytic solution of a two-dimensional hydrogen atom. I. Nonrelativistic theory". In: *Phys. Rev. A* 43 (3 Feb. 1991), pp. 1186–1196. URL: https://link.aps.org/doi/10.1103/PhysRevA.43.1186.

[272] A Mokhtari et al. "Direct femtosecond mapping of trajectories in a chemical reaction". In: *Nature* 348.6298 (1990), pp. 225–227.

[273] Chris Sparrow et al. "Simulating the vibrational quantum dynamics of molecules using photonics". In: *Nature* 557.7707 (2018), pp. 660–667.

[274] James E Hudson et al. "Absolute electron-impact ionization cross sections for a range of C1 to C5 chlorocarbons". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 34.15 (July 2001), pp. 3025–3039. URL: https://doi.org/10.1088/0953-4075/34/15/309.

[275] *Wikipedia page 'Atomic radii of the elements (data page)' collates multiple data sets*. URL: https://www.pnas.org/content/early/2008/11/24/0808245105.

[276] Dulal C Ghosh and Raka Biswas. "Theoretical calculation of absolute radii of atoms and ions. Part 1. The atomic radii". In: *International Journal of Molecular Sciences* 3.2 (2002), pp. 87–113.

[277] Russell D. Johnson III. *NIST Computational Chemistry Comparison and Benchmark Database*. Release 22. May 2022. URL: http://cccbdb.nist.gov.

[278] Michael Beverland et al. "Lower bounds on the non-clifford resources for quantum computations". In: *Quantum Science and Technology* 5 (2020), p. 035009.

[279] Craig Gidney and Austin G. Fowler. *Flexible layout of surface code computations using AutoCCZ states*. 2019. arXiv: 1905.08916 [quant-ph]. URL: https://arxiv.org/abs/1905.08916.

[280] Daniel Litinski. "A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery". In: *Quantum* 3 (Mar. 2019), p. 128. URL: https://doi.org/10.22331/q-2019-03-05-128.

[281] Austin G. Fowler et al. "Surface codes: Towards practical large-scale quantum computation". In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. URL: https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[282] Bálint Koczor. "Exponential Error Suppression for Near-Term Quantum Devices". In: *Phys. Rev. X* 11 (3 Sept. 2021), p. 031057. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.031057.

[283] William J. Huggins et al. “Virtual Distillation for Quantum Error Mitigation”. In: *Phys. Rev. X* 11 (4 Nov. 2021), p. 041036. URL: https://link.aps.org/doi/10.1103/PhysRevX.11.041036.

[284] Hamza Jnane et al. *Multicore Quantum Computing*. 2022. arXiv: 2201.08861 [quant-ph]. URL: https://arxiv.org/abs/2201.08861.

[285] Ramil Nigmatullin et al. “Minimally complex ion traps as modules for quantum communication and computing”. In: *New Journal of Physics* 18.10 (Oct. 2016), p. 103028. URL: https://doi.org/10.1088/1367-2630/18/10/103028.

[286] Nicholas Vence, Robert Harrison, and Predrag Krsti ć. “Attosecond electron dynamics: A multiresolution approach”. In: *Phys. Rev. A* 85 (3 Mar. 2012), p. 033403. URL: https://link.aps.org/doi/10.1103/PhysRevA.85.033403.

[287] Heiko Bauke and Christoph H. Keitel. “Accelerating the Fourier split operator method via graphics processing units”. In: *Computer Physics Communications* 182.12 (2011), pp. 2454–2463. URL: https://www.sciencedirect.com/science/article/pii/S0010465511002414.

[288] Nikitas Stamatopoulos et al. “Option Pricing using Quantum Computers”. In: *Quantum* 4 (July 2020), p. 291. URL: https://doi.org/10.22331/q-2020-07-06-291.

[289] Shouvanik Chakrabarti et al. “A Threshold for Quantum Advantage in Derivative Pricing”. In: *Quantum* 5 (June 2021), p. 463. URL: https://doi.org/10.22331/q-2021-06-01-463.

