Quantum computational chemistry methods for early-stage quantum computers

Yordan Stefanov Yordanov

Cavendish Laboratory, Department of Physics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Homerton College   July 2021

This dissertation is dedicated to my parents who introduced me to science and have always supported me.

Declaration

The work presented in this thesis was carried out at the Thin Film Magnetism and Quantum Information group at the Cavendish Laboratory, University of Cambridge, between October 2018 and June 2021. I state that this dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specified in the text. I further state that no part of my dissertation has already been submitted, or, is being concurrently submitted for any degree, diploma or other qualification at the University of Cambridge or any other University or similar academic institution. The thesis is fewer than 60, 000 words long and does not exceed the prescribed word limit of the Degree Committee.

Yordan Stefanov Yordanov
July 2021

Acknowledgements

Foremost I would like to express my immense gratitude towards my supervisor Prof Crispin Barnes for his continuous support and motivation, both academically and personally. His expertise in quantum mechanics and quantum information has guided me during all the years of Ph.D. research. His belief in giving students academic freedom has allowed me to explore areas of great interest to me and has shaped me into an independent researcher. I could not have imagined a better mentor for my Ph.D. degree.

I am grateful to my colleagues from the *Thin Film Magnetism and Quantum Information* group at the Cavendish Laboratory for providing me with an academically inspiring and friendly work environment. In particular I would like to thank: Dr David Arvidsson-Shukur, who has acted as my second supervisor, and who has been a close friend of mine outside the office; Aleksander Lasek with whom I shared the academic conference experience and learnt how to ski; Dr Hugo Lepage, who has always brought fun and cheers to the office; Dr Echo Zhang, who has been supplying the office with exotic sweets and teas; Niall Devlin who has been a source of great stories; and Jacob Chevalier-Drori, with whom I worked on a number of papers, and who tried to teach me proper maths.

I also want to express my thanks to the members of the Hitachi Cambridge Laboratory for assisting me in my research. In particular I wish to thank: Dr Aleksey Andreev, with whom I worked on many interesting problems during my first year of research; Dr Thierry Ferrus, with whom I have had many useful discussion and who have provided me with a link to industry; and Dr Normann Mertig who introduced me to the field of Quantum Machine Learning.

Last but not least, I wish to thank my close friends: Katerina Naydenova for standing by my side since the first year of my undergraduate study, and for proofreading many of my manuscripts; and Alexander Sklqrov for teaching me proper programming.

Abstract

One of the first practical applications of quantum computers is expected to be molecular modelling. Performing this task would profoundly affect areas such as chemistry, materials science and drug synthesis. Modelling of molecules, which are classically intractable, can be achieved with just over 30 qubits, whereas state of the art quantum computers already have more than 50 qubits. The Variational Quantum Eigensolver (VQE) algorithm and VQE based protocols, are promising candidates to enable this task on emerging Noisy Intermediate-Scale Quantum (NISQ) computers. These protocols require short quantum circuits and short coherence times, and are particularly resilient to quantum errors. Nevertheless, there is still a significant gap between the accuracy and the coherence times of current NISQ computers, and the hardware requirements of VQE protocols to simulate practically interesting molecules. In this thesis, I present my contribution to narrowing this gap by developing VQE protocols for molecular modelling that are less demanding on quantum hardware.

The VQE relies on the Rayleigh-Ritz variational principle to estimate the eigenvalues of a Hamiltonian operator, by minimizing its expectation value with respect to a trial quantum state, prepared by an ansatz. A major challenge for the practical realisation of VQE protocols on NISQ computers is to construct an ansatz that: (1) can accurately approximate the eigenstates of the Hamiltonian; (2) is easy to optimize; and (3) can be implemented by a shallow circuit, within the capabilities of a NISQ computer. The most widely used, unitary coupled cluster (UCC), type of ansätze mathematically correspond to a product of unitary evolutions of fermionic excitation operators. Owing to their fermionic structure, UCC ansätze preserve the symmetries of electronic wavefunctions, and thus are accurate and easy to optimize. Nevertheless, UCC ansätze are implemented by high depth circuits, which severely limit the size of the molecules that can be reliably simulated on NISQ computers. In this thesis, I begin by constructing efficient quantum circuits to perform evolutions of fermionic excitation operators. The circuits are optimized in the number of two-qubit entangling gates, which are the current bottleneck of NISQ computers. Compared to the standard circuits used to implement evolutions of fermionic excitation operators, the circuits derived in this thesis reduce the number of two-qubit entangling gates by more than 70% on average. As an intermediate result, I also derive efficient circuits to perform evolutions of

qubit excitation operators (excitation operators that account for qubit, rather than fermionic commutation relations).

Even with the fermionic-excitation-evolution circuits derived here, UCC ansätze still require very long circuits, with a particularly large number of two-qubit entangling gates. In this thesis, I consider the use of alternative VQE ansätze, based on evolutions of qubit excitation operators. Due to not accounting for fermionic anticommutation, evolutions of qubit excitation operators can be performed by circuits that require asymptotically fewer two-qubit entangling gates. Furthermore, qubit excitation operators preserve many of the physical properties of fermionic excitation operators. Performing a number of classical numerical VQE simulations for small molecules, I show that qubit-excitation-based ansätze can approximate molecular electronic wavefunctions almost as accurately as fermionic-excitation-based ansätze. Hence, I argue that evolutions of qubit excitation operators are more suitable to construct molecular ansätze than evolutions of fermionic excitation operators, especially in the era of NISQ computers.

Motivated by the advantage of qubit-excitation-based ansätze, I introduce the qubit-excitation-based adaptive variational quantum eigensolver (QEB-ADAPT-VQE). The QEB-ADAPT-VQE belongs to a family of ADAPT-VQE protocols for molecular modelling that grow a problem-tailored ansatz by iteratively appending unitary operators sampled from a predefined finite-size pool of operators. The operator at each iteration is sampled based on an ansatz-growing strategy, which aims to achieve the lowest estimate for the Hamiltonian expectation value at each iteration. In this way, ADAPT-VQE protocols construct shallow-circuit, few-parameter ansätze tailored specifically to the molecular systems of interest. In the case of the QEB-ADAPT-VQE, the operator pool is defined by a set of evolutions of single and double qubit excitation operators. I benchmark the performance of the QEB-ADAPT-VQE, by performing classical numerical simulations. I demonstrate that it can construct ansätze that are several orders of magnitude more accurate, and require significantly shallower circuits, than standard UCC ansätze. I also compare the QEB-ADAPT-VQE against the original fermionic-ADAPT-VQE, which utilizes a pool of fermionic excitation evolutions, and the qubit-ADAPT-VQE, which utilizes a pool of Pauli-string evolutions. I demonstrate that, in terms of circuit efficiency and convergence speed, the QEB-ADAPT-VQE systematically outperforms the qubit-ADAPT-VQE, which to my knowledge was the previous most circuit-efficient, scalable VQE protocol for molecular modeling. The QEB-ADAPT-VQE protocol, therefore represents a significant improvement in the field of VQE protocols for molecular modelling and brings us closer to achieving practical quantum advantage.

Lastly, I outline a modified version of the QEB-ADAPT-VQE, the excited-QEB-ADAPT-VQE, designed to estimate energies of excited molecular states. The excited-QEB-ADAPT-VQE is more robust to initial simulation conditions, at the expense of increased computational complexity.

Table of contents

List of Publications

Chapter 2 of the thesis is based on my published article "Efficient quantum circuits for quantum computational chemistry" [Phys. Rev. A 102, 062612(2020)]. Some of the circuit constructions used in this chapter are derived in another of my publications, "Implementation of a general single-qubit positive operator-valued measure on a circuit-based quantum computer" [Phys. Rev. A 100, 062317(2019)].

Chapters 3 and 4 are largely based on my preprint article "Iterative qubit-excitation based variational quantum eigensolver" [arXiv:2011.10540(2020)]. This article is currently under review by Nature Communications Physics. I note that due to a suggestion by the reviewers of Nature Communications Physics, this article will be renamed to "Qubit-excitation based ADAPT-VQE".

Chapter 5 is based on my preprint article "Molecular excited state VQE simulations with iteratively constructed qubit-excitation-based ansätze" [arXiv:submit/3788731(2021)].

List of figures

List of tables

Nomenclature

Roman Symbols

O Hermitian operator corresponding to an observable

e Natural number $e \approx 2.72..$

o eigenvalue of O

$a^\dagger, a$ Fermionic ladder operators

$b^\dagger, b$ Qubit ladder operators

$F(\theta)$ Unitary evolution of a fermionic excitation operator (fermionic evolution)

H Hamiltonian operator

I The identity operator

$Q(\theta)$ Unitary evolution of a qubit excitation operator (qubit evolution)

U Unitary operator

n_E Energy level of an electron orbital

n_q Number of qubits

N_e Number of electrons

N_n Number of nuclei

N_{so} Number of spin-orbitals

q_i The state of qubit i

E_k The k^{th} energy eigenvalue of a Hamiltonian operator H

$R_x(\theta)_i$ Single qubit x-rotation by θ, acting on qubit i

$R_y(\theta)_i$ Single qubit y-rotation by θ, acting on qubit i

$R_z(\theta)_i$ Single qubit z-rotation by θ, acting on qubit i

X_i Pauli-X gate, acting on qubit i

Had_i Hadamard gate, acting on qubit i

Y_i Pauli-Y gate, acting on qubit i

Z_i Pauli-Z gate, acting on qubit i

$cR_x(\theta)_{i,\{kl..\}}$ Single qubit x-rotation by θ, acting on qubit i, and controlled by the state of qubits $\{k,l..\}$

$cR_y(\theta)_{i,\{kl..\}}$ Single qubit y-rotation by θ, acting on qubit i, and controlled by the state of qubits $\{k,l..\}$

$cR_z(\theta)_{i,\{kl..\}}$ Single qubit z-rotation by θ, acting on qubit i, and controlled by the state of qubits $\{k,l..\}$

$\mathscr{E}$ The environment

$\hbar$ Reduced Planck constant

m_e The mass of an electron

r,R Distance

Greek Symbols

π Number Pi$\approx 3.14..$

ψ Wavefunction

Ψ Wavefunction

Φ Slater determinant

ϕ A single electron spin-orbital wavefunction

θ A real valued parameter

$\vec{\theta}$ A vector of independent real valued parameters

δ_{ij} Kronecker delta

ε Energy threshold

ε_0 Vacuum permittivity

Ξ Fermionic excitation operator

Ω Qubit excitation operator

Acronyms / Abbreviations

VQE   Variational Quantum Eigensolver

MBQC   Measurement Based Quantum Computer

NISQ   Noisy Intermediate-Scale Quantum

$CNOT_{ij}$ controlled-NOT quantum gate, with control qubit i and target qubit j

CP_{ij} controlled-Phase quantum gate, acting on qubit i and j

$SWAP_{ij}$ SWAP quantum gate

CC   Coupled cluster method

CI   Configuration interaction method

FCI   Full configuration interaction

BK   Bravyi-Kitaev qubit encoding

JW   Jordan-Wigner qubit encoding

UCC   unitary coupled cluster

QEB-ADAPT-VQE Qubit-excitation-based adaptive VQE

Chapter 1

Introduction and Theoretical Background

1.1 Motivation

I would like to begin by motivating this thesis with respect to the current state of quantum computation. In the beginning of the last century a series of crises in physics arose due to classical physics failing to predict phenomena like the black body radiation [1], and the photoelectric effect [2]. These crises provoked the gradual development of the theory of quantum mechanics, in the 1920s, by Schrödinger, Bohr, Heisenberg, Born, Einstein and others. Quantum mechanics has been hugely successful in explaining and predicting a variety of physical phenomena, ranging from the nuclear processes inside stars to superconductivity [3] and photoluminescence [4].

Quantum mechanics predicts the outcomes of quantum experiments in a probabilistic form. The outcome probabilities can be accurately calculated within the framework of the 'postulates of quantum mechanics'[1]. There are multiple formulations of these postulates; in this thesis I refer to the following formulation:

1. The state of a quantum mechanical system is represented by a complex vector (statevector) $|\psi\rangle$ in the Hilbert space[2] of the system. The statevector must also be normalizable such that $\langle\psi|\psi\rangle = 1$.

---

[1]The origin of the postulates is a subject of the field of 'Interpretations of Quantum Mechanics'. Adopting a particular interpretation [5–8] is not required for the studies presented in the thesis.

[2]The Hilbert space, named after David Hilbert, is a complex vector space with an inner product $\langle f, g \rangle$, such that $|f| = \langle f, f \rangle^{\frac{1}{2}}$.

2. To each observable property of the quantum system, there is a corresponding Hermitian operator O on the Hilbert space.

3. The outcome of a measurement of an observable, with corresponding Hermitian operator O, is one of the eigenvalues, $\{o_k\}$, of O. The immediate repetition of the measurement, yields the same measurement outcome.

4. The probability to obtain o_k in the measurement above, for a system in state $|\psi\rangle$, is $P_k = |\langle o_k|\psi\rangle|^2$. Hence, the expectation value of the observable is $\langle\psi|O|\psi\rangle$.

5. The evolution of $|\psi\rangle$ is described by the time-dependent Schrödinger equation $i\hbar\frac{\partial}{\partial t}|\psi\rangle = H|\psi\rangle$, where H is the Hamiltonian operator of the system. For a closed system the evolution is a unitary process.

6. The state of a composite quantum system is a vector in the tensor product of the Hilbert spaces of the constituent systems.

   Despite being mathematically prosaic, the postulates of quantum mechanics lead to interesting non-classical principles. Here I will remark on two of these principles that are central to quantum computation. First, since the Schrödinger equation is linear, any linear combination of its solutions is also a solution. This implies that states of quantum systems, much like waves in classical physics, can be added (or 'superposed') together, leading to the fundamental quantum mechanical principle known as 'quantum superposition'. Second, the principle of quantum superposition together with postulate 6 imply that a composite quantum system can exist in a superposition of product states of its constituent systems. This leads to the quantum principle of 'entanglement', where the states of two or more systems can be 'entangled', so that the states cannot be described independently and a measurement of any of the systems would infer information about the measurement probabilities of the other system(s).

   Beginning in the 1980s, with Paul Benioff's proposition of a quantum mechanical Turing machine [9, 10], a number of physicists and computer scientists, suggested that the principles of superposition and entanglement can be used to realize 'quantum computing' that can be superior to classical computing in performing certain tasks. One of the main mechanisms to achieve quantum advantage lies in utilizing a superposition state to encode and thus simultaneously process multiple inputs to a computational problem. In this way, for certain computational tasks, a quantum speed-up is achieved in comparison to processing the inputs classically in a sequential manner. In 1982 Richard Feynman suggested an application for quantum computaters to simulate quantum systems more efficiently than classical computers [11]. In the following decade a number of quantum algorithms, for various computational

problems were developed, which demonstrated reduced asymptotic complexities as compared to their classical counterparts. Some of the best known examples are the Deutch-Jozsa algorithm to determine if a function is balanced or constant in a constant number of steps [12], Grover's algorithm to find a unique input to a black-box-function in quadratically fewer steps than a classical computer [13], and Shor's algorithm to factor integers [14] in almost exponentially fewer steps than the best known classical algorithm.

In order to implement quantum algorithms, physicists will have to construct a quantum computer, a quantum device that could store, evolve unitarily in a controlled manner and measure a quantum state. The state of a quantum computer is represented by a finite register of n quantum bits, called *qubits* that, unlike classical bits, can exist in any superposition state of $|0\rangle$ and $|1\rangle$. In the prevailing *quantum-circuit* model of a quantum computer the qubit-register state is evolved by circuits of *quantum gates*, where each quantum gate corresponds to a unitary operation on the Hilbert space of the qubits. At the end, the outcome is obtained by individually measuring each qubit in the $\{|0\rangle, |1\rangle\}$ basis.

Despite the success of developing quantum algorithms, the process of constructing large and accurate quantum computers has been slow and full of technical obstacles owing to quantum errors and scalability problems. A major challenge is to isolate the quantum computer state from uncontrolled interactions with the environment. Such interactions destroy entanglement and superposition in a quantum computer over time, causing the quantum state to 'decohere', thus limiting the maximum run-time for tasks that can be performed reliably. Imperfect qubit control is another source of quantum errors, which even if not of decohering nature can still corrupt the outcome and limit significantly the capabilities of quantum computers. Error correction protocols allow quantum errors to be corrected faster than they are introduced, under the condition that the error rate is below a certain threshold [15–18]. However, such protocols require large numbers of qubits and will likely be difficult to implement for practically useful applications in the near future.

Nevertheless, there has been significant progress during the last decade, in constructing ever larger and more accurate quantum computers. Existing prototypes, constructed by IBM [19], and Google [20], can control more than 50 qubits and perform qubit operations with error rates on the order of 10^{-2}. In 2018 Preskill termed such quantum computers, characterized by 50-100 qubits and imperfect qubit control, as Noisy Intermediate-Scale Quantum (NISQ) computers [21]. In 2019 Google demonstrated [20] that their quantum computer can perform a task that, albeit not practically useful, is difficult for a classical

computer [3]. It is believed that in the foreseeable future emerging NISQ computers will also be able to perform useful classically intractable tasks.

A tasty fruit for NISQ computers is to perform molecular modelling[4], as modelling of classically intractable molecules can be done with just over 30 qubits [23]. Modelling molecules may help to explain and predict bio-chemical reactions [24], and help design new compounds for applications in the chemical and the drug industries [25, 26]. A promising algorithm to perform this task on NISQ computers is the Variational Quantum Eigensolver. The VQE is a hybrid quantum-classical algorithm that utilizes the Rayleigh Ritz variational principle to determine the lowest eigenvalue of a Hamiltonian operator, by optimizing a trial state, generated by an *ansatz*[5]. In particular the VQE can be used to solve the electronic structure problem [23, 27], and find the energies of the ground and low excited states of a molecule. By utilizing both a quantum and a classical computer, the VQE is less quantum-hardware intensive at the expense of requiring more quantum measurements and classical post-processing as compared to purely quantum algorithms for eigenvalue determination, e.g. the Quantum Phase Estimation (QPE) algorithm [28]. Moreover, due to its variational nature, the VQE is also particularly resilient to quantum errors [23].

The major component of the VQE is the ansatz, used to approximate the molecular electronic wavefunction. The ansatz mathematically corresponds to a parametrized unitary operator, and is implemented by an *ansatz circuit*. A central challenge for the successful implementation of the VQE algorithm on NISQ computers is to construct an ansatz that: (1) has the variational flexibility required to approximate the ground state wavefunction with sufficient accuracy;(2) can be efficiently optimized[6]; and (3) can be performed by a shallow quantum circuit, which does not exceed the capabilities of the quantum computer.

The last few years have seen significant efforts towards the development of protocols for molecular modelling based on the VQE, with the goal of constructing better ansätze. Despite these efforts and the success of the VQE protocols in modelling small molecules on real quantum computers [29–31], there is still a large gap between the capabilities of state of the art NISQ computers and the hardware requirements of VQE protocols for modelling larger molecules of practical interest. In this thesis, I pursue the goal of narrowing this gap

---

[3]Shortly after Ref. [20] was published, the researchers at IBM showed that the task performed by Google's quantum computer can be simulated on a classical supercomputer in just a few days [22].

[4]Molecular modelling encompasses all methods, theoretical and computational, used to model the behaviour of molecules.

[5]In the context of computational chemistry an *ansatz* is a parameterized unitary operator that acts on an initial reference state to generate a trial state.

[6]The classical optimization of an ansatz is efficient if it is done in a number of steps that scales polynomially with the number of qubits.

by improving existing techniques for molecular VQE ansatz construction, and developing
new VQE protocols for molecular modelling.

1.2   Quantum computing

1.2.1   Universal quantum computers

The goal of quantum computing is to use the principles of quantum superposition and
entanglement to perform computation superior to classical computing. Mathematically
quantum computing corresponds to unitarily transforming a statevector, representing a
quantum state, from an initial to a final point in a Hilbert space, followed by a projection
(measurement) onto orthonormal basis of statevectors. Hence, a quantum computer is a
quantum hardware device that can store, evolve unitarily, and measure a quantum state.

Analogously to a classical computer, the state of a quantum computer is represented
by a register of n qubits[7]. However, unlike a classical bit, the state of a qubit can be any
superposition of $|0\rangle$ and $|1\rangle$, which is represented by a statevector in a 2-dimensional Hilbert
space. Hence, the state of an n-qubit register can be represented by a statevector in an
2^n-dimensional Hilbert space. A universal, programmable quantum computer should be able
to perform or to approximate to an arbitrary accuracy, any unitary operation on the Hilbert
space of its qubit register. There are a number of quantum computer models, distinguished
by the manner in which a general unitary operation is performed.

Throughout this thesis I will consider the *quantum-circuit* model of a quantum computer,
since the majority of quantum computer prototypes, including NISQ computers, are based
on it. A quantum-circuit computer performs a general unitary operation, on the state of its
qubit register, as a circuit (sequence) of quantum gates, where each gate corresponds to an
elementary unitary operation, which can be directly performed by the quantum computer
hardware. In Sec. 1.2.3 the operation of the quantum-circuit model of a quantum computer
is described, and examples of common quantum gates are presented.

Another promising model of a universal quantum computer is the *measurement-based*[8]
quantum computer (MBQC) [35, 36]. The MBQC uses two registers of target and auxiliary
qubits. Computation starts by preparing the two qubit registers in a entangled state, e.g.
a cluster state [35] or a graph state [37]. Then, the state of the target qubits is evolved
by performing adaptively single qubit measurements on the auxiliary register. Although
the operations of MBQCs and quantum-circuit computers are conceptually different, both

---

[7]There are proposals, particularly in the field of quantum cryptography, to use d-dimensional quantum
systems, called 'qudits', instead of qubits [32–34].

[8]Also known as 'One-way' or 'Cluster-state' quantum computer.

quantum computer models can be realized by the same quantum hardware systems. MBQCs are suitable for systems where measurements are more easily performed than controlled unitary evolutions.

Other models of, not necessarily universal, quantum computers are adiabatic quantum computers, where computation is decomposed into a continuous transformation of an initial Hamiltonian to a final Hamiltonian [38], and topological quantum computers, where computation is based on braiding anyon quasiparticles on a 2-dimensional lattice [39].

1.2.2   Linear-algebraic representation of qubit states and operations

Before I proceed with a description of the quantum-circuit model of a quantum computer, let me describe the adopted linear-algebraic representation of qubit states and qubit operations.

Density matrices

A general quantum state ψ is described by a density matrix ρ_ψ. A density matrix is a Hermitian matrix that corresponds to a classical distribution of pure quantum states $|\phi_i\rangle$

$$\rho_\psi = \sum p_i |\phi_i\rangle\langle\phi_i| \equiv \begin{bmatrix} p_1 & 0 & \dots \\ 0 & p_2 & \dots \\ \vdots & & \ddots \end{bmatrix}, \tag{1.1}$$

where $\{p_i\}$ are (real positive) probabilities that sum to 1, and $\{|\phi_i\rangle\}$ is a complete basis of pure quantum states. If $p_k = 1$ and $p_{i \neq k} = 0$, then $\rho_\psi = |\phi_k\rangle\langle\phi_k|$ and ψ is the pure quantum state $|\phi_k\rangle$. A pure quantum state can be represented by a statevector in the Hilbert space of the corresponding quantum system.

Density matrices and statevectors are different formalisms used to describe quantum states. Using density matrices to represent qubit states is useful when studying quantum decoherence processes, which involve non-unitary evolutions. Studying decoherence processes, though, is not central to this thesis. Therefore, unless otherwise stated, the qubit states considered in this thesis will be assumed to be pure, and will be represented by statevectors in the Hilbert spaces of the corresponding qubit registers.

Single-qubit states and the Bloch sphere

The qubit is 2-level quantum system, defined by states $|0\rangle$ and $|1\rangle$. A general (pure) qubit state is represented by a statevector in a 2 dimensional Hilbert space, which can be described

by a point on the surface of the Bloch sphere[9] as shown in Fig. 1.1. Hence, any single-qubit operation (a unitary operation on the state of one qubit) can be represented as a combination of two rotations around the x, y or z axes of the Bloch sphere. For example a general single-qubit state can be obtained as

$$|q\rangle = R_z(\theta_2)R_y(\theta_1)|0\rangle = \cos\frac{\theta_1}{2}|0\rangle + \sin\frac{\theta_1}{2}e^{i\theta_2}|1\rangle, \qquad (1.2)$$

where $R_y(\theta_1)$ and $R_z(\theta_1)$ are single-qubit rotations around the y and z axes of the Bloch sphere, respectively.
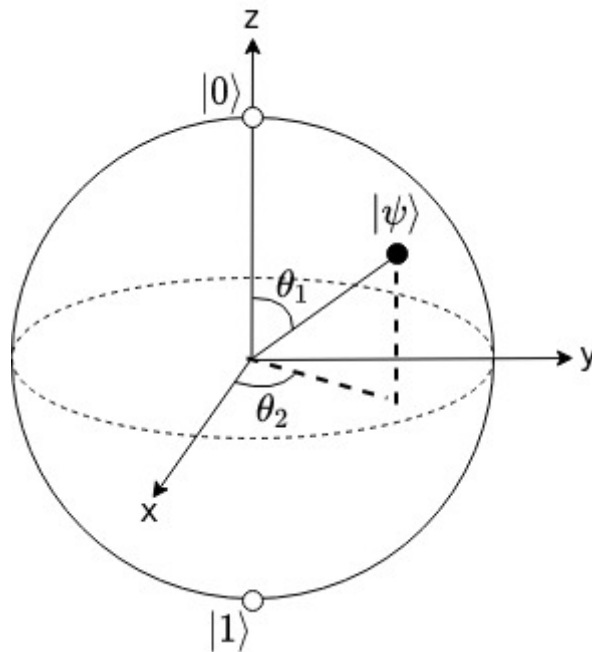
Fig. 1.1 The Bloch sphere geometrically represents the state of a qubit. The θ_1 and θ_2 angles define a one-to-one correspondence between a point of the surface on the Bloch sphere and a general pure qubit state $|q\rangle = \cos\frac{\theta_1}{2}|0\rangle + \sin\frac{\theta_1}{2}e^{i\theta_2}|1\rangle$.

Multi-qubit states

Using the Bloch sphere to visualize single-qubit states is convenient, however, it is not applicable to multi-qubit states. An n-qubit state is represented by a statevector in an 2^n-dimensional Hilbert space. Qubit statevectors are usually expressed in the computational basis. For a single qubit the computational basis states are $|0\rangle$ and $|1\rangle$, and the corresponding

---

[9]A mixed qubit state would lie within the sphere.

statevectors are

$$|0\rangle \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } |1\rangle \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{1.3}$$

For n qubits, the computational basis is the tensor product of the bases of all n qubits, such that the n-qubits basis state with binary value i corresponds to the statevector in the Hilbert space with $(i+1)^{\text{th}}$ non-zero component; e.g.

$$|00..\rangle \equiv \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{pmatrix}, |10..\rangle \equiv \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ \vdots \end{pmatrix}, \dots |1..11\rangle \equiv \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix} \tag{1.4}$$

Qubit operations

An n-qubit operation can be represented by a $2^n \times 2^n$ matrix acting on an n-qubits statevector. For example the two single-qubit rotations in Eq. 1.2 can be expressed by the unitary matrices

$$R_y(\theta) = \begin{bmatrix} \cos\frac{\theta_1}{2} & -\sin\frac{\theta_1}{2} \\ \sin\frac{\theta_1}{2} & \cos\frac{\theta_1}{2} \end{bmatrix} \text{ and } R_z(\theta_2) = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta_2} \end{bmatrix}. \tag{1.5}$$

Hence, the general single qubit state in Eq. 1.2 can be expressed as

$$|q\rangle = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta_2} \end{bmatrix} \begin{bmatrix} \cos\frac{\theta_1}{2} & -\sin\frac{\theta_1}{2} \\ \sin\frac{\theta_1}{2} & \cos\frac{\theta_1}{2} \end{bmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \cos\frac{\theta_1}{2} \\ e^{i\theta_2}\sin\frac{\theta_1}{2} \end{pmatrix} \tag{1.6}$$

I will frequently use matrix representations to describe and visualize the actions of various quantum circuits and quantum gates.

1.2.3 Quantum-circuit model of a quantum computer
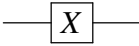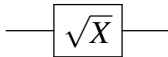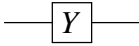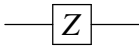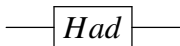
Universal quantum gate sets

The quantum-circuit model of a quantum computer is an analogue to the classical digital computer, where computation is performed by circuits of logic gates, with the difference that classical logic gates can be irreversible, whereas quantum gates are reversible unitary operations. Similarly to the case of classical computers, constructing a quantum computer that can directly implement any multi-qubit quantum gate operation is impossible.

Therefore, a universal quantum computer requires a finite set of basic quantum gates that can be combined to express or to approximate with an arbitrary accuracy any possible unitary operation on the Hilbert space of the qubit register. Such sets are referred to as *universal quantum-gate sets*. DiVizenzo showed [40, 41] that 1-qubit and 2-qubit quantum gates are sufficient to form a universal quantum-gate set[10]. An example of a minimal universal quantum-gate set is $\{Had, R_z(\frac{\pi}{8}), CNOT\}$ [43], where *Had* is the Hadamard gate, $R_z(\frac{\pi}{8})$ is 1-qubit rotation by $\frac{\pi}{8}$ around the *z*-axis of the Bloch sphere, and *CNOT* is a 2-qubit controlled-NOT operation (see Sec. 1.2.3). In general, universal quantum-gate sets consist of one entangling 2-qubit gate, and a set of (at least two) 1-qubit gates that can approximate any single-qubit gate.

Basic quantum gates
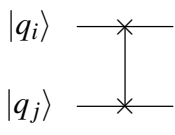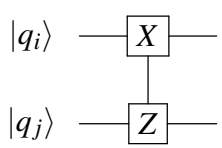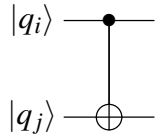
Here I summarize some of the basic 1-qubit and 2-qubit quantum gates that are commonly considered as fundamental building blocks of a quantum circuit. A list of the 1-qubit quantum gates is given in Table 1.2, and a list of the 2-qubit gates in Table 1.4.

Quantum gate	Symbol	Matrix
Pauli-X (X)	$-\boxed{X}-$	$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$
Pauli-X ($\sqrt{X}$)	$-\boxed{\sqrt{X}}-$	$\begin{bmatrix} \frac{1+i}{2} & \frac{1-i}{2} \\ \frac{1-i}{2} & \frac{1+i}{2} \end{bmatrix}$
Pauli-Y (Y)	$-\boxed{Y}-$	$\begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$
Pauli-Z (Z)	$-\boxed{Z}-$	$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$
Hadamard (*Had*)	$-\boxed{Had}-$	$\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$

---
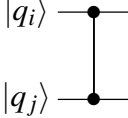
[10]In fact DiVizenzo showed that 2-qubit quantum gates, based on the exchange interaction, alone are sufficient for universal quantum computing, if logical qubits are encoded in the state of 3 physical qubits [42].

x-Pauli rotation by θ $\left(R_x(\theta)\right)$ $\boxed{R_x(\theta)}$ $\begin{bmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}$

y-Pauli rotation by θ $\left(R_y(\theta)\right)$ $\boxed{R_y(\theta)}$ $\begin{bmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}$

z-Pauli rotation by θ $\left(R_z(\theta)\right)$ $\boxed{R_z(\theta)}$ $\begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{bmatrix}$

Table 1.2 List of basic single-qubit quantum gates

Quantum gate	Symbol	Matrix
$\mathrm{SWAP}_{ij}{}^{\alpha}$	$\begin{array}{c}\lvert q_i\rangle \;\text{——}\times\text{——}\\ \lvert q_j\rangle \;\text{——}\times\text{——}\end{array}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
cross-resonance (RZX_{ij})	$\begin{array}{c}\lvert q_i\rangle \;\text{——}\boxed{X}\text{——}\\ \lvert q_j\rangle \;\text{——}\boxed{Z}\text{——}\end{array}$	$\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 0 & -i & 0 \\ 0 & 1 & 0 & i \\ -i & 0 & 1 & 0 \\ 0 & i & 0 & 1 \end{bmatrix}$
controlled-NOT (CNOT_{ij})	$\begin{array}{c}\lvert q_i\rangle \;\text{——}\bullet\text{——}\\ \lvert q_j\rangle \;\text{——}\oplus\text{——}\end{array}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$

controlled-Phase (CP_{ij})

$$|q_i\rangle \quad\quad\quad\quad\quad\quad\quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$
$$|q_j\rangle$$

Table 1.4 List of basic 2-qubit quantum gates. The subscripts i, j define the qubits on which the quantum gate acts. For the *CP* and the *CNOT* the first subscript, i, defines the control qubit and the second, j, the target.

It should be noted that most existing quantum computers do not directly perform all of the quantum gates listed here. As remarked in Sec. 1.2.3, a much smaller set of quantum gates is sufficient for universal quantum computing. For example, the only parameter-controlled quantum gate that some quantum computers perform is the $R_z(\theta)$, whereas $R_x(\theta)$ and $R_y(\theta)$ are obtained as combinations of H, $X^{\frac{1}{2}}$ and $R_z(\theta)$. For example

$$HadR_z(\theta)Had = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}\begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{bmatrix}\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = e^{i\frac{\theta}{2}}\begin{bmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix} = e^{i\frac{\theta}{2}}R_x(\theta).$$
$$(1.7)$$

Also most quantum computer interfaces feature the CNOT gate as a basic 2-qubit entangling gate, because its controlled-NOT logic makes it convenient for designing quantum circuits. However, in most cases the CNOT gate is synthesized as a combination of single-qubit gates and another 2-qubit gate, which is native to the quantum computer hardware. For example, the native 2-qubit gate in the IBM's superconducting quantum computers is the cross-resonance gate [44].

In this thesis, I assume the use of a quantum computer that can perform general single-qubit quantum gates and CNOT gates, and will regard these gates as fundamental building blocks of any quantum circuit.

Quantum circuits

Let me formally introduce what a quantum circuit is with the example circuit in Fig. 1.2.
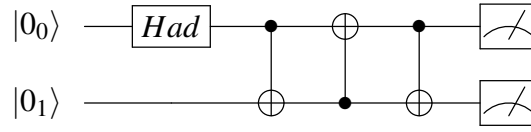
Fig. 1.2 Example quantum circuit

Each horizontal wire, corresponds to the state of a qubit. Unless otherwise stated, the qubits are assumed to be in state $|0\rangle$ initially. Quantum circuits are read from left to right, and the order of the quantum gates corresponds to the time order in which they are applied to the qubits state. For example, the circuit in Fig. 1.2, represents the state of 2 qubits, transformed by a Hadamard gate acting on qubit 0, followed by three *CNOT*s. The three *CNOT*s act to swap the state of the qubits, hence they are equivalent to a *SWAP* gate. The circuit in Fig. 1.2 transforms the 2-qubit state as

$$Had_0 : |00\rangle \rightarrow \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big)|0\rangle$$

$$CNOT_{01} :\rightarrow \frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big)$$

$$CNOT_{10} :\rightarrow |0\rangle\frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big)$$

$$CNOT_{01} :\rightarrow |0\rangle\frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big)$$

An important parameter that characterizes a quantum circuit is the circuit depth[11]. The circuit depth is defined as the longest path, in terms of number of basic qubit gates, from an input (preparation) to an output (measurement) moving forward (in time) only, along qubit wires and multi-qubit gates. For example, the quantum circuit in Fig. 1.3 has a depth of 6, as the longest path from an input to an output goes through the 4 *CNOT*s, the *Had* gate on qubit 0 and the $R_z(\theta)$ gate on qubit 2, but not through the *Z* gate on qubit 2. Assuming that any basic quantum gate is performed at some integer time step, then the circuit depth is the smallest number of time steps required to execute the circuit. Gates that do not act on common qubits can be performed simultaneously, e.g. the *Had* gate on qubit 0 and the *Z* gate on qubit 2 in Fig. 1.3, whilst gates that act on a common qubit must be performed sequentially, e.g. the *CNOT* gates in Fig. 1.3.

---

[11]In literature the circuit depth is also referred to as the 'circuit length'.
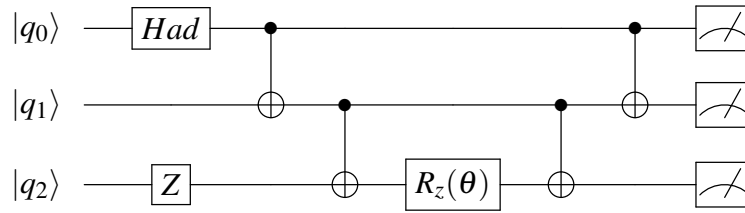
Fig. 1.3 Example quantum circuit 2

Composite and controlled quantum gates

As we proceed, composite and abstract multi-qubit quantum gates will be introduced. A symbol to depict a U-quantum-gate, corresponding to a multi-qubit unitary operator U, acting on the state of qubits i to j, is given in Fig. 1.4a. Also we can depict a controlled-U-quantum-gate, controlled by the state of qubit i as in Fig. 1.4b. Table 1.5 presents the quantum circuit symbols and the matrix representations of controlled single-qubit-rotation gates.

(a)   (b)

Fig. 1.4 **a)** Quantum gate corresponding to a general unitary operator U. **b)** Quantum gate corresponding to a general unitary operator U controlled by the state of qubit i.

Quantum gate	Symbol	Matrix
controlled-$R_x(\theta)\left(cR_x(\theta)_{i,\{j\}}\right)$	$\begin{array}{c}\lvert q_i\rangle \longrightarrow\bullet\longrightarrow \\[4pt] \lvert q_j\rangle \boxed{R_x(\theta)} \end{array}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\frac{\theta}{2} & 0 & -i\sin\frac{\theta}{2} \\ 0 & 0 & 1 & 0 \\ 0 & -i\sin\frac{\theta}{2} & 0 & \cos\frac{\theta}{2} \end{bmatrix}$
controlled-$R_y(\theta)\left(cR_y(\theta)_{i,\{j\}}\right)$	$\begin{array}{c}\lvert q_i\rangle \longrightarrow\bullet\longrightarrow \\[4pt] \lvert q_j\rangle \boxed{R_y(\theta)} \end{array}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\frac{\theta}{2} & 0 & -\sin\frac{\theta}{2} \\ 0 & 0 & 1 & 0 \\ 0 & \sin\frac{\theta}{2} & 0 & \cos\frac{\theta}{2} \end{bmatrix}$
controlled-$R_z(\theta)\left(cR_z(\theta)_{i,\{j\}}\right)$	$\begin{array}{c}\lvert q_i\rangle \longrightarrow\bullet\longrightarrow \\[4pt] \lvert q_j\rangle \boxed{R_z(\theta)} \end{array}$	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{i\theta} \end{bmatrix}$

Table 1.5 Controlled single-qubit-rotation gates.

1.3 Physical realization of quantum computers

Understanding the computational capabilities and limitations of NISQ computers is central to the development of algorithms for these devices. This section presents the guiding principles in physically realizing a quantum computer, describes the effects of quantum errors on the performance of quantum computers, and outlines the current most promising physical systems used to construct NISQ computers, and their respective limitations.

1.3.1 Guiding principles in physically representing a qubit

The elementary unit of a quantum computer is the qubit, a 2-state quantum system. Examples of 2-state quantum systems in nature are the spin-states of $\frac{1}{2}$-spin fermionic particles and the polarization-states of photons. Such systems can directly represent a qubit, and indeed spin

[45, 46] and photon-polarization [47, 48] qubits are frequently used in experimental devices for quantum computing and quantum encryption [47, 49].

A physical system representing a qubit should be: (1) stable to external perturbations, thus protecting the state of the qubit against quantum errors; and (2) controllable, so that the qubit state can be prepared, evolved and measured in a desired way. Generally these two requirements are opposing and can only be partially met. For example, a nuclear spin is a good qubit [50, 51] owing to its weak coupling to the external environment; a nuclear spin in a superposition state of being aligned and anti-aligned to a magnetic field can remain in this state for days. However, also due to its weak coupling, a nuclear spin is difficult to measure and to couple with other qubits.

Finding a delicate balance between qubit stability and qubit control is central to constructing a good quantum computer. Sometimes this involves considering multi-level quantum systems to represent qubits, e.g. electronic shell states [52], and superconducting Cooper pair states [53].

1.3.2 Quantum errors

A key concept in understanding how good a physical system is to represent qubits and construct a quantum computer, are *quantum errors* [17, 54]. A quantum error is considered any process that corrupts the desired evolution of the quantum computer state [17]. Quantum errors come in two types: (1) incoherent quantum errors, also termed as *decoherence*, corresponding to uncontrolled non-unitary evolutions of the quantum computer state, due to interactions with the environment; and (2) coherent errors, corresponding to uncontrolled unitary evolutions of the quantum computer state.

Decoherence and qubit decoherence times

Mathematically decoherence is a process where a unitary interaction between a quantum system $\mathscr{S}$ and the environment $\mathscr{E}$ results in a non-unitary evolution of $\mathscr{S}$. This non-unitary evolution transforms the state of $\mathscr{S}$ from a pure to a mixed state, thus destroying its quantum coherence. This process can also be viewed as a leak of quantum information from $\mathscr{S}$ to $\mathscr{E}$.

The most common reason for the decoherence of qubits are random fluctuations in the environment, such as lattice vibrations [55], radiation and external electromagnetic fields [56], that interact with the physical systems representing the qubits. Decoherence can also arise due to imperfect qubit control by coupling together the qubit state and the apparatus controlling it.

The decoherence of a qubit is characterized by the *relaxation time T_1*, and the *dephasing time T_2*, called together *decoherence times*. To define T_1 and T_2, I use the density matrix formalism. The relaxation time, T_1, is defined as the time scale, over which a qubit in an initial pure state $|1\rangle$, if not acted upon, decays to a mixed state of an equal classical mixture of $|1\rangle$ and $|0\rangle$. The decay of the qubit state is given as

$$\rho_{T_1}(t) = \frac{1}{2}\left(1 - e^{-\frac{t}{T_1}}\right)|0\rangle\langle 0| + \frac{1}{2}\left(1 + e^{-\frac{t}{T_1}}\right)|1\rangle\langle 1| \equiv \frac{1}{2}\begin{bmatrix} 1 - e^{-\frac{t}{T_1}} & 0 \\ 0 & 1 + e^{-\frac{t}{T_1}} \end{bmatrix} \quad (1.8)$$

Similarly, the dephasing time, T_2, is defined as the time scale over which a qubit in the initial pure state $|+\rangle \equiv \frac{|0\rangle + |1\rangle}{\sqrt{2}}$, if not acted upon, decays to a mixed state of an equal classical mixture of $|+\rangle$ and $|-\rangle \equiv \frac{|0\rangle - |1\rangle}{\sqrt{2}}$. The decay of the qubit state is given as

$$\rho_{T_2}(t) = \frac{1}{2}\left(1 - e^{-\frac{t}{T_2}}\right)|-\rangle\langle -| + \frac{1}{2}\left(1 + e^{-\frac{t}{T_2}}\right)|+\rangle\langle +| \equiv \frac{1}{2}\begin{bmatrix} 1 & e^{-\frac{t}{T_2}} \\ e^{-\frac{t}{T_2}} & 1 \end{bmatrix} \quad (1.9)$$

The relaxation and dephasing times are important parameters of the quantum computer hardware, since they limit the maximum duration of a computational task that can be performed reliably on a quantum computer. The values of T_1 and T_2 can be of the same order for some systems, e.g. superconducting qubits [44], and differ by orders of magnitude for other systems, e.g. spin-qubits [45].

Coherent quantum errors

Common sources of coherent quantum errors are random quantum fluctuations such as spin flips in spin-based qubits [57, 57], or spontaneous emissions/absorptions of photons in qubits represented by multi-energy-level systems [58]. Imperfect qubit control can result in quantum errors as well. For example, in superconducting qubits, performing a single-qubit gate means driving the qubit with a pulse from an arbitrary waveform generator [59, 60]. The amplitude, frequency, and duration of the pulse are real-valued parameters, thus are subject to some amount of error.

Coherent quantum errors are not as dangerous as decoherence, since they do not destroy the coherence of the qubit state. Coherent quantum errors can be mitigated by various *error mitigation* methods [61].

Quantum gate fidelity

As a result of both coherent and incoherent quantum errors the operation of a real quantum gate deviates from its intended unitary operation U. Generally, the operation of an imperfect noisy quantum gate is not necessarily unitary, and is described by a linear completely-positive trace-preserving map Λ. By Kraus' theorem [62], the action of a completely-positive map, such as Λ, on a quantum state, represented by a density matrix ρ, can be expressed as

$$\Lambda(\rho) = \sum_k \hat{A}_k^\dagger \rho \hat{A}_k, \tag{1.10}$$

where $\{A_k\}$ are linear Kraus operators satisfying $\sum_k A_k^\dagger A_k = I$. The quality of such an imperfect quantum gate can be quantified by its fidelity.

Fidelity is a measure of the "closeness" of two quantum states, represented by density matrices ρ_1 and ρ_2, and is defined as

$$\mathscr{F}(\rho_1, \rho_2) = \left(\mathrm{tr}\sqrt{\sqrt{\rho_1}\rho_2\sqrt{\rho_1}} \right)^2. \tag{1.11}$$

The value of the fidelity scales from 1, indicating that the states are the same, to 0 indicating that the states are orthogonal. However, the fidelity concept can be extended to describe the quality of a quantum gate as the closeness between U and Λ [63]:

$$\mathscr{F}(U, \Lambda) = \frac{1}{2^n} \sum_{i=0}^{2^n-1} \langle \phi_i | U^\dagger \Lambda(|\phi_i\rangle\langle\phi_i|) U |\phi_i\rangle, \tag{1.12}$$

where $\{|\phi_i\rangle\}$ is a set of basis states[12], e.g. the computational basis states, and n is the number of qubits.

1.3.3   Quantum-computer hardware parameters

The performance of a real quantum computer is determined by the set of hardware parameters defined below:

1. **Qubit decoherence times**: the relaxation time T_1, and the dephasing time T_2, defined in Sec. 1.3.2, limit the maximum duration of a computational task that can be performed reliably.

---

[12]In practice we would measure the fidelity using just a small set of basis states, since using the full set of basis states would be computationally intractable.

2. **Single- and 2-qubit quantum gate errors**: a quantum gate error (infidelity) is defined as one minus the quantum gate fidelity. Applying multiple imperfect gates results in an error, which approximately grows exponentially with the circuit depth. Hence, quantum gate errors limit the number of gates and the circuit depth.

3. **Readout errors**: the readout or the measurement error is the most error-prone process for most quantum hardware systems (see Table 1.6). However, this error is usually much smaller than the error accumulated by the multiple quantum gates in a quantum circuit. Hence, the measurement error is rarely a limiting factor in quantum-circuit computers.

4. **Quantum gate times**: a gate time is the time required to perform a quantum gate. Gate times together with qubit decoherence times limit the maximum depth of a quantum circuit. However, typically the limitation imposed by the quantum gate errors is more severe (see Sec. 1.3.4).

5. **Qubit connectivity**: qubit connectivity is a measure of how well the qubits of a quantum computer are connected to each other. Performing a 2-qubit gate between two qubits that are not connected, requires additional *SWAP* gates to bring the states of the qubits next to each other. This results in a larger quantum circuit. Qubit connectivity is often intrinsically restricted by the physical system used to construct the quantum computer. Qubit connectivity becomes increasingly important as the size (the number of qubits) of the quantum computer increases.

1.3.4 Physical systems

This section summarizes the three most promising types of physical systems, used to represent qubits and construct scalable quantum computers: superconducting qubits, trapped ion qubits and spin qubits. The typical best values for the hardware parameters defined in the previous section, for each of the three types of systems, are summarized in Table 1.6.

Apart from the systems described here, there are numerous other suggestions of physical systems to perform quantum computing, e.g. neutral atoms in optical lattices [64], linear optical systems [65, 66] and majorana quasiparticles [67]. However, the implementations of these systems are either in the early experimental stage or still at the theoretical stage.

Superconducting qubits

Superconducting qubits [60, 71, 72, 68] are implemented on superconducting electronic circuits, based on the Josephson junction [73]. The basic charge carriers in the superconducting

Qubit type	Dec. times T_1 / T_2 ,[s]	Single-qubit gate error	2-qubit gate error	Gate time,[s]	Readout error	Max num. of qubits
Supercond.	$10^{-4} - 10^{-5}$	$10^{-3} - 10^{-4}$	$10^{-2} - 10^{-3}$	$10^{-7} - 10^{-8}$	10^{-2}	65
Trapped ion	> 0.1	$10^{-3} - 10^{-4}$	$10^{-2} - 10^{-3}$	10^{-4}	$10^{-3} - 10^{-4}$	53
Spin	$> 10^{-5}$ ∗	10^{-3}	$10^{-1} - 10^{-2}$	10^{-7}	$10^{-1} - 10^{-2}$	~ 4

Table 1.6 Typical best values for quantum-computer hardware parameters. These values are obtained from Refs. [44, 68] for superconducting qubits, Refs. [18, 69, 52] for ion trapped qubits, and Refs. [45, 70] for spin qubits. * For some types of spin qubits T_1 can be on the order of minutes.

circuit are Cooper pairs, which can form a superconducting condensate at low temperatures. At every point in the circuit the condensate wavefunction is well-defined by a complex amplitude, and macroscopic quantum effects, such as discretization of magnetic flux, charge and phase can be measured. Hence, qubits can be represented by the discrete quantum states corresponding to different integer values of, for example, magnetic flux quanta and number of Cooper pairs (charge) in the superconducting circuit. Single-qubit gates can be implemented by microwave pulses tuned to the energy gaps between the energy levels of the qubits [59, 74], and 2-qubit gates can be implemented by coupling qubits by intermediate electrical coupling circuits [74].

Superconducting qubits are used to construct the current largest universal NISQ computers, e.g. the 53-qubit computer of Google, used in their quantum supremacy paper [20], a 62-qubit computer constructed by the researchers at the University of Science and Technology of China [75], a 65-qubit computer claimed by IBM [19], and as well as the non-universal D-Wave quantum annealers that operate with thousands of qubits [76]. Compared to other qubit technologies superconducting qubits feature fast and accurate quantum gates (see Table 1.6), and are relatively easy to produce, making the scaling (increasing the number of qubits) of superconducting quantum computers easier. However, superconducting qubits have short decoherence times and suffer from limited quantum gate connectivity.

Trapped ion qubits

Trapped ion qubits [52, 77, 78, 69] are implemented on ions suspended in free space by electromagnetic confinement. Qubits can be represented by the electronic states of each ion. For example, a qubit can be represented by the hyperfine levels of an electronic ground state (hyperfine qubit), or by the ground and first excited state (optical qubit). These qubits are

very stable and have decoherence times on the order of thousands of years, for hyperfine qubits, and seconds, for optical qubits [77]. Lasers can be used to induce coupling between internal qubit states to perform single-qubit quantum gates, or induce coupling between external motional states to entangle qubits and perform 2-qubit quantum gates such as the *CNOT* [79].

Trapped ion qubits are intrinsically more stable (longer decoherence times) and allow for a better qubit connectivity than superconducting qubits. Therefore, in the foreseeable future trapped-ion-based quantum computers may surpass superconducting quantum computers. Currently the largest trapped-ion-based quantum computers are a 53-qubit device presented in Ref. [80], a 32-qubit device constructed by IonQ [81], and an 8-qubit device constructed by Honeywell [82].

Spin qubits

Some of the earliest propositions to construct a quantum computer, were based on representing qubits as the spin states of $\frac{1}{2}$-spin fermionic particles, such as electrons [83], holes trapped in semiconductors [84], and atomic nuclei [50]. Single-qubit quantum gates on spin-qubits can be performed by localized magnetic fields that can address the spin-states of individual qubits, or by electric fields if the spin-qubits are subject to a magnetic field gradient. Qubits can be coupled by applying a gate voltage to control the exchange coupling between adjacent spins [46], thus performing a 2-qubit quantum gate corresponding to a power of *SWAP*. Spin-qubits have been implemented in a variety of physical systems, such as graphene [85], silicon [86, 70] and GaAs [45] quantum dots, and vacancies in a diamond [87].

The advantages of spin-qubits are their relatively small size, on the scale of a few nanometers, and their potential implementation in silicon based systems, which are compatible with the electronics manufacturing industry. However, spin-qubits are still highly experimental. Up to date, large spin-qubit-based quantum computers have not been constructed. Intel has announced [88] plans to construct a large scale spin-qubit-based quantum computer.

1.3.5 NISQ computers

Current and near future NISQ computers will likely correspond to superconducting and ion trapped quantum computers. For both types of systems, the qubit decoherence times and the gate times permit quantum circuits of depths of thousands of quantum gates. However, these numbers are greatly reduced by the quantum gate errors. The best current NISQ computers can support quantum circuits of a depth of just tens of quantum gates, without loosing all

useful coherence. Quantum gate errors are dominated by the 2-qubit gate errors that are about an order of magnitude larger than the single-qubit errors (see Table. 1.6). For example, one of the most accurate IBM devices, "ibmq-ourense", accessible via IBM's cloud [44], has an average $CNOT$ error (infidelity) of 7×10^{-3} and an average single-qubit gate error of 3×10^{-4}. Therefore, the 2-qubit gates are expected to be the major bottleneck of current and near future NISQ computers.

The goal of this thesis is to develop VQE protocols for molecular modelling, specifically designed for implementation on current and near future NISQ computers. I will consider the $CNOT$ count and the $CNOT$ depth [13] of the quantum circuit required to implement such a protocol, as the primary cost metrics of the efficiency of the protocol.

1.4   Classical computational chemistry

Quantum mechanics underpins the laws of all chemistry. Therefore, it is natural to assume that quantum mechanics can be used to simulate the behaviour of molecules and chemical compounds. However, exact simulation of a chemical system is a classically intractable problem as the size of the electronic wavefunction grows exponentially with the number of particles. Hence, storing and processing the wavefunction on a classical computer would require an exponentially large memory and an exponential number of steps, respectively. In 1982 Feynman [11] proposed a solution to this problem; using a quantum computer, which can efficiently[14] store and process a quantum multi-particle wavefunction.

This section begins by formulating the electronic structure problem, a central problem in molecular modelling. Then, I present some standard classical computational methods to solve the electronic structure problem, and discuss why these methods cannot be exact and computationally tractable at the same time.

1.4.1   The electronic structure problem

The electronic structure problem is central in chemistry, and particularly in molecular modelling. This problem corresponds to finding the electronic structures, the states of motion of electrons, and associated energies, corresponding to the low-lying energy states of a molecule. Mathematically, the electronic structure problem can be defined as finding the low-lying eigenvalues and eigenvectors of a molecular Hamiltonian, H_m.

---

[13]The $CNOT$ depth of a quantum circuit is defined as the longest path, in terms of $CNOT$ gates, between the input and the output of the circuit.

[14]By "efficient" I mean any algorithm that requires run time (number of steps) and memory that scale polynomially with the size (number of particles) of the problem.

The Hamiltonian, H_m, of a molecule with N_n nuclei and N_e electrons is

$$H_m = -\sum_i \frac{\hbar^2}{2m_e}\nabla_i^2 - \sum_k \frac{\hbar^2}{2M_k}\nabla_k^2 - \sum_{i,k} \frac{e^2}{4\pi\varepsilon_0}\frac{Z_k}{|\vec{R}_k - \vec{r}_i|} + \frac{1}{2}\sum_{i\neq j}\frac{e^2}{4\pi\varepsilon_0}\frac{1}{|\vec{r}_j - \vec{r}_i|} + \frac{1}{2}\sum_{k\neq l}\frac{e^2}{4\pi\varepsilon_0}\frac{Z_k Z_l}{|\vec{R}_k - \vec{R}_l|},$$
(1.13)

where $\vec{r}_i$ denotes the position of the i^{th} electron and $\vec{R}_k$, M_k and Z_k denote the position, mass and atomic number of the k^{th} nucleus. The sums run over all electrons and nuclei. The first two sums in the H_m expression correspond to the kinetic energies of the electrons and the nuclei, respectively. The last three sums correspond to the electron-nuclear, electron-electron and nuclear-nuclear Coulomb interactions, respectively.

Let us work in atomic units, where the unit of mass is the electron mass, $m_e \approx 9.109 \times 10^{-31}$kg, the unit of length is the Bohr length, $a_0 = \frac{4\pi\hbar^2}{m_e e^2} \approx 0.529 \times 10^{-10}$m, and the unit of energy is the Hartree energy, 1 Hartree $= \frac{e^2}{4\pi a_0} = 27.211$eV. In atomic units the Hamiltonian can be expressed as

$$H_m = -\sum_i \frac{\nabla_i^2}{2} - \sum_k \frac{\nabla_k^2}{2M_k'} - \sum_{i,k} \frac{Z_k}{|\vec{R}_k - \vec{r}_i|} + \frac{1}{2}\sum_{i\neq j}\frac{1}{|\vec{r}_j - \vec{r}_i|} + \frac{1}{2}\sum_{k\neq l}\frac{Z_k Z_l}{|\vec{R}_k - \vec{R}_l|},$$
(1.14)

where $M_k' = \frac{M_k}{m_e}$.

Nuclei are over 1000 times heavier than electrons, hence the Born-Oppenheimer approximation can be applied, where the nuclei are treated as static classical point-like charges. In this way the kinetic nuclear term in the expression for the molecular Hamiltonian in Eq. 1.14 goes to zero, and the Coulomb nuclei-nuclei term is constant. Then, for a given nuclear configuration (fixed values of $\{R_k\}$), we need to find the solutions (eigenstates) of the electronic Hamiltonian

$$H_e = -\sum_i \frac{\nabla_i^2}{2} - \sum_{i,k} \frac{Z_k}{|\vec{R}_k - \vec{r}_i|} + \frac{1}{2}\sum_{i\neq j}\frac{1}{|\vec{r}_j - \vec{r}_i|},$$
(1.15)

where we have omitted the constant nuclear-nuclear interaction terms. Throughout the rest of the thesis, I will drop the subscript e and will denote H_e as H. Our goal will be to find the eigenvalues $\{E_i\}$ and eigenvectors $|E_i\rangle$ corresponding to the ground state and lowest excited states of H. We want to estimate the energies to within *chemical accuracy* of approximately 10^{-3} Hartree. Chemical accuracy is the typical error of thermochemical experiments. Also, if the energies are known to this accuracy then at room temperature chemical reaction rates can be calculated with a precision to within the correct order of magnitude, using Eyring's equation [24]

$$k \propto e^{-\Delta E/k_B T},$$
(1.16)

where k is the reaction rate, k_B is the Boltzmann constant, T is the temperature and ΔE is the energy barrier between reactant and product states.

1.4.2 First and second quantized, and grid-based and set-basis computational chemistry methods

Many-particle quantum problems can be described in either the first or the second quantization formalism. Hence, we distinguish between first and second quantized methods to solve the electronic structure problem. In first quantized methods variables such as particle position and momentum are expressed by operators and quantized, and the exchange antisymmetry of electrons is accounted for by the wavefunction. In second quantized methods fields rather than variables are quantized, and the symmetry of the particles is accounted for by the properties of the operators that are applied to the wavefunction.

In addition to first and second quantized methods, we also distinguish between grid-based and basis-set methods. Grid-based methods directly store the electronic wavefunction by evaluating it on a discretized spatial grid. Whereas, basis-set methods project the Hamiltonian onto a basis of electronic orbital functions and the electron wavefunction is expressed as a superposition of Slater determinants.

First quantized, grid-based methods

In first-quantized grid-based methods [89], the wavefunction is antisymmetrized to account for the electronic exchange antisymmetry, and is considered in the position representation $|\boldsymbol{x}\rangle$. An N_e electron wavefunction can be expressed as

$$|\psi\rangle = \int_{\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}} \psi(\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}) \mathscr{A}\left(|\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}\rangle\right) d\boldsymbol{x}_1 .. d\boldsymbol{x}_{N_e}, \tag{1.17}$$

where $\boldsymbol{x}_i \equiv (\vec{r}_i, \sigma_i) \equiv (x_i, y_i, z_i, \sigma_i)$ are the spatial and the spin coordinates of the i^{th} electron, $\mathscr{A}$ denotes the antisymmetrization operator, and $\psi(\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}) = \langle \boldsymbol{x}_1 .. \boldsymbol{x}_{N_e} | \psi \rangle$.

The system is simulated by evaluating $|\psi\rangle$ on a discretized spatial grid. If each spatial axis is disretized in N_p points, the wavefunction can be expressed as

$$|\psi\rangle = \sum_{\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}} \psi(\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}) \mathscr{A}\left(|\boldsymbol{x}_1 .. \boldsymbol{x}_{N_e}\rangle\right), \tag{1.18}$$

where $|\boldsymbol{x}_i\rangle = |x_i\rangle |y_i\rangle |z_i\rangle |\sigma_i\rangle$ with $x_i, y_i, z_i \in [0, 1, \ldots N_p - 1]$ and $\sigma_i \in [0, 1]$. This wavefunction is represented by $2^{N_e} N_p^{3N_e}$ complex amplitudes, hence the memory required to store it on a classical computer is exponential in the number of electrons N_e.

Grid-based methods are useful when the Born-Oppenheimer approximation is not appropriate or when considering chemical dynamics [90, 91]. However, these methods quickly become intractable for classical computers as the number of electrons increases.

First quantized, basis-set methods

In basis-set methods the Hamiltonian of the system, H, is projected onto N_{so} single-electron basis wavefunctions $\{\phi_k(\boldsymbol{x}_i)\}$, where $\boldsymbol{x}_i$ is as usual the spatial-spin coordinate of the i^{th} electron. The basis wavefunctions $\{\phi_k(\boldsymbol{x}_i)\}$ usually approximate molecular (or atomic) spin-orbitals. The choice of $\{\phi_k(\boldsymbol{x}_i)\}$ is discussed in Sec. 1.4.4.

Again, the exchange symmetry of the electrons should be accounted for by their wavefunction. Hence, a wavefunction of N_e electrons can be written as an antisymmetrized product of N_e basis wavefunctions $\{\phi_{k_0}(\boldsymbol{x}_i), \ldots \phi_{k_{N_e-1}}(\boldsymbol{x}_i)\}$, which can be represented as a Slater determinant:

$$\Phi_{k_0,\ldots k_{N_e-1}}(\boldsymbol{x}_0,\boldsymbol{x}_1 \ldots \boldsymbol{x}_{N_e-1}) = \frac{1}{\sqrt{N_e!}} \begin{vmatrix} \phi_{k_0}(\boldsymbol{x}_0) & \phi_{k_1}(\boldsymbol{x}_0) & \ldots & \phi_{k_{N_e-1}}(\boldsymbol{x}_0) \\ \phi_{k_0}(\boldsymbol{x}_1) & \phi_{k_1}(\boldsymbol{x}_1) & \ldots & \phi_{k_{N_e-1}}(\boldsymbol{x}_1) \\ \vdots & \vdots & \vdots & \vdots \\ \phi_{i_0}(\boldsymbol{x}_{N_e-1}) & \phi_{k_1}(\boldsymbol{x}_{N_e-1}) & \ldots & \phi_{k_{N_e-1}}(\boldsymbol{x}_{N_e-1}) \end{vmatrix} \quad (1.19)$$

Exchanging the position of two electrons, corresponds to exchanging two rows of the Slater determinant, which changes the signs of the wavefunction as required. Typically the number of considered spin orbitals N_{so} is larger than the number of electrons N_e. Hence, a general N_e electron wavefunction will be represented by a superposition of Slater determinants:

$$\psi(\boldsymbol{x}_0,\boldsymbol{x}_1 \ldots \boldsymbol{x}_{N_e-1}) = \sum_{k_0=0}^{N_{so}-1} \sum_{k_1=k_0+1}^{N_{so}-1} \ldots \sum_{k_{N_e-1}=i_{N_e-2}+1}^{N_{so}-1} c_{k_0,\ldots k_{N_e-1}} \Phi_{k_0,\ldots k_{N_e-1}}(\boldsymbol{x}_0,\boldsymbol{x}_1 \ldots \boldsymbol{x}_{N_e-1})$$

$$(1.20)$$

The grid-based methods discussed in the section above store the electronic wavefunction directly without assuming any prior knowledge about its form. In contrast, basis-set methods exploit such knowledge by projecting the Hamiltonian on the spin orbitals $\{\phi_k(\boldsymbol{x}_i)\}$, thus reducing the resources needed to store the electronic wavefunction. A general electronic wavefunction written in terms of Slater determinants can be represented by $2^{N_e}N_e!$ complex amplitudes. Accounting for electron conservation, this number can be reduced further. As we will see next, the information encoded by a Slater determinant can be even further reduced by second quantization, hence first-quantization basis-set methods are rarely used.

Second quantization, grid-based methods

Second-quantized grid-based methods are rarely used in either classical or quantum computational chemistry. They are discussed in Ref. [92].

Second quantization, basis-set methods

Second-quantized basis-set methods follow straightforwardly from first-quantized basis-set methods. Again the Hamiltonian is projected onto a basis of N_{so} molecular spin-orbital wavefunctions $\{\phi_k(\boldsymbol{x}_i)\}$. However, in second-quantization formalism the wavefunction does not account for the fermionic anticommutation, so the Slater determinant in Eq. 1.19 can be simply represented as the *occupation number vector*

$$|\boldsymbol{f}\rangle = |f_0 f_1 \ldots f_{N_{so}-1}\rangle \equiv \Phi_{k_0,\ldots k_{N_e}}(\boldsymbol{x}_0, \boldsymbol{x}_1 \ldots \boldsymbol{x}_{N_e-1}), \tag{1.21}$$

where $f_k = 1$ if $k \in \{k_0, \ldots k_{N_e-1}\}$, indicating that spin orbital ϕ_k is occupied by an electron, and $f_k = 0$ otherwise. The space spanned by all $2^{N_{so}}$ such vectors is known as the Fock space, and the vectors $\{|\boldsymbol{f}\rangle\}$ form an orthonormal basis for this space. Throughout this thesis, I will use occupation number vectors predominantly to represent Slater determinants, so I will refer to them as Slater determinants.

Electrons are excited (de-excited) into a spin orbital ϕ_i by the fermionic creation (annihilation) operator $a_i^\dagger$ (a_i). I refer to $a_i^\dagger$ and a_i together as fermionic ladder operators. In second quantization formalism the exchange symmetry of the electrons is accounted for by the fermionic ladder operators $a_i^\dagger$ and a_i, therefore they obey the anticommutation relations

$$\{a_i, a_j^\dagger\} = a_i a_j^\dagger + a_j^\dagger a_i = \delta_{ij} \text{ and } \{a_i^\dagger, a_j^\dagger\} = \{a_i, a_j\} = 0. \tag{1.22}$$

The action of the fermionic ladder operators on the Slater determinants is

$$a_i |f_0 f_1 \ldots f_i \ldots f_{N_{so}}\rangle = (-1)^{\sum_{r=0}^{i-1} f_r} \delta_{f_i 1} |f_0 f_1 \ldots f_i \oplus 1 \ldots f_{N_{so}}\rangle \text{ and} \tag{1.23}$$

$$a_i^\dagger |f_0 f_1 \ldots f_i \ldots f_{N_{so}}\rangle = (-1)^{\sum_{r=0}^{i-1} f_r} \delta_{f_i 0} |f_0 f_1 \ldots f_i \oplus 1 \ldots f_{N_{so}}\rangle, \tag{1.24}$$

where $\oplus$ denotes addition modulo 2. The phase term $(-1)^{\sum_{r=0}^{i-1} f_r}$ accounts for the exchange symmetry of the electrons.

In second quantization the electronic Hamiltonian H in Eq. 1.15 is represented in terms of fermionic ladder operators as

$$H = \sum_{k,i} h_{ki} a_k^\dagger a_i + \frac{1}{2} \sum_{l,k,j,i} h_{lkji} a_l^\dagger a_k^\dagger a_j a_i, \qquad (1.25)$$

where

$$h_{ki} = \int \phi_k^*(\boldsymbol{x}) \left(-\frac{\nabla^2}{2} - \sum_p \frac{Z_r}{|\vec{r} - \vec{R}_p|} \right) \phi_i(\boldsymbol{x}) d\boldsymbol{x} \quad \text{and} \qquad (1.26)$$

$$h_{klji} = \int \frac{\phi_l^*(\boldsymbol{x}_1) \phi_k^*(\boldsymbol{x}_2) \phi_j(\boldsymbol{x}_2) \phi_i(\boldsymbol{x}_1)}{|\vec{r}_1 - \vec{r}_2|} d\boldsymbol{x}_1 d\boldsymbol{x}_2. \qquad (1.27)$$

The indices i, j, k, l run over all spin orbitals $\{\phi_i\}$, the index p runs over all nuclei, $\boldsymbol{x} \equiv (\vec{r}, \sigma)$ denotes spatial and spin coordinates, and R_p is the position of the p^{th} nucleus. The first integrals correspond to the electronic kinetic terms and the electron-nuclei Coulomb interaction terms of the Hamiltonian, and the second integrals correspond to the electron-electron Coulomb interaction terms. The number of fermionic operator terms in the Hamiltonian expression scales as $O(N_{so}^4)$

The Slater determinants $|\boldsymbol{f}\rangle$ form a complete orthonormal basis for the electronic Hamiltonian in Eq. 1.25. Therefore, the eigenstates of the electronic Hamiltonian can be expressed as linear combinations of Slater determinants, as

$$|\psi\rangle = \sum_f c_f |\boldsymbol{f}\rangle, \qquad (1.28)$$

where c_f are real coefficients. If the molecular spin orbital wavefunctions $\{\phi_i\}$ form a complete basis for the single-electron states of the molecular Hamiltonian, then these solutions are exact. In Sec. 1.4.4 we will see how to choose a basis of spin orbital wavefunctions. If $|\psi\rangle$ includes all $\binom{N_{so}}{N_e}$ Slater determinants that correspond to a N_e-electron wavefunction, then the wavefunction $|\psi\rangle$ is called the full configuration interaction (FCI). However, the number of all Slater determinants scales exponentially with N_e, making classical FCI calculations intractable for large systems.

Second-quantized basis-set methods are the most widely used type of methods in classical computational chemistry. Also they translate the most easily to quantum computational chemistry. Throughout the rest of the thesis, we will be concerned with second-quantized basis-set methods only.

1.4.3 Classical computational chemistry methods

This section briefly describes three standard classical second-quantized basis-set methods used to approximate the ground state wavefunction of an electronic Hamiltonian: the Hartree-Fock (HF) method, the configuration interaction (CI) method and the coupled cluster (CC) method. The purpose of this summary is to present the limitations of these classical methods in balancing between computational cost and accuracy, and to serve as motivation for developing quantum computational chemistry methods.

Hartree-Fock

The HF method [93, 94] is a variational method that represents the electronic wavefunction as only a single Slater determinant. Typically, the method considers a number of N_{so} spin orbitals, that is larger than the number of electrons N_e. Since, only one Slater determinant is considered, N_e of the spin orbitals will be occupied and $N_{so} - N_e$ will be unoccupied (also referred to as virtual orbitals).

The HF method operates by minimizing the expectation value for the ground state energy. This is done by varying the spatial parts of the N_e occupied spin-orbitals, subject to the condition that they form an orthonormal basis [94]. Typically, the input to the HF method is a set of single-electron atomic spin orbitals (orbitals localized around each atom), and the output is a set of single-electron molecular orbitals [93].

The HF method works well for systems where the true electronic wavefunction is dominated by a single Slater determinant. However, the HF method does not account for static and dynamic correlations. Dynamic correlations arise from the Coulomb interaction term, and result in small contributions from many Slater determinants that have occupied virtual orbitals [95]. Static correlations occur in excited states, transition states and near the dissociation limit [95, 96], and result in an electronic wavefunction where two or more Slater determinants are dominant. The HF method performs poorly for strongly correlated systems. Typically the Slater determinant generated by the HF method is used as an initial reference state for more accurate methods.

Configuration interaction method

The configuration interaction (CI) method [93, 97] includes additional determinants to account for the correlation effects of the electron wavefunction. The CI creates a correlated wavefunction by considering excitations above a reference state, usually chosen to be the HF

state. The correlated wavefunction is represented as

$$|\psi_{CI}\rangle = \left(I + \sum_{ki} c_{k,i} a_k^\dagger a_i + \sum_{l>k,j>i} c_{lkji} a_l^\dagger a_k^\dagger a_j a_i + \dots\right)|\psi_{HF}\rangle, \tag{1.29}$$

where c are parameters to be optimized. The CI method optimizes these parameters based on the Rayleigh-Ritz variational principle, which states that the expectation value of a Hamiltonian with respect to a parametrized wavefunction, is always greater or equal to the lowest eigenvalue of the Hamiltonian.

If the CI includes all excited state determinants, then the optimized wavefunction is the FCI wavefunction. However, this would result in exponentially large number of parameters and wavefunction size, which would make the CI method classically intractable, except for small molecules. Usually, the CI considers only single and double excitations (CISD), and sometimes triple excitations (CISDT).

The CI method is effective at recovering dynamic correlation, but less effective at recovering static correlation. The truncated CI method also is not size-extensive [96].

Coupled cluster method

Similarly to the CI method, the CC method also includes additional determinants as excitations above a reference HF state. However, the trial wavefunction is represented as a parametrized product instead:

$$|\psi_{CC}\rangle = \prod_{k,i}\left(I + c_{ki} a_k^\dagger a_i\right) \prod_{l>k,j>i}\left(I + c_{lkji} \hat{a}_l^\dagger a_k^\dagger a_j \hat{a}_i\right) \dots |\psi_{HF}\rangle = e^{\sum_p \Xi_p}|\psi_{HF}\rangle \tag{1.30}$$

with

$$\begin{aligned}
\Xi_1 &= \sum_{k\in virt, i\in occ} \theta_{ki} a_k^\dagger a_i, \\
\Xi_2 &= \sum_{k,l\in virt,\ j,i\in occ} \theta_{lkji} a_l^\dagger a_k^\dagger a_j a_i, \\
&\vdots
\end{aligned} \tag{1.31}$$

where θ are variational parameters, *occ* (*virt*) denote occupied (virtual) spin-orbitals, and Ξ_p denotes a sum of all fermionic excitation operators of order p. If all excitation operators Ξ_p, for $p \in [0, \frac{N_{so}}{2}]$ are considered, the FCI wavefunction is recovered. However, this would result in an exponential number of parameters to be optimized. Therefore, usually only single and double excitations are used, resulting in the CCSD method.

Even with single and double excitations only, the CCSD wavefunction, in general, includes all determinants [23], and would require exponential memory size to be stored directly. Therefore, the CC method does not store the wavefunction directly, but rather coupled non-linear equations are derived, whose solutions approximate the electronic ground state wavefunction [93]. The time complexity of solving these equations scales as $O(N_e^2(N_{so} - N_e)^4)$, and the required memory as $O(N_{so}^4)$.

The CC method converges faster than the CI method, and is also size-extensive. However, the wavefunction constructed by the CC does not obey the Rayleigh-Ritz principle [93]. Also the CC method does not perform well for strongly correlated systems [98]. Section 1.5.5 describes the unitary version of the CC method, the UCC, that solves the issues above, and which can be efficiently implemented on a quantum computer.

1.4.4   Spin-orbital basis sets

Section 1.4.3 presented basis-set methods that project the electron Hamiltonian onto a basis of spin-orbital wavefunctions $\{\phi_i\}$, and represent the electron wavefunction as a linear combination of Slater determinants of $\{\phi_i\}$. This section summarizes some standard spin-orbital basis sets.

The true spin-orbitals of a molecule can be obtained by numerically solving the Schrödinger equation using grid-based methods. However, this is computationally intractable for large molecules. Instead, a basis of spin-orbital wavefunctions can be obtained by considering approximations of the spatial form of the atomic spin-orbitals of the individual nuclei.

Such standard approximations are the Slater-type orbitals (STOs) [93, 23]. The form of a STO is

$$R_{n_E}^{STO}(r) \propto \left(\zeta r\right)^{n_E - 1} e^{-\zeta r},\tag{1.32}$$

where n_E is the energy level, and ζ is a fitting parameter. The STO functions do not display oscillatory behaviour like the true atomic orbitals. Therefore, an atomic orbital can be better approximated as a linear combination of several STOs, each with a different parameter ζ, thus increasing the radial flexibility of the approximate function. A representation where an atomic orbital is represented by n STOs is referred to as an $n - \zeta$ representation. STOs display many features of true atomic orbitals and can approximate them accurately. However, STOs make the evaluation of the two-electron integrals in Eq. 1.27 computationally expensive.

To simplify the calculation of these integrals, Gaussian-type orbital (GTO) functions [99] can be used instead. The spatial form of the GTOs is

$$R_{n_E l}^{GTO}(r) \propto \left(\sqrt{\alpha_{n_E l}}r\right)^l e^{-\alpha_{n_E l}r^2},\tag{1.33}$$

where l is the angular momentum quantum number of the orbital and $\alpha_{n_E l}$ is a fitting parameter. Owing to the exponential r^2 dependence, GTOs are more localized than STOs, and do not approximate the atomic orbitals as well as the STOs. Therefore, more GTOs are required to describe an atomic orbital. However, this is compensated by the two-electron integrals in Eq. 1.27 being more easily calculated.

Typically, STOs are approximated as linear combinations of GTOs, and the STOs are used as a basis of atomic orbital wavefunctions. The number of basis wavefunctions determines the accuracy, run-time and memory requirements of classical chemistry algorithms. In the case of quantum chemistry algorithms, as we will see in Sec. 1.5.1, the number of basis functions determines the number of required qubits. Below are presented several different spin-orbital bases, depending on the number of STOs and GTOs considered.

STO-nG basis sets

Some of the simplest spin-orbital basis-sets, are the STO-nG basis sets [100]. In the STO-nG, each spin-orbital is represented by one STO, which is approximated by n-GTOs. In STO-nG basis sets, only the spin-orbitals that participate in the HF states, and those of similar energies are considered. These basis sets are called minimal basis sets, since only the spin-orbitals sufficient to contain the electrons in the neutral atoms are used.

Split-valence basis sets

More complex and accurate basis sets are the so called split-valence basis sets. Split-valence basis sets are STO-nGs, where the valence orbitals (the highest-energy occupied orbitals in the HF state) are represented by n STOs ($n - \zeta$ representation). An example of a split-valence basis set is the $6 - 31G$ basis [101].

Correlation-consistent basis sets

Correlation-consistent polarization valence $n\zeta$ (cc-pVnZ) bases [102] can achieve additional accuracy. These bases sets include additional virtual spin-orbitals that can recover the correlation energy. The core occupied spin-orbitals have $1 - \zeta$ STO representations, whilst the valence and virtual spin-orbitals have $n - \zeta$ STO representations. The occupied spin-orbitals are generated by the HF method, while the virtual spin-orbitals can be generated by correlated calculations on atoms [27].

1.5   Quantum computational chemistry

Section 1.4 defined the electronic structure problem and presented standard classical computational chemistry methods to solve this problem. In this section, we look into how the electronic structure problem can be mapped and solved on a quantum computer. The considered techniques correspond to second-quantized basis-set methods. Quantum first-quantized methods exist [23, 103], however they are rarely used and I do not discuss them in the thesis.

   I begin by presenting three different qubit encodings for the spin-orbitals and the fermionic ladder operators. I then describe the quantum phase estimation (QPE) and the variational quantum eigensolver (VQE) algorithms to find the ground state energy of an electronic Hamiltonian.

1.5.1   Qubit encoding methods

The Jordan Wigner encoding

A qubit encoding method in the context of quantum computational chemistry is a map from the Fock space of electronic Slater determinants to the Hilbert space of qubit states. The Jordan-Wigner (JW) encoding is the most straightforward qubit encoding, where the occupation of the spin-orbital ϕ_i is represented by the state of qubit i, such that qubit state $|0_i\rangle$ $\left(|1_i\rangle\right)$ indicates that ϕ_i is unoccupied (occupied). Hence, there is a one to one correspondence between Slater determinants and computational basis qubit states:

$$|f_0 f_1 \ldots f_{N_{so-1}}\rangle \rightarrow |q_0\rangle |q_1\rangle \ldots |q_{N_{so-1}}\rangle$$
$$q_i = f_i \in \{0,1\} \tag{1.34}$$

The fermionic ladder operators $a_i^\dagger$ and a_i are mapped to quantum gate operations as

$$a_i^\dagger = b_i^\dagger Z_{i-1} \ldots Z_0 \ \text{ and} \tag{1.35}$$

$$a_i = b_i Z_{i-1} \ldots Z_0, \tag{1.36}$$

with

$$b_i^\dagger = \frac{(X_i - iY_i)}{2} \ \text{ and} \tag{1.37}$$

$$b_i = \frac{(X_i + iY_i)}{2}, \tag{1.38}$$

where I refer to $b_i^\dagger$ and b_i as *qubit ladder operators*. The operators, $b_i^\dagger$ and b_i, act to change the occupation of spin-orbital ϕ_i, whereas the products (strings) of Pauli-z gate operators

act as the exchange phase factor $(-1)^{\sum_{k=0}^{i} f_k}$ in Eqs. (1.23) and (1.24), accounting for the electron antisymmetry. The action of the product of Pauli-z gate operators is also referred to as computing the parity of the (qubit) state.

Using the JW encoding of the fermionic ladder operators [Eqs. (1.36) and (1.38)], the electron Hamiltonian in Eq. (1.25) can be expressed as a linear combination of products of Pauli operators as

$$H = \sum_j h_j P_j = \sum_j h_j \prod_{i=0}^{N_{so}-1} \sigma_i^j, \tag{1.39}$$

where h_j are a real scalar coefficients (not to be confused with the one- and two-electron integrals h_{ki} and h_{lkji} in Eqs. (1.26) and (1.27)), P_j is a product of Pauli operators, (or a *Pauli string* for a shorthand), and σ_i^j is one of the Pauli operators $\{I_i, X_i, Y_i, Z_i\}$ acting on qubit i. The number of Pauli string terms, in the Hamiltonian quantum gate-expression, scales as the number of fermionic terms, in the Hamiltonian expression in Eq. (1.25), $O(N_{so}^4)$.

Within the JW encoding it is easy to see the first advantage of using quantum computers to solve chemistry problems: an N_{so} qubit state can represent a superposition of $2^{N_{so}}$ Slater determinants, hence it can store the full FCI wavefunction. On the other hand, directly storing the $2^{N_{so}}$ complex amplitudes of the FCI wavefunctions on a classical computer would require $O(2^{N_{so}})$ classical bits of memory.

Parity encoding

In the JW encoding the occupation number is stored locally (one needs to measure the state of one qubit to find the occupation number), while the parity is stored non-locally (one needs to measure $O(N_{so})$ qubits to determine the parity).

Parity encoding is an alternative encoding where the occupation number is stored non-locally and the parity locally. The correspondence between Slater determinants and qubit states is

$$|f_0 f_1 \dots f_{N_{so-1}}\rangle \rightarrow |q_0\rangle |q_1\rangle \dots |q_{N_{so-1}}\rangle$$
$$q_i = \bigoplus_{j=0}^{i} f_j, \tag{1.40}$$

where $\bigoplus$ denotes summation modulus 2. The transformed fermionic ladder operators are described in Ref. [104]. The parity encoding has not yet found a practical use.

Bravyi-Kitaev encoding

Bravyi-Kitaev encoding is midway between the JW and the parity encoding in the sense that the occupation number and parity are both stored partially locally. The correspondence between Slater determinants and qubit states is

$$|f_0 f_1 \ldots f_{N_{so}-1}\rangle \rightarrow |q_0\rangle |q_1\rangle \ldots |q_{N_{so}-1}\rangle$$
$$q_i = \bigoplus_{j=0}^{i} \beta_{ij}^{(N_{so})} f_j, \tag{1.41}$$

where $\beta_{ij}^{(N_{so})}$ are elements of the BK matrix for N_{so} qubits.

The BK matrix is defined recursively as

$$\beta^{(1)} = \begin{bmatrix} 1 \end{bmatrix} \tag{1.42}$$

$$\beta^{(2^{x+1})} = \begin{bmatrix} \beta^{(2^x)} & \mathbf{0} \\ \mathbf{A} & \beta^{(2^x)} \end{bmatrix}, \tag{1.43}$$

where $\mathbf{0}$ is a $2^x \times 2^x$ matrix of zeros and $\mathbf{A}$ is a $2^x \times 2^x$ matrix of zeros except the bottom row, which is filled with ones. For example, for 2 and 4 qubits the BK matrix is

$$\beta^{(2)} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \text{ and } \beta^{(4)} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} \tag{1.44}$$

If the number of orbitals (qubits) N_{so} is not a power of 2, then the BK matrix is created by taking the first N_{so} rows of the BK matrix for the next power, $\lceil \log N_{so} \rceil$, of 2.

The quantum gate operation representation of the fermionic ladder operators in the BK encoding is considerably more complicated than in the case of the JW encoding, and it is explained in detail in Ref. [104]. The BK encoding represents a fermionic ladder operator with $O(\log N_{so})$ Pauli operators, acting on $O(\log N_{so})$ qubits, whereas the JW and the parity encoding represent a fermionic ladder operator with $O(N_{so})$ Pauli operators, acting on $O(N_{so})$ qubits. A thorough comparison between the JW and the BK methods is presented in Ref. [105].

1.5.2   Plane wave spin-orbital basis sets

Section 1.4.4 described some of the standard STO spin-orbital basis sets, used in classical computational chemistry. While these bases can be used in quantum computational chemistry methods to obtain an accurate description of molecular electronic wavefunctions with relatively few spin-orbitals, they also lead to a number of terms in the Hamiltonian that scales as $O(N_{so}^4)$, irrespective of the qubit encoding. The number of terms in the Hamiltonian is proportional to the number of quantum computer measurements required to evaluate the expectation value of H (see Sec. A).

This number can be reduced by the plane wave and dual plane wave spin-orbital basis sets [106]. The plane wave basis functions are

$$\phi_v = \frac{1}{\sqrt{V}} \exp\left(\frac{2\pi i v r}{L}\right), \tag{1.45}$$

for a plane wave with wave vector corresponding to the v^{th} harmonic of a cell with size L and volume V. The dual plane wave basis functions are obtained as the Fourier transform of the plane wave basis functions [92]. The plane wave and the dual plane wave bases diagonalize the kinetic and potential operators in the electron Hamiltonian, reducing the number of H terms from $O(N_{so}^4)$ to $O(N_{so}^3)$ and to $O(N_{so}^2)$, respectively. However, to achieve the same accuracy as STO bases, plane wave bases need 10 to 100 times more basis functions. This, results in requiring more qubits, and correspondingly more quantum gates. Hence, plane wave bases are not ideal for NISQ computers.