[290] Fischer Black and Myron Scholes. “The Pricing of Options and Corporate Liabilities”. In: *Journal of Political Economy* 81.3 (1973), pp. 637–654. eprint: https://doi.org/10.1086/260062. URL: https://doi.org/10.1086/260062.

[291] Robert C. Merton. “Theory of Rational Option Pricing”. In: *The Bell Journal of Economics and Management Science* 4.1 (1973), pp. 141–183. URL: http://www.jstor.org/stable/3003143.

[292] C. G. JUNG. “Good and Evil in Analytical Psychology1”. In: *Journal of Analytical Psychology* 5.2 (1960), pp. 91–100. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1465-5922.1960.00091.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1465-5922.1960.00091.x.

[293] Yanwu Gu et al. “Noise-resilient phase estimation with randomized compiling”. In: *arXiv preprint arXiv:2208.04100* (2022).

[294] Kh P Gnatenko, HP Laba, and VM Tkachuk. “Energy levels estimation on a quantum computer by evolution of a physical quantity”. In: *Physics Letters A* 424 (2022), p. 127843.

[295] George EP Box and David A Pierce. “Distribution of residual autocorrelations in autoregressive-integrated moving average time series models”. In: *Journal of the American statistical Association* 65.332 (1970), pp. 1509–1526.

[296] Greta M Ljung and George EP Box. “On a measure of lack of fit in time series models”. In: *Biometrika* 65.2 (1978), pp. 297–303.

[297] David Ramirez, Javier Via, and Ignacio Santamaria. "A generalization of the magnitude squared coherence spectrum for more than two signals: definition, properties and estimation". In: *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE. 2008, pp. 3769–3772.

[298] Sheida Malekpour, John A Gubner, and William A Sethares. "Measures of generalized magnitude-squared coherence: Differences and similarities". In: *Journal of the Franklin Institute* 355.5 (2018), pp. 2932–2950.

[299] Ignacio Santamaria and Javier Via. "Estimation of the magnitude squared coherence spectrum based on reduced-rank canonical coordinates". In: *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*. Vol. 3. IEEE. 2007, pp. III–985.

[300] Monson H Hayes. *Statistical digital signal processing and modeling*. John Wiley & Sons, 1996.

[301] David J Luitz, Nicolas Laflorencie, and Fabien Alet. "Many-body localization edge in the random-field Heisenberg chain". In: *Physical Review B* 91.8 (2015), p. 081103.

[302] Rahul Nandkishore and David A Huse. "Many-body localization and thermalization in quantum statistical mechanics". In: *Annu. Rev. Condens. Matter Phys.* 6.1 (2015), pp. 15–38.

[303] Jutho Haegeman et al. "Elementary Excitations in Gapped Quantum Spin Systems". In: *Phys. Rev. Lett.* 111 (8 Aug. 2013), p. 080401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.111.080401.

[304] Ying Li and Simon C. Benjamin. "Efficient Variational Quantum Simulator Incorporating Active Error Minimization". In: *Phys. Rev. X* 7 (2 June 2017), p. 021050. URL: https://link.aps.org/doi/10.1103/PhysRevX.7.021050.

[305] Ming-Cheng Chen et al. "Demonstration of Adiabatic Variational Quantum Computing with a Superconducting Quantum Coprocessor". In: *Phys. Rev. Lett.* 125 (18 Oct. 2020), p. 180501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.180501.

[306] Bálint Koczor and Simon C. Benjamin. "Quantum natural gradient generalized to noisy and nonunitary circuits". In: *Phys. Rev. A* 106 (6 Dec. 2022), p. 062416. URL: https://link.aps.org/doi/10.1103/PhysRevA.106.062416.

[307] Barnaby van Straaten and Bálint Koczor. "Measurement Cost of Metric-Aware Variational Quantum Algorithms". In: *PRX Quantum* 2 (3 Aug. 2021), p. 030324. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.030324.