1.5.3   Quantum phase estimation

The Quantum phase estimation (QPE) [28, 107] algorithm is a quantum algorithm to estimate the eigenvalue (or phase) of an eigenvector of a unitary operator. The QPE is used as a subroutine in a number of quantum algorithms, e.g. Shor's algorithm [14]. In the context of quantum computational chemistry the QPE can be used to estimate the eigenvalues of an electronic Hamiltonian [Eq. (1.15)]. The QPE is described by the steps below [17] and the circuit in Fig. 1.5

1. Prepare two qubit registers: (1) a target qubit register of N_{so} qubits to represent the electron wavefunction, initialized in state $|\psi_0\rangle = \sum_i c_i |E_i\rangle$, which should have non-zero overlap with the true FCI ground state wavefunction ($c_o \neq 0$), and (2) an ancilla qubit register of n_{out} qubits to store the output value for the energy eigenvalue, initialized in an equal superposition of all computational basis states, $1/\sqrt{2^{n_{out}}} \sum_{x=0}^{2^{n_{out}}-1} |x\rangle$.
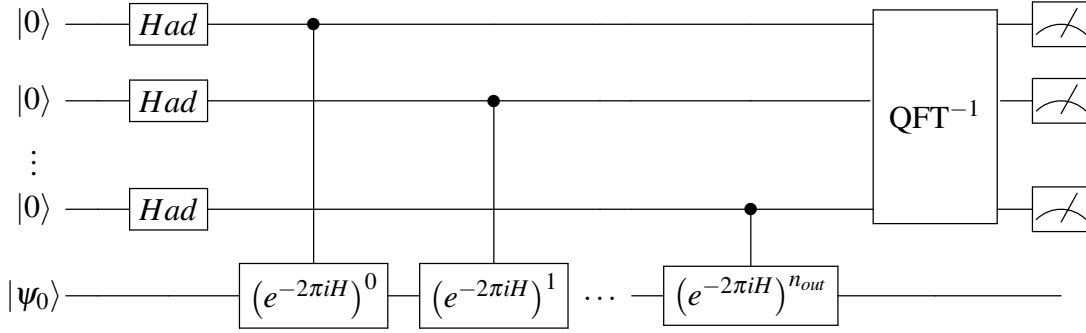
Fig. 1.5 Quantum phase estimation circuit. The state $|\psi_0\rangle$ is a N_{so} qubit state representing the electron wavefunction. H is the electron Hamiltonian [Eq. (1.25)], and n_{out} is the number of ancilla qubits initialized in state $|0\rangle$. QFT^{-1} denotes the inverse Fourier transform. Upon measurement of the ancilla qubits an estimate for an eigenvalue of H is obtained.

2. Apply the controlled gates, shown in Fig. 1.5, that evolve the two qubit register states as

$$\frac{1}{\sqrt{2^{n_{out}}}} \sum_{x=0}^{2^{n_{out}}-1} |x\rangle \sum_i c_i |E_i\rangle \rightarrow \frac{1}{\sqrt{2^{n_{out}}}} \sum_{x=0}^{2^{n_{out}}-1} \sum_i e^{-2\pi i x E_i} |x\rangle c_i |E_i\rangle \tag{1.46}$$

3. Apply the inverse quantum Fourier transform, QFT^{-1} [108], to the ancilla qubit register to record the phase into the ancilla register

$$\frac{1}{\sqrt{2^{n_{out}}}} \sum_{x=0}^{2^{n_{out}}-1} \sum_i e^{-2\pi i x E_i} |x\rangle c_i |E_i\rangle \rightarrow \frac{1}{\sqrt{2^{n_{out}}}} \sum_i c_i |\text{bin}(E_i)\rangle |E_i\rangle, \tag{1.47}$$

where $\text{bin}(E_i)$ is the value of E_i as a binary bit string of length n_{out}.

4. Measure the ancilla register and obtain the eigenvalue E_i with probability $|c_i|^2$. The state of the target qubit register is collapsed to the corresponding eigenstate $|E_i\rangle$.

The complexity of the QPE is determined by the depth of the circuits that implement the controlled unitary evolutions of H. Generaly, the complexity is polynomial in the number of spin-orbitals, N_{so}, and the number of electrons, N_e, thus the QPE is exponentially faster than the classical FCI method. Some of the most efficient implementations of the QPE have complexities of $O(N_{so}^5)$ [109] and $O(N_e^2 N_{so}^3)$ [110].

The success probability and precision of the QPE are determined by the number of ancilla qubits [17]. For a success probability p and a precision $\varepsilon_{pe} = \frac{1}{2^n}$ to the n^{th} bit

$$n_{out} = n + \left\lceil \log_2 \left(2 + \frac{1}{2p} \right) \right\rceil. \tag{1.48}$$

The QPE evolves the target register with H for times $\{2\pi, 4\pi, \ldots, 2^{n_{out}}\pi\}$. Thus, the total evolution time is approximately $T = 2^{n_{out}+1}\pi$. For a success probability $p = 0.5$, $n_{out} = n+2$, and so $T = 2^{n+3}\pi = 8\pi/\varepsilon_{pe}$. Hence, the target state is transformed by $16\pi/\varepsilon_{pe} = 2^{n+4}$ controlled unitary evolutions e^{-iH}. On a quantum-circuit computer e^{-iH} cannot be performed exactly, so it should be approximated using, for example, Trotter decomposition [111]. This introduces an additional error ε_U to the finite precision ε_{pe}. There are also errors arising due to quantum gate operations, referred to as circuit synthesis errors ε_{cs}. Hence, the total error of the QPE is upper bounded by $\varepsilon_{pe} + \varepsilon_U + \varepsilon_{cs}$ [112]. There has been considerable work done in minimizing and achieving balance between these three errors, for example, by optimal state initialization [113], and efficient Hamiltonian simulation [111, 114, 115].

However, to perform the multiple controlled e^{-iH} evolutions, the QPE requires long quantum circuits with large numbers of quantum gates, on the order of millions to billions for practical applications [116, 117]. Such circuits greatly exceed the capabilities of NISQ computers. Therefore, the QPE is considered to be a long-term method, which requires fault-tolerant quantum computers, or at least quantum computers with enough qubits for error correction.

1.5.4 Variational Quantum Eigensolver

A promising candidate to solve the electronic structure problem on NISQ computers is the variational quantum eigensolver (VQE) [23, 29, 118]. The VQE is an algorithm that determines the lowest eigenvalue, E_0, of a Hamiltonian operator H. The VQE relies on the Rayleigh-Ritz variational principle, which states that

$$\langle \psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle \equiv E(\vec{\theta}) \geq E_0, \tag{1.49}$$

to iteratively minimize the expectation value $E(\vec{\theta})$ with respect to the parametrized trial state $|\psi(\vec{\theta})\rangle$, and thus obtain an estimate for the ground state energy E_0. A single VQE iteration is described by the three steps below, and is illustrated by Fig. 1.6:

1. On a quantum computer prepare the trial state $|\psi(\vec{\theta})\rangle$

2. On a quantum computer measure the expectation value $E(\vec{\theta})$

3. On a classical computer, process the measurement data, and update the variational parameters $\vec{\theta}$ for the next iteration using a classical optimizer

The iterations continue until $E(\vec{\theta})$ converges.
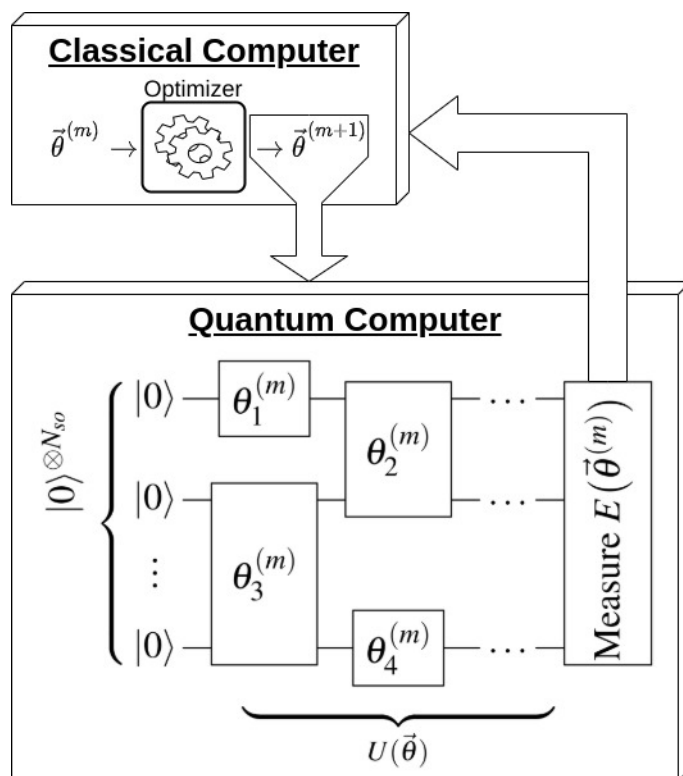
Fig. 1.6 A diagram summarizing a single VQE iteration

Since classical computers cannot efficiently store and process the electronic wavefunction, the VQE uses a quantum computer to do so. The variational parameters at the end of each iteration are updated by a classical computer though. This allows the VQE to store intermediate results on the classical computer, and restart the state of the quantum computer after each iteration. Hence, the VQE exchanges the long coherence times (long quantum circuits) required by the QPE, for multiple quantum measurements and additional classical processing. As we will see later, this results in a polynomial overhead in the asymptotic complexity of the VQE as compared to the best QPE implementations. Nevertheless, the VQE remains exponentially faster than exact classical methods like the FCI. Owing to its use of both a quantum and a classical computer, the VQE is considered a *hybrid quantum-classical algorithm*.

Preparing the trial wavefunction

The trial state for the VQE, $|\psi(\vec{\theta})\rangle$, is prepared by a parametrized quantum circuit, corresponding to a unitary operation $U(\vec{\theta})$, acting on an initial reference state $|\psi_0\rangle$, such that

$|\psi(\vec{\theta})\rangle = U(\vec{\theta})|\psi_0\rangle$. I refer to the unitary operator $U(\vec{\theta})$ as the *ansatz* and to the circuit that implements $U(\vec{\theta})$ as the *ansatz circuit*, respectively[15].

Hamiltonian measurement

Once the trial state is prepared, the expectation value of the Hamiltonian H should be measured. In the context of computational chemistry, H is an electronic Hamiltonian, with second-quantized form given by Eq. (1.25). In Sec. 1.5.1 we saw that with the JW encoding the electronic Hamiltonian can be expressed in terms of quantum gate operations as

$$H = \sum_j h_j P_j = \sum_j h_j \prod_{i=0}^{N_{so}-1} \sigma_i^j. \tag{1.50}$$

Using a different encoding like the BK, also results in a similar expression. Therefore, the VQE needs to measure the expectation value

$$E(\vec{\theta}) = \langle\psi(\vec{\theta})| \sum_j h_j \prod_{i=0}^{N_{so}-1} \sigma_i^j |\psi(\vec{\theta})\rangle = \sum_j h_j \langle\psi(\vec{\theta})| \prod_{i=0}^{N_{so}-1} \sigma_i^j |\psi(\vec{\theta})\rangle \tag{1.51}$$

at the end of each iteration. In practice the expectation values of all Pauli strings in the expression for H above are measured on a quantum computer, then they are multiplied with the corresponding coefficients, h_j, and added together on a classical computer, to obtain $E(\vec{\theta})$.

The expectation value of a Pauli string, $\langle\psi(\vec{\theta})| \prod_{i=0}^{N_{so}-1} \sigma_i |\psi(\vec{\theta})\rangle$, can be conveniently measured on a quantum computer by preparing the state $|\psi(\vec{\theta})\rangle$, performing a multi-qubit basis rotation as described in Appendix A, and measuring just a single qubit in the Z-basis. If a single Pauli string of length $l(\leq N_{so})$ is measured at a time, the circuit performing the basis rotations requires $(l-1)$ *CNOT*s and at most l single-qubit gates, and has a depth of $O(\log l)$. Therefore, the depth and gate count of the circuit required to measure $E(\vec{\theta})$ is dominated by the ansatz circuit, which generates $|\psi(\vec{\theta})\rangle$.

The number of measurements required to estimate the expectation value of each Hamiltonian term to a precision ε scales as $O(1/\varepsilon^2)$ [119]. Therefore, using a STO-based spin-orbital basis where the Hamiltonian is expressed by $O(N_{so}^4)$ terms, results in $O(N_{so}^4/\varepsilon^2)$ measurements required to obtain $E(\vec{\theta})$. As mentioned in Sec. 1.5.2 using doubly plane wave spin-orbital bases, the number of H terms can be reduced to $O(N_{so}^2)$, thus reducing the number of measurements to $O(N_{so}^2/\varepsilon^2)$. The number of measurements can be additionally decreased by a constant factor by grouping together and simultaneously measuring mutually

---

[15]In literature *ansatz* is often used interchangeably for any of the two terms

commuting Hamiltonian terms [120, 121]. Some works have also proposed applying unitary transformations, resulting in an increased circuit length, to group the Hamiltonian terms even further, thus reducing the number of measurements from $O(N_{so}^4/\varepsilon^2)$ to $O(N_{so}^3/\varepsilon^2)$ [122].

In this thesis, I consider circuit efficiency as a primary cost metric. Therefore, I assume the use of STO-based spin-orbitals and Hamiltonian term groupings that do not increase the circuit length. This implies that measuring the expectation value of a molecular electronic Hamiltonian will require a number of quantum computer measurements, which scale as $O(N_{so}^4)$.

Classical optimization

The classical optimization subroutine is crucial for the success of the VQE. The classical optimizer should be able to minimize the multi-variable Hamiltonian expectation value, and it also should be robust to noise due to quantum errors. Discussion of the details of various optimization methods is beyond the scope of this thesis. A thorough comparison of various classical optimizers for the VQE is presented in Ref. [123].

I only mention that classical optimizers are broadly divided into two groups: gradient-descent and direct-search methods. As the name suggests, gradient-descent methods rely on calculating the gradient vector of the minimized function with respect to the variational parameters. In the case of finding the ground state energy of a Hamiltonian with the VQE, the energy gradient vector can be explicitly calculated by the classical or the quantum computer (see Appendix C.2). Gradient-descent optimizers are faster than direct-search optimizers, however the former are more susceptible to noise and are more likely to get stuck in a local minima.

1.5.5 VQE ansatz

The ansatz, $U(\vec{\theta})$, is the most critical component of the VQE. First, the variational flexibility of the ansatz determines how accurately the VQE can approximate the ground state, and its corresponding energy. Second, as remarked above the circuits required to measure the expectation values of the Hamiltonian terms are dominated by the ansatz circuit. Hence, both the run-time of the quantum subroutine and the accumulated quantum errors are determined almost entirely by the depth and the gate count of the ansatz circuit. Third, the success and complexity of the classical optimization are determined by the number of ansatz parameters, $\vec{\theta}$, and the dependence of the energy expectation value on these parameters.

Therefore, successfully performing molecular modelling with the VQE on a NISQ computer requires an ansatz that is: (1) variationally flexible enough to approximate accurately

the ground state of H, (2) implemented by a shallow circuit, especially one with as few *CNOT* gates as possible, and (3) easy to optimize. These desired qualities are satisfied, to various levels, by ansätze that lie in between two extremes: chemically-inspired and hardware-efficient ansätze.

Hardware-efficient ansätze

Hardware-efficient ansätze correspond to universal unitary transformations implemented as periodic sequences of parametrized one- and two-qubit gates that are easy to implement on the available quantum hardware [124–128]. These ansätze aim to prepare a variationally flexible ansatz state using as few quantum gates as possible. Hence, they are particularly suitable for small quantum computers that have imperfect quantum gates and limited qubit connectivity. Hardware efficient ansätze have been successfully used to simulate a number of small molecules with the VQE [129].

However, as hardware-efficient ansätze lack chemically motivated structure, two major problems arise. First, in order for the ansatz to be able to accurately approximate an arbitrary state (in an exponentially large Hilbert space) a large, potentially exponential, number of variational parameters is required [130]. This problem can be partially alleviated by constructing a hardware-efficient ansatz that conserves the number of electrons [126]. However, even in this case, an exponentially large Hilbert space needs to be searched. The second problem is that using a hardware-efficient ansatz with randomly initialized variational parameters, makes the energy gradients along most parameters close to zero [131]. This problem becomes exponentially worse with the number of qubits and the circuit depth. In some scenarios, this is known as the barren-plateau problem [132–134]. Because of these problems, the classical optimization of hardware-efficient ansätze is considered to be intractable for large systems, and VQE protocols, using such ansätze, are considered to be unscalable for practical applications.

Chemically-inspired ansätze

The majority of chemically inspired ansätze come from adapting classical computational chemistry methods to run on a quantum computer. The most widely used type of ansätze are the unitary coupled cluster (UCC) ansätze, which as the name suggests are motivated by the CC (Sec. 1.4.3) method. A general UCC ansatz corresponds to a unitary evolution of a sum of fermionic excitation operators:

$$U_{UCC}(\vec{\theta}) = \exp\Big(\sum_{k,i}^{N_{so}-1} \theta_{ki}(a_k^\dagger a_i - a_i^\dagger a_k) + \sum_{l,k,j,i}^{N_{so}-1} \theta_{lkji}(a_l^\dagger a_k^\dagger a_j a_i - a_j^\dagger a_i^\dagger a_l a_k) + \ldots \Big) \quad (1.52)$$

The exponential above cannot be implemented exactly on a quantum-circuit computer. Therefore, in practice the UCC ansätze are approximated to one or a few Trotter steps [135, 136] as

$$U_{UCC}(\vec{\theta}) = \prod_{k,i}^{N_{so}-1} e^{\theta_{ki}(a_k^\dagger a_i - a_i^\dagger a_k)} \prod_{l,k,j,i}^{N_{so}-1} e^{\theta_{kkji}(a_l^\dagger a_k^\dagger a_j a_i - a_j^\dagger a_i^\dagger a_l a_k)} \dots \ , \qquad (1.53)$$

where each exponential of a fermionic excitation operator can be implemented directly as a quantum circuit (see Appendix B). In this thesis, I will refer to an exponential of a fermionic excitation operator as *fermionic evolution* for shorthand.

As in the case of the CC method, UCC ansätze are typically truncated to include first and second order fermionic excitation operators only. The most commonly used UCC ansatz is the UCC Singles and Doubles (UCCSD) [29, 119, 137–140], which includes single and double fermionic excitation operators above the HF state:

$$U_{UCCSD}(\vec{\theta}) = \exp\left(\sum_{k\in virt, i\in occ} \theta_{ki}(a_k^\dagger a_i - a_i^\dagger a_k) + \sum_{l,k\in virt, j,i\in occ} \theta_{kkji}(a_l^\dagger a_k^\dagger a_j a_i - a_j^\dagger a_i^\dagger a_l a_k) \right)$$

$$(1.54)$$

The UCCSD has been successfully used to experimentally implement the VQE for various small molecules [29–31].

Other examples of UCC ansätze include: the generalized UCCSD (GUCCSD), which includes generalized single and double fermionic excitation operators [135], the k-UpCCSD, corresponding to repeated layers of selected UCCSD excitation operators [135], and the Bogoliubov-UCC [139], a quasiparticle variant of the UCC, which includes pairing-terms present in superconductivity. These UCC ansätze are *fixed ansätze* that depend only on the number of spin-orbitals and the number of electrons of the system of interest. By and large, the various fixed UCC ansätze offer different trade-offs between ansatz size (number of included fermionic excitation evolutions) and accuracy. Recently a number of VQE protocols, e.g. the fermionic-ADAPT-VQE [141] and the Pruning-VQE [142], were proposed that instead of using fixed UCC ansätze, iteratively construct problem-tailored UCC ansätze. I further discuss such iterative VQE protocols in Chapter 4.

Owing to their physically-motivated fermionic structure, UCC ansätze generate states that conserve the symmetries of electronic wavefunctions; e.g. electron-number-conservation, spin-conservation and time-symmetry. This results in small search spaces as compared to hardware-efficient ansätze, which makes UCC ansätze easy to optimize. In addition, these symmetries are invariant under certain types of noise, which makes UCC ansätze resilient to these types of noise [143, 128]. Furthermore, even relatively simple UCC ansätze, like

the single-Trotterized UCCSD, are highly accurate for weakly correlated systems, such as molecules near their equilibrium configuration [29, 30, 135, 136].

However, UCC ansätze require long quantum circuits that use particularly large numbers of *CNOT*s. Typically, UCC ansätze consist of $O(N_{so}^4)$ number of fermionic evolutions, resulting in UCC ansatz circuits that require $O(N_{so}^4 \log N_{so}) - O(N_{so}^5)$ *CNOT*s and have depths that roughly scale as $O(N_{so}^3 \log N_{so}) - O(N_{so}^4)$ [23]. The depths and *CNOT* counts of UCC ansatz circuits significantly limit the size of the molecules that can be simulated on NISQ computers. The *CNOT* complexity of the UCC ansatz circuits stems from (1) the *CNOT* cost of the circuits that implement fermionic evolutions, and (2) the large number of fermionic evolutions, included in the ansatz, which are not necessarily required for the approximation of the electronic wavefunction. Moreover, as we will see in Chapter 3, simple UCC ansätze, like the UCCSD, cannot accurately approximate strongly correlated systems. Hence, higher-order excitations and/or multiple-step Trotterization [135, 136] are required, which additionally increases the depth of the ansatz circuit.

In Chapter 2, I construct *CNOT*-efficient quantum circuits to implement fermionic evolutions that significantly reduce the *CNOT* counts of UCC ansatz circuits.

Chapter 2

Quantum circuits for unitary coupled cluster ansätze

Chapter summary

The ansatz is the backbone of any VQE protocol. The most widely used ansätze for molecular modelling with VQE protocols are UCC ansätze. UCC ansätze correspond to products of unitary evolutions of fermionic excitation operators. Due to their fermionic excitation structure these ansätze are accurate and easy to optimize. However, they are implemented by long circuits, with large $CNOT$ counts and $CNOT$ depths, that severely limit the size of the molecular systems that can be simulated on NISQ computers. In this Chapter, I construct circuits, optimized in terms of $CNOT$ count and $CNOT$ depth, to implement unitary evolutions of single and double fermionic excitation operators. The circuits can be used to build $CNOT$ efficient UCC ansätze circuits, and can be directly incorporated within existing VQE protocols to lower their quantum hardware requirements.

Section 2.1 presents some preliminary circuit identities and circuit constructions. Section 2.2 describes the canonical quantum circuits used to implement evolutions of fermionic excitation operators. In Sec. 2.3, I construct $CNOT$ efficient circuits to implement unitary evolutions of single and double *qubit excitation operators*, excitation operators that obey *qubit commutation relations*. Then, in Sec. 2.4 I expand the functionality of these circuits to account for the parity of the state, and thus implement unitary evolutions of fermionic excitation operators instead. The resulting circuits require up to 2 (8) times fewer $CNOT$s in the case of single (double) fermionic excitation evolutions than the canonical circuits.

The multi-qubit controlled rotation circuits in Sec. 2.1.3 are derived in my paper [144]. Sections 2.3 and 2.4 are based on my paper [145].

2.1 Preliminary circuit identities and circuit constructions

2.1.1 Two-qubit circuit identities

Here I summarize some commonly used 2-qubit circuit identities.
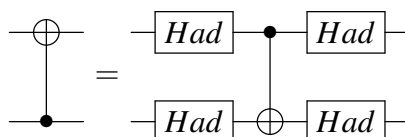
Fig. 2.1 A circuit identity exchanging the control and target qubits of a *CNOT* gate.
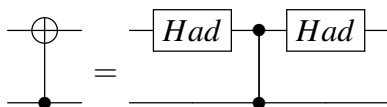
Fig. 2.2 A circuit identity for transformation between a *CNOT* and a *CP*. Both gates are equivalent in terms of *CNOT* cost.
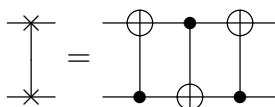
Fig. 2.3 A circuit to construct a *SWAP* gate with 3 *CNOT*s.
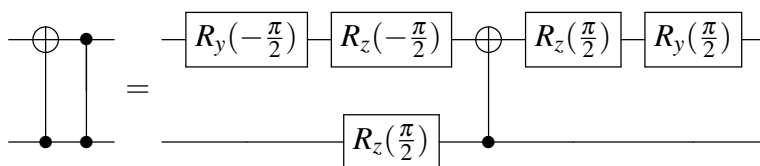
Fig. 2.4 A circuit identity to replace a *CNOT* next to a *CP* with a circuit with a *CNOT* cost of 1.

2.1.2 *CNOT* staircases

An important circuit construction in quantum computational chemistry is the *CNOT staircase* that can be used to compute the parity of a qubit state and record it into the state of a single

qubit. *CNOT* staircases are used to measure the expectation values of the Hamiltonian terms (see Appendix A), and to construct circuits implementing evolutions of fermionic excitation operators (see Sec. 2.2).

The parity of an *n*-qubit state can be computed by a staircase of $(n-1)$ *CNOT* gates such as the one in Fig. 2.5a. The action of the *CNOT* staircase on the qubit state can be expressed as

$$\prod_{i=n-2}^{0} CNOT_{i,i+1} |q_0\rangle |q_1\rangle \ldots |q_{n-1}\rangle = |q_0\rangle |q_1 \oplus q_0\rangle \ldots |(q_{n-1} \oplus (q_{n-2} \oplus (\ldots \oplus (q_1 \oplus q_0)\ldots))\rangle.$$
(2.1)

Hence, the parity of all *n* qubits is recorded in the state of the $(n-1)^{\text{th}}$ qubit.

The *CNOT* staircase circuit in Fig. 2.5.a has a depth of $(n-1)$, equal to the number of *CNOT*s. The depth can be reduced by rearranging the *CNOT*s. If we construct 2 half-staircases of $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$ *CNOT*s, respectively, which compute the parities of the first and the second half of the qubits, and then connect the two end qubits of each half-staircase with a *CNOT*, the parity of all *n* qubits will be recorded on the target qubit of the last *CNOT*. This rearrangement is visualized in Fig. 2.5b. The result is that the depth is decreased by a factor of approximately 2. This procedure of halving the staircase depth can be repeated approximately $\log_2 n$ times with each half-staircase, until the $(n-1)$ *CNOT*s are packed optimally into a circuit of depth $O(\log n)$.

(a)

(b)
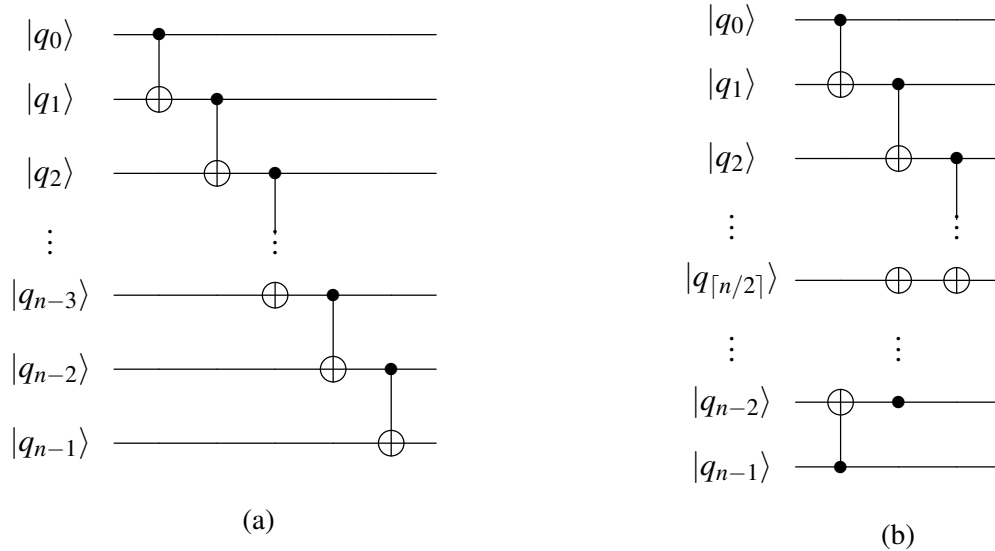
Fig. 2.5 **a)** A *CNOT* staircase to compute and record the parity of the *n*-qubits state on qubit $n-1$. **b)** A rearranged *CNOT* staircase with reduced circuit depth of $\lceil n/2 \rceil$, and parity recorded on qubit $\lceil n/2 \rceil$.

2.1.3   Multi-qubit-controlled single-qubit rotation

Here, I derive a method to construct circuits that implement multi-qubit-controlled rotations, without using ancilla qubits. These circuits are used later in the chapter (Sec. 2.3) to construct circuits implementing evolutions of qubit excitation operators.

Lets denote a $R_y(\theta)$ rotation on a target qubit 0, controlled by qubits $\{1,..,n\}$ in state $|1_1..1_n\rangle$, as $cR_y(\theta)_{0,\{1,..,n\}}$. If there are no control qubits, $cR_y(\theta)_{0,\{\}} \equiv R_y(\theta)_0$. The n-qubit-controlled rotation $cR_y(\theta)_{0,\{1,..,n\}}$ can be decomposed into two opposite half-way rotations, controlled by $(n-1)$ qubits as

$$cR_y(\theta)_{0,\{1,..,n\}} = CNOT_{1,0} cR_y \left(-\frac{\theta}{2} \right)_{0,\{2,..,n\}} CNOT_{1,0} cR_y \left(\frac{\theta}{2} \right)_{0,\{2,..,n\}}, \qquad (2.2)$$

or equivalently, as

$$cR_y(\theta)_{0,\{1,..,n\}} = cR_y \left(\frac{\theta}{2} \right)_{0,\{2,..,n\}} CNOT_{1,0} cR_y \left(-\frac{\theta}{2} \right)_{0,\{2,..,n\}} CNOT_{1,0}. \qquad (2.3)$$

By decomposing the controlled rotations further in the same way, the overall operation can be reduced to $CNOT$s and single-qubit rotations. Implementing naively Eqs. (2.2) or (2.3), results in a circuit with $2^{n+1} - 2$ $CNOT$s. However, for $n > 2$, I alternately combine Eqs. (2.2) and (2.3) to cancel adjacent $CNOT$s, as shown in Fig. 2.6, and obtain a circuit with 2^n $CNOT$s. This number of $CNOT$s is optimal and cannot be reduced further, because there are 2^n distinct control states and at least 2^n controlled gates are required to "identify" the correct control state.
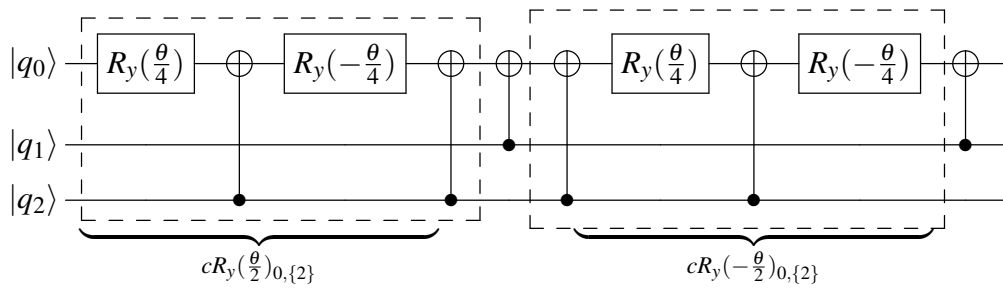
Fig. 2.6 A circuit to implement the 2-qubit controlled rotation $cR_y(\theta)_{0,\{1,2\}}$. The first half-way rotation $cR_y(\frac{\theta}{2})_{0,\{2\}}$ is implemented as in Eq. (2.2), and the second half-way rotation $cR_y(-\frac{\theta}{2})_{0,\{2\}}$, as in Eq. (2.3). This allows the two middle $CNOT$s between qubits 0 and 1 to be cancelled (note that $CNOT$ gates with the same target qubit commute).

R_z- and R_x-controlled rotations can be obtained by additional single-qubit rotations on the target qubit. Similarly the control state can be modified by additional operations on

the control qubits. Lastly, *CP* gates can be used instead of *CNOT* gates in Eqs. (2.2) and (2.3). In the following sections, I will use implementations with both *CNOT* and *CP* gates, depending on which favours adjacent gate cancellation.

2.2 Canonical fermionic-evolution circuits

This section describes the canonical circuits used to implement evolutions of fermionic excitation operators [116, 23]. I consider evolutions of single and double fermionic excitations operators only, since these are sufficient to recover the FCI wavefunction [146, 147]. The majority of UCC ansätze, e.g. the UCCSD and the k-UpCCSD, are also based only on single and double excitation operators.

Single and double fermionic excitation operators are defined, respectively, by the skew-Hermitian operators

$$\Xi_{ik} = a_i^\dagger a_k - a_k^\dagger a_i \quad \text{and} \tag{2.4}$$

$$\Xi_{ijkl} = a_i^\dagger a_j^\dagger a_k a_l - a_k^\dagger a_l^\dagger a_i a_j \tag{2.5}$$

where $a_i^\dagger$ and a_i are the fermionic ladder operators, which satisfy the anticommutation relations

$$\{a_i, a_j^\dagger\} = a_i a_j^\dagger + a_i^\dagger a_j = \delta_{ij} \text{ and } \{a_i^\dagger, a_j^\dagger\} = \{a_i, a_j\} = 0. \tag{2.6}$$

Unitary evolutions of single and double fermionic excitation operators are expressed, respectively, by the unitaries

$$F_{ik}(\theta) = e^{\theta \Xi_{ik}} \quad \text{and} \tag{2.7}$$

$$F_{ijkl}(\theta) = e^{\theta \Xi_{ijkl}}, \tag{2.8}$$

The angle θ is a variational parameter, whose physical meaning can be interpreted, e.g., as the duration of the evolution. Throughout the thesis, I refer to a unitary evolution generated by a fermionic excitation operators as a *fermionic evolution* for shorthand.

As shown in Sec. 1.5.1, within the Jordan-Wigner encoding, a and $a^\dagger$ can be written in terms of quantum gate operators as

$$a_i = b_i \prod_{r=0}^{i-1} Z_r = \frac{1}{2}(X_i + iY_i) \prod_{r=0}^{i-1} Z_r \quad \text{and} \tag{2.9}$$

$$a_i^\dagger = b_i^\dagger \prod_{r=0}^{i-1} Z_r = \frac{1}{2}(X_i - iY_i) \prod_{r=0}^{i-1} Z_r, \tag{2.10}$$

with the qubit ladder operators

$$b_i^\dagger \equiv \frac{1}{2}(X_i - iY_i) \text{ and} \tag{2.11}$$

$$b_i \equiv \frac{1}{2}(X_i + iY_i). \tag{2.12}$$

Substituting Eqs. (2.9) and (2.10), into Eqs. (2.7) and (2.8), for $l > k > j > i$, single and double fermionic evolutions can be re-expressed in terms of quantum gate operators as

$$F_{ik}(\theta) = \exp\left[i\frac{\theta}{2}(X_iY_k - Y_iX_k) \prod_{r=i+1}^{k-1} Z_r\right] \text{ and} \tag{2.13}$$

$$F_{ijkl}(\theta) = \exp\left[i\frac{\theta}{8}(X_iY_jX_kX_l + Y_iX_jX_kX_l + Y_iY_jY_kX_l + Y_iY_jX_kY_l \right.$$
$$\left. - X_iX_jY_kX_l - X_iX_jX_kY_l - Y_iX_jY_kY_l - X_iY_jY_kY_l) \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right], \tag{2.14}$$

respectively. The Pauli string terms in each of the exponentials above commute. Hence, single and double fermionic evolutions can be expressed as products of 2 and 8 exponentials of individual Pauli strings (Pauli string evolutions), respectively:

$$F_{ik}(\theta) = \exp\left[i\frac{\theta}{2}X_iY_k \prod_{r=i+1}^{k-1} Z_r\right] \exp\left[-i\frac{\theta}{2}Y_iX_k \prod_{r=i+1}^{k-1} Z_r\right] \text{ and} \tag{2.15}$$

$$F_{ijkl}(\theta) = \exp\left[i\frac{\theta}{8}X_iY_jX_kX_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right] \exp\left[i\frac{\theta}{8}Y_iX_jX_kX_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right]$$
$$\times \exp\left[i\frac{\theta}{8}Y_iY_jY_kX_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right] \exp\left[i\frac{\theta}{8}Y_iY_jX_kY_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right]$$
$$\times \exp\left[-i\frac{\theta}{8}X_iX_jY_kX_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right] \exp\left[-i\frac{\theta}{8}X_iX_jX_kY_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right]$$
$$\times \exp\left[-i\frac{\theta}{8}Y_iX_jY_kY_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right] \exp\left[-i\frac{\theta}{8}X_iY_jY_kY_l \prod_{r=i+1}^{j-1} Z_r \prod_{p=k+1}^{l-1} Z_p\right] \tag{2.16}$$

Exponentials of individual Pauli strings can be directly implemented as quantum circuits. For example, the first exponential in Eq. 2.15, $\exp\left[i\frac{\theta}{2}X_iY_k\prod_{r=i+1}^{k-1}Z_r\right]$, can be implemented by the circuit in Fig. 2.7 (see Appendix B). The two *CNOT* staircases together with the $R_z(\theta)$ rotation in between them, are referred to as a *CNOT staircase construction*. This construction implements $\exp\left[i\frac{\theta}{2}\prod_{r=i}^{k}Z_r\right]$. The *Had* and $R_x(\pm\frac{\pi}{2})$ gates, on both sides of the *CNOT* staircase construction, act as qubit basis rotations to transform the Z_i and the Z_k operators to X_i and Y_k operators, respectively. Similarly, by sandwiching a *CNOT* staircase construction between single-qubit rotations that transform individual Pauli-Z operators to Pauli-Y or Pauli-X operators, a circuit for any exponential of a Pauli string can be constructed. Constructing circuits to implement exponentials of Pauli strings is a standard procedure, which is described in more detail in Appendix B.
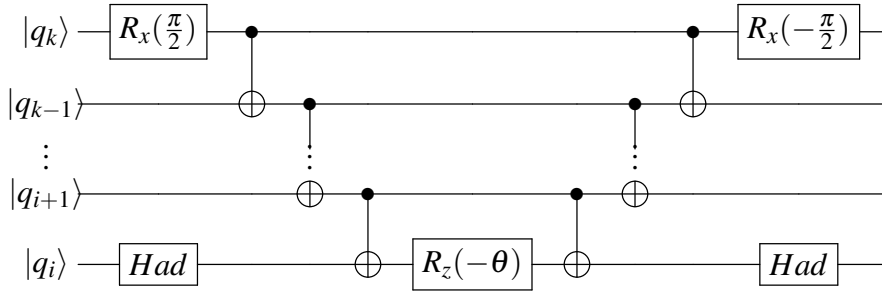
Fig. 2.7 A circuit implementing the exponential $\exp\left[i\frac{\theta}{2}X_iY_k\prod_{r=i+1}^{k-1}Z_r\right]$.

A single fermionic evolution $F_{ik}(\theta)$ [Eq. (2.15)] can be implemented by the circuit in Fig. 2.8 that contains 2 *CNOT* staircase constructions (resulting in $2\times 2 = 4$ *CNOT* staircases), and has a *CNOT* count and a *CNOT* depth of $2(k-i)$. A double fermionic evolution $F_{ijkl}(\theta)$ [Eq. (2.16)] can be implemented by a circuit that contains 8 *CNOT* staircase constructions (resulting in $8\times 2 = 16$ *CNOT* staircases), and has a *CNOT* count and a *CNOT* depth of $8(l+j-k-i)$ (the circuit for a double fermionic evolution is shown in Fig. E.1 in Appendix E due to its length). If the *CNOT*s of the staircases are rearranged as described in Sec. 2.1.2, the scaling of the *CNOT* depth with the number of qubits participating in the fermionic evolutions can be exponentially reduced.

2.3   Qubit-evolution circuits

The first step in constructing *CNOT*-efficient fermionic evolution circuits, is to construct circuits that implement unitary evolutions of qubit excitation operators. I shall refer to a unitary evolution of a qubit excitation operator as a *qubit evolution*.
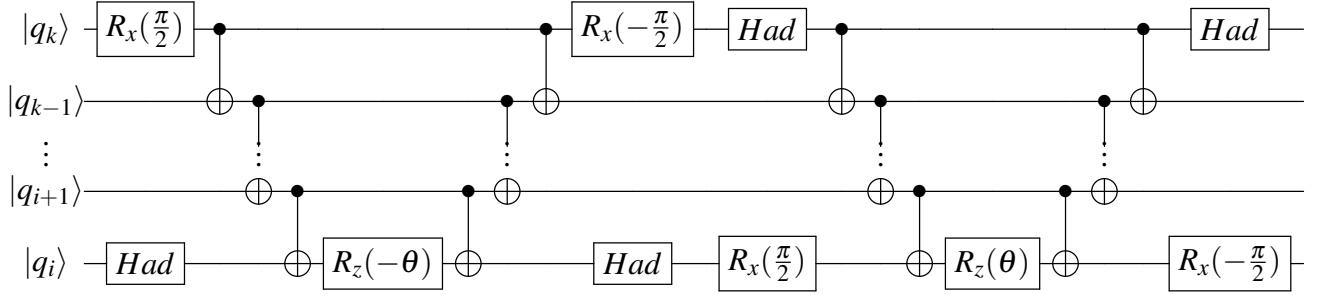
Fig. 2.8 A canonical circuit to implement a single fermionic evolution [Eq. (2.15)]. The vertical dots denote *CNOT* staircases on qubits q_{i+1} to q_{k-1}.

Single and double qubit excitation operators are generated by the qubit annihilation and creation operators, b and $b^{\dagger}$ [Eqs. (2.11) and (2.12)], and are given, respectively, by the skew-Hermitian operators

$$\Omega_{ik} = b_i^{\dagger} b_k - b_k^{\dagger} b_i \quad \text{and} \tag{2.17}$$

$$\Omega_{ijkl} = b_i^{\dagger} b_j^{\dagger} b_k b_l - b_k^{\dagger} b_l^{\dagger} b_i b_j. \tag{2.18}$$

The operators b and $b^{\dagger}$ satisfy the commutation relations

$$\{b_i, b_i^{\dagger}\} = I, \ \ [b_i, b_j^{\dagger}] = 0 \ \text{if} \ i \neq j, \ \text{and} \ [b_i, b_j] = [b_i^{\dagger}, b_j^{\dagger}] = 0 \ \text{for all} \ i, j. \tag{2.19}$$

Interestingly, these relations are neither bosonic nor fermionic. Some authors have referred to them as parafermionic [148]. However, as already mentioned, I shall refer to them as *qubit commutation relations* instead.

Single and double qubit evolutions are expressed, respectively, by the unitaries

$$Q_{ik}(\theta) = e^{\theta \Omega_{ik}} = \exp\left[i\frac{\theta}{2}(X_i Y_k - Y_i X_k)\right] \quad \text{and} \tag{2.20}$$

$$Q_{ijkl}(\theta) = e^{\theta \Omega_{ijkl}} = \exp\Big[i\frac{\theta}{8}(X_i Y_j X_k X_l + Y_i X_j X_k X_l + Y_i Y_j Y_k X_l + Y_i Y_j X_k Y_l$$
$$- X_i X_j Y_k X_l - X_i X_j X_k Y_l - Y_i X_j Y_k Y_l - X_i Y_j Y_k Y_l)\Big]. \tag{2.21}$$

These unitary operations are similar to the single and double fermionic evolutions [Eqs. (2.13) and (2.14)], apart from lacking in the exponential the products of Pauli-Z operators that compute the parity of the qubit state, and account for the anticommutation relations of the fermionic ladder operators.

2.3.1 Single-qubit-evolution circuit

A single qubit evolution $Q_{ik}(\theta)$ [Eq. (2.20)] is a 2-qubit operation that acts to continuously exchange states $|0_i 1_k\rangle$ and $|1_i 0_k\rangle$ as θ is varied, and acts trivially on all other computational basis states. Its action on a 2-qubit state can be represented by the matrix

$$Q_{01}(\theta) \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{2.22}$$

This action also corresponds to an evolution of an exchange interaction [126, 149]. A single qubit evolution can be implemented optimally by a circuit with only 2 *CNOT*s. I present a variant of such a circuit in Fig. 2.9.
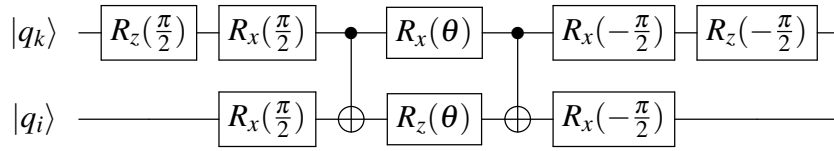
Fig. 2.9 A *CNOT* optimal circuit to implement a single qubit evolution [Eq. (2.20)].

However, for the purpose of constructing a double-qubit-evolution circuit, and consequently a double-fermionic-evolution circuit, I construct a single-qubit-evolution circuit as a *partial SWAP gate*. A *SWAP* gate can be expressed by 3 *CNOT*s as shown in Fig. 2.3, and has a matrix representation as given in Table 1.4. I modify the *SWAP* circuit, by replacing the middle *CNOT* with a controlled rotation $cR_y(-2\theta)_{0,\{1\}}$ so that we obtain the circuit in Fig. 2.10.a. The action of this circuit is

$$CNOT_{0,1} cR_y(-2\theta)_{0,\{1\}} CNOT_{0,1} =$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos\theta & \sin\theta \\ 0 & 0 & -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{2.23}$$

which is equivalent to a single qubit evolution. The controlled rotation $cR_y(-2\theta)_{0,\{1\}}$ can be implemented as in Eq. (2.3), using 2 *CP*s instead of *CNOT*s. This allows one of the *CP*s to be removed by applying the circuit identity in Fig. 2.4. Then, by converting the remaining *CP* to a *CNOT* (Fig. 2.2), and cancelling adjacent single-qubit gates, the explicit basic-quantum-gate circuit in Fig. 2.10.b can be obtained. This circuit uses 3 *CNOT*s, one

more *CNOT* than the circuit in Fig. 2.9. However, the partial *SWAP* circuit in Fig. 2.10 has parametrised qubit rotations acting on one qubit only. As we will see next, this is an advantage when expanding the circuit to implement a double qubit evolution.

(a)

(b)

Fig. 2.10 **a)** A circuit implementing a single qubit evolution $Q_{ik}(\theta)$ [Eq. (2.20)]. **b)** An explicit basic-quantum-gate circuit for **a)** obtained by implementing the controlled-$R_y(-2\theta)$ rotation as in Eq. (2.3), and using the circuit identity in Fig. 2.4, to reduce the *CNOT* cost by 1.

2.3.2 Double-qubit-evolution circuit

A double qubit evolution $Q_{ijkl}(\theta)$ [Eq. (2.20)] is a qubit operation that continuously exchanges states $|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle$ and $|0_i\rangle|0_j\rangle|1_k\rangle|1_l\rangle$ as θ is varied, and acts trivially on all other computational basis states. Its action on a 4-qubit state can be represented by the matrix

$$Q_{ijkl}(\theta) \equiv \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \cos\theta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\sin\theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sin\theta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cos\theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (2.24)$$

This operation can be implemented by expanding the partial *SWAP* circuit, which implements a single qubit evolution (Fig. 2.10). Now it should exchange the states $|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle$ and $|0_i\rangle|0_j\rangle|1_k\rangle|1_l\rangle$, instead of the states $|1_i\rangle|0_k\rangle$ and $|0_i\rangle|1_k\rangle$. To ensure this, it must act non-trivially only if the parities of both qubit pairs $\{k,l\}$ and $\{i,j\}$ are even. To perform this parity-controlled partial *SWAP* operation, I construct the circuit in Fig. 2.11.

In order to understand better the logic of this circuit, let us consider the transformation of state $|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle$. The first two *CNOT*s, between qubits l and k, and qubits j and i, compute and encode the parities of the two respective qubit pairs on qubits k and i, respectively (0 for even parity):

$$|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle \rightarrow |0_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle$$

Then qubits k and i are used as control qubits, with a control state $|0_i\rangle|0_k\rangle$, for a controlled partial *SWAP* operation (the dotted rectangle in Fig. 2.11), between qubits l and j:

$$|0_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle \rightarrow \cos\theta|0_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle + \sin\theta|0_i\rangle|0_j\rangle|0_k\rangle|1_l\rangle$$

Lastly, the last two *CNOT*s between qubits k and l, and qubits i and j, respectively, reverse the parity-encoding action of the first two *CNOT*s:

$$\cos\theta|0_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle + \sin\theta|0_i\rangle|0_j\rangle|0_k\rangle|1_l\rangle \rightarrow \cos\theta|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle + \sin\theta|0_i\rangle|0_j\rangle|1_k\rangle|1_l\rangle$$

For all initial states except $|0_i\rangle|0_j\rangle|1_k\rangle|1_l\rangle$ and $|1_i\rangle|1_j\rangle|0_k\rangle|0_l\rangle$, the controlled partial SWAP operation will act trivially, hence the overall action of the circuit in Fig. 2.11 will be trivial as well.
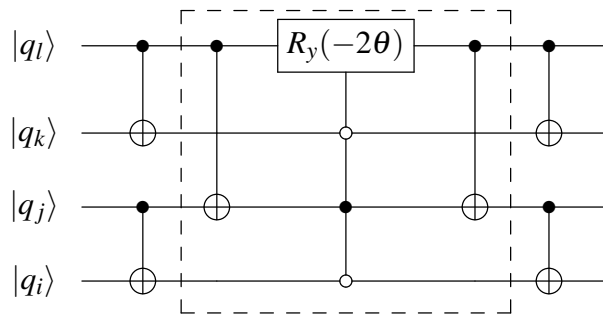

Fig. 2.11 A circuit to implement a double qubit evolution $Q_{ijkl}(\theta)$ [Eq. (2.21)]. The explicit basic-quantum-gate circuit is given in Fig. 2.12b. The dotted rectangle denotes a partial *SWAP* gate between qubits l and j controlled by qubits k and i.

The controlled $cR_y(-2\theta)_{l,\{kji\}}$ rotation in Fig. 2.11 is implemented by the circuit in Fig. 2.12a. The circuit identity in Fig. 2.4 is again applied to reduce the *CNOT* cost by one. In this way, after some single-qubit cancellation, the explicit basic-quantum-gate circuit in Fig. 2.12b is obtained. This circuit has a *CNOT* count of 13 and *CNOT* depth of 11.

I note that a circuit for an operation equivalent to a double qubit evolution was suggested in Refs. [31, 150]. This circuit also has a *CNOT* count of 13, but is less compact and has a higher *CNOT* depth of 13.

2.4 Efficient fermionic-evolution circuits

As already mentioned, the expressions for single and double fermionic evolutions [Eqs. (2.13) and (2.14)] differ from those for single and double qubit evolutions [Eqs. (2.20) and (2.21)] only by the additional products of Pauli-Z operators in their exponents. These products compute the parity of the corresponding qubit state, and account for the fermionic anticommutation relations [Eq. (2.6)]: in a single (double) fermionic excitation, the Pauli-Z products change the sign before the parameter θ if the parity of qubit state $|q_{i+1}\rangle \dots |q_{k-1}\rangle$ $\big(|q_{i+1}\rangle \dots |q_{j-1}\rangle |q_{k+1}\rangle \dots |q_{l-1}\rangle\big)$ is odd. The fermionic evolutions can be expressed in terms of qubit evolutions as

$$F_{ik}(\theta) = \begin{cases} Q_{ik}(\theta) \text{ if } Par\big(|q_{i+1}\rangle \dots |q_{k-1}\rangle\big) = 0 \\ \\ Q_{ik}(-\theta) \text{ if } Par\big(|q_{i+1}\rangle \dots |q_{k-1}\rangle\big) = 1 \end{cases} , \quad \text{and} \quad (2.25)$$

$$F_{ijkl}(\theta) = \begin{cases} Q_{ijkl}(\theta) \text{ if } Par\big(|q_{i+1}\rangle \dots |q_{j-1}\rangle |q_{k+1}\rangle \dots |q_{l-1}\rangle\big) = 0 \\ \\ Q_{ijkl}(-\theta) \text{ if } Par\big(|q_{i+1}\rangle \dots |q_{j-1}\rangle |q_{k+1}\rangle \dots |q_{l-1}\rangle\big) = 1 \end{cases} . \quad (2.26)$$

I expand the circuits for qubit evolutions in Figs. 2.10 and 2.11 to implement fermionic evolutions, by sandwiching the controlled-$R_y(-2\theta)$ rotation in each of the circuits between two *CNOT* staircases (Figs. 2.13 and 2.14). In this way the *CNOT* staircases compute the parities of the relevant qubit states, as given by Eqs. (2.25) and (2.26), and change the sign before the corresponding excitation parameter θ, thus recovering the action of a fermionic evolution. Compared to the canonical circuits for fermionic evolutions, (see Sec. 2.2), the circuits outlined here use only 2 *CNOT* staircases, instead of 4 or 16, per single or double
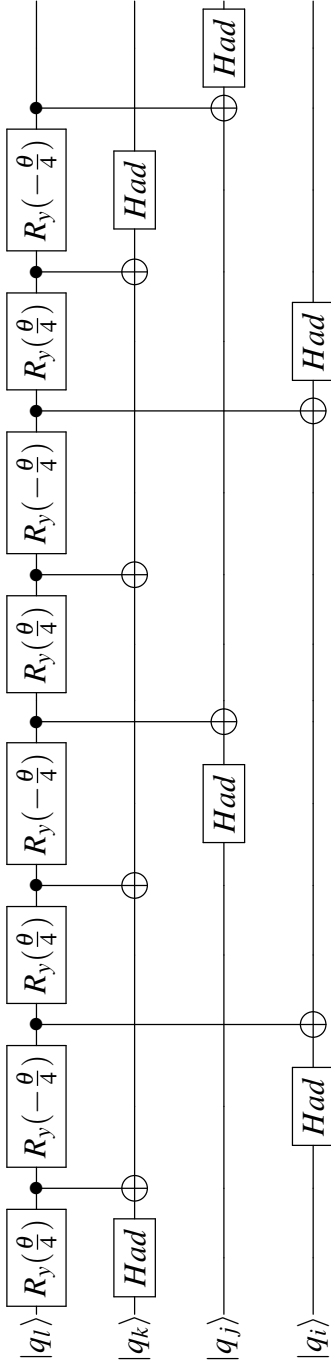
(a) An explicit basic-quantum-gate circuit to implement the controlled rotation $cR_y(2\theta)_{l,\{jkl\}}$ used in the circuit for a double qubit evolution in Fig. 2.11. The circuit is obtained with the method described in Sec. 2.1.3. The directions of the *CNOT* are reversed, as shown in Fig. 2.1.
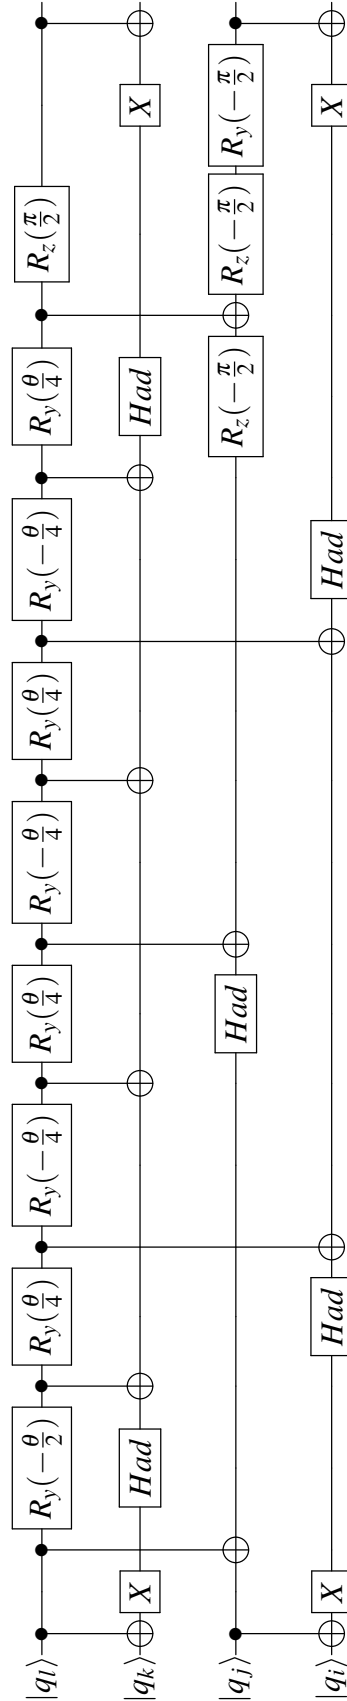
(b) An explicit basic-quantum-gate circuit to implement a double qubit evolution $Q_{ijkl}(\theta)$ [Eq. (2.21)].

Fig. 2.12

fermionic evolution, respectively. This construction can be expanded to any order of a fermionic evolution.

2.4.1 Single-fermionic-evolution circuit

Figure 2.13 shows the circuit for a single fermionic evolution, obtained by modifying the circuit for a single qubit evolution in Fig. 2.10a. The parity of $|q_{i+1}\rangle \dots |q_{k-1}\rangle$ is computed and recorded in qubit $(k-1)$ by a staircase of *CNOT*s. Conditioned on qubit $(k-1)$ being in state $|1\rangle$ (odd parity), the two *CP* gates between qubits k and $k-1$, reverse the direction of the $R_y(-2\theta)$ rotation: $R_y(-2\theta) \to R_y(2\theta)$. The controlled $R_y(-2\theta)$ rotation is implemented as in the single-qubit-evolution circuit (Fig. 2.12a) so that the circuit identity in Fig. 2.4 is again applied to reduce the number of *CNOT*s (*CP*s) by 1. A second *CNOT* staircase corrects for the parity computing action of the first one.

Let me define the total number of qubits involved in the single fermionic evolution as $n_{\text{sf}} \equiv k - i + 1$. For $n_{\text{sf}} \geq 2$, the circuit in Fig. 2.13 has a *CNOT* count of $(2n_{\text{sf}} - 1)$ and a *CNOT* depth of $\max[5, 2n_{\text{sf}} - 3]$. For $n_{\text{sf}} = 2$, the circuit equals that in Fig. 2.10b. Additionally, the *CNOT*s in the two staircases can be rearranged as described in Sec. 2.1.2, so that the *CNOT* depth of the circuit scales as $O(\log n_{sf})$, instead of $O(n_{sf})$.

Compared to the canonical circuit (Fig. 2.8) for a single fermionic evolution, the circuit constructed here, requires up to 2 times fewer *CNOT*s and is up to twice as shallow.
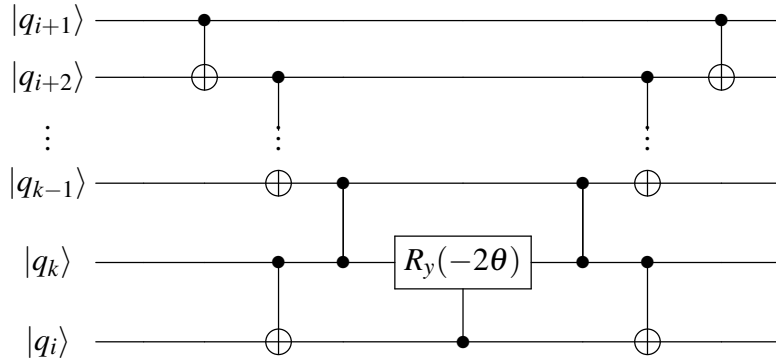
Fig. 2.13 A *CNOT*-efficient circuit to implement single fermionic evolution $F_{ik}(\theta)$ [Eq. (2.13)].

2.4.2 Double fermionic-evolution circuit

Figure 2.14 shows the circuit for a double fermionic evolution, obtained by modifying the double-qubit-evolution circuit in Fig. 2.11. The parity of $|q_{i+1}\rangle \dots |q_{j-1}\rangle |q_{k+1}\rangle \dots |q_{l-1}\rangle$ is

computed and recorded on qubit $(l-1)$ by a *CNOT* staircase. Conditioned on qubit $(l-1)$ being in state $|1\rangle$, the two *CP* gates between qubits l and $(l-1)$ reverse the direction of the controlled $R_y(-2\theta)$ rotation. The controlled $R_y(-2\theta)$ rotation is implemented as in Fig. 2.12a, so that the circuit identity, from Fig. 2.4, can be applied again to reduce the number of *CNOT*s by 1. A second *CNOT* staircase corrects for the parity computing action of the first one.

Let me define the total number of qubits participating in the double fermionic evolution as $n_{\mathrm{df}} \equiv l + j - k - i + 2$. The circuit in Fig. 2.14 has a *CNOT* count of $2n_{\mathrm{df}} + 5$ and a *CNOT* depth of $\max[13, 2n_{\mathrm{df}} - 1]$ for $n_{\mathrm{df}} \geq 5$. For $n_{\mathrm{df}} = 4$ the circuit equals that in Figs. 2.11. Again, the *CNOT*s in the two staircases can be rearranged as described in Sec. 2.1.2, so that the *CNOT* depth scaling with n_{df} is reduced from $O(n_{\mathrm{df}})$ to $O(\log n_{df})$.

Compared to the canonical double-fermionic-evolution circuit (Sec. 2.2) the circuit constructed here requires up to 8 times fewer *CNOT*s, and is up to 8 times shallower, in terms of *CNOT*s.
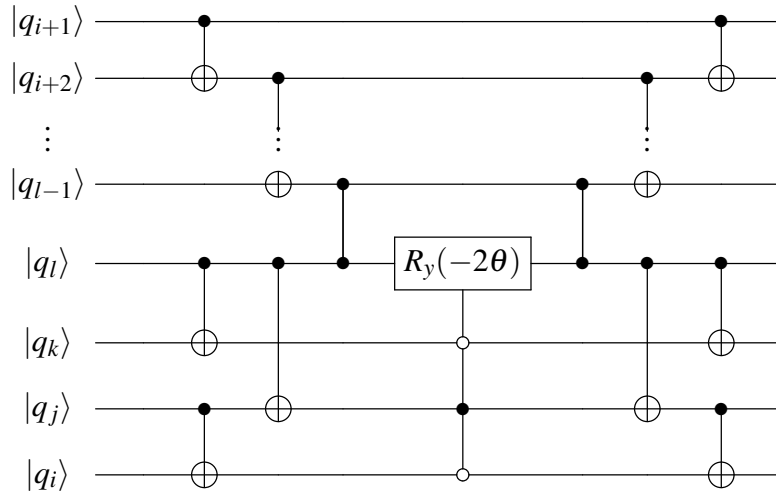

Fig. 2.14 A *CNOT*-efficient circuit to implement a double fermionic evolution $F_{ijkl}(\theta)$ [Eq.(2.14)].

2.5   Conclusion

Single and double fermionic evolutions are the building blocks of the majority of UCC ansätze. In this Chapter, I presented *CNOT*-efficient quantum circuits to implement single and double fermionic evolutions in the Jordan Wigner encoding.

First, I constructed circuits that implement single and double qubit evolutions, evolutions of qubit excitation operators that satisfy qubit commutation relations. Next, I expanded the functionality of the qubit evolution circuits to account for qubit state parity, in order to implement fermionic evolutions instead. The resulting circuits for single and double fermionic evolutions consist of 2 *CNOT* staircases, as compared to 4 and 16 *CNOT* staircases for the canonical single- and double-fermionic-evolution circuits, respectively. Hence, the single-(double-)fermionic-evolution circuit constructed here, has a *CNOT* count and a depth that scale by a factor of 2 (8) less than the *CNOT* count and the depth of the canonical single- (double-)fermionic-evolution circuit. The exact *CNOT* counts for the circuits are summarised in the table below.

Evolution	Canonical circuit	Efficient circuit
single fermionic F_{ik}	$4(k-i) = 4n_{\text{sf}} - 4$	$2(k-i) + 1 = 2n_{\text{sf}} - 1$
double fermionic F_{ijkl}	$16(l+j-k-i) = 16n_{\text{df}} - 16$	$2(l+j-k-i) + 9 = 2n_{\text{df}} + 5$

These circuits can be readily integrated in existing VQE protocols, which use UCC ansätze, to reduce their quantum hardware requirements. The *CNOT* gate reduction achieved by constructing a UCC ansatz with our circuits is bounded between 50%, for single fermionic evolutions, and 87% for double fermionic evolutions. However, for typical UCC ansätze, e.g. the UCCSD, where double fermionic evolutions dominate, the minimum *CNOT* count reduction is about 70%. As the system size (N_{so}) increases and the number of double fermionic evolutions in the ansatz becomes more dominant, the *CNOT* count reduction increases and converges to 87%. For example, for a UCCSD ansatz for BeH$_2$, in the STO-3G basis ($N_{so} = 14$), the gate reduction is about 81%.

The ansatz circuits constructed here correspond to *local* circuit optimizations, in the sense that individual parts of the ansatz are optimized independently from one another. *Global circuit optimizations* on the other hand consider the ansatz circuit as a whole. Examples of global circuit optimizations for VQE ansätze, include the various optimization methods deployed by IBM's compiler, Qiskit [151], with gate count reductions of up to 30% for UCC ansätze, and a method deployed by the compiler of Cambridge Quantum Computing (CQC)[152] that achieves gate count reduction of nearly 60%. Shortly after we published the results presented in this Chapter, the researchers at CQC presented another, competitive global optimization method that achieves a *CNOT* count reduction of up to 75% [153].

Lastly, I remark that the fermionic evolution circuits presented in this chapter correspond to the JW encoding. The *CNOT* counts of these circuits, scale linearly with the system size, as $O(N_{so})$. Using the BK encoding (Sec. 1.5.1), fermionic evolution circuits can be

constructed with a number of $CNOT$s that scales logarithmically, as $O(\log N_{so})$. However, the BK encoding gate-count reduction becomes significant only for large systems. For example, the BK gate-count reduction (with respect to the canonical fermionic evolution circuits in the JW enconding) is at maximum about 25% for systems of up to 50 spin-orbitals (qubits) [105]. Therefore, using the JW encoding together with the fermionic evolution circuits constructed here will be more advantageous than using the BK encoding for small ($N_{so} < 50$) and perhaps medium size ($50 < N_{so} < 100$) electronic systems.

In the next chapter I propose the use of a new type of ansätze for molecular modelling with the VQE that consist of qubit instead of fermionic evolutions. Qubit evolutions act on a fixed number of qubits, so they are implemented by fixed size circuits, whose gate count does not scale with the system size. Hence, qubit evolutions are potentially more circuit-efficient than fermionic evolutions, irrespectively of the qubit encoding method.

Chapter 3

Qubit unitary coupled cluster ansätze

Chapter summary

In second-quantized basis-set methods, the exchange symmetry of electrons is accounted for by the anticommutation relations of the fermionic ladder operators that act on the electronic wavefunction. In the context of molecular modelling with VQE protocols, the fermionic exchange symmetry is accounted for by the fermionic-operator-representation of the electronic Hamiltonian [Eq. (1.25)], and consequently translated to its quantum-gate-operator representation [Eq. (1.39)] by the respective qubit encoding of the fermionic ladder operators. The electronic wavefunction is expressed as a linear combination (superposition) of Slater determinants, where each Slater determinant is represented by a computational basis state of the qubit register, which is independent of the particles' statistics. Therefore, the structure of the ansatz, which generates the trial qubit state of a molecular VQE protocol, is not restricted by an exchange symmetry conservation.

In Chapter 2, I introduced qubit evolutions as a means to derive $CNOT$-efficient circuits to implement fermionic evolutions, which are used to construct UCC ansätze. In this chapter, I consider using qubit evolutions to construct molecular VQE ansätze, thus disposing of the fermionic structure of UCC ansätze.

Qubit excitation operators act on a fixed number of qubits, as opposed to fermionic excitation operators that act on $O(N_{so})$ $\left(O(\log N_{so})\right)$ qubits in the JW (BK) encoding to account for fermionic exchange-symmetry. Hence, qubit evolutions are implemented by fixed-sized quantum circuits that require $O(1)$ $CNOT$s, as opposed to fermionic evolutions that are implemented by circuits that require $O(N_{so})$ $\left(O(\log N_{so})\right)$ $CNOT$s in the JW (BK) encoding. Moreover, qubit evolutions capture many of the physical features of fermionic evolutions, like time-symmetry and particle conservation. This and the fact that qubit evolutions are implemented by constant size circuits, suggest that ansätze, built as series of

qubit evolutions, might be able to approximate electronic wavefunctions almost as accurately as UCC ansätze, while requiring asymptotically fewer *CNOT*s. In this Chapter, I investigate this hypothesis using classical numerical VQE simulations. Throughout the rest of the thesis, I will refer to qubit-evolution-based ansätze as Q-UCC. A similar terminology is used in Ref. [154], which also studies Q-UCC ansätze.

In Sec. 3.1 I formally define Q-UCC ansätze. In Sec. 3.2 I perform a comparison between Q-UCCSD and UCCSD ansätze based on numerical results obtained with classical VQE simulations.

3.1 Qubit unitary coupled cluster ansätze

A general Q-UCC ansatz can be written as

$$
U_{Q-UCC} = \exp\Big(\sum_{k,i} \theta_{k,i} \Omega_{ki} + \sum_{l,k,j,i} \theta_{l,k,j,i} \Omega_{lkji} + \dots \Big) =
$$
$$
\exp\Big(\sum_{k,i} \theta_{k,i} \big(b_k^\dagger b_i - b_i^\dagger b_k \big) + \sum_{l,k,j,i} \theta_{l,k,j,i} \big(b_l^\dagger b_k^\dagger b_j b_i - b_j^\dagger b_i^\dagger b_l^\dagger b_k^\dagger \big) + \dots \Big) \tag{3.1}
$$

where the sum goes over all qubit excitation operators $\{\Omega_{ki}\}$, $\{\Omega_{lkji}\}$, …, acting on qubits $\{0, \dots, N_{so} - 1\}$.

For comparison a general UCC ansatz can be written as

$$
U_{UCC} = \exp\Big(\sum_{k,i} \theta_{k,i} \Xi_{ki} + \sum_{l,k,j,i} \theta_{l,k,j,i} \Xi_{lkji} + \dots \Big) =
$$
$$
\exp\Big(\sum_{k,i} \theta_{k,i} \big(a_k^\dagger a_i - a_i^\dagger a_k \big) + \sum_{l,k,j,i} \theta_{l,k,j,i} \big(a_l^\dagger a_k^\dagger a_j a_i - a_j^\dagger a_i^\dagger a_l a_k \big) + \dots \Big) =
$$
$$
\exp\Big(\sum_{k,i} \theta_{k,i} \big(b_k^\dagger b_i - b_i^\dagger b_k \big) \prod_{p=i+1}^{k-1} Z_p - \sum_{l,k,j,i} \theta_{l,k,j,i} \big(b_l^\dagger b_k^\dagger b_j b_i - b_j^\dagger b_i^\dagger b_l^\dagger b_k^\dagger \big) \prod_{p=i+1}^{j-1} Z_p \prod_{p'=k+1}^{l-1} Z_{p'} + \dots \Big),
$$
$$
\tag{3.2}
$$

where again the sum goes over all fermionic excitation operators $\{\Xi_{ki}\}$, $\{\Xi_{lkji}\}$, … acting on qubits $\{0, \dots, N_{so} - 1\}$.

The only difference between the two ansätze are the computationally expensive products of Pauli-Z operators, which account for the exchange symmetry of the fermionic ladder operators, $a^\dagger$ and a, in the case of the UCC. However, irrespective of this, both the sum of qubit excitation operators and the sum of fermionic excitation operators in the exponentials in Eqs. (3.1) and (3.2), respectively, correspond to a universal N_{so}-body interaction [23, 93].

Hence, both the general Q-UCC and the general UCC, by acting on the same initial qubit state, span the same Hilbert space; if the initial state is a computational basis state of Hamming weight[1] N_e (equal to the number of electrons), then the spanned Hilbert space is the space of real valued qubit states with Hamming weight N_e, which contains the exact FCI wavefunction. In fact, it was shown that only single and double excitation operators, and their commutators, are sufficient to exactly approximate the FCI wavefunction [146, 147]. The lack of the products of Pauli-Z operators in the expression for the Q-UCC affects only its parameter dependence.

Even so, similarly as in the case of UCC ansätze, in practice we would like to truncate Q-UCC ansätze to single and double qubit excitation operators only, and to approximate the Q-UCC ansätze using one or a few Trotterization steps [111]. In this way Q-UCC ansätze can be represented as products of single and double qubit evolutions and implemented by the quantum circuits derived in Chapter 2. However, when only single and double excitation operators are included, and a Trotter approximation applied, it is not apparent if the spaces spanned by Q-UCC and UCC ansätze are equivalent, and if the two types of ansätze can approximate electronic wavefunctions equally well.

3.2 Classical numerical simulations

In this section, I compare the Q-UCCSD and the UCCSD ansätze (a Q-UCC and a UCC ansätze that include only single and double excitation operators above the Hartree-Fock state) in their ability to approximate electronic wavefunctions, by performing noiseless classical numerical VQE simulations for small molecules: LiH, chain H_6 and BeH_2. LiH and BeH_2 have been simulated with the VQE on real quantum computers [124, 29], and are often used in the field of quantum computational chemistry to classically benchmark various VQE protocols [155, 141, 154, 156]. H_6 corresponds to a chain of 6 hydrogen atoms. The ground state of this molecule is metastable and it does not occur in nature. Here I used it, similarly as in Ref. [156], as a prototype of a molecule with a strongly correlated ground state.

The numerical simulations presented in this thesis are performed by a custom in-house software, which I call the Cavendish Quantum Computational Chemistry (CQCC) package[2], designed to simulate and study the performance of VQE protocols. For the noiseless simulations presented in the thesis the expectation values of operators, including H, are calculated as products of their matrix representations and the corresponding trial state statevector. The statevector is calculated by a statevector simulator designed specifically for the quick compu-

---

[1]The Hamming weight of a multi-qubit computational basis state, is the number of qubits in state $|1\rangle$.

[2]The code is available at https://github.com/JordanovSJ/VQE

tation of statevectors, generated by qubit/fermionic-evolution-based ansätze (see Appendix C.1). The classical optimization subroutine of the VQE, unless otherwise stated, is performed by the gradient descent Broyden–Fletcher–Goldfarb–Shanno (BFGS) [157] method. For the majority of simulations presented in the thesis, the BFGS is supplied an explicit function, suggested in Ref. [141], to calculate the energy-gradient vector of the ansatz (see Appendix C.2). Additional computational details about the CQCC package and the classical numerical simulations presented in the thesis are given in Appendix C.

The simulation results presented here, are obtained for all three molecules (LiH, H_6 and BeH_2) in the STO-3G (see Sec. 1.4.4) spin-orbital basis, without assuming frozen orbitals[3]. For LiH, H_6 and BeH_2 the STO-3G basis includes 12, 12 and 14 spin-orbitals, respectively, represented by 12, 12 and 14 qubits. All simulation are performed with the JW qubit encoding.

3.2.1   Q-UCCSD versus UCCSD

Analogously to the UCSSD, let us define the Q-UCCSD as a Q-UCC ansatz that contains all single and double qubit excitation operators above the Hartree Fock state. For both the UCCSD and the Q-UCCSD, let us consider a single-step Trotter approximation. Hence, the Q-UCCSD is given by the product of single and double qubit evolutions,

$$
U_{Q-UCCSD} = \prod_{k \in virt, i \in occ} \exp\left(\theta_{ki} \Omega_{ki}\right) \prod_{k,l \in virt, j, i \in occ} \exp\left(\theta_{lkji} \Omega_{lkji}\right) =
$$

$$
\prod_{k \in virt, i \in occ} \exp\left(\theta_{ki}\left(b_k^\dagger b_i - b_i^\dagger b_k\right)\right) \prod_{k,l \in virt, j, i \in occ} \exp\left(\theta_{lkji}\left(b_l^\dagger b_k^\dagger b_j b_i - b_j^\dagger b_i^\dagger b_l b_k\right)\right). \quad (3.3)
$$

and the UCCSD by the product of single and double fermionic evolutions

$$
U_{UCCSD} = \prod_{k \in virt, i \in occ} \exp\left(\theta_{ki} \Xi_{ki}\right) \prod_{k,l \in virt, j, i \in occ} \exp\left(\theta_{lkji} \Xi_{lkji}\right) =
$$

$$
\prod_{k \in virt, i \in occ} \exp\left(\theta_{ki}\left(a_k^\dagger a_i - a_i^\dagger a_k\right)\right) \prod_{k,l \in virt, j, i \in occ} \exp\left(\theta_{lkji}\left(a_l^\dagger a_k^\dagger a_j a_i - a_j^\dagger a_i^\dagger a_l a_k\right)\right). \quad (3.4)
$$

I compare the Q-UCCSD and the UCCSD by using the VQE with each of the two ansätze (let us refer to these as the UCCSD-VQE and the Q-UCCSD-VQE methods) to obtain energy dissociation curves (energy as function of bond distance) for LiH, H_6 and BeH_2. The simulation results are presented in Fig. 3.1.

---

[3]A frozen spin-orbital is a spin-orbital assumed to be fully occupied or unoccupied
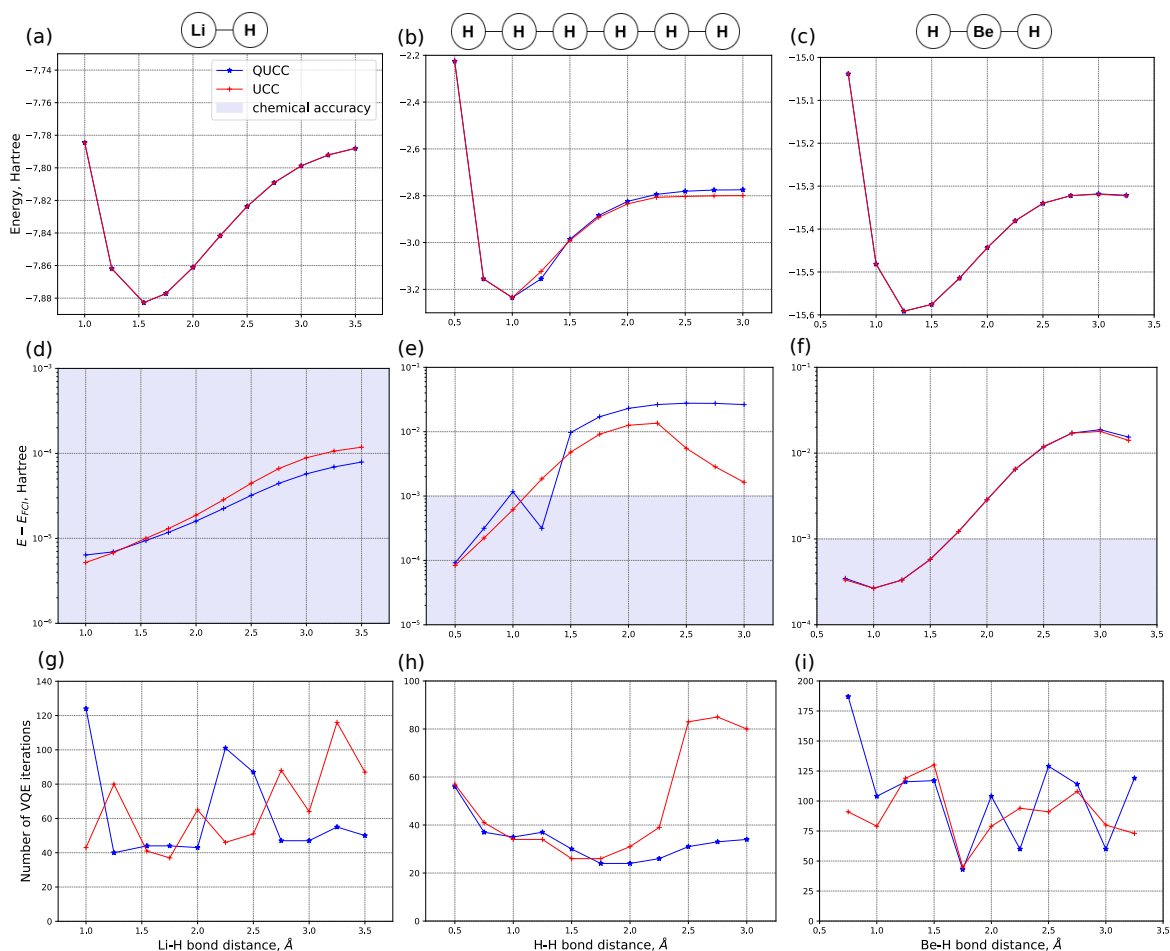
Fig. 3.1 Energy dissociation curves for LiH, H_6 and BeH_2 in the STO-3G basis, obtained with the UCCSD-VQE and the Q-UCCSD-VQE. **a,b,c** Absolute value for the ground state energy estimation. **d,e,f** Error in the estimated value for the ground state energy with respect to the exact FCI energy. **g,h,i** Number of VQE iterations required to optimize the parameters of the ansätze.

First, we observe that the energy dissociation curves obtained by Q-UCCSD-VQE and UCCSD-VQE for LiH, H_6 and BeH_2 in Figs. 3.1.a,b,c respectively, are close to indistinguishable. In Figs. 3.1.d,e,f the FCI energy [4] is subtracted from the absolute values of the energy estimates, thus plots of the energy errors as function of the bond distance are obtained for the Q-UCCSD-VQE and the UCCSD-VQE. For BeH_2 the error plots (Fig. 3.1.f) for the two types of ansätze are almost identical, and for LiH (Fig. 3.1.d) the Q-UCCSD-VQE is slightly more accurate on average. However, for H_6 there is a significant difference in the error plots (Fig. 3.1.e) at large bond distances, with the UCCSD-VQE being more accurate

---

[4]For small molecules, the exact FCI energy can be obtained by diagonalizing the matrix representation of the electronic Hamiltonian, and getting the lowest eigenvalue.

in this case. Nevertheless, for H_6 at bond distances where chemical accuracy is achieved, the two ansätze actually perform similarly. These results are not conclusive, but signify that even with only single and double excitations included, the Q-UCCSD can approximate electronic wavefunction almost as accurately as the UCCSD. Similar results on the accuracy of the Q-UCCSD, as compared to that of the UCCSD were obtained in Ref. [154].

We are also interested in how easy to optimize the Q-UCCSD is, in comparison to the UCCSD. The ease of optimization of an ansatz can be quantified by the number of VQE-iterations required for convergence. Plots of the iteration numbers as function of the bond distance, for the Q-UCCSD-VQE and the UCCSD-VQE, are given in Figs. 3.1g,h,i. On average the number of iterations required for convergence is roughly the same for each method. The only consistent difference in the number of VQE iterations appears in the case of H_6: at bond distances where the UCCSD is more accurate, it consistently requires more VQE iterations to converge, which is expected. Overall, it can be concluded that the Q-UCCSD is as easy to optimize as the UCCSD.

Despite the close equivalence in terms of accuracy and number of iterations required for convergence of the two types of ansätze, the Q-UCCSD has lower gate counts due to it being constructed by qubit evolutions. The numbers of fermionic and qubit evolutions included in the UCCSD and the Q-UCCSD, respectively, scale as $O((N_{so} - N_e)^2 N_e^2)$. Hence, the *CNOT* count scaling of the Q-UCCSD is $O((N_{so} - N_e)^2 N_e^2)$, and the *CNOT* count scaling of the UCCSD is $O((N_{so} - N_e)^2 N_e^2 N_{so})$ [5]. The table below summarizes the exact gates counts of the two ansätze for each molecule[6].

	UCCSD	Q-UCCSD
LiH, single-qubit gates	6186	5768
LiH, *CNOT* gates	3496	2280
H_6, single-qubit gates	8217	7677
H_6, *CNOT* gates	4593	3033
BeH_2, single-qubit gates	13678	12900
BeH_2, *CNOT* gates	8347	5158

Table 3.1 UCCSD and Q-UCCSD gates counts for LiH, H_6 and BeH_2 in the STO-3G basis. The two ansätze include all single and double qubit/fermionic evolutions above the Hartree-Fock state.

The *CNOT* count reduction from the UCCSD to the Q-UCCSD is not so significant for these three molecules (about 30% for each molecule). However, this reduction will grow

---

[5]The *CNOT* count of UCCSD would scale as $O((N_{so} - N_e)^2 N_e^2 \log N_{so})$ in the BK encoding
[6]The gate counts are obtained for the qubit and fermionic evolution circuits derived in the Chapter 2.

linearly with the system size, N_{so}. Thus, for large molecules we can expect that Q-UCC ansätze will be more advantageous than UCC ansätze.

I finally remark that both the Q-UCCSD and the UCCSD ansätze fail to achieve chemical accuracy at most bond distances in the cases of H_6 and BeH_2 (Fig. 3.1d). The likely reason for this is that the ansätze fail to recover the correlation effects that become stronger as the dissociation distance is increased [95]. I will come back to this issue in Chapter 4 when iterative VQE protocols are introduced to construct problem tailored ansätze.

3.3 Conclusion

In this Chapter, I proposed the use of Q-UCC ansätze for molecular modelling that are constructed as products of qubit evolutions. This proposal was based on the fact that qubit evolutions are performed by quantum circuits that require asymptotically fewer *CNOT*s than the standard fermionic evolutions, and a presumption that Q-UCC ansätze can approximate electronic wavefunctions as accurately as UCC ansätze.

I investigated this presumption with classical numerical VQE simulations for LiH, H_6 and BeH_2 using the Q-UCCSD and the UCCSD ansätze. For LiH and BeH_2, the two ansätze could achieve nearly identical accuracies in the estimation of the ground state energies. However, in the case of the more strongly correlated H_6, at large bond distances, the UCCSD was noticeably more accurate than the Q-UCCSD. These results indicated that the UCCSD can be more accurate than the Q-UCCSD for strongly correlated states. Nonetheless, the slight advantage of the UCCSD in terms of accuracy, is likely to be off-set by the circuit-efficiency of the Q-UCCSD, when it comes to simulations on NISQ computers. However, these results are inconclusive and additional investigation is required. In terms of ease of optimization, the two ansätze required approximately the same number of VQE iterations for convergence on average. These results suggest that the Pauli-Z strings that account for the anticommutation of the fermionic ladder operators, and are present in the expressions for UCC ansätze, play little role in the variational flexibility of an ansatz. Similar conclusions were derived in Refs. [154, 156].

Comparing the Q-UCCSD and the UCCSD is not a thorough comparison of qubit and fermionic evolutions in their ability to approximate electronic wavefunctions. For instance, both the Q-UCCSD and the UCCSD have redundantly high numbers of variational parameters, whose additional variational flexibility may partially compensate for the lack of fermionic features of qubit evolutions. In the next chapter, I introduce iterative VQE protocols that construct problem tailored ansätze, using significantly fewer parameters than fixed ansätze,

such as the Q-UCCSD and the UCCSD. There, I perform a comparison between problem-tailored iteratively-constructed Q-UCC and UCC ansätze.

Chapter 4

The qubit-excitation based adaptive VQE

Chapter summary

Most ansätze used for molecular modelling with the VQE, such as the various UCC ansätze, e.g. the UCCSD [29, 119], the GUCCSD [135] and the k-upCCSD [135], are *fixed ansätze*, whose structures are determined only by the number of spin-orbitals (qubits), N_{so}, and by the number of electrons, N_e. Generally, even simple fixed UCC ansätze, like the UCCSD, are accurate for weakly correlated systems and are easy to optimize due to their fermionic structure. In Chapter 3 we saw that this also holds for the Q-UCCSD ansatz.

However, in the cases of H_6 and BeH_2 (Fig. 3.1.d), both the UCCSD and the Q-UCCSD failed to achieve chemical accuracy for long bond distances where correlation effects become stronger. In order to approximate a strongly correlated wavefunction, fixed UCC and Q-UCC ansätze need to include higher order excitations and/or use multiple-step Trotterzation, both of which result in longer ansatz circuits and more variational parameters. Fixed UCC and Q-UCC ansätze also include redundant excitation terms that do not contribute to the approximation of the electronic wavefunction, and unnecessarily increase the length of the ansatz circuit and the number of variational parameters.