[308] IBM Quantum. *https://quantum.ibm.com/*. 2021.

[309] Qiskit contributors. *Qiskit: An Open-source Framework for Quantum Computing*. 2023.

[310] Jonathan Foldager and Balint Koczor. "Can shallow quantum circuits scramble local noise into global white noise?" In: *arXiv preprint arXiv:2302.00881* (2023).

[311] J. M. Foster and S. F. Boys. "Quantum Variational Calculations for a Range of CH_2 Configurations". In: *Rev. Mod. Phys.* 32 (2 Apr. 1960), pp. 305–307. URL: https://link.aps.org/doi/10.1103/RevModPhys.32.305.

[312] Henry F. Schaefer. "Methylene: A Paradigm for Computational Quantum Chemistry". In: *Science* 231.4742 (1986), pp. 1100–1107. URL: http://www.jstor.org/stable/1696782 (visited on 12/17/2022).

[313] James F. Harrison. "Methylene: A Personal Perspective". In: *Advances in the Theory of Atomic and Molecular Systems: Conceptual and Computational Advances in Quantum Chemistry*. Ed. by Piotr Piecuch et al. Dordrecht: Springer Netherlands, 2009, pp. 33–43. URL: https://doi.org/10.1007/978-90-481-2596-8_2.

[314] C. David Sherrill et al. "Structures and vibrational frequencies in the full configuration interaction limit: Predictions for four electronic states of methylene using a triple-zeta plus double polarization (TZ2P) basis". In: *The Journal of Chemical Physics* 108.3 (1998), pp. 1040–1049. eprint: https://doi.org/10.1063/1.475465. URL: https://doi.org/10.1063/1.475465.

[315] Libor Veis and Jiří Pittner. "Quantum computing applied to calculations of molecular energies: CH2 benchmark". In: *The Journal of Chemical Physics* 133.19 (2010), p. 194106. eprint: https://doi.org/10.1063/1.3503767. URL: https://doi.org/10.1063/1.3503767.

[316] Libor Veis and Jiří Pittner. "Adiabatic state preparation study of methylene". In: *The Journal of Chemical Physics* 140.21 (2014), p. 214111. eprint: https://doi.org/10.1063/1.4880755. URL: https://doi.org/10.1063/1.4880755.

[317] Kenji Sugisaki et al. "A quantum algorithm for spin chemistry: a Bayesian exchange coupling parameter calculator with broken-symmetry wave functions". In: *Chem. Sci.* 12 (6 2021), pp. 2121–2132. URL: http://dx.doi.org/10.1039/D0SC04847J.

[318] Isaiah Shavitt. "Geometry and singlet-triplet energy gap in methylene: a critical review of experimental and theoretical determinations". In: *Tetrahedron* 41.8 (1985), pp. 1531–1542. URL: https://www.sciencedirect.com/science/article/pii/S0040402001963938.

[319] B. Foxen et al. "Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms". In: *Phys. Rev. Lett.* 125 (12 Sept. 2020), p. 120504. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.120504.

[320] Kenji Sugisaki et al. "Quantum Chemistry on Quantum Computers: A Method for Preparation of Multiconfigurational Wave Functions on Quantum Computers without Performing Post-Hartree–Fock Calculations". In: *ACS Central Science* 5.1 (2019). PMID: 30693335, pp. 167–175. eprint: https://doi.org/10.1021/acscentsci.8b00788. URL: https://doi.org/10.1021/acscentsci.8b00788.

[321] Seunghoon Lee et al. *Is there evidence for exponential quantum advantage in quantum chemistry?* 2022. URL: https://arxiv.org/abs/2208.02199.

[322] Adam D Bookatz. "QMA-complete problems". In: *arXiv preprint arXiv:1212.6312* (2012).

[323] Toby S Cubitt, David Perez-Garcia, and Michael M Wolf. "Undecidability of the spectral gap". In: *Nature* 528.7581 (2015), pp. 207–211.