To solve these issues, recently, a number of works [141, 156, 158, 159, 155, 142, 160] suggested VQE protocols that instead of using fixed ansätze, iteratively construct ansätze tailored to the system of interest. Here I refer to these as *iterative VQE protocols*.

The ADAPT-VQE protocols [141, 156] are a family of iterative VQE protocols that grow a problem tailored ansatz by iteratively appending together parametrized unitary operators, which I refer to as *ansatz elements*, sampled from a predefined *ansatz element pool*. The

sampling is based on an *ansatz-growing strategy* that aims to minimize the energy estimate $E(\vec{\theta})$, after a VQE optimization at each iteration. In this way, an arbitrarily accurate ansatz can be constructed that has fewer variational parameters, and a correspondingly shallower circuit than a fixed ansatz of the same accuracy. Similarly to the motivation behind the VQE, the motivation behind ADAPT-VQE protocols is to further reduce the circuit depth at the expense of more quantum computer measurements and classical processing. Generally, the number of additional quantum computer measurements required by ADAPT-VQE protocols, scales polynomially with the number of ansatz-constructing iterations[1]. The fermionic-ADAPT-VQE protocol [141] was the first iterative VQE protocol. It grows a problem tailored UCC ansatz by iteratively appending spin-complement pairs of single and double fermionic evolutions[2]. The fermionic excitation operator, defining the spin-complement pair of fermionic evolutions at each ansatz-growing-iteration, is sampled from an ansatz element pool of single and double fermionic excitation operators. The ansatz-growing strategy is based on an energy gradient hierarchy. The fermionic-ADAPT-VQE was demonstrated to achieve chemical accuracy for various molecules, using ansätze with several times fewer variational parameters and correspondingly shallower circuits, than the UCCSD ansatz. In the follow-up work [156], the qubit-ADAPT-VQE was proposed, utilizing an ansatz element pool of more rudimentary, but more variationally flexible evolutions of individual Pauli strings. Due to this pool, the qubit-ADAPT-VQE constructs even shallower ansatz circuits than the fermionic-ADAPT-VQE, but requires more variational parameters and ansatz-growing iterations to achieve a given accuracy.

The Evolutionary-VQE (EVQE), proposed in Ref. [158], is another type of an iterative VQE protocol, suitable for optimizations outside the field of quantum computational chemistry, that constructs its ansatz using basic quantum-circuit constructions. In addition, the EVQE also accounts for hardware noise in constructing its ansatz. In Refs. [159, 155] the iterative qubit coupled cluster (IQCC) protocol was suggested, where instead of growing an ansatz, the Hamiltonian is transformed iteratively by being "dressed" with operators from a pool of Pauli string evolutions. The IQCC can use arbitrary-shallow constant-depth ansatz circuits, however it needs to evaluate a number of Hamiltonian terms that is exponential in the number of variational parameters. Finally, I mention the adaptive-pruning VQE, suggested in Ref. [142], that starts with a UCC ansatz and iteratively removes irrelevant parts of it.

---

[1]In the context of ADAPT-VQE protocols, I shall refer to the parameter-optimizing iterations of a single VQE run as *VQE-iterations*, and the iterations of an ADAPT-VQE protocol as *ansatz-growing iterations*; a single ansatz-growing iteration involves a complete VQE run, which involves multiple VQE iterations

[2] The spin-complementary of unitary operator, acting on some spin orbitals, is another unitary operator, that acts on the opposite-spin orbitals.

Motivated by the similarity of the Q-UCCSD and the UCCSD ansätze in approximating electronic wavefunctions, and the success of iterative VQE protocols, in particular the ADAPT-VQE, here I introduce the qubit-excitation-based adaptive VQE (QEB-ADAPT-VQE) protocol. As the name suggests, the QEB-ADAPT-VQE is an ADAPT-VQE protocol that grows a problem-tailored ansatz by iteratively appending qubit evolutions. The QEB-ADAPT-VQE, and in particular its ansatz-growing strategy, are explained in details in Sec. 4.1. Owing to the utilization of qubit evolutions, the QEB-ADAPT-VQE presumably should construct more circuit-efficient ansätze than the fermionic-ADAPT-VQE without acquiring as large overhead in the number of variational parameters and ansatz-constructing iterations as the qubit-ADAPT-VQE. In Sec. 4.2 I test this hypothesis and benchmark the performance of the QEB-ADAPT-VQE by estimating the ground state energies of small molecules.

In Sec. 4.3 I use my simulation results to compare iteratively constructed Q-UCC and UCC ansätze.

Most of the results presented in this Chapter are based on my manuscript [161].

4.1 The QEB-ADAPT-VQE protocol

4.1.1 Protocol description

Here, I describe the three preparation components and the iterative loop of the QEB-ADAPT-VQE protocol.

Before we start constructing the ansatz, first, we transform the molecular Hamiltonian H to a 2^{nd} quantized form (see Sec. 1.4), by calculating the one- and two-electron integrals h_{ki} [Eq. (1.26)] and h_{lkji} [Eq. (1.27)], respectively. Then, we map H to a quantum-gate-operation representation, as described in Sec. 1.5.1. Since the QEB-ADAPT-VQE uses qubit evolutions, which act on a fixed number of qubits irrespectively of the qubit encoding method, we naturally use the simpler JW encoding. These Hamiltonian transformations are a standard step in every VQE protocol.

Second, we need to define an ansatz element pool. The default choice is to include all unique single and double qubit evolutions, $\{Q_{ki}(\theta)\}$ and $\{Q_{lkji}(\theta)\}$, respectively, acting on the N_{so} spin-orbitals (qubits), so that $i, j, k, l \in \{0, N_{so} - 1\}$. Let us denote this pool as $\mathbb{P}_{\text{full}}(Q, N_{so})$. The size of this pool is $|\mathbb{P}_{\text{full}}(Q, N_{so})| = \binom{N_{so}}{2} + 3\binom{N_{so}}{4}$. Note, that $\mathbb{P}_{\text{full}}(Q, N_{so})$ will include evolutions, where $i < j < k < l$ is not true. Alternatively, a smaller ansatz element pool can be defined that includes fewer elements, e.g. only evolutions of single and double qubit excitation operators above the Hartree-Fock state. Smaller pools will decrease the computational cost (see Sec. 4.1.2), but might result in a less efficient ansatz construction.

Third, we choose an initial reference state $|\psi_0\rangle$. For a guaranteed and faster convergence, $|\psi_0\rangle$ should have a significant overlap with the unknown ground state. For ground state simulations near equilibrium configuration the Hartree-Fock state is a standard choice, because the Hartree Fock Slater determinant is likely to be dominant in the ground state.

Finally, we can begin to construct the ansatz. We set the iteration number to $m = 1$, the initial ansatz to the identity $U(\vec{\theta}) \rightarrow U^{[0]} = I$, and initiate the iterative loop of the QEB-ADAPT-VQE. Below, I describe the six steps of the m^{th} iteration of the QEB-ADAPT-VQE. Afterwards I explain what is the function of each step.

1. Prepare the state $|\psi^{[m-1]}\rangle = U^{[m-1]}(\vec{\theta}^{[m-1]})|\psi_0\rangle$, using the ansatz $U^{[m-1]}(\vec{\theta}^{[m-1]})$, with values for the variational parameters $\vec{\theta}^{[m-1]}$ as determined in the previous iteration.

2. For each qubit evolution $Q_p(\theta_p) = e^{\theta_p \Omega_p} \in \mathbb{P}_{\text{full}}(Q, N_{so})$, which can be either a single or a double, measure the energy gradient

$$
\begin{aligned}
\frac{\partial}{\partial \theta_p} E^{[m-1]}\big(Q_p(\theta_p)\big)\Big|_{\theta_p=0} &= \frac{\partial}{\partial \theta_p}\langle\psi^{[m-1]}|Q_p^\dagger(\theta_p)HQ_p(\theta_p)|\psi^{[m-1]}\rangle\Big|_{\theta_p=0} = \\
&\frac{\partial}{\partial \theta_p}\langle\psi^{[m-1]}|e^{-\theta_p\Omega_p}He^{\theta_p\Omega_p}|\psi^{[m-1]}\rangle\Big|_{\theta_p=0} = \\
&-\langle\psi^{[m-1]}|\Omega_p H|\psi^{[m-1]}\rangle + \langle\psi^{[m-1]}|H\Omega_p|\psi^{[m-1]}\rangle = \\
&\langle\psi^{[m-1]}|[H,\Omega_p]|\psi^{[m-1]}\rangle. \quad (4.1)
\end{aligned}
$$

3. Identify the set of n_{qe} qubit evolutions, $\mathbb{Q}^{[m]}(n_{qe})$, with largest energy-gradient magnitudes, measured in the previous step. For each $Q_p(\theta_p) \in \mathbb{Q}^{[m]}(n_{qe})$:

    (a) Run the VQE to find

    $$
    \min_{\vec{\theta}^{[m-1]},\theta_p} E(\vec{\theta}^{[m-1]},\theta_p) = \min_{\vec{\theta}^{[m-1]},\theta_p}\langle\psi_0|U^{[m-1]\dagger}(\vec{\theta}^{[m-1]})Q_p^\dagger(\theta_p)HQ_p(\theta_p)U^{[m-1]}(\vec{\theta}^{[m-1]})|\psi_0\rangle
    $$

    (b) Find the energy reduction $\Delta E_p^{[m]} = E^{[m-1]} - \min_{\vec{\theta}^{[m-1]},\theta_p} E(\vec{\theta}^{[m-1]},\theta_p)$

    (c) Save the (re)optimized values of $\vec{\theta}^{[m-1]} \cup \{\theta_p\}$ as $\vec{\theta}_p^{[m]}$

4. Identify the largest energy reduction $\Delta E^{[m]} \equiv \Delta E_{p'}^{[m]} = \max\big(\{\Delta E_p^{[m]}\}\big)$, and its corresponding qubit evolution $Q^{[m]}(\theta^{[m]}) \equiv Q_{p'}(\theta_{p'})$.

    If $\Delta E^{[m]} < \varepsilon$, where $\varepsilon > 0$ is an energy-reduction threshold:

    (a) Exit

Else:

(a) Append $Q^{[m]}(\theta^{[m]})$ to the ansatz: $U^{[m]}(\vec{\theta}^{[m]}) = Q^{[m]}(\theta^{[m]})U^{[m]}(\vec{\theta}^{[m-1]})$

(b) Set $E^{[m]} = E^{[m-1]} - \Delta E_p^{[m]}$

(c) Set the values of the new set of variational parameters, $\vec{\theta}^{[m]} = \vec{\theta}^{[m-1]} \cup \{\theta_{p'}\}$, to $\vec{\theta}_{p'}^{[m]}$

5. (Optional) If the ground state of the system of interest is known, *a priori*, to have the same spin as $|\psi_0\rangle$, additionally append to the ansatz the spin-complement of $Q^{[m]}(\theta^{[m]})$, $\overline{Q}^{[m]}(\overline{\theta}^{[m]})$, unless $Q^{[m]}(\theta^{[m]}) \equiv \overline{Q}^{[m]}(\overline{\theta}^{[m]})$:

$$U^{[m]}(\vec{\theta}^{[m]}) = \overline{Q}^{[m]}(\overline{\theta}^{[m]})Q^{[m]}(\theta^{[m]})U^{[m-1]}(\vec{\theta}^{[m-1]}).$$

6. Enter the $(m+1)^{\text{th}}$ iteration by returning to step 1

The iterative loop of the QEB-ADAPT-VQE is visualized on the diagram in Fig. 4.1. The loop begins by preparing the trial state $|\psi^{[m-1]}\rangle$, obtained in the previous, $(m-1)^{\text{th}}$ iteration.

To identify a suitable qubit evolution to append to the ansatz, first the gradient of the energy expectation value, with respect to the variational parameter of each qubit evolution in $\mathbb{P}_{\text{full}}(Q, N_{so})$ is calculated in step 2. The gradients are evaluated at $\theta_p = 0$, because of the presumption that $|\psi_0\rangle$ is close to the true ground state, which suggests that the optimized value of θ_p is close to 0. As shown in Eq. (4.1), each of the gradients can be calculated by measuring, on a quantum computer, the expectation value of the commutator of H and the corresponding qubit excitation operator Ω_p, with respect to $|\psi^{[m-1]}\rangle$. Steps 1 and 2 are identical to the original fermionic-ADAPT-VQE.

The gradients calculated in step 2, roughly indicate how much each qubit excitation can decrease $E^{[m-1]}$. However, the largest gradient does not necessarily correspond to the largest energy reduction, optimized over all variational parameters. In step 3 the set of n_{qe} qubit evolutions, $\mathbb{Q}^{[m]}(n_{qe})$, with the largest energy gradient magnitudes is identified. The larger n_{qe} is, the more likely is $\mathbb{Q}^{[m]}(n_{qe})$ to contain the qubit evolution that reduces $E^{[m-1]}$ the most. For each qubit evolution $Q_p(\theta_p)$ in $\mathbb{Q}^{[m]}(n_{qe})$, the VQE is run with ansatz $Q_p(\theta_p)U^{[m-1]}(\vec{\theta}^{[m]})$ to calculate how much adding $Q_p(\theta_p)$ to the the ansatz reduces the energy estimate. Step 3 is not present in the original fermionic-ADAPT-VQE protocol [156], which directly grows its ansätze by the ansatz element with largest energy gradient. I introduce this step as additional means to reduce the ansatz circuit depth. The performance of the QEB-ADAPT-VQE for different values of n_{qe} is studied in Sec. 4.1.3.

Step 4 picks the qubit excitation, $Q^{[m]}\big(\theta^{[m]}\big)$, with the largest contribution to the energy reduction, $\Delta E^{[m]}$. If $|\Delta E^{(m)}|$ is below some real positive energy-reduction threshold ε, the iterative loop is exited. If instead $|\Delta E^{[m]}| > \varepsilon$, $Q^{[m]}\big(\theta^{[m]}\big)$ is appended to the ansatz.
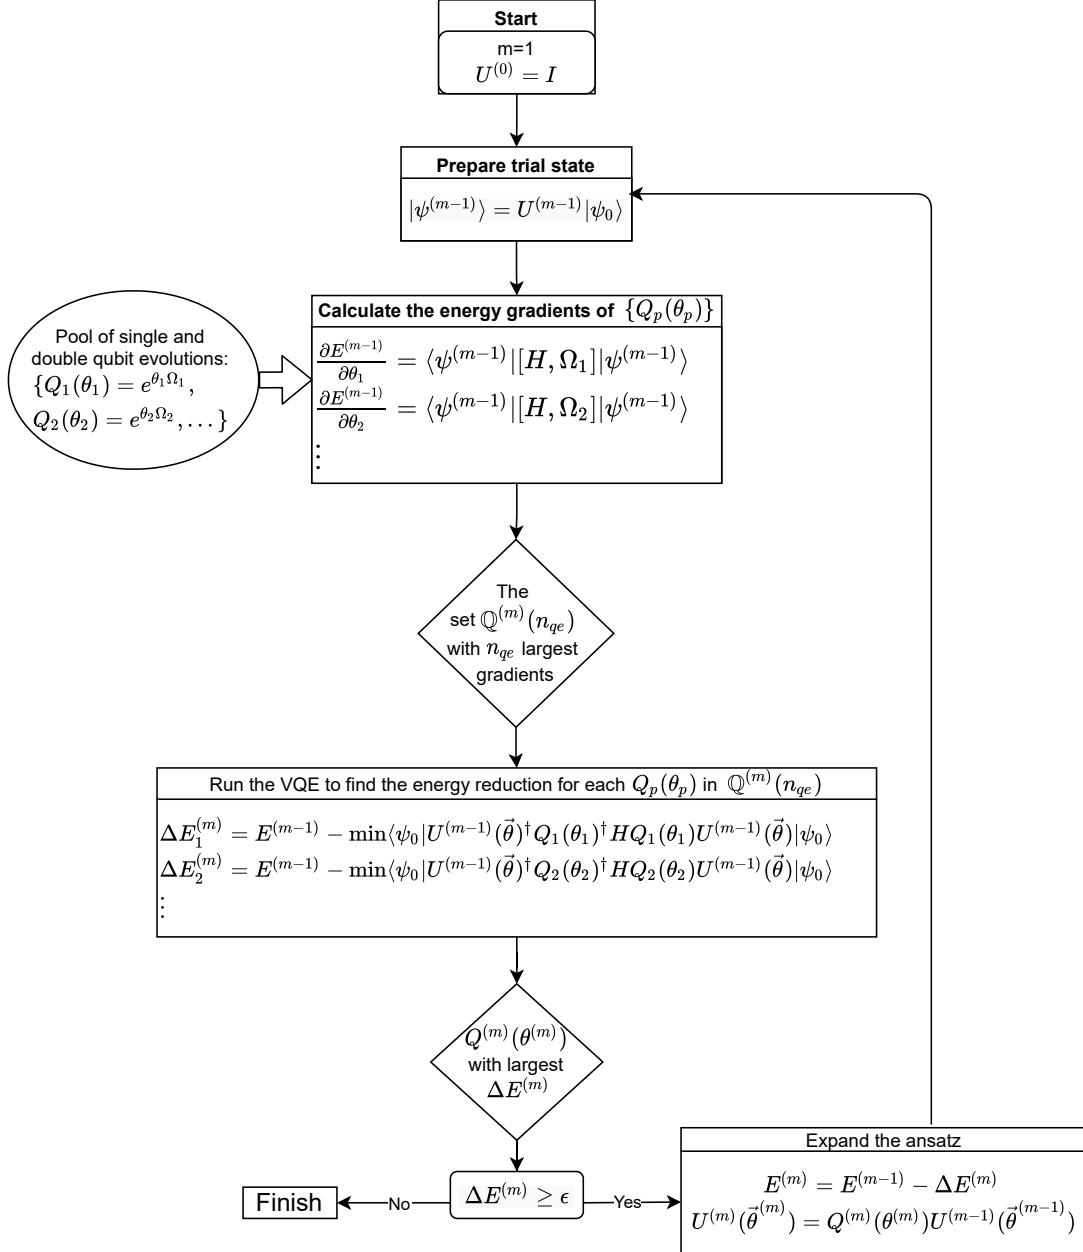

Fig. 4.1 Schematic diagram of the ansatz constructing loop of the QEB-ADAPT-VQE protocol (see Sec. 4.1.1). The optional step 5 is not depicted.

Lastly, if we know *a priori* the spin of the true ground state of the simulated electronic system, $|\psi_0\rangle$ can be chosen to have the same spin. Then, it can be assumed that qubit evolutions would come in spin-complement pairs, so that the ansatz remains spin-symmetric and conserves the spin of the state. Hence, step 5 appends the spin-complementary of $Q^{[m]}(\theta^{[m]})$, $\overline{Q}^{[m]}(\overline{\theta}^{[m]})$ to the ansatz. As demonstrated in Sec. 4.1.4, performing step 5 significantly reduces (up to 50%) the number of ansatz-constructing iterations. One might suggest that the variational parameters θ_p and $\overline{\theta_p}$ must be equal in magnitude, so that $Q^{[m]}(\theta^{[m]})$ and $\overline{Q}^{[m]}(\overline{\theta}^{[m]})$ contribute equally and the ansatz remains spin-symmetric. However, as shown in Sec. 4.1.4 in practice this is not the case. Setting $|\theta_p| = |\overline{\theta_p}|$ might reduce the number of variational parameters in some cases, but results in constructing a longer ansatz circuit.

4.1.2 Analysis of the computational complexity of the QEB-ADAPT-VQE

Here, I comment on the computational complexity of the QEB-ADAPT-VQE in terms of number of quantum computer measurements and run time. Both of these complexities are determined by steps 2 and 3.

Given that H is represented by up to $O(N_{so}^4)$ Pauli strings [Eq. (1.39)], calculating each gradient in step 2 would require $O(N_{so}^4)$ quantum computer measurements. Since $|\mathbb{P}(\tilde{A}, N_{so})| \propto N_{so}^4$, the complexity of step 2, in terms of quantum computer measurements is $O(N_{so}^8)$. Step 2 is completely parallelizable so if multiple quantum computers are available, its time complexity can be arbitrarily reduced down to the time required to evaluate the expectation value of a Pauli string term, which is proportional to the ansatz circuit depth, scaling as $O(m/N_{so})$ (a N_{so}-qubits circuit of $O(m)$ qubit evolutions).

Using a gradient-descent minimizer, like the BFGS [157], optimizing ansatz $U^{(m)}(\vec{\theta}^{[m]})$, which has $O(m)$ variational parameters (at the m^{th} ansatz-constructing iteration), would require $O(m^2)$ VQE iterations[3]. Therefore, each VQE run in step 3 would require $O(m^2 N_{so}^4)$ quantum measurements. Hence, the overall complexity of step 3 in terms of measurements would be $O(n_{qe} m^2 N_{so}^4)$. This complexity is a worst case estimate, assuming that at each ansatz-constructing iteration, all parameters $\vec{\theta}^{[m]}$ are initialized at zero. In fact though, we initiate $\vec{\theta}^{[m]}$ as $\vec{\theta}^{[m-1]} \cup 0$, so we will need fewer VQE-iterations to optimize the new ansatz, $U^{(m)}(\vec{\theta}^{[m]})$. However, the complexity also can be higher if we use a direct search minimizer, e.g. the Nelder-Mead [163], which is likely to be the case in practice, when noisy quantum

---

[3]Most gradient-descent optimizers require a number of iterations thaat is quadratic in the number of parameters [157, 162].

hardware is used. A detailed discussion of the various minimizers that can be used for the VQE, and their complexities is provided in Ref. [123]. Again, if multiple quantum devices are available, each of the n_{qe} VQE runs can be executed in parallel. Hence, the run time complexity of step 3 is lower bounded by the run time of a single VQE run. With the BFGS optimizer a single VQE run will involve $O(m^2)$ VQE-iterations, where the run time of one VQE-iteration will scale linearly with the depth of the ansatz circuit as $O(m/N_{so})$ ($O(m)$ quantum gates, assumed to be evenly distributed over the N_{so} qubits). So the time complexity of a single VQE run in step 3 would be $O(m^3/N_{so})$.

Let us denote as $|U(\vec{\theta})|$ the number of qubit evolutions required to construct an ansatz to achieve some desired accuracy. Then, overall the QEB-ADAPT-VQE would require $O\big(|U(\vec{\theta})|(N_{so}^8 + n_{qe}|U(\vec{\theta})|^2 N_{so}^4)\big)$ quantum computer measurements, and its run-time complexity would be lower bounded by $O(|U(\vec{\theta})|^4/N_{so})$. The size of the ansatz $|U(\vec{\theta})|$ depends on the desired accuracy, and also is problem specific. Therefore, it is difficult to predict how it would scale with N_{so}. For strongly correlated states, achieving chemical accuracy might require an ansatz that consists of as many as $O(N_{so}^4)$ qubit evolutions [4]. However, for weakly correlated states, the scaling of $|U(\vec{\theta})|$ with N_{so} is likely to be lower. Assuming the worst case scenario, the time complexity of the QEB-ADAPT-VQE will be lower-bounded by $O(N_{so}^{15})$ and it will require $O(n_{qe}N_{so}^{16})$ quantum computer measurements. For comparison, the UCCSD-VQE has a time complexity of $O(N_{so}^{11})$, assuming maximum parallelization, and requires $O(N_{so}^{12})$ quantum computer measurements. However, as we will see in Sec. 4.2.1, the ansätze constructed by the QEB-ADAPT-VQE are implemented by much shallower circuits than the UCCSD.

4.1.3 QEB-ADAPT-VQE dependence on n_{qe}

Here I investigate the performance of the QEB-ADAPT-VQE for different values of n_{qe}, the number of qubit evolutions considered in step 3. As we increase n_{qe}, we increase the chance to pick at each iteration the qubit evolution that, added to the ansatz, achieves the largest energy reduction. Following this greedy strategy is no guarantee for an optimal ansatz, since qubit evolutions do not commute in general. Nevertheless, we do expect, on average, to construct a more circuit-efficient ansatz by increasing n_{qe} up to some saturation value.

To test this presumption, and to pick a suitable value for n_{qe}, I perform classical numerical simulations to obtain energy convergence plots for the ground states of LiH, H_6 and BeH_2 in the STO-3G basis. The simulations for the three molecules are performed for bond distances

---

[4]I base this statement of the fact that the UCCSD, which consists of $O(N_{so}^4)$ fermionic evolutions, cannot achieve chemical accuracy in some cases.

$r_{Li-H} = 3\text{Å}$, $r_{H-H} = 3\text{Å}$ and $r_{Be-H} = 3\text{Å}$, away from equilibrium configurations[5], where correlation effects are stronger, and the effect of increasing n_{qe} should be more evident. The simulation results are presented in Fig. 4.2.
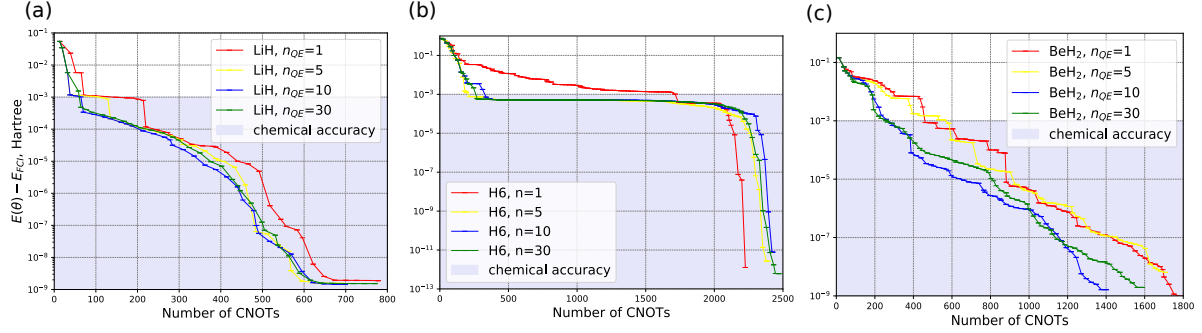

Fig. 4.2 Energy convergence plots obtained with the QEB-ADAPT-VQE with different values of n_{qe} for the ground states of LiH, H_6 and BeH_2 in the SO-3G basis, at bond distances of $r_{Li-H} = 3\text{Å}$, $r_{H-H} = 3\text{Å}$ and $r_{Be-H} = 3\text{Å}$. The QEB-ADAPT-VQE is terminated for $\varepsilon = 10^{-12}$ Hartree.

The table below summarizes the average (over number of qubit evolutions) *CNOT* count reductions, with respect to $n_{qe} = 1$, for each molecule and different values of n_{qe}:

	$n_{qe} = 5$	$n_{qe} = 10$	$n_{qe} = 30$
LiH	16%	20%	16%
BeH_2	3%	26%	22%
H_6	15%	11%	13%

Table 4.1 Average *CNOT* count reduction for QEB-ADAPT-VQE($n_{qe} > 1$) as compared to QEB-ADAPT-VQE($n_{qe} = 1$).

For LiH (Fig. 4.2.a) the QEB-ADAPT-VQE clearly constructs ansatz circuits with fewer *CNOT*s as n_{qe} is increased above 1. For BeH_2 (Fig. 4.2.c) a significant *CNOT* count reduction is obtained for $n_{qe} = 10$ and $n_{qe} = 30$, but not for $n_{qe} = 5$. For H_6 (Fig. 4.2.b) the average *CNOT* reduction is about the same for $n_{qe} = 5$, $n_{qe} = 10$ and $n_{qe} = 30$, but strangely the ansatz constructed by the QEB-ADAPT-VQE for $n_{qe} = 1$ is the most *CNOT*-efficient for accuracies higher than 10^{-4} Hartree. Also, for all three molecules, we observe no further *CNOT* reduction for $n_{qe} = 30$ as compared to $n_{qe} = 10$. Actually for $n_{qe} = 30$ the *CNOT* reduction is a bit lower. As noted above these inconsistencies can be explained by the fact that

---

[5]The equilibrium configurations of LiH and BeH_2 correspond to $r_{Li-H} = 1.546\text{Å}$ and $r_{Be-H} = 1.316\text{Å}$, respectively.

the greedy strategy to obtain the lowest estimate for $E(\vec{\theta})$ at each iteration is no guarantee for constructing an optimal ansatz, because qubit evolutions do not commute.

Nevertheless, there is a clear advantage in terms of *CNOT* count, in performing step 3 of the QEB-ADAPT-VQE for $n_{qe} > 1$. Despite the associated overhead in the number of quantum computer measurements with increasing n_{qe}, this is justified as long as the bottleneck of NISQ computers is the quantum gate fidelity. Furthermore, we can expect the *CNOT* count reduction for $n_{qe} > 1$ to increase for larger molecules, because the QEB-ADAPT-VQE will have to consider a larger pool of qubit evolutions. The results in Table 4.1 also indicate that for the three considered molecules, 10 is a good value for n_{qe}, around which the *CNOT* reduction saturates.

4.1.4   Appending spin-complement qubit evolutions

Here I compare the performance of the QEB-ADAPT-VQE when the optional step 5 is (1) not performed, (2) performed as described in Sec. 4.1.1, assigning a variational parameter to each qubit evolution, so that $\theta^{(m)}$ and $\overline{\theta}^{(m)}$ are independent, and (3) performed, by assigning one variational parameter to a spin-complement pair of qubit evolutions, so that $|\theta^{(m)}| = |\overline{\theta}^{(m)}|$.

The comparison is made with energy convergence plots for the ground states of LiH and BeH$_2$ in the STO-3G basis, at equilibrium bond distances of $r_{Li-H} = 1.546$Å and $r_{Be-H} = 1.316$Å, respectively. Both molecules have zero spin ground states, so we expect performing step 5 to significantly reduce the number of ansatz-constructing-iterations. The energy convergence plots are presented in Fig. 4.3.

From the plots in Figs. 4.3.a and 4.3.b we see that the QEB-ADAPT-VQE constructs ansätze with approximately the same *CNOT* counts whether step 5 is performed or not (blue and green plots, respectively). The reason for this is that even without enforcing appending of spin-complement evolutions, the QEB-ADAPT-VQE is likely to construct an ansatz where spin-complement qubit evolutions are appended at adjacent or nearby positions in the ansatz, since they equally reduce $E(\vec{\theta})$. This can be observed in the ansätze constructed by the QEB-ADAPT-VQE for LiH and BeH$_2$ when step 5 is not performed, which are included up to the 32$^{\text{th}}$ ansatz element in Tables D.1 and D.3 in Appendix D.1. Hence, as demonstrated by the blue and green plots in Figs. 4.3.c and 4.3.d, by performing step 5 and "guessing" that the next appended qubit evolution will be the spin-complementary of the one appended in step 4, the number of ansatz-constructing iterations is reduced by up to a factor of 2.

On the other hand, if we assign one parameter for each spin-complement pair of qubit evolutions (the red plots in Fig. 4.3), the number of variational parameters is reduced in the case of LiH (Fig. 4.3.e). However, no parameter reduction is achieved in the case of BeH$_2$ (Fig. 4.3.f). Moreover, as seen in Figs. 4.3.a and 4.3.b, the constructed ansätze, on average,
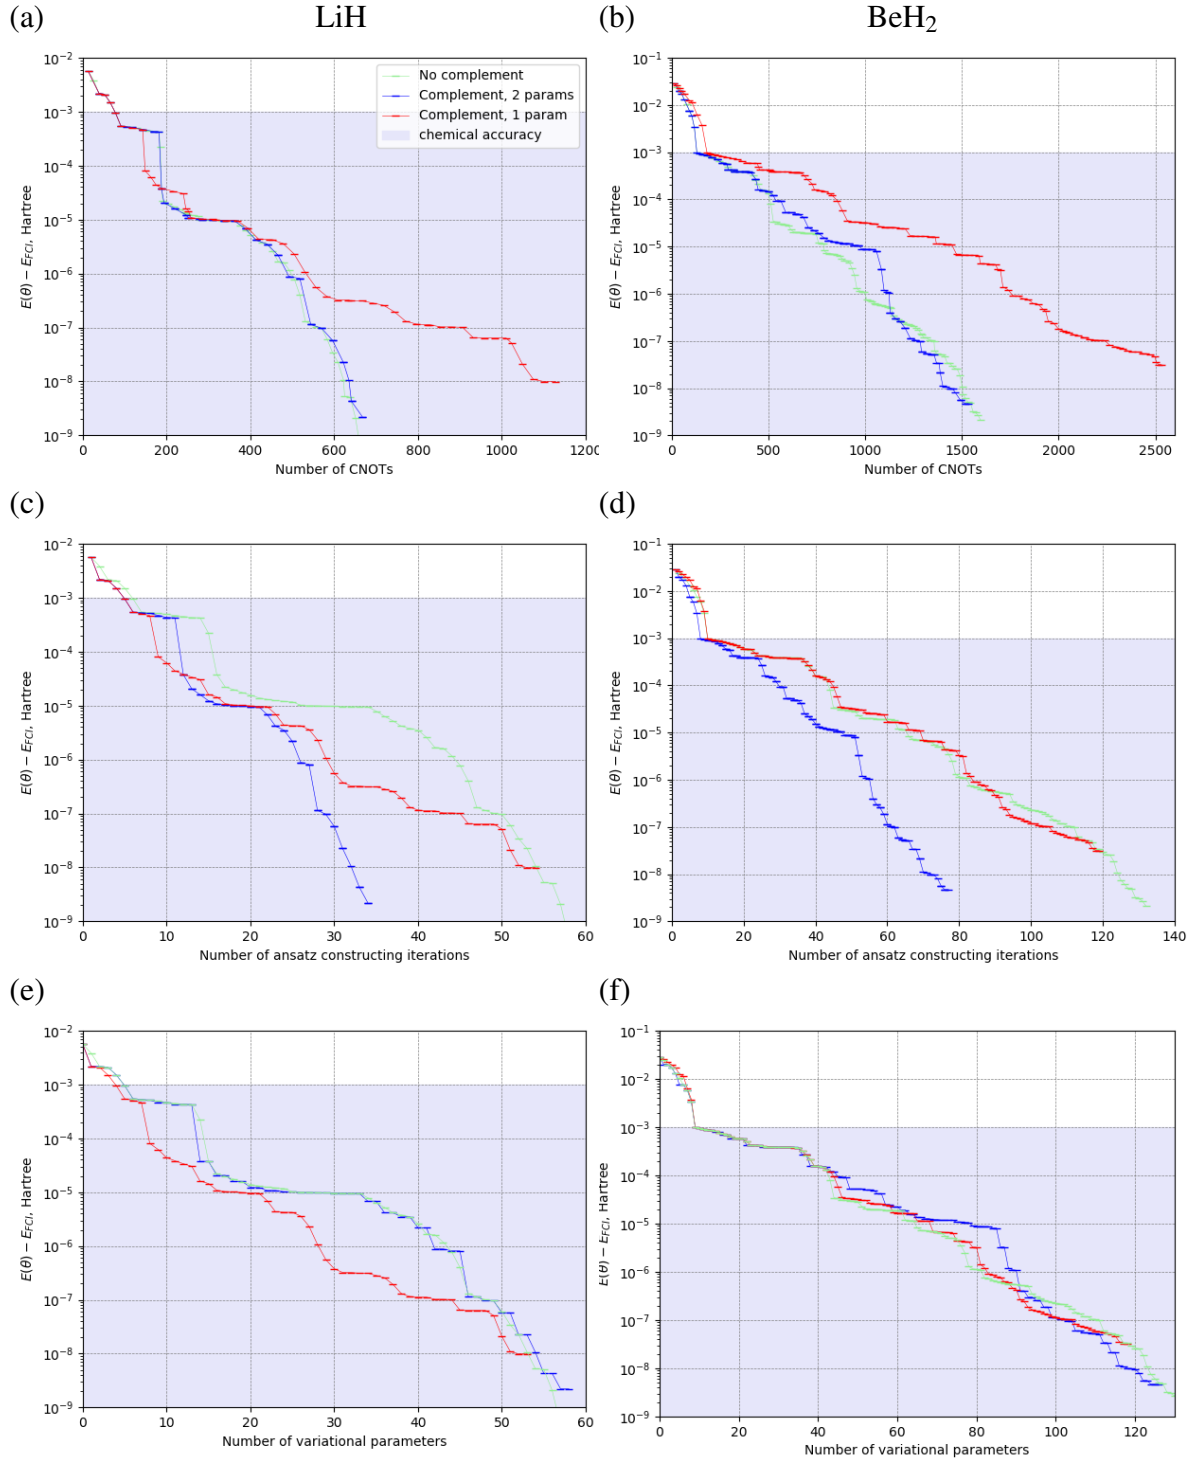
Fig. 4.3 Energy convergence plots for the ground states of LiH and BeH$_2$ in the STO-3G basis, at equilibrium bond distances of $r_{Li-H} = 1.546$Å and $r_{Be-H} = 1.316$Å, respectively. **Green** plots are obtained with the QEB-ADAPT-VQE when step 5 is not performed. **Blue** plots are obtained with the QEB-ADAPT-VQE, when step 5 is performed and each qubit evolution is assigned an independent parameter. **Red** plots are obtained with the QEB-ADAPT-VQE, when step 5 is performed and one parameter is assigned per spin-complement pair of qubit evolutions. For all plots $n_{qe} = 1$, and $\varepsilon = 10^{-12}$ Hartree. **a,b** Accuracy as function of *CNOT* count. **c,d** Accuracy as function of number of ansatz-constructing iterations. **e,f** Accuracy as function of number of variational parameters.

require significantly more *CNOT*s for both molecules, especially as the circuit sizes increase. This approach does not work well, because of two reasons. First, a pair of spin-complement qubit evolutions would have parameters equal in magnitude only if the two qubit evolutions commute. This issue can be solved by checking if the pair of qubit evolutions commute. Second, however even if we know that the two spin-complement qubit evolutions commute and they can share parameters equal in magnitude, the correct relative sign between the parameters would depend on the parity of the qubit state, which qubit evolutions do not account for.

Unlike qubit evolutions, fermionic evolutions account for the parity of the qubit state and can efficiently form pairs of spin-complement fermionic evolutions. Spin-complement pairs of fermionic evolutions are utilized as ansatz elements by the fermionic-ADAPT-VQE, thus as we will see in Sec. 4.2.2 it requires up to half as many parameters than the QEB-ADAPT-VQE.

4.2 Benchmarking the QEB-ADAPT-VQE

In this section, I benchmark the QEB-ADAPT-VQE against other standard methods to obtain estimates for ground state energies. In Sec. 4.2.1 I compare the QEB-ADAPT-VQE against the UCCSD-VQE, by obtaining energy dissociation plots. In Sec. 4.2.2 I compare the QEB-ADAPT-VQE against the fermionic-ADAPT-VQE [141] and the qubit-ADAPT-VQE [156], by obtaining energy convergence plots. All simulations in this section are performed for LiH, H_6 and BeH_2 in the STO-3g orbital basis set, with no frozen orbitals assumed. For all VQE optimizations, the BFGS optimizer is used. All *CNOT* counts are obtained assuming the use of the *CNOT*-efficient circuits derived in Chapter 2.

4.2.1 Dissociation curves

Figure 4.4 shows energy dissociation curves for LiH, H_6 and BeH_2, obtained with the QEB-ADAPT-VQE for $n_{qe} = 10$ and energy-reduction thresholds $\varepsilon_4 = 10^{-4}$ Hartree, $\varepsilon_6 = 10^{-6}$ Hartree and $\varepsilon_8 = 10^{-8}$ Hartree. Dissociation curves obtained with the Hartree-Fock method, the FCI method, and the VQE, using an untrotterized UCCSD ansatz (UCCSD-VQE) are also included for comparison. Unlike, in Sec. 3.2.1 the UCCSD here only includes spin-conserving single and double fermionic evolutions[6] for a fairer comparison to the QEB-ADAPT-VQE.

---

[6]The number of included single-fermionic excitations is $\frac{1}{2}(N_{so} - N_e)N_e$, and the number of double-fermionic excitation is $\frac{1}{8}(N_{so} - N_e)\left(\frac{N_{so} - N_e}{2} - 1\right)N_e\left(\frac{N_e}{2} - 1\right) + \frac{1}{16}(N_{so} - N_e)^2 N_e^2$.

Figures 4.4a,b,c show the absolute values for the ground-state energy estimates. All methods except the HF, produce similar energy estimates that cannot be clearly distinguished. In Figs. 4.4d,e,f the exact FCI energy is subtracted in order to differentiate better the different methods and their corresponding errors.
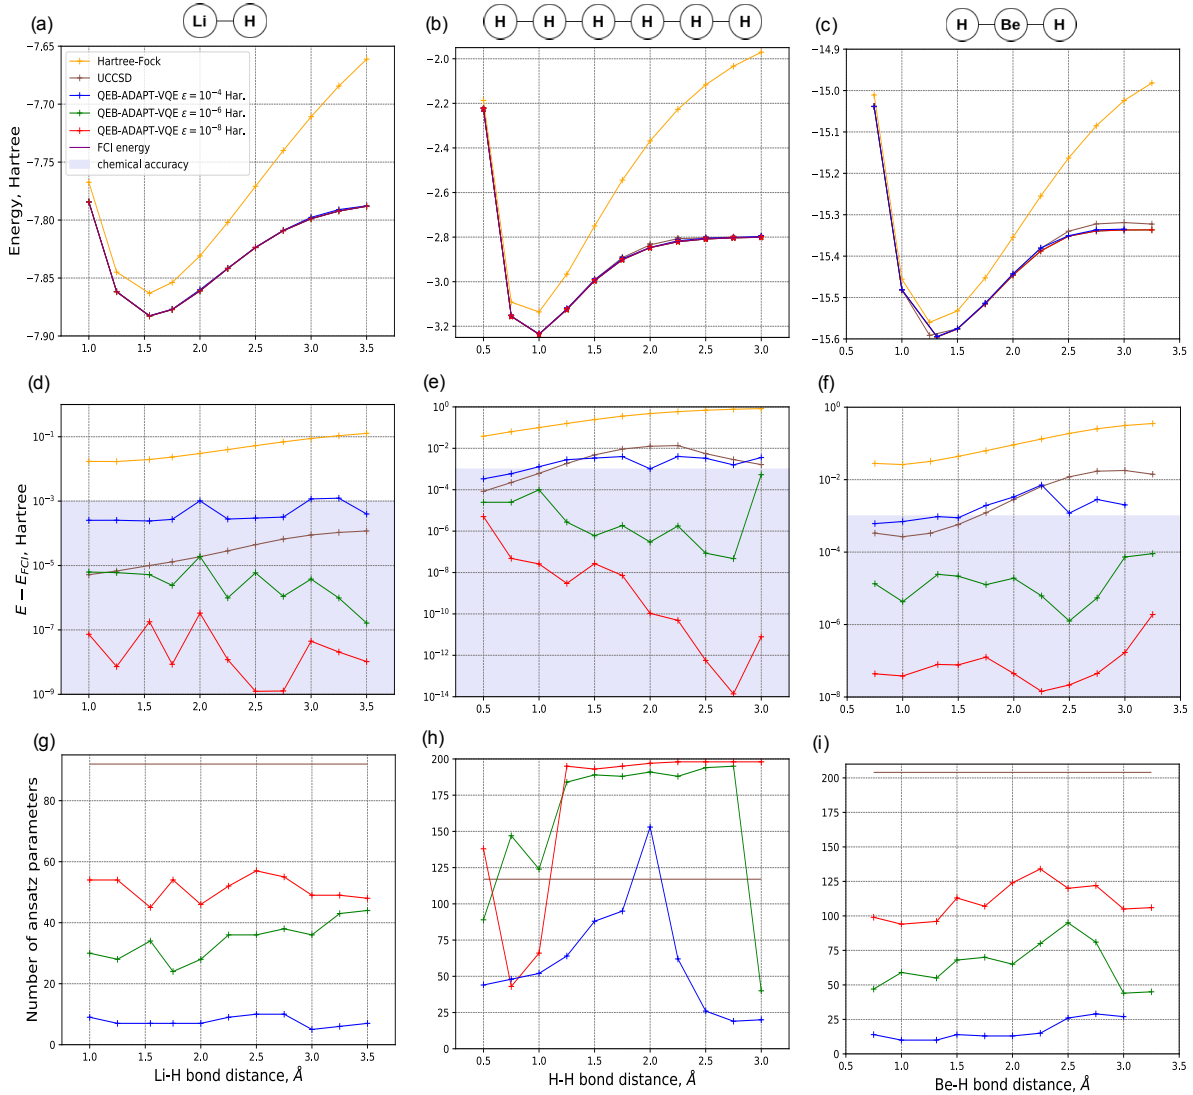
Fig. 4.4 Energy dissociation curves for LiH, H_6 and BeH_2 in the STO-3G spin-orbital basis, obtained with the QEB-ADAPT-VQE, the UCCSD-VQE and the Hartree-Fock. **a,b,c** Estimated ground state energy. **d,e,f** Error of the ground-state-energy estimate with respect to the exact FCI energy. **g,h,i** Number of ansatz parameters. The UCCSD ansätze considered here contain only spin-conserving fermionic evolutions: for LiH, H_6 and BeH_2 the UCCSD has 92, 117 and 204 parameters/fermionic evolutions, respectively.

As HF does not account for static and dynamic correlations, it is unable to achieve chemical accuracy at any bond-distance for any of the molecules.

The UCCSD-VQE achieves chemical accuracy over all bond distances for LiH (Fig. 4.4d) and over bond distances close to equilibrium configuration for H_6(Fig. 4.4e) and BeH_2(Fig. 4.4f). However, as expected the UCCSD-VQE fails to achieve chemical accuracy for bond distances away from the equilibrium configurations of H_6 and BeH_2, where the ground states become more strongly correlated.

The QEB-ADAPT-VQE for ε_4, similarly to the UCCSD-VQE, struggles to achieve chemical accuracy for strongly correlated ground states. However, for ε_6 and ε_8 the QEB-ADAPT-VQE achieves chemical accuracy over all investigated bond distances for all three molecules. This indicates that the QEB-ADAPT-VQE can successfully construct ansätze to accurately approximate strongly correlated states.

However, the real strength of the QEB-ADAPT-VQE, similarly to other ADAPT-VQE protocols, is not just in constructing accurate ansätze, but in constructing accurate problem-tailored ansätze with few variational parameters, and corresponding shallow ansatz circuits. Figures 4.4g.h.i show plots of the number of variational parameters used by the ansatz of each method as function of bond distance. In the cases of LiH (Fig. 4.4g) and BeH_2 (Fig. 4.4i), the ansätze constructed by the QEB-ADAPT-VQE for ε_6 and ε_8 are not only more accurate than the UCCSD, but also have significantly fewer parameters. However, in the case of H_6 the QEB-ADAPT-VQE on average requires more parameters than the UCCSD. The reason for this is that H_6 is more strongly correlated than LiH and BeH_2. Thus, even an optimally constructed ansatz requires more variational parameters than the UCCSD, to accurately approximate the ground state of H_6.

An interesting point to note are the abrupt changes in the number of variational parameters used by the QEB-ADAPT-VQE for H_6 at bond distances of around 1 Å, 2 Å, and 2.75 Å. The likely reason for these changes are molecular structure transformations, such as isomerizations, where different eigenstates of H become lowest in energy (energy level crossings).

4.2.2 Comparison to the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE

In this section, I compare the QEB-ADAPT-VQE against the fermionic-ADAPT-VQE [141] and the qubit-ADAPT-VQE [156] with energy convergence plots for LiH, H_6 and BeH_2.

All convergence plots are terminated for $\varepsilon = 10^{-12}$ Hartree[7]. For a fair comparison the QEB-ADAPT-VQE is performed for $n_{qe} = 1$. Before I present the energy convergence plots, I briefly describe the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE protocols.

The fermionic-ADAPT-VQE and the qubit-ADAPT-VQE

The fermionic-ADAPT-VQE is the original ADAPT-VQE protocol that, as already mentioned, grows a problem tailored UCC ansatz by iteratively appending spin-complement pairs of single and double fermionic evolutions. Unlike the QEB-ADAPT-VQE, the fermionic-ADAPT-VQE assigns one parameter per spin-complement pair. This is viable because fermionic evolutions account for the parity of the qubit state. In this way, the correct relative sign between the variational parameters of the two spin-complement fermionic evolutions is "automatically" obtained. Let us denote a spin-complement pair of a single (double) fermionic evolution as $F'_{ki}(\theta_{ki})$ $\left(F'_{lkji}(\theta_{lkji})\right)$. Then,

$$F'_{ki}(\theta_{ki}) = \overline{F}_{ki}(\theta_{ki})F_{ki}(\theta_{ki}) = e^{\theta_{ki}\overline{\Xi}_{ki}}e^{\theta_{ki}\Xi_{ki}} \quad \text{and} \tag{4.2}$$

$$F'_{lkji}(\theta_{lkji}) = \overline{F}_{lkji}(\theta_{lkji})F_{lkji}(\theta_{lkji}) = e^{\theta_{lkji}\overline{\Xi}_{lkji}}e^{\theta_{lkji}\Xi_{lkji}}, \tag{4.3}$$

where the overline denotes the spin-complement of a fermionic evolution, or a fermionic excitation operator. In the original definition of the fermionic-ADAPT-VQE in Ref. [141], the protocol utilizes an ansatz element pool defined by all unique spin-complement pairs of single and double fermionic evolutions, $\{F'_{ki}(\theta_{ki})\}$ and $\{F'_{lkji}(\theta_{lkji})\}$, respectively, that act on N_{so} spin-orbitals. The ansatz-growing strategy of the ADAPT-VQE, is to append at each iteration the ansatz element with largest energy gradient magnitude. This is equivalent to the ansatz-growing strategy of the QEB-ADAPT-VQE for $n_{qe} = 1$. For the simulations presented in this thesis, I implement the fermionic-ADAPT-VQE exactly as in its original paper [141].

The qubit-ADAPT-VQE is similar to the fermionic-ADAPT-VQE, except for its ansatz element pool. The ansatz element pool of the qubit-ADAPT-VQE consists of unitary evolutions of individual Pauli strings (I will refer to these as Pauli string evolutions). In the original paper of the qubit-ADAPT-VQE [156], the authors suggested several different pools of Pauli string excitations. For the simulations presented in this section, I use a pool that includes all unique XY-Pauli strings that have lengths of 2 or 4 and have an odd number of Ys. This pool consists of $O(N_{so}^4)$ Pauli string evolutions, which can be combined to obtain all qubit evolutions in the ansatz element of the QEB-ADAPT-VQE (see Chapter 2).

---

[7]In the original papers of the fermionic-ADAPT-VQE [141] and the qubit-ADAPT-VQE [156] an energy-gradient threshold is used for termination instead.

Hence, the comparison between the QEB-ADAPT-VQE and qubit-ADAPT-VQE, in terms of ansatz-circuit efficiency will be fair.

I note that the authors of Ref. [156] actually proved that the qubit-ADAPT-VQE can construct an ansatz that exactly approximates an electronic wavefunction, using a reduced ansatz element pool of only $2N_{so} - 2$ Pauli string evolutions. This reduced pool can decrease the number of quantum computer measurements required to evaluate the energy gradients at each ansatz-constructing iteration from $O(N_{so}^8)$ to $O(N_{so}^5)$. However, this reduced ansatz element pool will also result in a slower and less circuit-efficient ansatz construction, so a comparison to the QEB-ADAPT-VQE would not be fair.

Results

The energy convergence plots for LiH, H_6 and BeH_2, obtained with the QEB-ADAPT-VQE, the ADAPT-VQE and qubit-ADAPT-VQE are presented in Fig. 4.5. LiH and BeH_2 are simulated at equilibrium bond distances of $r_{Li-H} = 1.546$Å and $r_{Be-H} = 1.316$Å, and H_6 at bond distance $r_{H-H} = 1.5$Å. Figure D.1 in Appendix D includes similar plots for LiH, H_6 and BeH_2 at bond distances of $r_{Li-H} = 3$Å, $r_{H-H} = 3$Å and $r_{Be-H} = 3$Å instead.

The plots compare the three protocols in terms of three cost metrics, required to achieve a specific accuracy: (1) the number of ansatz-constructing iterations; (2) the number of variational parameters; and (3) the *CNOT* count of the ansatz circuit. The number of ansatz-constructing iterations and the number of variational parameters[8] determine the total number of required quantum computer measurements and the total run time (see the analysis of the complexity of the QEB-ADAPT-VQE in Sec. 4.1.2, which is similar for the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE). The *CNOT* count of the ansatz circuit is approximately proportional to its depth. Hence, the *CNOT* count can be used as a measure of the run time of the quantum subroutine of the VQE, which also reflects the resulting quantum error. Since the *CNOT*s are the current bottleneck of NISQ computers, I consider the *CNOT* count as a primary cost metric.

First, we observe, in Figs. 4.5a,b,c that the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE perform very similarly in terms of number of ansatz-constructing iterations, especially in the cases of LiH and BeH_2. This indicates that the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE use approximately the same number of qubit and fermionic evolutions, respectively, in constructing their ansätze. The level of this similarity is remarkable, but not unexpected, since as we saw in Sec. 3.2.1 UCCSD and Q-UCCSD ansätze perform

---

[8] For the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE methods, the number of iterations is the same as the number of variational parameters. In the case of the QEB-ADAPT-VQE, one or two single-parameter qubit evolutions are added per iteration. Thus, the number of iterations is less than the number of variational parameters for the QEB-ADAPT-VQE.
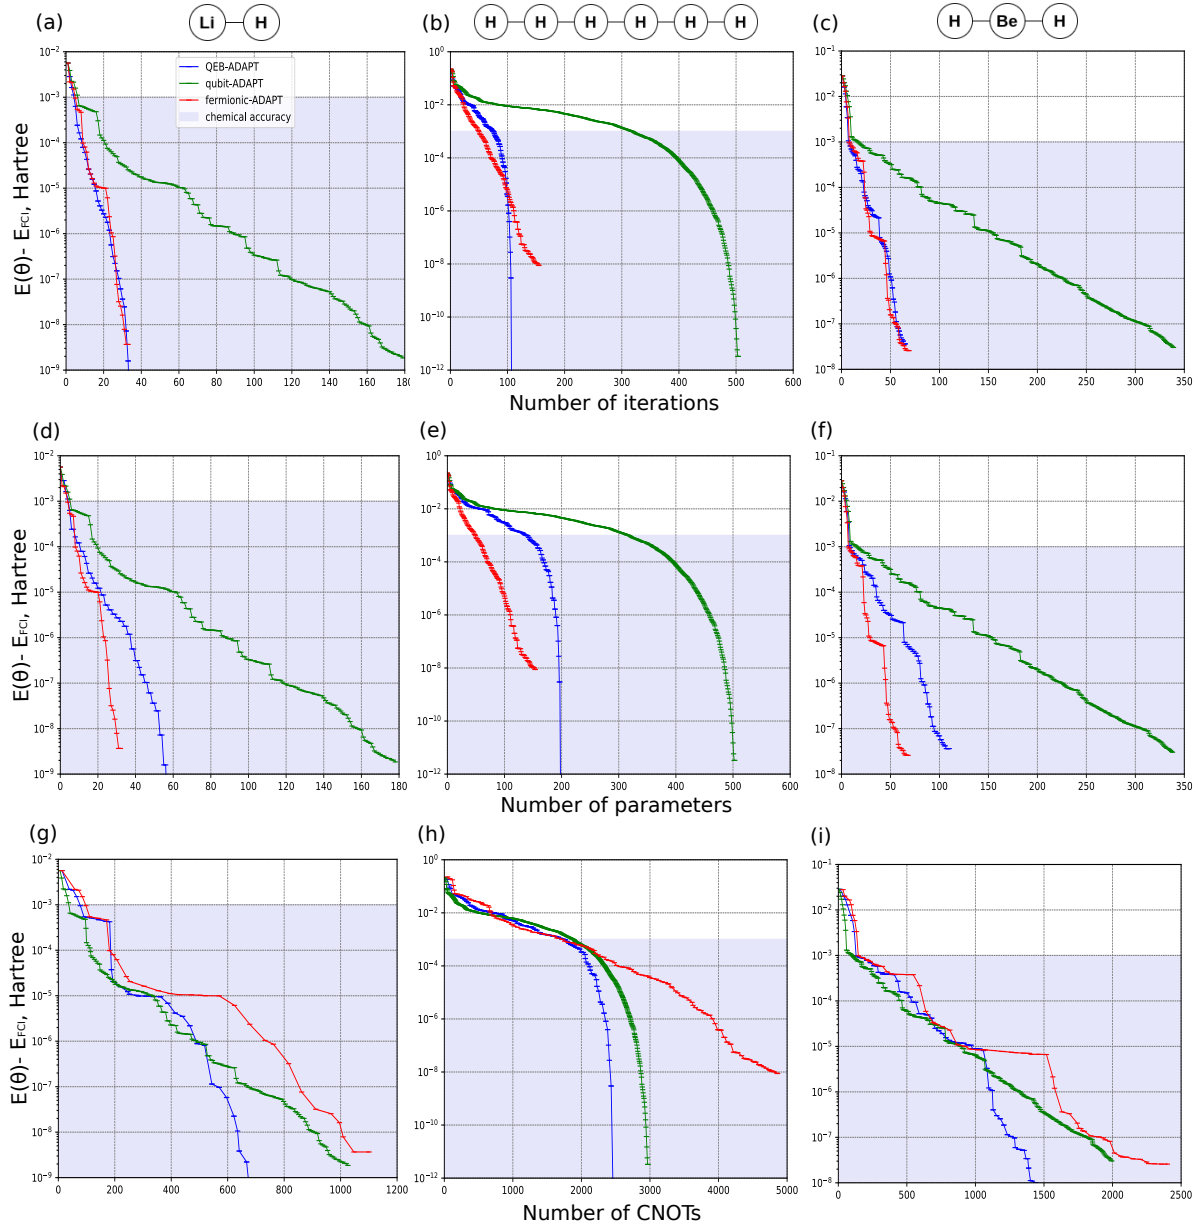
Fig. 4.5 Energy convergence plots for the ground states of LiH, H_6 and BeH_2 in the STO-3G orbital basis set, at bond distances $r_{\text{Li-H}} = 1.546\text{Å}$, $r_{H-H} = 1.5\text{Å}$ and $r_{\text{Be-H}} = 1.316\text{Å}$. **Blue** plots are obtained with the QEB-ADAPT-VQE for $n_{qe} = 1$. **Red** plots are obtained with the fermionic-ADAPT-VQE. **Green** plots are obtained with the qubit-ADAPT-VQE. All convergence plots are terminated for energy-reduction threshold of $\varepsilon = 10^{-12}$ Hartree. **a,b,c** Accuracy as function of ansatz-constructing iterations. **d,e,f** Accuracy as function of ansatz parameters (the number of parameters is the same as the number of ansatz-constructing iterations for the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE). **g,h,i** Accuracy as function of the *CNOT* count of the ansatz circuit. The *CNOT* counts are calculated assuming the use of the circuits derived in Chapter 2.

similarly in approximating electronic wavefunctions. Moreover, the QEB-ADAPT-VQE systematically outperforms the fermionic-ADAPT-VQE in terms of $CNOT$ count (Figs. 4.5g,h,i), because qubit evolutions are implemented by simpler circuits than fermionic evolutions.

While the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE require similar numbers of iterations, the QEB-ADAPT-VQE requires up to twice as many variational parameters (Fig. 4.5d,e,f). This difference is due to the fact that the QEB-ADAPT-VQE assigns one parameter to each qubit evolution in its ansatz, whereas the fermionic-ADAPT-VQE assigns one parameter to a spin-complement pair of fermionic evolutions.

Figures 4.5a,b,c,d,e,f show that the QEB-ADAPT-VQE converges faster, requiring systematically fewer ansatz-constructing iterations and variational parameters than the qubit-ADAPT-VQE. The reason for this is that the qubit-ADAPT-VQE utilizes more rudimentary ansatz elements. As seen from Eqs. (2.20) and (2.21) single and double qubit evolutions correspond to products of 2 and 8, mutually commuting evolutions of individual Pauli strings, respectively. Hence, using qubit evolutions the QEB-ADAPT-VQE is able to construct ansätze more rapidly than the qubit-ADAPT-VQE.

In terms of $CNOT$ count (Figs. 4.5g,h,i), the qubit-ADAPT-VQE is more efficient than the QEB-ADAPT-VQE at low accuracies. However, at higher accuracies, and corresponding larger ansätze, the QEB-VQE-ADAPT starts to systematically outperform the qubit-ADAPT-VQE in terms of $CNOT$-efficiency. This surprising result can be attributed to the fact that qubit evolutions allow for the local circuit optimizations introduced in Chapter 2, whereas Pauli string evolutions, albeit more variationally flexible, do not allow for any local circuit optimizations[9].

Lastly, I note that the observations made here are supported by the additional convergence plots in Fig. D.1 in Appendix D.

4.3   Comparison of iteratively constructed UCC and Q-UCC ansätze

In Sec. 3.2.1 I presented simulation results that compared the UCCSD and the Q-UCCSD ansätze. Generally, the results indicated that the two ansätze can approximate electronic wavefunctions comparably well. However, the results also indicated that the UCCSD might be more accurate for strongly correlated states.

Here I further investigate the capabilities of qubit and fermionic evolutions to construct molecular electronic ansätze, by comparing iteratively-constructed UCC and Q-UCC ansätze.

---

[9]A Pauli string evolution is implemented by a single $CNOT$-staircase construction (as shown in Appendix B), whose $CNOT$ count cannot be reduced.

The comparison between the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE in the previous Sec. 4.2.2 is not a fair comparison between the two types of evolutions: the QEB-ADAPT-VQE assigns one variational parameter per a qubit evolution in its ansatz, whereas the fermionic-ADAPT-VQE assigns one variational parameter per a spin-complement pair of fermionic evolutions. In this section, I compare the QEB-ADAPT-VQE for $n_{qe} = 1$ and step 5 not implemented, to the fermionic-ADAPT-VQE when it grows its ansatz by appending individual fermionic evolutions (instead of complement-pairs of fermionic evolutions). In this way the two protocols differ only in using a pool of qubit evolutions, and a pool fermionic evolutions, respectively. The ansätze constructed by the two ADAPT-VQE protocols contain only qubit/fermionic evolutions that contribute to the approximation of the electronic wavefunction. Hence, comparing these iteratively-constructed UCC and Q-UCC ansätze is a more comprehensive comparison of the two types of evolutions than comparing the fixed UCCSD and Q-UCCSD ansätze, which contain a redundant number of qubit and fermionic evolutions, respectively.

Figure 4.6 shows energy convergence plots for the ground states of LiH, H_6 and BeH_2 at bond distances of $r_{\text{Li-H}} = 1.546$Å, $r_{\text{H-H}} = 1.5$Å and $r_{\text{Be-H}} = 1.316$Å, respectively. The plots correspond to UCC and Q-UCC ansätze constructed by the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE as explained above, for energy-reduction threshold of $\varepsilon = 10^{-12}$ Hartree. The plots for the two types of ansätze are similar for each molecule. However, there is a clear tendency the UCC to be slightly more accurate per number of ansatz elements. This advantage is most evident for the more strongly correlated H_6 (Fig. 4.6.b), where the UCC requires up to 20% fewer evolutions than the Q-UCC to achieve a given accuracy. These observations further indicate that UCC ansätze can approximate strongly correlated states a bit better than Q-UCC ansätze.
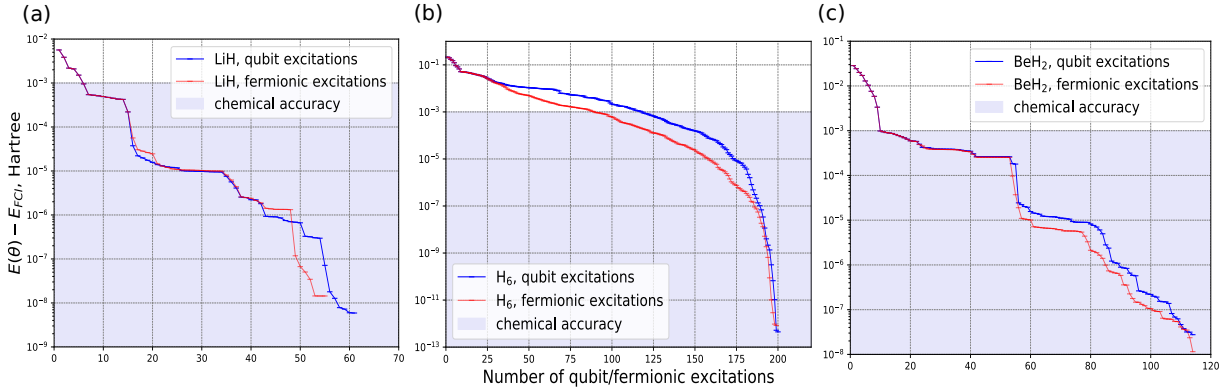
Fig. 4.6 Energy convergence plots for the ground states of LiH, H_6 and BeH_2 in the STO-3G basis at bond distances of $r_{\text{Li-H}} = 1.546\text{Å}$, $r_{\text{H-H}} = 1.5\text{Å}$ and $r_{\text{Be-H}} = 1.316\text{Å}$, respectively. **Blue** plots are obtained with the QEB-ADAPT-VQE when step 5 is not implemented and $n_{qe} = 1$. **Red** plots are obtained with the fermionic-ADAPT-VQE, when using ansatz element pool of independent single and double fermionic evolutions. All convergence plots are terminated for $\varepsilon = 10^{-12}$ Hartree.

To further test this observation, in Fig. 4.7 I include energy convergence plots, similar to those in Fig. 4.6, but for bond distances of $r_{\text{Li-H}} = 3\text{Å}$, $r_{\text{H-H}} = 3\text{Å}$ and $r_{\text{Be-H}} = 3\text{Å}$. At these larger bond distances the ground states of the LiH, and BeH_2 are more strongly correlated. So we expect to see larger difference in the accuracies achieved by the UCC and the Q-UCC ansätze per number of fermionic and qubit evolutions, respectively.
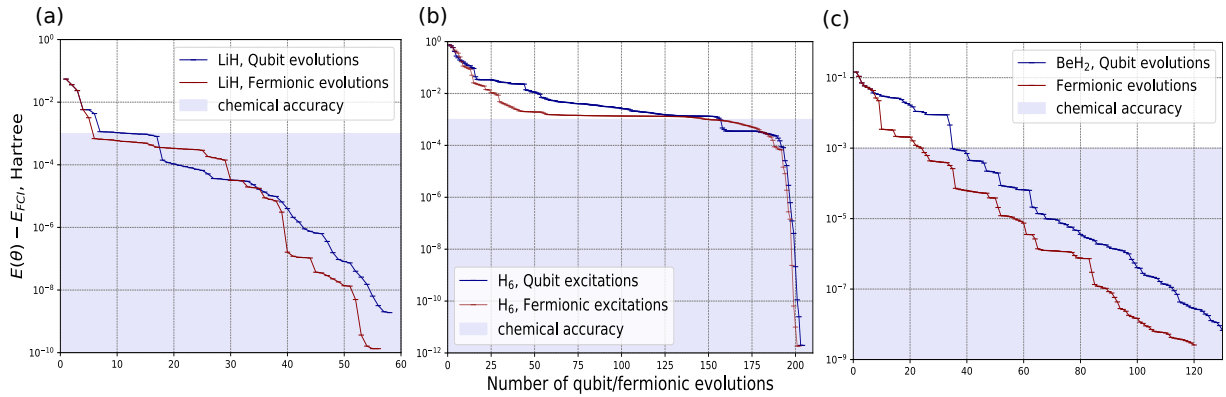
Fig. 4.7 Energy convergence plots, similar to those in Fig. 4.6, for the ground states of LiH, H_6 and BeH_2 in the STO-3G basis at bond distances of $r_{\text{Li-H}} = 3\text{Å}$, $r_{\text{H-H}} = 3\text{Å}$ and $r_{\text{Be-H}} = 3\text{Å}$, respectively.

In the case of LiH for $r_{\text{Li-H}} = 3\text{Å}$ (Fig. 4.7a) the difference between the convergence plots for the Q-UCC and the UCC is similar to that for $r_{\text{Li-H}} = 1.546\text{Å}$ (Fig. 4.6c). However, in the

case of BeH$_2$ for $r_{\text{Be-H}} = 3$Å (Fig. 4.7c) the difference between the convergence plots for the Q-UCC and the UCC is noticeably larger than for $r_{\text{Be-H}} = 1.316$Å (Fig. 4.6c), with the UCC requiring on average about 20% fewer ansatz elements to achieve a given accuracy, than the Q-UCC. These results further indicate that UCC ansätze are more accurate per number ansatz elements when approximating strongly correlated states.

As a side point it is interesting to note that when the fermionic-ADAPT-VQE is implemented with a pool of independent single and double fermionic evolutions (Figs. 4.6 and 4.7) it is able to converge, albeit more slowly, to higher final accuracy than when it is implemented with a pool of spin-complement pairs of single and double fermionic evolutions (Fig. 4.5). This is due to the fact that the pool of independent fermionic excitation is more variationally flexible.

4.4   Conclusion

In this chapter, I utilized qubit evolutions to introduce the qubit-excitation based adaptive variational quantum eigensolver (QEB-ADAPT-VQE). The QEB-ADAPT-VQE simulates molecular electronic wavefunctions with a problem-tailored ansatz, grown iteratively by appending single and double qubit evolutions. I benchmarked the performance of the QEB-ADAPT-VQE with classical numerical simulations for LiH, H$_6$ and BeH$_2$.

First, I compared the QEB-ADAPT-VQE against the UCCSD-VQE and demonstrated that the ansätze constructed by the QEB-ADAPT-VQE can achieve much higher accuracies than a fixed ansatz such as the UCCSD, and at the same time require multiple times fewer variational parameters and correspondingly shallower circuits.

Next, I compared the QEB-ADAPT-VQE to the original fermionic-ADAPT-VQE, and its more circuit-efficient cousin, the qubit-ADAPT-VQE. Compared to the fermionic-ADAPT-VQE, the QEB-ADAPT-VQE requires up to twice as many variational parameters. However, the QEB-ADAPT-VQE requires asymptotically fewer *CNOT*s, owing to its use of qubit evolutions that are implemented by simpler circuits than the fermionic evolutions used by the fermionic-ADAPT-VQE.

The simulations also showed that the qubit-ADAPT-VQE is more *CNOT*-efficient than the QEB-ADAPT-VQE in achieving low accuracies that correspond to small ansatz circuits. However, for higher accuracies and correspondingly larger ansatz circuits, the QEB-ADAPT-VQE systematically outperformed the qubit-ADAPT-VQE in terms of *CNOT*-efficiency. The primary reason for this is that qubit evolutions allow for the local circuit optimizations derived in Chapter 2, whilst the more rudimentary Pauli string evolutions, utilized by the qubit-ADAPT-VQE, do not. However, in practice we are only interest in reaching chemical

accuracy. Therefore, one might question what is the usefulness of constructing more *CNOT*-efficient ansätze with the QEB-ADAPT-VQE for accuracies higher than chemical accuracy. Although the numerical results presented here are not sufficient to draw a general conclusion, they indicate that the *CNOT*-efficiency of the QEB-ADAPT-VQE becomes more evident for larger ansatz circuits. Therefore, for larger molecules, the QEB-ADAPT-VQE will likely be able to reach chemical accuracy using fewer *CNOT*s than the qubit-ADAPT-VQE. Additionally, as demonstrated in Sec. 4.1.3, the QEB-ADAPT-VQE can construct even more circuit-efficient ansätze by increasing the parameter n_{qe}, at the expense of more quantum computer measurements. The simulation results also demonstrated that in terms of convergence speed, the QEB-ADAPT-VQE requires fewer variational parameters, and correspondingly fewer ansatz-constructing iterations, than the qubit-ADAPT-VQE.

These results imply that the QEB-ADAPT-VQE is more circuit-efficient and converges faster than the qubit-ADAPT-VQE, which to my knowledge was the previously most circuit-efficient, scalable VQE protocol for molecular modelling. I do remark though, that in my comparison of the QEB-ADAPT-VQE and the qubit-ADAPT-VQE, I ignored the fact that the latter protocol can use a reduced ansatz element pool of $O(N_{so})$ Pauli string evolutions. Using such a reduced pool would decrease the number of quantum computer measurements required by the qubit-ADAPT-VQE to measure the energy gradients at each ansatz-constructing iteration from $O(N_{so}^8)$ to $O(N_{so}^5)$. However, using a reduced pool would also result in a slower and less efficient ansatz construction. Moreover, as discussed in Sec. 4.1.2, the complexity of a single ansatz-constructing iteration, for both the QEB-VQE-ADAPT and the qubit-ADAPT-VQE, might turn out to be dominated by the complexity of the VQE run, which on average can be as large as $O(N_{so}^{12})$.

I also note that in theory, hardware-efficient ansätze and the ansätze of the IQCC protocol [159, 155] might be implemented by shallower circuits than the ansätze constructed by the QEB-ADAPT-VQE. However, hardware-efficient ansätze and the IQCC are unlikely to be scalable for large systems: as noted in Sec. 1.5.5, the optimization of hardware-efficient ansätze is likely to become intractable for large systems; and the IQCC requires evaluating a number of expectation values, exponential in the number of variational parameters.

In addition to outlining and benchmarking the QEB-ADAPT-VQE, in this Chapter I also further compared qubit and fermionic evolutions in constructing molecular electronic ansätze. I used the QEB-ADAPT-VQE and a slightly modified version of the fermionic-ADAPT-VQE to compare iteratively-constructed UCC and Q-UCC ansätze. The results indicated that the Q-UCC and the UCC ansätze can achieve similar accuracies per number of ansatz elements when approximating weakly correlated states. However, it was also observed a tendency the UCC ansätze to be more accurate in approximating strongly correlated states. In the

cases of H_6 and BeH_2 the UCC ansätze required up to 20% fewer ansatz elements than the Q-UCC ansätze to achieve a given accuracy. It was also found that qubit evolutions cannot efficiently form single-parameter spin-complement pairs, whereas fermionic evolutions can. This feature of fermionic evolutions originates from the fact that they account for the qubit state parity. Grouping fermionic evolutions in spin-complement pairs, as utilized by the fermionic-ADAPT-VQE, can be used to nearly halve the number of variational parameters when the ansatz must respect spin conservation. However, in the era of NISQ computers, when the ansatz-circuit depth, and in particular the *CNOT* depth, is a primary cost factor, these advantages of UCC ansätze are off-set by the circuit-efficiency of Q-UCC ansätze.

So far, I have investigated VQE protocols to find molecular ground state energies. In the next Chapter, I consider excited states.

Chapter 5

Estimating excited state energies

Chapter Summary

Finding the energies of excited states is important for calculating molecular spectral properties, e.g. photodissociation rates [164] and absorption bands [165]. In Chapter 4, I introduced the QEB-ADAPT-VQE protocol for molecular modelling, and benchmarked its performance in estimating molecular ground state energies. In this chapter, I propose and benchmark a modified version of the QEB-ADAPT-VQE protocol designed to estimate the energies of excited states. I call this modified version the "excited-QEB-ADAPT-VQE", or just the e-QEB-ADAPT-VQE for shorthand.

Molecular ground states near equilibrium bond configuration typically have a single dominant Slater determinant, which corresponds to the Hartree-Fock state. Thus, when dealing with ground states we can straightforwardly choose the Hartree-Fock state as an initial reference state. However, excited states are generally more statically correlated than ground states, and might have two or more dominant Slater determinants. In addition, often excited states have non-zero spin, which is difficult to deduce *a priori*. Hence, choosing *a priori* an initial reference state that has a significant overlap with an unknown excited state is more difficult than for a ground state. The e-QEB-ADAPT-VQE features a modified ansatz-growing strategy designed to be less dependent on the choice of an initial reference state.

The Chapter is organised as follows: In Sec. 5.1 I describe how the energies of excited states can be calculated with the VQE. In Sec. 5.2 I outline the e-QEB-ADAPT-VQE, and in Sec. 5.3, I benchmark its performance by estimating excited state energies for LiH and BeH_2. Lastly in Sec. 5.4 I compare iteratively constructed Q-UCC and UCC ansätz for excited states.

5.1   Finding excited state energies with the VQE

Two methods to calculate excited state energies with the VQE are the quantum subspace expansion [166–168], and the overlap-based method [169, 170]. Here I use the more straightforward overlap-based method.

As the name suggests the overlap-based method works by including the overlaps between previously found, lower energy, eigenstates of the electronic Hamiltonian H and the state we currently want to approximate in the cost function of the VQE. When the cost function is being minimized, the additional overlap terms force the newly found state to be orthogonal to the previously found eigenstates. In this way, we obtain a new eigenstate. For example, after we find the ground state $|E_0\rangle$, we modify H as

$$H \to H_1 = H + \alpha_0 |E_0\rangle\langle E_0|, \tag{5.1}$$

where α_0 is a real positive scalar coefficient. If $\alpha_0 + E_0 > E_1$, where E_1 is the energy of the first excited state of H, the lowest eigenvalue of the modified Hamiltonian H_1 will be shifted from E_0 to E_1. In order to make sure that $\alpha_0 + E_0 > E_1$ is satisfied, prior to knowing E_1, we can choose an arbitrary value for α_0, which is large compared to the energy scale of the problem[1]. Running the VQE for H_1 will output an estimate for E_1, and the optimized ansatz will generate an approximation for $|E_1\rangle$:

$$E_1 = \min_{\vec{\theta}} \langle \psi_0 | U^\dagger(\vec{\theta}) H_1 U(\vec{\theta}) | \psi_0 \rangle = \min_{\vec{\theta}} \langle \psi_0 | U^\dagger(\vec{\theta}) \big(H + \alpha_0 |E_0\rangle\langle E_0| \big) U(\vec{\theta}) | \psi_0 \rangle. \tag{5.2}$$

After E_1 is found the same procedure can be repeated multiple times to find the next eigenvalues of H. The k^{th} eigenvalue of H, E_k, can be estimated as

$$E_k = \min_{\vec{\theta}} \langle \psi_0 | U^\dagger(\vec{\theta}) H_k U(\vec{\theta}) | \psi_0 \rangle = \min_{\vec{\theta}} \langle \psi_0 | U^\dagger(\vec{\theta}) \big(H + \sum_{r=0}^{k-1} \alpha_r |E_r\rangle\langle E_r| \big) U(\vec{\theta}) | \psi_0 \rangle. \tag{5.3}$$

The first Hamiltonian term in Eq. (5.3) is calculated as usual by evaluating the expectation values of the Pauli string terms in the quantum-gate operator expression for H [Eq. (1.39)]. Each overlap term in Eq. (5.3), $\langle E_r | U(\vec{\theta}) | \psi_0 \rangle$, can be evaluated: (1) either directly , which would require a circuit with total depth equal to the depth of $U(\vec{\theta})$ plus the depth of the ansatz generating $|E_r\rangle$; or (2) with the SWAP test [171], which would require twice as many, $2N_{so}$, qubits, but a circuit with the same depth as the circuit for $U(\vec{\theta})$. Given that the bottleneck

---

[1] For the simulation results presented in the thesis I use $\alpha_0 = 2|E_{HF}|$.

of NISQ computers is the circuit depth, using the SWAP test would be the better choice in practice.

I note that the overlap-based method is efficient and reliable for low-lying excited states. However, due to its recursive nature and build-up of errors in the approximations of the eigenstates and eigenenergies of H, the method is not suitable for high excited states. In such cases the quantum subspace expansion method is more suitable. Alternatively one can also use the method of *classical shadows* [172] to obtain an estimate for the overlap with a reduced number of quantum computer measurements.

5.2 The e-QEB-ADAPT-VQE

5.2.1 Protocol description

In this section, I describe the e-QEB-ADAPT-VQE protocol. The preparation part of the e-QEB-ADAPT-VQE is similar to that of the QEB-ADAPT-VQE. The difference is that now we need to add the additional overlap terms to H, as described in the previous section, to obtain H_k, a Hamiltonian to find the k$^{\text{th}}$ eigenvalue of H.

As I mentioned in the chapter summary, the e-QEB-ADAPT-VQE is designed to be less dependent on the initial reference state, $|\psi_0\rangle$. The only condition for $|\psi_0\rangle$ is to have Hamming weight equal to the number of electrons, N_e. Therefore, I arbitrarily choose $|\psi_0\rangle$ to be the Hartree Fock state again.

We begin constructing the ansatz by setting the iteration number to $m = 1$, the initial ansatz to the identity $U(\vec{\theta}) \rightarrow U^{[0]} = I$, and then initiate the ansatz-constructing loop of the e-QEB-ADAPT-VQE. Below, I describe the five steps of the m^{th} iteration. Afterwards I explain how and why the e-QEB-ADAPT-VQE is changed with respect to the QEB-ADAPT-VQE.

1. Prepare the state $|\psi^{[m-1]}\rangle = U^{[m-1]}\big(\vec{\theta}^{[m-1]}\big)|\psi_0\rangle$, using the ansatz $U^{[m-1]}\big(\vec{\theta}^{[m-1]}\big)$ with values for the variational parameters $\vec{\theta}^{[m-1]}$ as determined in the previous iteration.

2. For each qubit evolution $Q_p(\theta_p) = e^{\theta_p \Omega_p} \in \mathbb{P}_{\text{full}}(Q, N_{so})$:

   (a) Run a single-parameter VQE optimization to find

   $$\min_{\theta_p} E^{[m]}(\theta_p) = \min_{\theta_p}\langle\psi^{[m-1]}|Q_p^\dagger(\theta_p)H_k Q_p(\theta_p)|\psi^{[m-1]}\rangle.$$

   (b) Calculate the energy reduction $E^{[m-1]} - \min_{\theta_p} E^{[m]}(\theta_p)$

3. Identify the set of n_{qe} qubit evolutions, $\mathbb{Q}^{[m]}(n_{qe})$, with largest energy reductions, measured in the previous step. For $Q_p(\theta_p) \in \mathbb{Q}^{[m]}(n_{qe})$:

   (a) Run the VQE to find

   $$\min_{\vec{\theta}^{[m-1]},\theta_p} E(\vec{\theta}^{[m-1]},\theta_p) = \min_{\vec{\theta}^{[m-1]},\theta_p} \langle \psi_0 | U^{[m-1]^\dagger}(\vec{\theta}^{[m-1]}) Q_p^\dagger(\theta_p) H Q_p(\theta_p) U^{[m-1]}(\vec{\theta}^{[m-1]}) | \psi_0 \rangle.$$

   (b) Find the energy reduction $\Delta E_p^{[m]} = E^{[m-1]} - \min_{\vec{\theta}^{[m-1]},\theta_p} E(\vec{\theta}^{[m-1]},\theta_p)$

   (c) Save the (re)optimized values of $\vec{\theta}^{[m-1]} \cup \{\theta_p\}$ as $\vec{\theta}_p^{[m]}$

4. Identify the largest energy reduction $\Delta E^{[m]} \equiv \Delta E_{p'}^{[m]} = \max(\{\Delta E_p^{[m]}\})$, and its corresponding qubit evolution $Q^{(m)}(\theta^{[m]}) \equiv Q_{p'}(\theta_{p'})$.

   If $\Delta E^{[m]} < \varepsilon$, where $\varepsilon > 0$ is an energy-reduction threshold:

   (a) Exit

   Else:

   (a) Append $Q^{[m]}(\theta^{[m]})$ to the ansatz: $U^{[m]}(\vec{\theta}^{[m]}) = Q^{[m]}(\theta^{[m]}) U^{[m]}(\vec{\theta}^{[m-1]})$

   (b) Set $E^{[m]} = E^{[m-1]} - \Delta E_p^{[m]}$

   (c) Set the values of the new set of variational parameters, $\vec{\theta}^{[m]} = \vec{\theta}^{[m-1]} \cup \{\theta_{p'}\}$, to $\vec{\theta}_{p'}^{[m]}$

5. Enter the $(m+1)^{\text{th}}$ iteration by returning to step 1

The major difference between the QEB-ADAPT-VQE and the e-QEB-ADAPT-VQE is step 2. In step 2 of the e-QEB-ADAPT-VQE the individual energy reduction contribution of each qubit evolution $Q_p(\theta_p) \in \mathbb{P}_{\text{full}}(Q, N_{so})$ is calculated. Each of these energy reductions is calculated by a single-parameter VQE optimization performed to minimize the energy expectation value $\langle \psi^{[m-1]} | Q_p^\dagger(\theta_p) H Q_p(\theta_p) | \psi^{[m-1]} \rangle$. In this way, we get an indication about how much the energy expectation value will be reduced when each $Q_p(\theta_p)$ is added to the ansatz and the energy is minimized along the full set of parameters $\vec{\theta}^{[m-1]} \cup \{\theta_p\}$, at a cost of performing $O(N_{MO}^4)$ single-parameter VQE optimizations.

In the QEB-ADAPT-VQE (and the fermionic-ADAPT-VQE [141]) such an indication is obtained by measuring the energy gradients $\left\{ \frac{\partial}{\partial \theta_p} \langle \psi^{[m-1]} | Q_p^\dagger(\theta_p) H Q_p(\theta_p) | \psi^{[m-1]} \rangle \right\}$ (see Sec. 4.1). However, these energy gradients must be evaluated for some values of the respective parameters $\{\theta_p\}$. If we are approximating the ground state $|E_0\rangle$ and we have a

high level of confidence that $|\psi_0\rangle$ has a large overlap with $|E_0\rangle$, then the gradients can be evaluated conveniently for $\{\theta_p = 0\}$. But if we are approximating the excited state $|E_k\rangle$, where we do not know how close $|\psi_0\rangle$ and $|E_k\rangle$ are, measuring the energy gradients for $\{\theta_p = 0\}$, or any other arbitrary value, can be a poor indicator, leading to an inefficient and slow construction of the ansatz, and a possibility of finding a local energy minimum (see Sec. 5.2.3). Instead the e-QEB-ADAPT-VQE measures the individual energy reduction contribution for each qubit evolution, which can be obtained by the VQE, using a direct search minimizer [173], e.g. the Nelder Mead [163], irrespectively of the overlap of $|\psi_0\rangle$ and $|E_k\rangle$.

If we denote by N_H the number of terms in the Pauli string representation of H [Eq. (1.39)], calculating a single energy gradient requires measuring $2N_H$ expectation values. On the other hand, running a single-parameter VQE requires measuring γN_H expectation values, where γ is the number of function evaluations required for a single parameter minimization. Hence, the additional cost of the technique pursued here in comparison to using energy gradients as in the QEB-ADAPT-VQE, is a factor of $\gamma/2$ more quantum computer measurements in step 2. In practice, the value of γ would depend on the used optimizer, the desired accuracy and the noise levels. However, γ will not scale with the ansatz size. Also, once the ansatz reaches some critical size, so that $|\psi^{[m-1]}\rangle$ is close to $|E_k\rangle$, we can switch to the cheaper QEB-ADAPT-VQE.

In addition to step 2, the e-QEB-ADAPT-VQE differs from the QEB-ADAPT-VQE in not appending spin-complement qubit evolutions. Appending spin-complement qubit evolutions is advantageous when we know that the ansatz should respect spin-conservation. However, often excited states have non-zero spin that is difficult to deduce *a priori*.

5.2.2   The e-QEB-ADAPT-VQE dependence on n_{qe}

Here I compare the performance of the e-QEB-ADAPT-VQE for different values of n_{qe}, similarly as for the QEB-ADAPT-VQE in Sec. 4.1.3. Figure 5.1 shows energy convergence plots for the first excited states of LiH and BeH_2 in the STO-3G basis at bond distances of $r_{\text{Li-H}} = 1.546\text{Å}$ and $r_{\text{Be-H}} = 1.316\text{Å}$, respectively. For each molecule, plots are obtained with the e-QEB-ADAPT-VQE for $n_{qe} = 1$, $n_{qe} = 5$, $n_{qe} = 10$ and $n_{qe} = 20$.

Interestingly, the *CNOT* reduction for $n_{qe} > 1$ is much smaller than in the case of the QEB-ADAPT-ADAPT (see Sec. 4.1.3). At low accuracies ($>10^{-5}$ Hartree), for each molecule, the convergence plots are almost identical. Only at higher accuracies we can observe any significant *CNOT* reduction for $n_{qe} > 1$. The maximum *CNOT* reduction in the case of LiH is about 15% and in the case of BeH_2, about 10%. These results suggest that the individual energy reductions calculated in step 2 of the e-QEB-ADAPT-VQE, are a better indicator

about which ansatz element to add to the ansatz, than the energy-gradients calculated in step 2 of the QEB-ADAPT-VQE. Therefore, implementing step 3 of the e-QEB-ADAPT-VQE for $n_{qe} > 1$ is less necessary.

In Fig. 5.1 we also see that the plots for $n_{qe} = 5$, $n_{qe} = 10$ and $n_{qe} = 20$ are similar at all levels of accuracy. These results suggest that the saturation value of n_{qe} for the e-QEB-ADAPT-VQE will be lower than for the QEB-ADAPT-VQE. Therefore, the e-QEB-ADAPT-VQE might be suitable for ground state simulations as well, because the additional computational overhead of step 2, will be partially offset by using a smaller value for n_{qe}.
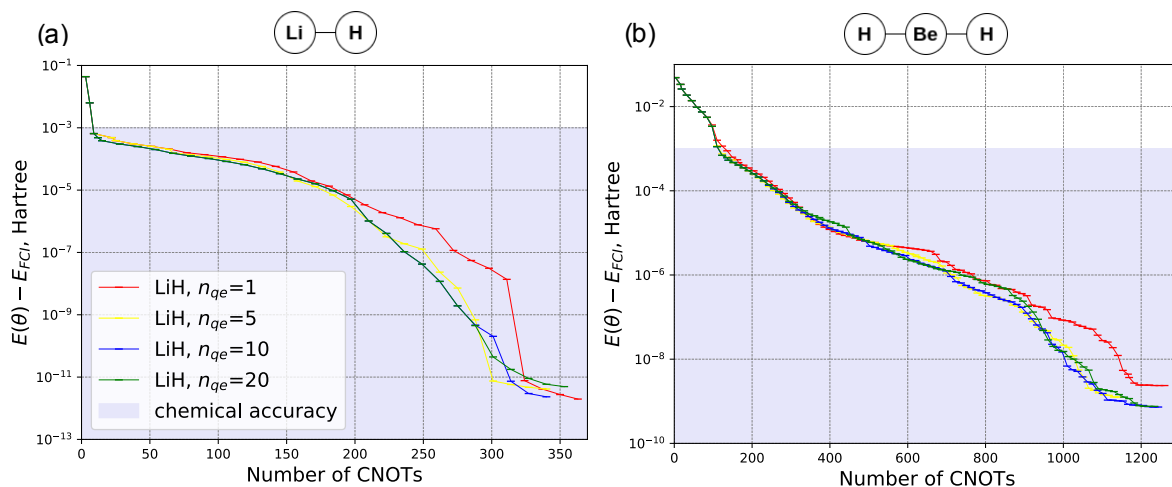

Fig. 5.1 Energy convergence plots for the first excited states of LiH and BeH_2 in the STO-3G basis at bond distances of $r_{Li\text{-}H} = 1.546$Å and $r_{Be\text{-}H} = 1.316$Å, respectively. The plots are obtained with the e-QEB-ADAPT-VQE for different values of n_{qe}. All convergence plots are terminated for $\varepsilon = 10^{-12}$ Hartree.

5.2.3   The QEB-ADAPT-VQE vs the e-QEB-ADAPT-VQE

Here I compare the performance of the e-QEB-ADAPT-VQE and the QEB-ADAPT-VQE in finding excited state energies. Figure 5.2 presents energy convergence plots for the first excited states of LiH and BeH_2 in the STO-3G basis at equilibrium bond distances of $r_{Li\text{-}H} = 1.546$Å and $r_{Be\text{-}H} = 1.316$Å, respectively.

In the case of LiH the QEB-ADAPT-VQE converges slower, requiring more than twice as many ansatz-constructing iterations, than the e-QEB-ADAPT-VQE. Since each ansatz-constructing iteration corresponds to a qubit evolution in the ansatz, the ansatz constructed by the QEB-ADAPT-VQE is roughly twice as large as the one constructed by the excited-QEB-ADAPT-VQE. As noted above, the reason for this is that energy-gradient-based ansatz-growing strategy is not suitable when there is no guarantee that the initial reference state $|\psi_0\rangle$

has a significant overlap with the target state $|E_k\rangle$. In the case of BeH$_2$, the QEB-ADAPT-VQE completely fails to converge to chemical accuracy, likely getting stuck in a local energy minimum. On the other hand, the e-QEB-ADAPT-VQE, being independent on the overlap of $|\psi_0\rangle$ and $|E_k\rangle$, converges successfully in the cases of both molecules.
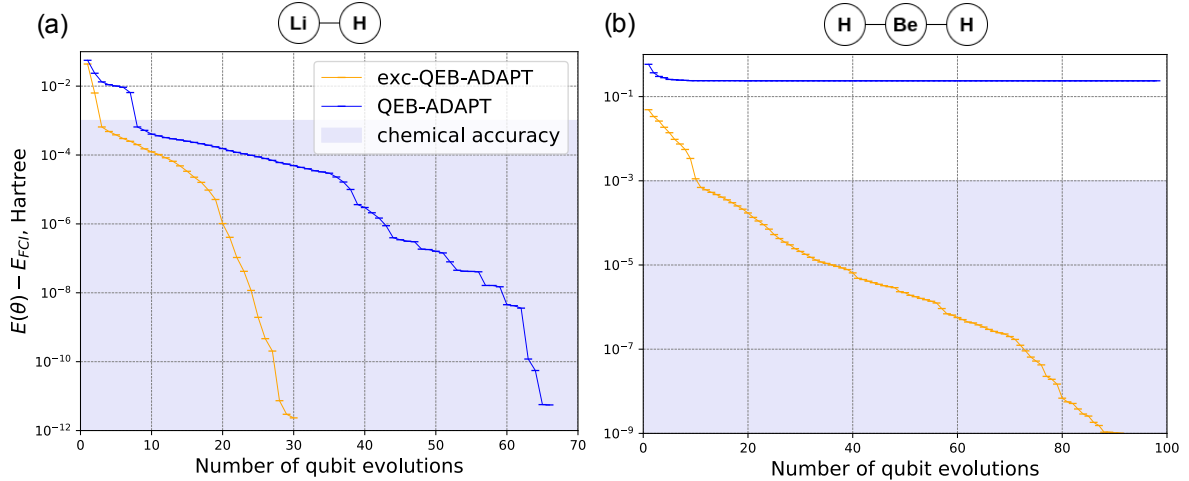

Fig. 5.2 Energy convergence plots for the first excited states of LiH and BeH$_2$ in the STO-3G basis at equilibrium bond distances of $r_{\text{Li-H}} = 1.546$Å and $r_{\text{Be-H}} = 1.316$Å, respectively. The plots are obtained with the QEB-ADAPT-VQE for $n_{qe} = 10$ and step 5 not implemented, and the e-QEB-ADAPT-VQE for $n_{qe} = 10$. The plots are terminated for energy-reduction threshold of $\varepsilon = 10^{-12}$ Hartree.

5.3 Benchmarking the e-QEB-ADAPT-VQE

5.3.1 First excited states

Here I benchmark the performance of the e-QEB-ADAPT-VQE, for $\varepsilon = 10^{-6}$ Hartree and $\varepsilon = 10^{-8}$ Hartree, by obtaining energy dissociation plots for the first excited states of LiH and BeH$_2$ in the STO-3G basis. I also include energy dissociation plots obtained with UCCSD-VQE and the GUCCSD-VQE, where the GUCCSD is the generalized UCCSD ansatz, which includes all unique single and double fermionic evolutions.

The energy dissociation curves for the two molecules are given in Figs. 5.3a and 5.3b. In Figs. 5.3c and 5.3d are plotted the errors of each method with respect to the FCI energy, as a function of bond distance. In the case of LiH, all methods, surprisingly even the UCCSD-VQE, achieve chemical accuracy. Figures. 5.3e and 5.3f show the number of variational parameters, which is also equivalent to the number of qubit/fermionic evolutions, used by the ansatz of each method. The ansätze constructed by the e-QEB-ADAPT-VQE for LiH are
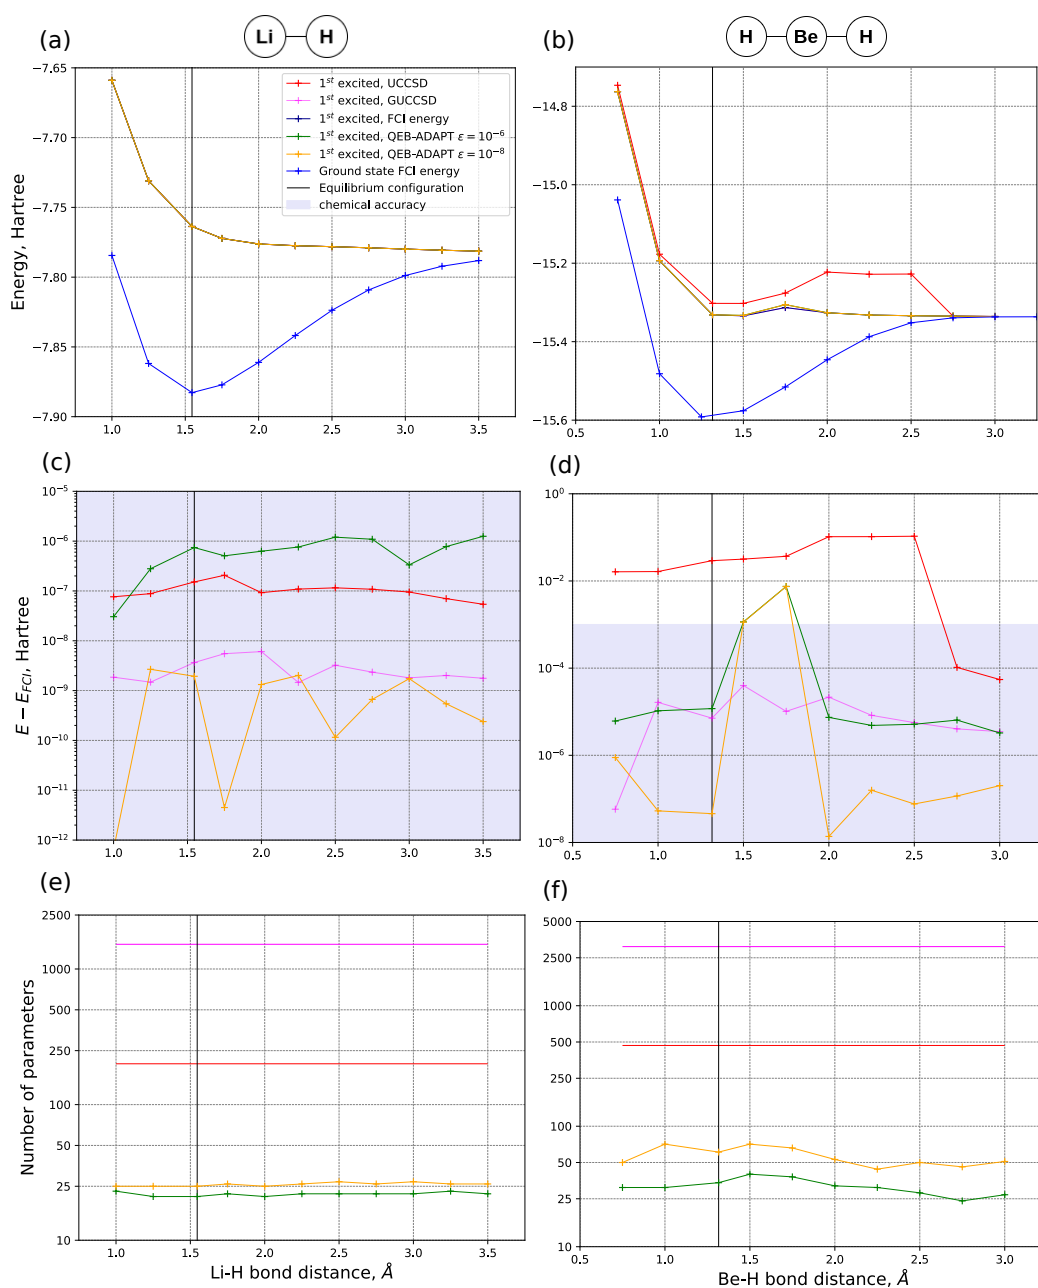
Fig. 5.3 Energy dissociation curves for LiH and BeH$_2$ in the STO-3G orbital basis. **a,b** Absolute values for the estimated first excited state energies. **c,d** Error in the estimated energy values with respect to the exact FCI energy, as function of bond distance. **e,f** Number of variational parameters, also equal to the number of ansatz elements, used by the ansatz of the corresponding method, as function of bond distance. The UCCSD includes non-spin conserving fermionic evolutions and has 200 and 468 parameters for LiH and BeH$_2$, respectively. The GUCCSD has 1551 and 3094 parameters for LiH and BeH$_2$, respectively.

extremely compact, consisting of at most 27 qubit evolutions, whereas the UCCSD and the GUCCSD consist of 200 and 1521 fermionic evolutions, respectively.

In the case of BeH_2 the results are more interesting. First, the UCCSD-VQE fails to achieve chemical accuracy for the majority of bond distances (Fig. 5.3d). This result is not surprising as the simple UCCSD ansatz is not suitable to approximate strongly correlated states such as excited states. Second, in Fig. 5.3d we see that the e-QEB-ADAPT-VQE fails to achieve chemical accuracy at bond distances $r_{Be-H} = 1.5$Å and $r_{Be-H} = 1.75$Å. I comment on the reason for this below. The only method that achieves chemical accuracy at all bond distances is the GUCCSD-VQE, owing to the high variational flexibility of the GUCCSD ansatz. Nevertheless, for the majority of bond distances the e-QEB-ADAPT-VQE constructs ansätze that are more accurate and consist of nearly 50 times fewer ansatz elements than the GUCCSD (Figs. 5.3d and 5.3f).

Finding the energy of a wrong excited state

As we saw in Fig. 5.3d, the e-QEB-ADAPT-VQE fails to achieve chemical accuracy for BeH_2 at bond distances $r_{Be-H} = 1.5$Å and $r_{Be-H} = 1.75$Å. To investigate the reason for this, let us consider Fig. 5.4, which depicts the FCI energies for the 10 lowest energy states of BeH_2 (in the STO-3G basis) as functions of bond distance.

The 9 excited states shown in Fig. 5.4 are ordered in three degenerate energy levels. The table below summarises the energy levels, and their corresponding degenerate excited states, at bond distances $r_{Be-H} = 1.5$Å and $r_{Be-H} = 1.75$Å:

$r_{Be-H} = 1.5$Å		$r_{Be-H} = 1.75$Å	
exc. states	energy, Hartree	exc. states	energy, Hartree
{1,2}	$E_I(1.5$Å$) \approx -15.3343$	{1,2}	$E_I(1.75$Å$) \approx -15.3131$
{3,4,5,6,7,8}	$E_{II}(1.5$Å$) \approx -15.3331$	{3,4,5}	$E_{II}(1.75$Å$) \approx -15.3059$
{9}	$E_{III}(1.5$Å$) \approx -15.3026$	{6,7,8,9}	$E_{III}(1.75$Å$) \approx -15.3056$

Table 5.1 Low lying excited energy levels of BeH_2 at bond distances of $r_{Be-H} = 1.5$ and $r_{Be-H} = 1.75$. Note that the ground state energy is not included here, and the counting of the excited states starts from 1, so that state 0 is the ground state.

For $r_{Be-H} = 1.5$Å we can see that $E_I(1.5$Å$)$ and $E_{II}(1.5$Å$)$ are very close. Upon inspection of the energy estimate of the e-QEB-ADAPT-VQE($\varepsilon = 10^{-8}$), for $r_{Be-H} = 1.5$Å, it turns out that the algorithm actually finds $E_{II}(1.5$Å$)$ with accuracy of 3×10^{-8} Hartree. Therefore, the problem is not a lack of accuracy, but converging to the wrong excited state.

Two possible reasons for convergence to a wrong excited state are: (1) the classical minimizers used by the e-QEB-ADAPT-VQE get stuck in a local minimum corresponding to $E_{II}(1.5\text{Å})$, and/or (2) the e-QEB-ADAPT-VQE fails to construct an ansatz that can approximate the degenerate excited states corresponding to energy level $E_I(1.5\text{Å})$. Since the e-QEB-ADAPT-VQE uses a combination of the BFGS and the Nelder-Mead optimization methods, where the Nelder-Mead is a direct search method that is unlikely to get stuck in local minima, the first possibility can be ruled out. Hence, we are left with the second one.
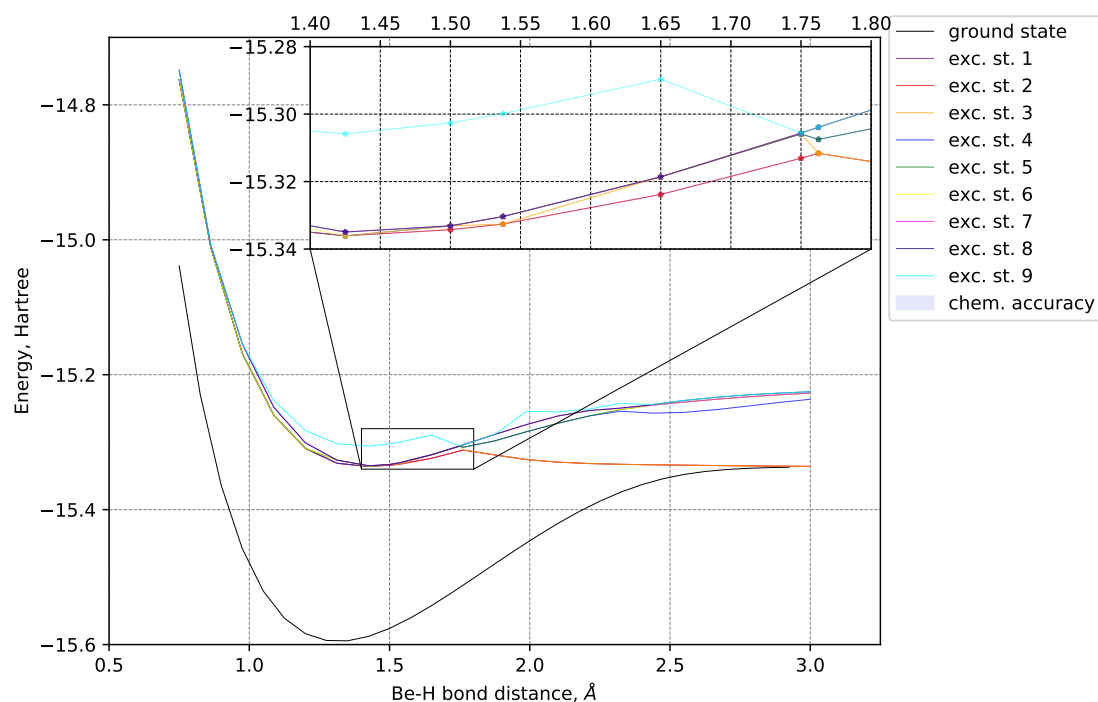
Fig. 5.4 FCI energies for the 10 lowest energy states of BeH_2 in the STO-3G basis. The numbers of the excited states correspond to their order at equilibrium bond distance of $r_{Be-H} = 1.316$.

Upon explicit inspection of the degenerate states corresponding to energy level $E_I(1.5\text{Å})$, it is found that they have strong static correlations, where two dominant Slater determinants contribute equally to the wavefuncion. On the other hand, for each of the degenerate states corresponding to energy level $E_{II}(1.5\text{Å})$ there is one dominant Slater determinant. Therefore, what happens is that the first few ansatz-constructing-iterations of the e-QEB-ADAPT-VQE (see Sec. 5.2) select qubit evolutions that map between the initial reference state, $|\psi_0\rangle$, and the dominant Slater determinant of one of the states corresponding to $E_{II}(1.5\text{Å})$, because these qubit evolutions decrease the estimate for the energy by the most. From that point on, the e-

QEB-ADAPT-VQE continues along a "wrong" path to construct an ansatz that approximates a state corresponding to $E_{II}(1.5\text{Å})$, instead of a state corresponding to $E_I(1.5\text{Å})$.

One might suggest that a possible solution would be to initialize $|\psi_0\rangle$ as an equal superposition of two Slater determinants. However, this would not work, because a qubit evolution that maps between one of the Slater determinants of $|\psi_0\rangle$ and the dominant Slater determinant of a state corresponding to $E_{II}(1.5\text{Å})$ will still decrease the energy estimate by the most.

In the case of $r_{Be-H} = 1.75\text{Å}$, the situation is slightly different. The three energy levels $E_I(1.75\text{Å})$, $E_{II}(1.75\text{Å})$ and $E_{III}(1.75\text{Å})$ are close, and the states corresponding to $E_I(1.75\text{Å})$ and $E_{II}(1.75\text{Å})$ have two dominant Slater determinants, whereas the states corresponding to $E_{III}(1.75\text{Å})$ have one dominant Slater determinant. Hence, the excited-QEB-ADAPT-VQE converges to an excited state corresponding to the energy level with highest energy, but weakest static correlations, $E_{III}(1.75\text{Å})$.

Overall the following conclusion can be made: A generic feature of the e-QEB-ADAPT-VQE is that it is more "willing" to construct ansätze for less statically correlated states (which have just one dominant Slater determinant). Therefore, in the case when (1) the lowest energy eigenvalues of an electronic Hamiltonian are ordered in two or more closely spaced energy levels, and (2) the lowest energy level corresponds to eigenstates that have stronger static correlations than the eigenstates above them, the e-QEB-ADAPT-VQE might converge to the less statically correlated eigenstates, instead of to the those with lowest energy.

5.3.2 Higher excited states

Here I further benchmark the performance of the e-QEB-ADAPT-VQE by obtaining higher excited state energies for LiH. Energy dissociation curves for the five lowest energy states of LiH in the STO-3G basis, obtained with the e-QEB-ADAPT-VQE for $n_{qe} = 10$ and $\varepsilon = 10^{-8}$ Hartree, are plotted in Fig. 5.5a. Figure 5.5b shows the corresponding errors.
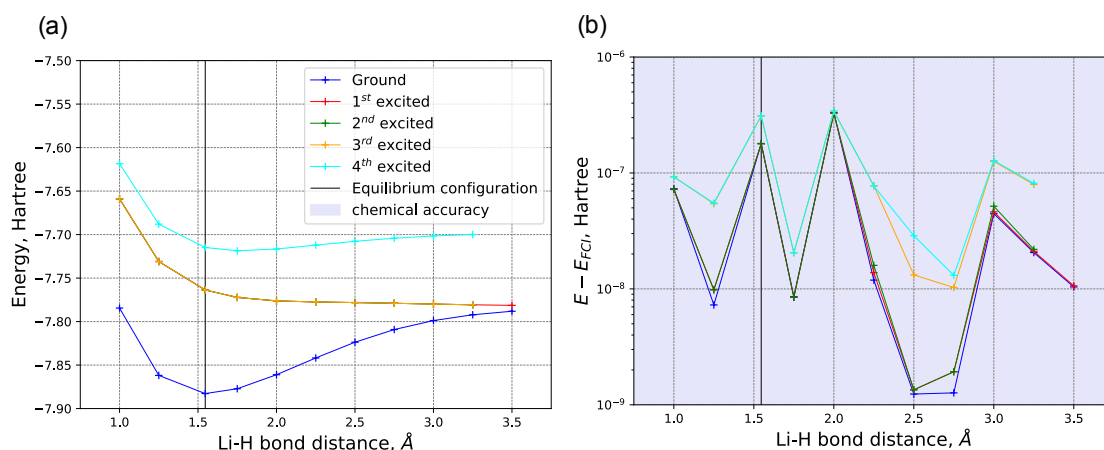
Fig. 5.5 **a** Energies for the 5 lowest energy states of LiH, in the STO-3G basis, obtained with e-QEB-ADAPT-VQE for $\varepsilon = 10^{-8}$ Hartree. Note that excited states $\{1, 2, 3\}$ are degenerate. **b** Errors in the estimated energy values with respect to the FCI energies.

The first important thing to verify in these plots is that the e-QEB-ADAPT-VQE successfully constructs accurate ansätze for all three degenerate states corresponding to the lowest excited energy level of LiH. Also, Fig. 5.5b demonstrates how the recursive error of the overlap-based method scales as more excited states are approximated. Assuming that the relative error of the excited-QEB-ADAPT-VQE for each excited state is approximately the same, then the absolute error in the estimate of the energy of excited state k, E_k, should increase approximately linearly with k. Although we do not have enough excited states to check if this linear dependence is obeyed, the increase in the error is evident.

5.4 Comparison of iteraively constructed Q-UCC and UCC ansätze for excited states

In this section, once again I perform a comparison between iteratively constructed Q-UCC and UCC ansätze, this time approximating excited states. For this purpose I construct both types of ansätze, using the excited-QEB-ADAPT-VQE($n_{qe} = 10$), for the first excited states of LiH and BeH$_2$ at bond distances of $r_{\text{Li-H}} = 1.546$Å and $r_{\text{Be-H}} = 1.316$Å, respectively. Energy convergence plots for these ansätze are presented in Fig. 5.6.

Fig. 5.6 Energy convergence plots for the first excited states of LiH and BeH$_2$ in the STO-3G basis at equilibrium bond distances of $r_{\text{Li-H}} = 1.546$Å and $r_{\text{Be-H}} = 1.316$Å, respectively. The plots are obtained with the e-QEB-ADAPT-VQE for $n_{qe} = 10$ and $\varepsilon = 10^{-12}$ Hartree.

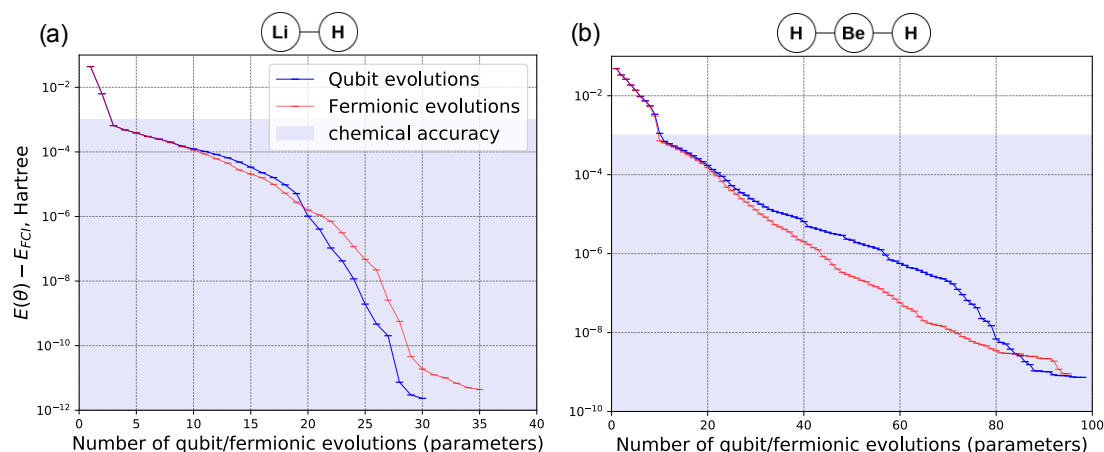We observe close similarity between the Q-UCC and the UCC ansätze. In the case of LiH (Fig. 5.6a) the Q-UCC ansatz is slightly more accurate per number of ansatz elements, consisting of up to 15% fewer ansatz elements than the UCC ansatz. This result does not agree with our expectation for the UCC to be more accurate for the more strongly correlated first excited state. However, in the case of BeH$_2$ (Fig. 5.6b), whose first excited state has strong static correlations, the UCC ansatz is more accurate per number of ansatz elements, consisting of up to 25% fewer ansatz elements than the Q-UCC ansatz. Nevertheless, since qubit evolutions are implemented by simpler circuits than fermionic evolutions, the Q-UCC ansatz will require fewer CNOTs even in the case of BeH$_2$.

5.5   Conclusion

In this chapter, I proposed a modified version of the QEB-ADAPT-VQE protocol, the excited-QEB-ADAPT-VQE, intended to estimate molecular excited state energies. Choosing an initial reference state that has a significant overlap with an unknown target excited state is not as straightforward as for a target ground state. Therefore, the e-QEB-ADAPT-VQE is designed to be less dependent on the choice of an initial reference state. This is achieved by an ansatz-growing strategy that does not rely on energy-gradient evaluations, but instead on individual-ansatz-element energy-reduction evaluations. The modified ansatz-growing strategy comes at a cost of up to a constant factor of more quantum computer measurements per ansatz-growing iteration, as compared to the QEB-ADAPT-VQE.

I benchmarked the performance of the e-QEB-ADAPT-VQE by constructing ansätze for the first excited states of LiH and BeH_2. Generally, it was found that the e-QEB-ADAPT-VQE can construct highly accurate ansätze that require few variational parameters, and are implemented by short ansatz circuits, similarly to the ansätze constructed by the QEB-ADAPT-VQE for ground states.

However, it was also found that the e-QEB-ADAPT-VQE might fail to find excited state energies in order of increasing energy, if the excited states are ordered in closely spaced degenerate energy levels, and the lower lying excited states are more statically correlated than the states above them. This issue derives from the fact that the e-QEB-ADAPT-VQE pursues the greedy strategy to achieve a lowest estimate for the energy at each ansatz-growing iteration, which results into the protocol being more likely to construct an ansatz for a less statically correlated state. In fact, this issue will be present in any iterative-VQE protocol that relies on such a greedy strategy, including all ADAPT-VQE protocols and the IQCC [155]. Nevertheless, this is unlikely to be a significant problem, because in practice one would be interested in finding the whole spectrum of low lying excited state energies, so the order in which they are found is not necessarily important.

Lastly, I used the e-QEB-ADAPT-VQE to compare iteratively constructed Q-UCC and UCC ansätze in approximating excited states. The results supported my previous observations in Chapters 3 and 4 that UCC ansätze exhibit a slight advantage in terms of accuracy per number of ansatz elements, when approximating strongly correlated states.

Chapter 6

Conclusion

In this thesis, I derived techniques to reduce the quantum circuit requirements of existing VQE protocols for molecular modelling, and developed a new protocol, the QEB-ADAPT-VQE. The QEB-ADAPT-VQE supersedes previous state of the art VQE protocols in terms of circuit efficiency and convergence speed. Below I summarise my findings and suggest future directions for research.

6.1 Summary

In Chapter 1 I outlined the current state of quantum computing. I started by summarising the different models of universal quantum computers. In particular, I considered the most widely adopted quantum-circuit-based model of a quantum computer, and discussed its various physical implementations and their corresponding limitations. The bottleneck in the performance of quantum-circuit computers is the number of quantum gates, in particular 2-qubit entangling gates, that can be performed reliably before quantum errors build up and destroy any useful amounts of quantum information. This has prompted the development of hybrid quantum-classical algorithms, in particular the prominent VQE algorithm, that require shallower quantum circuits and shorter coherence times than purely quantum algorithms.

The VQE is a hybrid quantum-classical algorithm that can estimate the lowest eigenvalue of a Hermitian operator exponentially faster than any classical algorithm. A promising, practically useful, application for emerging NISQ computers is molecular modelling using the VQE. In Chapter 1 we saw how the electronic structure problem, of finding the ground state energy and wavefunction of a molecule (or any other electron system), can be mapped to a quantum computer, and then solved by the VQE. The major component of the VQE, is an ansatz to generate a trial state to approximate the molecular electronic wavefunction. For a successful implementation of the VQE on NISQ computers the ansatz should be: (1)

variationally flexible in order to accurately approximate the ground state of the molecule; (2) easy to optimize; and (3) performed by a shallow quantum circuit. The most commonly used anätze for molecular modelling with the VQE are the UCC ansätze, which correspond to products of parametrized evolutions of fermionic excitation operators (fermionic evolutions). These ansätze are accurate and easy to optimize owing to their fermionic structure.

In Chapter 2 we saw what are the canonical circuits used to implement fermionic evolutions. Stacking such circuits to construct UCC ansätze results in ansatz circuits that, even for relatively small molecules, greatly exceed the current capabilities of NISQ computers. In Chapter 2 I derived alternative circuits to implement single and double fermionic evolutions, optimized in the number of $CNOT$ gates. For this purpose, first I constructed circuits to implement single and double qubit evolutions, evolutions of qubit excitation operators that satisfy qubit commutation relations. Then, I expanded the functionality of the qubit evolution circuits to account for fermionic anticommutation, thus obtaining $CNOT$-optimized fermionic evolution circuits. Compared to the canonical circuits, the circuits I derived reduce the $CNOT$ count by between 25% to 50%, and 73% to 87%, for single and double fermionic evolutions, respectively.

Fermionic evolutions are implemented by circuits that act on a number of qubits, and require a number of gates, which scale with the system size, defined by the number of molecular spin-orbitals, N_{so}. In the Jordan-Wigner qubit encoding these numbers scale as $O(N_{so})$, and in the Bravyi-Kitaev qubit encoding as $O(\log N_{so})$. On the other hand qubit evolutions are implemented by circuits that act on a fixed, $O(1)$, number of qubits, and consequently require a fixed, $O(1)$, number of gates, irrespectively of the qubit encoding. In Chapter 3, I investigated the use of ansätze corresponding to products of qubit evolutions, termed as Q-UCC, as a means to further reduce the circuit requirements of the VQE. I performed a comparison between the UCCSD and the Q-UCCSD ansätze by conducting classical numerical VQE simulations for LiH, H_6 and BeH_2. The simulation results for LiH and BeH_2 indicated that the two types of ansätze perform equivalently in approximating electronic wavefunctions. However, for H_6, which has a strongly correlated ground state, the results indicated a slight accuracy advantage for the UCCCSD.

Motivated by the results about the equivalence, in terms of accuracy, of the Q-UCCSD and the UCCSD ansätze, in Chapter 4 I proposed the qubit-excitation based adaptive (QEB-ADAPT) VQE protocol for molecular modelling. The QEB-ADAPT-VQE is a VQE protocol for molecular modelling, which constructs its own problem tailored ansätze by iteratively appending single and double qubit evolutions. I benchmarked the performance of the QEB-ADAPT-VQE with classical numerical simulations for LiH, H_6 and BeH_2. First, I compared the ansätze constructed by the QEB-ADAPT-VQE against the fixed UCCSD ansatz, and

demonstrated that the former ansätze are both more accurate and require multiple times fewer variational parameters, and correspondingly shallower circuits, than the UCCSD. Second, I compared the QEB-ADAPT-VQE against the original fermionic-ADAPT-VQE and its more circuit-efficient cousin, the qubit-ADAPT-VQE. The QEB-ADAPT-VQE was up to 4 times slower than the fermionic-ADAPT-VQE, but much faster than the qubit-ADAPT-VQE. For low accuracies the qubit-ADAPT-VQE was the most $CNOT$-efficient protocol. However, for higher accuracies, and correspondingly larger ansatz circuits, the QEB-ADAPT-VQE constructed systematically more $CNOT$-efficient ansätze than both the fermionic-ADAPT-VQE and qubit-ADAPT-VQE. These results imply that for large molecules the QEB-ADAPT-VQE is likely to be both faster and more circuit-efficient than the qubit-ADAPT-VQE, which to my knowledge was the previous most-circuit efficient practically-scalable VQE protocol for molecular modelling. Therefore, I believe that the QEB-ADAPT-VQE represents the current state of the art method to determine molecular state energies on a quantum computer.

In Chapter 5, I proposed a modified version of the QEB-ADAPT-VQE, the excited-QEB-ADAPT-VQE, designed to estimate energies of excited molecular states. Choosing an initial reference state, which has a large overlap with an unknown target excited state is rarely possible. Hence, the e-QEB-ADAPT-VQE, unlike the QEB-ADAPT-VQE, was designed to be less dependent on the choice of an initial reference state, at the cost of requiring up to a constant factor of more quantum computer measurements. The performance of the e-QEB-ADAPT-VQE was benchmarked with classical numerical simulations for the excited states of LiH and BeH$_2$. Generally, the e-QEB-ADAPT-VQE was found to construct ansätze, which achieve well below chemical accuracy over a range of molecular bond distances. This implied that the e-QEB-ADAPT-VQE is well suited to approximate excited states. Moreover, these ansätze required multiple times fewer parameters, and correspondingly shallower ansatz circuits, than the generalized UCCSD (GUCCSD) ansatz, which could achieve similar accuracies. However, it was also found that the e-QEB-ADAPT-VQE might fail to yield molecular state energies in order of increasing energy, if the molecular system of interest has energy levels that are closely spaced and the corresponding eigenstates have strong static correlations. Nevertheless, this issue is unlikely to be important in practice if we are interested in obtaining the whole spectrum of low lying excited states.

I also used the fermionic-ADAPT-VQE, the QEB-ADAPT-VQE and the excited-QEB-ADAPT-VQE to perform a comparison of iteratively constructed Q-UCC and UCC ansätze. The ansätze constructed by these ADAPT-VQE protocols consist of close to optimal numbers of ansatz elements (qubit/fermionic evolutions). Hence, comparing iteratively constructed Q-UCC and UCC ansätze corresponds to a more thorough comparison of qubit and fermionic evolutions, than the comparison of the Q-UCCSD and the UCCSD ansätze (Chapter 3), which

include redundant numbers of qubit and fermionic evolutions, respectively. In Chapter 4, I compared Q-UCC and UCC ansätze constructed by the QEB-ADAPT-VQE and a modified version of the fermionic-ADAPT-VQE, respectively, for the ground states of LiH, H_6 and BeH_2. The comparison indicated that on average the UCC ansätze are a bit more accurate per number of ansatz elements than the Q-UCC ansätze, when approximating strongly correlated states. In Chapter 5 a similar comparison was performed between Q-UCC and UCC ansätze constructed by the excited-QEB-ADAPT-VQE for the first excited states of LiH and BeH_2. In this case, it was again observed that there is a tendency for the UCC to be slightly more accurate when approximating strongly correlated states. Another advantage of fermionic evolutions over qubit evolutions was found to be that the former can efficiently form spin-complement pairs, owing to their anticommutation structure, whereas the latter cannot. This feature of fermionic evolutions allows the number of variational parameters of an UCC ansatz to be roughly halved when spin-conservation must be obeyed by the ansatz. However, in the era of NISQ computers, these two advantages of fermionic evolutions are outweighed by the asymptotically better circuit-efficiency of qubit evolutions.

6.2   Further Work

The main result of this thesis is the QEB-ADAPT-VQE protocol. It improves on previous VQE protocols for molecular modelling in terms of circuit-efficiency and convergence speed. Nevertheless, there are numerous opportunities to further improve the QEB-ADAPT-VQE. Below I suggest three such improvements.

First, the ansatz element pool of the QEB-ADAPT-VQE can be expanded to include non-symmetry-preserving terms as suggested in Ref. [174]. This expanded pool could improve the speed of convergence and boost the resilience to symmetry-breaking errors of the QEB-ADAPT-VQE. Second, the methods suggested in Ref. [142] can be incorporated into the QEB-ADAPT-VQE, to "prune" (remove) from the already constructed ansatz, ansatz elements that have little contribution to the energy reduction. These methods could further reduce the size of the constructed ansatz. Third, the ansatz growing strategy of the QEB-ADAPT-VQE could be entirely replaced by a generalized machine learning routine that incorporates both, growing and pruning of the ansatz at the same time. Such an upgrade could result in constructing close to optimal ansätze, in terms of circuit size and number of variational parameters. Moreover, a machine learning ansatz-growing routine could allow the QEB-ADAPT-VQE to "learn" how to construct ansätze for new molecules or excited states, based on previously constructed ansätze for ground states. This could remove the need to calculate the energy gradients or energy reductions for all ansatz elements at each iteration.

Another important direction of future work is to study how the QEB-ADAPT-VQE protocol performs in the presence of noise. For example, quantum noise causes fluctuations in the energy landscape, with respect to the parameters of an ansatz, which can produce inaccurate values for the energy gradients required to grow the ansatz of the QEB-ADAPT-VQE. This can result in constructing an inefficient or/and inaccurate ansatz. Because of this, the excited-QEB-ADAPT-VQE might turn out to be more suitable for high noise levels, since it does not rely on evaluating energy gradients. This and other questions related to the performance of the QEB-ADAPT-VQE in the presence of noise are being investigated by Zhenghao Li. This investigation will first include classical noisy simulation of the QEB-ADAPT-VQE, using the CQCC package. Then, depending on the results of the classical simulation, the QEB-ADAPT-VQE will be benchmarked on real cloud accessible quantum computers, offered by IBM, IonQ and Rigetti. As a longer-term goal, we also plan to incorporate error mitigatation [61] methods in the ADAPT-VQE framework.

Lastly, the optimized circuits for qubit and fermionic evolutions, derived in Chapter 2, correspond to local circuit optimization. In Chapter 2, I also referenced several global optimization methods. It would be interesting to investigate if these global optimization methods, especially those suggested in Refs. [152, 153], can be applied to reduce even further the size of the ansatz circuits constructed with the locally optimized qubit or fermionic evolution circuits derived in Chapter 2.

References

[1] Rodney Loudon. *The quantum theory of light*. OUP Oxford, 2000.

[2] D Ter Haar and A Einstein. 3 on a heuristic point of view about the creatidn and conversion of light?

[3] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:1175–1204, Dec 1957.

[4] H. J. Kimble, M. Dagenais, and L. Mandel. Photon antibunching in resonance fluorescence. *Phys. Rev. Lett.*, 39:691–695, Sep 1977.

[5] Jan Faye. Copenhagen interpretation of quantum mechanics. 2002.

[6] Lev Vaidman. Many-worlds interpretation of quantum mechanics. 2002.

[7] Peter R Holland. The de broglie-bohm theory of motion and quantum field theory. *Physics reports*, 224(3):95–150, 1993.

[8] Edward Nelson. Derivation of the schrödinger equation from newtonian mechanics. *Phys. Rev.*, 150:1079–1085, Oct 1966.

[9] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of statistical physics*, 22(5):563–591, 1980.

[10] Paul Benioff. Quantum mechanical hamiltonian models of turing machines. *Journal of Statistical Physics*, 29(3):515–546, 1982.

[11] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys*, 21(6/7), 1982.

[12] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[13] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996.

[14] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[15] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane. Quantum error correction and orthogonal geometry. *Phys. Rev. Lett.*, 78:405–408, Jan 1997.

[16] D. G. Cory, M. D. Price, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo. Experimental quantum error correction. *Phys. Rev. Lett.*, 81:2152–2155, Sep 1998.

[17] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[18] Laird Egan, Dripto M Debroy, Crystal Noel, Andrew Risinger, Daiwei Zhu, Debopriyo Biswas, Michael Newman, Muyuan Li, Kenneth R Brown, Marko Cetina, et al. Fault-tolerant operation of a quantum error-correction code. *arXiv preprint arXiv:2009.11482*, 2020.

[19] Guanyu Zhu and Andrew Cross. Hardware-aware approach for fault-tolerant quantum computation. *https://www.ibm.com/blogs*, 2020.

[20] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[21] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[22] Edwin Pednault, John A Gunnels, Giacomo Nannicini, Lior Horesh, and Robert Wisnieff. Leveraging secondary storage to simulate deep 54-qubit sycamore circuits. *arXiv preprint arXiv:1910.09534*, 2019.

[23] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. Quantum computational chemistry. *Rev. Mod. Phys.*, 92:015003, Mar 2020.

[24] Henry Eyring. The activated complex in chemical reactions. *The Journal of Chemical Physics*, 3(2):107–115, 1935.

[25] Yudong Cao, Jhonathan Romero, and Alán Aspuru-Guzik. Potential of quantum computing for drug discovery. *IBM Journal of Research and Development*, 62(6):6–1, 2018.

[26] Jacob D Durrant and J Andrew McCammon. Molecular dynamics simulations and drug discovery. *BMC biology*, 9(1):1–9, 2011.

[27] Trygve Helgaker, Poul Jorgensen, and Jeppe Olsen. *Molecular electronic-structure theory*. John Wiley & Sons, 2014.

[28] Kitaev A Yu. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.

[29] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P. Lanyon, Peter Love, Ryan Babbush, Alán Aspuru-Guzik, Rainer Blatt, and Christian F. Roos. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X*, 8:031022, Jul 2018.

[30] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5:4213, 2014.

[31] Yunseong Nam, Jwo-Sy Chen, Neal C Pisenti, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R Brown, Stewart Allen, Jason M Amini, Joel Apisdorf, et al. Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Information*, 6(1):1–6, 2020.

[32] Thomas Durt, Dagomir Kaszlikowski, Jing-Ling Chen, and L. C. Kwek. Security of quantum key distributions with entangled qudits. *Phys. Rev. A*, 69:032313, Mar 2004.

[33] Yuchen Wang, Zixuan Hu, Barry C Sanders, and Sabre Kais. Qudits and high-dimensional quantum computing. *Frontiers in Physics*, 8:479, 2020.

[34] Yu. I. Bogdanov, M. V. Chekhova, S. P. Kulik, G. A. Maslennikov, A. A. Zhukov., C. H. Oh, and M. K. Tey. Qutrit state engineering with biphotons. *Phys. Rev. Lett.*, 93:230503, Dec 2004.

[35] Michael A Nielsen. Cluster-state quantum computation. *Reports on Mathematical Physics*, 57(1):147–161, 2006.

[36] Robert Raussendorf, Daniel E. Browne, and Hans J. Briegel. Measurement-based quantum computation on cluster states. *Phys. Rev. A*, 68:022312, Aug 2003.

[37] M. Hein, J. Eisert, and H. J. Briegel. Multiparty entanglement in graph states. *Phys. Rev. A*, 69:062311, Jun 2004.

[38] Arnab Das and Bikas K. Chakrabarti. Colloquium: Quantum annealing and analog quantum computation. *Rev. Mod. Phys.*, 80:1061–1081, Sep 2008.

[39] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma. Non-abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80:1083–1159, Sep 2008.

[40] David Elieser Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 425(1868):73–90, 1989.

[41] David P. DiVincenzo. Two-bit gates are universal for quantum computation. *Phys. Rev. A*, 51:1015–1022, Feb 1995.

[42] David P DiVincenzo, Dave Bacon, Julia Kempe, Guido Burkard, and K Birgitta Whaley. Universal quantum computation with the exchange interaction. *nature*, 408(6810):339–342, 2000.

[43] P Oscar Boykin, Tal Mor, Matthew Pulver, Vwani Roychowdhury, and Farrokh Vatan. On universal and fault-tolerant quantum computing. *arXiv preprint quant-ph/9906054*, 1999.

[44] Ibm quantum experience. *https://quantum-computing.ibm.com/*, 2021.

[45] John M Nichol, Lucas A Orona, Shannon P Harvey, Saeed Fallahi, Geoffrey C Gardner, Michael J Manfra, and Amir Yacoby. High-fidelity entangling gate for double-quantum-dot spin qubits. *npj Quantum Information*, 3(1):1–5, 2017.

[46] E. A. Laird, J. M. Taylor, D. P. DiVincenzo, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent spin manipulation in an exchange-only qubit. *Phys. Rev. B*, 82:075403, Aug 2010.

[47] Xiongfeng Ma, Chi-Hang Fred Fung, and Hoi-Kwong Lo. Quantum key distribution with entangled photon sources. *Phys. Rev. A*, 76:012307, Jul 2007.

[48] Mercedes Gimeno-Segovia, Pete Shadbolt, Dan E. Browne, and Terry Rudolph. From three-photon greenberger-horne-zeilinger states to ballistic universal quantum computation. *Phys. Rev. Lett.*, 115:020502, Jul 2015.

[49] Artur K. Ekert, Bruno Huttner, G. Massimo Palma, and Asher Peres. Eavesdropping on quantum-cryptographical systems. *Phys. Rev. A*, 50:1047–1056, Aug 1994.

[50] Bruce E Kane. A silicon-based nuclear spin quantum computer. *nature*, 393(6681):133–137, 1998.

[51] Jarryd J Pla, Kuan Y Tan, Juan P Dehollain, Wee H Lim, John JL Morton, Floris A Zwanenburg, David N Jamieson, Andrew S Dzurak, and Andrea Morello. High-fidelity readout and control of a nuclear spin qubit in silicon. *Nature*, 496(7445):334–338, 2013.

[52] Andrew M Steane. The ion trap quantum information processor. *arXiv preprint quant-ph/9608011*, 1996.

[53] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A*, 76:042319, Oct 2007.

[54] A. A. Clerk, M. H. Devoret, S. M. Girvin, Florian Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Rev. Mod. Phys.*, 82:1155–1208, Apr 2010.

[55] N Bundaleski, P Soulisse, A Momeni, H Khemliche, and P Roncin. Decoherence in fast atom diffraction from surfaces. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 269(11):1216–1220, 2011.

[56] P. Mohanty and R. A. Webb. Decoherence and quantum fluctuations. *Phys. Rev. B*, 55:R13452–R13455, May 1997.

[57] S. Ulmer, C. C. Rodegheri, K. Blaum, H. Kracke, A. Mooser, W. Quint, and J. Walz. Observation of spin flips with a single trapped proton. *Phys. Rev. Lett.*, 106:253001, Jun 2011.

[58] O. Astafiev, Yu. A. Pashkin, Y. Nakamura, T. Yamamoto, and J. S. Tsai. Quantum noise in the josephson charge qubit. *Phys. Rev. Lett.*, 93:267007, Dec 2004.

[59] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Phys. Rev. Lett.*, 103:110501, Sep 2009.

[60] Rami Barends, J Kelly, A Megrant, A Veitia, D Sank, E Jeffrey, TC White, J Mutus, AG Fowler, B Campbell, et al. Logic gates at the surface code threshold: Superconducting qubits poised for fault-tolerant quantum computing. *arXiv preprint arXiv:1402.4848*, 2014.

[61] Suguru Endo, Simon C. Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Phys. Rev. X*, 8:031027, Jul 2018.

[62] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear algebra and its applications*, 10(3):285–290, 1975.

[63] Michael A Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, 303(4):249–252, 2002.

[64] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Neutral atom quantum register. *Phys. Rev. Lett.*, 93:150501, Oct 2004.

[65] Christoph Adami and Nicolas J Cerf. Quantum computation with linear optics. In *NASA International Conference on Quantum Computing and Quantum Communications*, pages 391–401. Springer, 1998.

[66] Hannes Pichler, Soonwon Choi, Peter Zoller, and Mikhail D Lukin. Universal photonic quantum computation via time-delayed feedback. *Proceedings of the National Academy of Sciences*, 114(43):11362–11367, 2017.

[67] Jacob F. Steiner and Felix von Oppen. Readout of majorana qubits. *Phys. Rev. Research*, 2:033255, Aug 2020.

[68] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11:369–395, 2020.

[69] VM Schäfer, CJ Ballance, K Thirumalai, LJ Stephenson, TG Ballance, AM Steane, and DM Lucas. Fast quantum logic gates with trapped-ion qubits. *Nature*, 555(7694):75–78, 2018.

[70] Xiao Xue, Benjamin D'Anjou, Thomas F. Watson, Daniel R. Ward, Donald E. Savage, Max G. Lagally, Mark Friesen, Susan N. Coppersmith, Mark A. Eriksson, William A. Coish, and Lieven M. K. Vandersypen. Repetitive quantum nondemolition measurement and soft decoding of a silicon spin qubit. *Phys. Rev. X*, 10:021006, Apr 2020.

[71] Andrew A Houck, Jens Koch, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Life after charge noise: recent results with transmon qubits. *Quantum Information Processing*, 8(2):105–115, 2009.

[72] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O'Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.*, 111:080502, Aug 2013.

[73] Brian David Josephson. Possible new effects in superconductive tunnelling. *Physics letters*, 1(7):251–253, 1962.

[74] Chad Tyler Rigetti. *Quantum gates for superconducting qubits*. Yale University, 2009.

[75] Ming Gong, Shiyu Wang, Chen Zha, Ming-Cheng Chen, He-Liang Huang, Yulin Wu, Qingling Zhu, Youwei Zhao, Shaowei Li, Shaojun Guo, et al. Quantum walks on a programmable two-dimensional 62-qubit superconducting processor. *Science*, 372(6545):948–952, 2021.

[76] D-wave systems. *https://www.dwavesys.com/*, 2021.

[77] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504, Aug 2016.

[78] Nicolai Friis, Oliver Marty, Christine Maier, Cornelius Hempel, Milan Holzäpfel, Petar Jurcevic, Martin B. Plenio, Marcus Huber, Christian Roos, Rainer Blatt, and Ben Lanyon. Observation of entangled states of a fully controlled 20-qubit system. *Phys. Rev. X*, 8:021012, Apr 2018.

[79] J. J. García-Ripoll, P. Zoller, and J. I. Cirac. Speed optimized two-qubit gates with laser coherent control techniques for ion trap quantum computing. *Phys. Rev. Lett.*, 91:157901, Oct 2003.

[80] Jiehang Zhang, Guido Pagano, Paul W Hess, Antonis Kyprianidis, Patrick Becker, Harvey Kaplan, Alexey V Gorshkov, Z-X Gong, and Christopher Monroe. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature*, 551(7682):601–604, 2017.

[81] Peter Chapman. Introducing the world's most powerful quantum computer. *https://ionq.com/posts*, 2020.

[82] The world's highest performing quantum computer is here. *https://www.honeywell.com/us/en/news*, 2020.

[83] Daniel Loss and David P. DiVincenzo. Quantum computation with quantum dots. *Phys. Rev. A*, 57:120–126, Jan 1998.

[84] Hannes Watzinger, Josip Kukučka, Lada Vukušić, Fei Gao, Ting Wang, Friedrich Schäffler, Jian-Jun Zhang, and Georgios Katsaros. A germanium hole spin qubit. *Nature communications*, 9(1):1–6, 2018.

[85] Björn Trauzettel, Denis V Bulaev, Daniel Loss, and Guido Burkard. Spin qubits in graphene quantum dots. *Nature Physics*, 3(3):192–196, 2007.

[86] F. A. Calderon-Vargas, George S. Barron, Xiu-Hao Deng, A. J. Sigillito, Edwin Barnes, and Sophia E. Economou. Fast high-fidelity entangling gates for spin qubits in si double quantum dots. *Phys. Rev. B*, 100:035304, Jul 2019.

[87] L Childress, MV Gurudev Dutt, JM Taylor, AS Zibrov, F Jelezko, J Wrachtrup, PR Hemmer, and MD Lukin. Coherent dynamics of coupled electron and nuclear spin qubits in diamond. *Science*, 314(5797):281–285, 2006.

[88] Jim Clarke. What intel is planning for the future of quantum computing: Hot qubits, cold control chips, and rapid testing. *https://spectrum.ieee.org/*, 2020.

[89] Ivan Kassal, Stephen P Jordan, Peter J Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, 2008.

[90] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos. Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients. *Rev. Mod. Phys.*, 64:1045–1097, Oct 1992.

[91] Dennis C Rapaport. *The art of molecular dynamics simulation*. Cambridge university press, 2004.

[92] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan. Low-depth quantum simulation of materials. *Phys. Rev. X*, 8:011044, Mar 2018.

[93] Trygve Helgaker, Sonia Coriani, Poul Jorgensen, Kasper Kristensen, Jeppe Olsen, and Kenneth Ruud. Recent advances in wave function-based methods of molecular-property calculations. *Chemical reviews*, 112(1):543–631, 2012.

[94] Crispin H. W. Barnes. Lecture notes in advance quantum condensed matter physics, - 2016.

[95] Christof Hattig, Wim Klopper, Andreas Koohn, and David P Tew. Explicitly correlated electrons in molecules. *Chemical reviews*, 112(1):4–74, 2012.

[96] Dmitry I Lyakh, Monika Musial, Victor F Lotrich, and Rodney J Bartlett. Multireference nature of chemistry: The coupled-cluster view. *Chemical reviews*, 112(1):182–243, 2012.

[97] Isaiah Shavitt. The method of configuration interaction. In *Methods of electronic structure theory*, pages 189–275. Springer, 1977.

[98] Hans Lischka, Dana Nachtigallova, Adelia JA Aquino, Peter G Szalay, Felix Plasser, Francisco BC Machado, and Mario Barbatti. Multireference approaches for excited states of molecules. *Chemical reviews*, 118(15):7293–7361, 2018.

[99] BI Dunlap and N Rosch. The gaussian-type orbitals density-functional approach to finite systems. In *Advances in Quantum Chemistry*, volume 21, pages 317–339. Elsevier, 1990.

[100] Warren J Hehre, Robert F Stewart, and John A Pople. Self-consistent molecular-orbital methods. i. use of gaussian expansions of slater-type atomic orbitals. *The Journal of Chemical Physics*, 51(6):2657–2664, 1969.

[101] RHWJ Ditchfield, W J_ Hehre, and John A Pople. Self-consistent molecular-orbital methods. ix. an extended gaussian-type basis for molecular-orbital studies of organic molecules. *The Journal of Chemical Physics*, 54(2):724–728, 1971.

[102] Thom H Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen. *The Journal of chemical physics*, 90(2):1007–1023, 1989.

[103] Nathan Wiebe and Chris Granade. Efficient bayesian phase estimation. *Phys. Rev. Lett.*, 117:010503, Jun 2016.

[104] Jacob T Seeley, Martin J Richard, and Peter J Love. The bravyi-kitaev transformation for quantum computation of electronic structure. *The Journal of chemical physics*, 137(22):224109, 2012.

[105] Andrew Tranter, Peter J Love, Florian Mintert, and Peter V Coveney. A comparison of the bravyi–kitaev and jordan–wigner transformations for the quantum simulation of quantum chemistry. *Journal of chemical theory and computation*, 14(11):5617–5630, 2018.

[106] Ryan Babbush, Craig Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Alexandru Paler, Austin Fowler, and Hartmut Neven. Encoding electronic spectra in quantum circuits with linear t complexity. *Phys. Rev. X*, 8:041015, Oct 2018.

[107] U. Dorner, R. Demkowicz-Dobrzanski, B. J. Smith, J. S. Lundeen, W. Wasilewski, K. Banaszek, and I. A. Walmsley. Optimal quantum phase estimation. *Phys. Rev. Lett.*, 102:040403, Jan 2009.

[108] Y. S. Weinstein, M. A. Pravia, E. M. Fortunato, S. Lloyd, and D. G. Cory. Implementation of the quantum fourier transform. *Phys. Rev. Lett.*, 86:1889–1891, Feb 2001.

[109] Ryan Babbush, Dominic W Berry, Ian D Kivlichan, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially more precise quantum simulation of fermions in second quantization. *New Journal of Physics*, 18(3):033032, 2016.

[110] Ryan Babbush, Dominic W Berry, Yuval R Sanders, Ian D Kivlichan, Artur Scherer, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially more precise quantum simulation of fermions in the configuration interaction representation. *Quantum Science and Technology*, 3(1):015006, 2017.

[111] Hale F Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.

[112] Markus Reiher, Nathan Wiebe, Krysta M Svore, Dave Wecker, and Matthias Troyer. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences*, 114(29):7555–7560, 2017.

[113] Jarrod R McClean, Ryan Babbush, Peter J Love, and Alán Aspuru-Guzik. Exploiting locality in quantum computation for quantum chemistry. *The journal of physical chemistry letters*, 5(24):4368–4380, 2014.

[114] Andrew M Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *arXiv preprint arXiv:1202.5822*, 2012.

[115] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019.

[116] James D Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure hamiltonians using quantum computers. *Molecular Physics*, 109(5):735–750, 2011.

[117] N Cody Jones, James D Whitfield, Peter L McMahon, Man-Hong Yung, Rodney Van Meter, Alán Aspuru-Guzik, and Yoshihisa Yamamoto. Faster quantum chemistry simulation on fault-tolerant quantum computers. *New Journal of Physics*, 14(11):115023, 2012.

[118] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Phys. Rev. X*, 6:031007, Jul 2016.

[119] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, 2018.

[120] Tzu-Ching Yen, Vladyslav Verteletskyi, and Artur F Izmaylov. Measuring all compatible operators in one series of single-qubit measurements using unitary transformations. *Journal of chemical theory and computation*, 16(4):2400–2409, 2020.

[121] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F Izmaylov. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *The Journal of chemical physics*, 152(12):124114, 2020.

[122] Ophelia Crawford, Barnaby van Straaten, Daochen Wang, Thomas Parks, Earl Campbell, and Stephen Brierley. Efficient quantum measurement of pauli operators in the presence of finite sampling error. *Quantum*, 5:385, 2021.

[123] Wim Lavrijsen, Ana Tudor, Juliane Müller, Costin Iancu, and Wibe de Jong. Classical optimizers for noisy intermediate-scale quantum devices. *arXiv preprint arXiv:2004.03004*, 2020.

[124] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.

[125] Abhinav Kandala, Kristan Temme, Antonio D Corcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta. Extending the computational reach of a noisy superconducting quantum processor. *arXiv preprint arXiv:1805.04492*, 2018.

[126] Marc Ganzhorn, Daniel J Egger, P Barkoutsos, Pauline Ollitrault, Gian Salis, Nikolaj Moll, M Roth, A Fuhrer, P Mueller, S Woerner, et al. Gate-efficient simulation of molecular eigenstates on a quantum computer. *Physical Review Applied*, 11(4):044092, 2019.

[127] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, Stefan Filipp, and Ivano Tavernelli. Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions. *Phys. Rev. A*, 98:022322, Aug 2018.

[128] Bryan T Gard, Linghua Zhu, George S Barron, Nicholas J Mayhall, Sophia E Economou, and Edwin Barnes. Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. *npj Quantum Information*, 6(1):1–9, 2020.

[129] Abhinav Kandala, Kristan Temme, Antonio D Córcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749):491–495, 2019.

[130] Lennart Bittel and Martin Kliesch. Training variational quantum algorithms is np-hard–even for logarithmically many qubits and free fermionic systems. *arXiv preprint arXiv:2101.07267*, 2021.

[131] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):1–6, 2018.

[132] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. Noise-induced barren plateaus in variational quantum algorithms. *arXiv preprint arXiv:2007.14384*, 2020.

[133] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost-function-dependent barren plateaus in shallow quantum neural networks. *arXiv preprint arXiv:2001.00550*, 2020.

[134] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *arXiv preprint arXiv:2011.00027*, 2020.

[135] Joonho Lee, William J Huggins, Martin Head-Gordon, and K Birgitta Whaley. Generalized unitary coupled cluster wave functions for quantum computation. *Journal of chemical theory and computation*, 15(1):311–324, 2018.

[136] Harper R Grimsley, Daniel Claudino, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. Is the trotterized uccsd ansatz chemically well-defined? *Journal of Chemical Theory and Computation*, 2019.

[137] Gaurav Harsha, Toru Shiozaki, and Gustavo E Scuseria. On the difference between variational and unitary coupled cluster theories. *The Journal of chemical physics*, 148(4):044107, 2018.

[138] Nicholas P. Bauman, Jaroslav Chládek, Libor Veis, Jiří Pittner, and Karol Kowalski. Variational quantum eigensolver for approximate diagonalization of down-folded hamiltonians using generalized unitary coupled cluster ansatz. *arXiv preprint arXiv:2011.01985*, 2020.

[139] Pierre-Luc Dallaire-Demers, Jonathan Romero, Libor Veis, Sukin Sim, and Alán Aspuru-Guzik. Low-depth circuit ansatz for preparing correlated fermionic states on a quantum computer. *Quantum Science and Technology*, 4(4):045005, 2019.

[140] Igor O Sokolov, Panagiotis Kl Barkoutsos, Pauline J Ollitrault, Donny Greenberg, Julia Rice, Marco Pistoia, and Ivano Tavernelli. Quantum orbital-optimized unitary coupled cluster methods in the strongly correlated regime: Can quantum algorithms outperform their classical equivalents? *The Journal of Chemical Physics*, 152(12):124107, 2020.

[141] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature communications*, 10(1):1–9, 2019.

[142] Sukin Sim, Jonathan Romero, Jerome F Gonthier, and Alexander A Kunitsa. Adaptive pruning-based optimization of parameterized quantum circuits. *arXiv preprint arXiv:2010.00629*, 2020.

[143] George S Barron, Bryan T Gard, Orien J Altman, Nicholas J Mayhall, Edwin Barnes, and Sophia E Economou. Preserving symmetries for variational quantum eigensolvers in the presence of noise. *arXiv preprint arXiv:2003.00171*, 2020.

[144] Yordan S Yordanov and Crispin HW Barnes. Implementation of a general single-qubit positive operator-valued measure on a circuit-based quantum computer. *Physical Review A*, 100(6):062317, 2019.

[145] Yordan S Yordanov and Crispin HW Barnes. Efficient quantum circuits for quantum computational chemistry. *arXiv preprint arXiv:2005.14475*, 2020.

[146] David A. Mazziotti. Exact two-body expansion of the many-particle wave function. *Phys. Rev. A*, 102:030802, Sep 2020.

[147] Marcel Nooijen. Can the eigenstates of a many-body hamiltonian be represented exactly using a general two-body cluster expansion? *Phys. Rev. Lett.*, 84:2108–2111, Mar 2000.

[148] L-A Wu and DA Lidar. Qubits as parafermions. *Journal of Mathematical Physics*, 43(9):4506–4525, 2002.

[149] Hugo V. Lepage, Aleksander A. Lasek, David R. M. Arvidsson-Shukur, and Crispin H. W. Barnes. Entanglement generation via power-of-swap operations between dynamic electron-spin qubits. *Phys. Rev. A*, 101:022329, Feb 2020.

[150] Qingfeng Wang, Ming Li, Christopher Monroe, and Yunseong Nam. Resource-optimized fermionic local-hamiltonian simulation on quantum computer for quantum chemistry. *arXiv preprint arXiv:2004.04151*, 2020.

[151] IBM. Qiskit. 2021.

[152] Alexander Cowtan, Silas Dilkes, Ross Duncan, Will Simmons, and Seyon Sivarajah. Phase gadget synthesis for shallow circuits. *arXiv preprint arXiv:1906.01734*, 2019.

[153] Alexander Cowtan, Will Simmons, and Ross Duncan. A generic compilation strategy for the unitary coupled cluster ansatz. *arXiv preprint arXiv:2007.10515*, 2020.

[154] Rongxin Xia and Sabre Kais. Qubit coupled cluster singles and doubles variational quantum eigensolver ansatz for electronic structure calculations. *Quantum Science and Technology*, 2020.

[155] Robert A Lang, Ilya G Ryabinkin, and Artur F Izmaylov. Iterative qubit coupled cluster method with involutory linear combinations of pauli products. *arXiv preprint arXiv:2002.05701*, 2020.

[156] Ho Lun Tang, Edwin Barnes, Harper R Grimsley, Nicholas J Mayhall, and Sophia E Economou. qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansatze on a quantum processor. *arXiv preprint arXiv:1911.10205*, 2019.

[157] Roger Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2013.

[158] Arthur G Rattew, Shaohan Hu, Marco Pistoia, Richard Chen, and Steve Wood. A domain-agnostic, noise-resistant evolutionary variational quantum eigensolver for hardware-efficient optimization in the hilbert space. *arXiv preprint arXiv:1910.09694*, 2019.

[159] Ilya G Ryabinkin, Robert A Lang, Scott N Genin, and Artur F Izmaylov. Iterative qubit coupled cluster approach with efficient screening of generators. *Journal of Chemical Theory and Computation*, 16(2):1055–1063, 2020.

[160] Jie Liu, Zhenyu Li, and Jinlong Yang. An efficient adaptive variational quantum solver of the schrodinger equation based on reduced density matrices. *arXiv preprint arXiv:2012.07047*, 2020.

[161] Yordan S Yordanov, V Armaos, Crispin HW Barnes, and David RM Arvidsson-Shukur. Iterative qubit-excitation based variational quantum eigensolver. *arXiv preprint arXiv:2011.10540*, 2020.

[162] Dewi Retno Sari Saputro and Purnami Widyaningsih. Limited memory broyden-fletcher-goldfarb-shanno (l-bfgs) method for the parameter estimation on geographically weighted ordinal logistic regression model (gwolr). 1868(1):040009, 2017.

[163] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 01 1965.

[164] BT Draine and Frank Bertoldi. Structure of stationary photodissociation fronts. *arXiv preprint astro-ph/9603032*, 1996.

[165] Yutaka Toyozawa. Theory of line-shapes of the exciton absorption bands. *Progress of Theoretical Physics*, 20(1):53–81, 1958.

[166] Jarrod R McClean, Zhang Jiang, Nicholas C Rubin, Ryan Babbush, and Hartmut Neven. Decoding quantum errors with subspace expansions. *Nature Communications*, 11(1):1–9, 2020.

[167] Jarrod R. McClean, Mollie E. Kimchi-Schwartz, Jonathan Carter, and Wibe A. de Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A*, 95:042308, Apr 2017.

[168] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X*, 8:011021, Feb 2018.

[169] Oscar Higgott, Daochen Wang, and Stephen Brierley. Variational quantum computation of excited states. *Quantum*, 3:156, 2019.

[170] Tyson Jones, Suguru Endo, Sam McArdle, Xiao Yuan, and Simon C. Benjamin. Variational quantum algorithms for discovering hamiltonian spectra. *Phys. Rev. A*, 99:062304, Jun 2019.

[171] Lukasz Cincio, Yiğit Subaşı, Andrew T Sornborger, and Patrick J Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, 2018.

[172] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

[173] Christian Kokail, Christine Maier, Rick van Bijnen, Tiff Brydges, Manoj K Joshi, Petar Jurcevic, Christine A Muschik, Pietro Silvi, Rainer Blatt, Christian F Roos, et al. Self-verifying variational quantum simulation of lattice models. *Nature*, 569(7756):355–360, 2019.

[174] Alexandre Choquette, Agustin Di Paolo, Panagiotis Kl Barkoutsos, David Sénéchal, Ivano Tavernelli, and Alexandre Blais. Quantum-optimal-control-inspired ansatz for variational quantum algorithms. *arXiv preprint arXiv:2008.01098*, 2020.

[175] Jarrod McClean, Nicholas Rubin, Kevin Sung, Ian David Kivlichan, Xavier Bonet-Monroig, Yudong Cao, Chengyu Dai, Eric Schuyler Fried, Craig Gidney, Brendan Gimby, et al. Openfermion: the electronic structure package for quantum computers. *Quantum Science and Technology*, 2020.

[176] Justin M Turney, Andrew C Simmonett, Robert M Parrish, Edward G Hohenstein, Francesco A Evangelista, Justin T Fermann, Benjamin J Mintz, Lori A Burns, Jeremiah J Wilke, Micah L Abrams, et al. Psi4: an open-source ab initio electronic structure program. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(4):556–565, 2012.

[177] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.

[178] Cleve Moler and Charles Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1):3–49, 2003.

Appendix A

Pauli string measurements on a quantum computer

The expectation value of a Pauli string (a product of Pauli operators) with respect to a qubit state $|\psi\rangle$ can be directly measured on a quantum computer by a basis rotation followed by a single-qubit measurement in the Z (computational) basis.

Let us first consider the simplest case of measuring the expectation value of a Z operator with respect to a normalized single-qubit state $|\psi_1\rangle = a|0\rangle + b|1\rangle$. The matrix representation of Z is

$$Z \equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{A.1}$$

The two eigenstates of Z are $|0\rangle$ with expectation value $+1$, and $|1\rangle$ with expectation value -1. Hence, the expectation value of Z is the probability to measure $|0\rangle$ minus the probability to measure $|1\rangle$:

$$\langle\psi_1|Z|\psi_1\rangle = |\langle\psi_1|0\rangle|^2 - |\langle\psi_1|1\rangle|^2 = |a|^2 - |b|^2. \tag{A.2}$$

The expectation values of an X or a Y Pauli operator can be measured similarly, by an additional x or y basis rotation, respectively. An x or a y basis rotation can be implemented by a *Had* gate or a $R_x(\frac{\pi}{2})$ rotation gate, respectively.

Let us now consider the case of measuring the expectation value of a single Pauli-Z operator acting on qubit 0 , Z_0, with respect to a 2-qubit state $|\psi_2\rangle$. Now there are two eigenvectors corresponding to each eigenvalue of Z_0: states $|00\rangle$ and $|01\rangle$ where qubit 0 is in state $|0\rangle$ with eigenvalue $+1$, and states $|10\rangle$ and $|11\rangle$ where qubit 0 is in state $|1\rangle$ with eigenvalue -1. Then, the expectation value is

$$\langle\psi_2|Z_0|\psi_2\rangle = |\langle\psi_2|00\rangle|^2 + |\langle\psi_2|10\rangle|^2 - |\langle\psi_2|01\rangle|^2 - |\langle\psi_2|11\rangle|^2 \tag{A.3}$$

This can be extended to the case of an n-qubit state, and operator Z_i, acting on qubit i, as

$$\langle \psi_n | Z_i | \psi_n \rangle = \sum |\langle \psi_n | \ldots 0_i \ldots \rangle|^2 - \sum |\langle \psi_n | \ldots 1_i \ldots \rangle|^2, \tag{A.4}$$

where each sum runs over all 2^{n-1} computational basis states with qubit i in state $|0\rangle$ or $|1\rangle$, respectively. The two sums of n-qubit measurement probabilities are equivalent to the single-qubit measurement probabilities

$$\sum |\langle \psi_n | \ldots 0_i \ldots \rangle|^2 = |\langle \psi_n | 0_i \rangle|^2 \quad \text{and} \tag{A.5}$$

$$\sum |\langle \psi_n | \ldots 1_i \ldots \rangle|^2 = |\langle \psi_n | 1_i \rangle|^2. \tag{A.6}$$

Therefore, the expectation value of a single Pauli-Z operator with respect to an n-qubit state can be obtained by single-qubit measurements on the i^{th} qubit only:

$$\langle \psi_n | Z_i | \psi_n \rangle = |\langle \psi_n | 0_i \rangle|^2 - |\langle \psi_n | 1_i \rangle|^2 \tag{A.7}$$

Let us now consider the more interesting case of measuring the expectation value of two Z operators with respect to a 2-qubit state, $|\psi_2\rangle$. The tensor product of the two Z operators is

$$Z_0 \otimes Z_1 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{A.8}$$

Again $Z_0 \otimes Z_1$ has two eigenvalues; $+1$, corresponding to states with even parity, $\{|00\rangle, |11\rangle\}$, and -1, corresponding to states with odd parity, $\{|10\rangle, |01\rangle\}$. Hence,

$$\langle \psi_2 | Z_0 \otimes Z_1 | \psi_2 \rangle = |\langle \psi_2 | 00 \rangle|^2 + P |\langle \psi_2 | 11 \rangle|^2 - |\langle \psi_2 | 10 \rangle|^2 - |\langle \psi_2 | 01 \rangle|^2. \tag{A.9}$$

This expectation value cannot be obtained directly by measurements on one qubit in the Z basis only, because the parity of the state is stored non-locally in the two qubits. However, if we apply a *CNOT* gate with control qubit 0 and target qubit 1, the parity of the state will be recorded on qubit 1. Thus, the expectation value will be the probability to measure qubit 1 in $|0\rangle$ minus the probability to measure qubit 1 in $|1\rangle$:

$$\langle \psi_2 | Z_0 \otimes Z_1 | \psi_2 \rangle = |\langle 0_1 | CNOT_{01} | \psi_2 \rangle|^2 - |\langle 1_1 | CNOT_{01} | \psi_2 \rangle|^2 \tag{A.10}$$
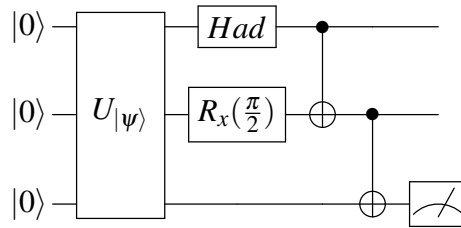
Fig. A.1 Circuit to measure the expectation value $\langle \psi | X_0 \otimes Y_1 \otimes Z_2 | \psi \rangle$. The unitary U_ψ prepares the state $|\psi\rangle$.

The eigenvalues of a string of l Pauli-Z operators, $Z_0 \otimes Z_1 \otimes \ldots \otimes Z_{l-1}$, are also either $+1$ or -1, corresponding to even and odd parity eigenstates, respectively. Therefore, the expectation value of $Z_0 \otimes Z_1 \otimes \ldots \otimes Z_{l-1}$, with respect to an n-qubit state $|\psi_n\rangle$ is the measurement probability that qubits $\{0, \ldots, l-1\}$ have even parity minus the measurement probability that qubits $\{0, \ldots, l-1\}$ have odd parity.

We can use a staircase of $(l-1)$ *CNOT*s to record the parity of qubits $\{0, \ldots, l-1\}$ on, say, qubit $l-1$ (see Sec. 2.1.2). In this way the expectation value $\langle \psi_n | Z_0 \otimes Z_1 \otimes \ldots \otimes Z_{l-1} | \psi_n \rangle$ can be obtained by determining the measurement probabilities for qubit $(l-1)$ only:

$$\langle \psi_n | Z_{q_0} \otimes Z_{q_1} \otimes \ldots \otimes Z_{q_{l-1}} | \psi_n \rangle = |\langle \psi_n | \bigotimes_{k=0}^{l-2} CNOT_{k,k+1} | 0_{l-1} \rangle|^2 - |\langle \psi_n | \bigotimes_{k=0}^{l-2} CNOT_{k,k+1} | 1_{l-1} \rangle|^2$$

$$(A.11)$$

The expectation value of any string of Pauli operators can be measured in a similar way, by additional single qubit basis rotations, applied before the *CNOT* staircase (see the example in Fig. A.1).

Appendix B

Pauli-string-exponential circuit

An exponential of a Pauli string can be directly implemented as a quantum circuit. To see how, let us first consider the case of constructing a circuit for $\exp\left[i\theta Z_0\right]$. The matrix representation of the action of this exponential on a single qubit state is

$$\exp\left[i\theta Z_0\right] = \sum_{k=0}^{\infty} \frac{(i\theta)^k}{k!} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}^k = \begin{bmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{bmatrix}, \tag{B.1}$$

which corresponds, up to a global phase, to a single-qubit z-rotation gate, $R_z(-2\theta)$.

Let us now consider $\exp\left[i\theta Z_0 Z_1\right]$. The matrix representation of the action of this exponential on the state of 2 qubits is

$$\exp\left[i\theta Z_0 Z_1\right] = \sum_{k=0}^{\infty} \frac{(i\theta)^k}{k!} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}^k = e^{i\theta} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-2i\theta} & 0 & 0 \\ 0 & 0 & e^{-2i\theta} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{B.2}$$

The odd parity states, $|10\rangle$ and $|01\rangle$, pick up a phase difference of $e^{-2i\theta}$ with respect to the even parity states, $|00\rangle$ and $|11\rangle$. This unitary action can be implemented by the circuit in Fig. B.1. The first *CNOT* computes and records the parity of the 2-qubit state on qubit 1. Hence, the $R_z(-2\theta)$ gate on qubit 1 adds a positive phase $e^{i\theta}$ to the states that had initial even parity, and negative phase $e^{-i\theta}$ to the states that had initial odd parity. The second *CNOT* recovers the original states, but with an additional phase of $e^{\pm i\theta}$ depending on their parity.

Following from Eq. (B.2), it is straightforward to see that the action of an exponential of a string of l Pauli Z operators, $\exp\left[i\theta Z_0 \dots Z_{l-1}\right]$, is to add a positive phase $e^{i\theta}$ to states of
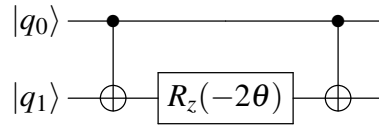
Fig. B.1 A circuit to implement the exponential $\exp\left[i\theta Z_0 Z_1\right]$

even parity, and a negative phase $e^{-i\theta}$ to states of odd parity:

$$\exp\left[i\theta Z_0 \ldots Z_{l-1}\right] = \sum_{x=0}^{2^l - 1} \exp\left[i(-1)^{Par(x)}\theta\right]|x\rangle\langle x|, \tag{B.3}$$

where $Par(x)$ denotes the parity of the binary representation of x. Therefore, the circuit in Fig. B.1 can be expanded as in Fig. B.2 to implemented $\exp\left[i\theta Z_0 \ldots Z_{l-1}\right]$. This type of a circuit, consisting of two *CNOT* staircases and a R_z rotation sandwiched in between them, is referred to as a *CNOT staircase construction*. The first sequence (staircase) of $(l-1)$ *CNOT*s computes and records the parity of the l-qubit state on qubit $l-1$. Hence, similarly to the case of 2-qubits, the $R_z(-2\theta)$ rotation adds a positive phase $e^{i\theta}$ to the states of even initial parity, and negative phase $e^{-i\theta}$ to the states of odd initial parity. The second *CNOT* staircase recovers the original states, with their corresponding phase shifts.
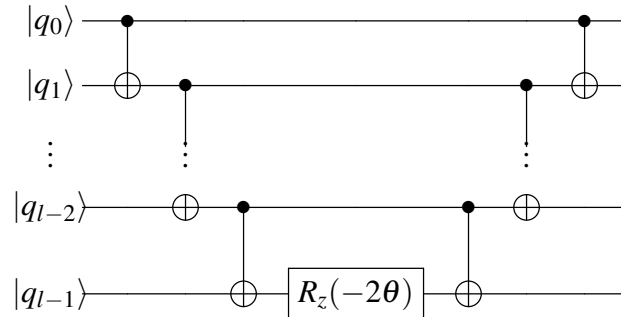
Fig. B.2 A circuit to implement the exponential $\exp\left[i\theta Z_0 \ldots Z_{l-1}\right]$. This type of a circuit is referred to as a *CNOT* staircase construction.

An exponential of an arbitrary Pauli string can be expressed in terms of an exponential of a Z-Pauli string as

$$\exp\left[i\theta X_0 Y_1 Z_2 \ldots\right] = \sum_{k=0}^{\infty} \frac{(i\theta)^k}{k!}(X_0 Y_1 Z_2 \ldots)^k =$$

$$\sum_{k=0}^{\infty} \frac{(i\theta)^k}{k!} Had_0 R_x(\frac{\pi}{2})_1 (Z_0 Z_1 Z_2 \ldots)^k R_x(\frac{-\pi}{2})_1 Had_0 =$$

$$Had_0 R_x(\frac{\pi}{2})_1 \exp\left[i\theta Z_0 Z_1 Z_2 \ldots\right] R_x(\frac{-\pi}{2})_1 Had_0. \tag{B.4}$$

Hence, a circuit implementing an exponential of an arbitrary Pauli string can be constructed by additional single-qubit basis rotations to the circuit in Fig. B.2. A circuit to implement the exponential $\exp\left[i\theta X_0 Y_1 Z_2\right]$ is given as an example in Fig. B.3

Fig. B.3 A circuit to implement the exponential $\exp\left[i\theta X_0 Y_1 Z_2\right]$.

Appendix C

The Cavendish Quantum Computational Chemistry package

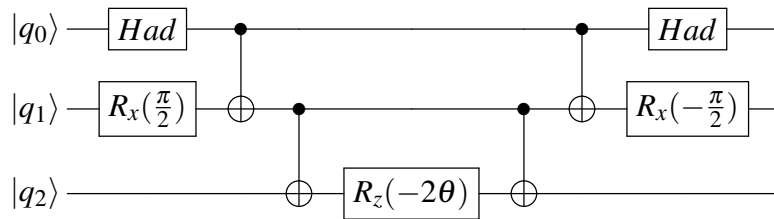The Cavendish Quantum Computational Chemistry (CQCC) package[1] is an in-house software that incorporates existing quantum computational packages, and a fast custom statevector simulator, into a platform intended to simulate and study of the performance of VQE protocols for electronic system modelling. The CQCC includes:

1. Transforming Hamiltonians of electronic systems to second-quantized form, and mapping them to quantum gate operator or matrix representation

2. Constructing custom ansätze in the form of QASM strings and unitary matrices

3. Constructing custom ansatz elements, e.g. qubit/fermionic evolutions and Pauli string evolutions, for ADAPT-VQE protocols

4. Running the VQE, given an electronic Hamiltonian, an ansatz and a classical optimizer.

5. Implementation of ADAPT-VQE protocols with a custom ansatz element pool, and an ansatz-growing strategy

6. Evaluating expectation values of unitary operators, using a fast custom statevector simulator, a quantum computer shot simulator, implemented by Qiskit [151], or one of IBM's cloud accessible quantum computing devices [44]

The various Hamiltonian transformations are implemented using the *openfermion*[175] and *psi4*[176] packages .

---

[1] Available at: https://github.com/JordanovSJ/VQE

The classical optimizers are implement by the *scipy*[177] package. For the majority of simulations presented in the thesis the gradient-descent BFGS [157] or the direct-search Nelder-Mead [163] methods are used. Optionally, gradient descent optimizers are supplied with a function that calculates the energy gradient vector of the ansatz. This function is described in Sec. C.2.

The CQCC offers three methods to evaluate the various expectation values required by VQE protocols: (1) a custom-built fast statevector simulator (see Sec. C.1), which explicitly evaluates the qubit statevector, and then calculates the exact expectation value as a matrix product; (2) a shot simulator, implemented by Qiskit [151], which simulates the operation of a quantum computer; and (3) a cloud connection to one of IBM's quantum computers. The simulations presented in the thesis are noiseless and are obtained with the statevector simulator.

C.1 The CQCC statevector simulator

Calculating qubit statevectors is the most computationally expensive and time consuming process in numerical classical VQE simulations. The statevector simulator, used by the CQCC, is specifically designed to efficiently calculate statevectors that are generated by ansätze that correspond to products of exponentials of (fermionic/qubit) excitation operators.

Let us consider an ansatz $U(\vec{\theta}) = \prod_{i=N_U}^{1} e^{\theta_i A_i}$, where $\{A_i\}$ are qubit/fermionic excitation operators. Excitation operators are anti-Hermitian operators so they satisfy

$$A_i^\dagger = -A_i. \tag{C.1}$$

Additionally, excitation operators satisfy the relation

$$A_i^3 = -A_i. \tag{C.2}$$

Let $|\psi(\vec{\theta})\rangle$ is a state generated by $U(\vec{\theta})$, acting on an initial reference state $|\psi_0\rangle$:

$$|\psi(\vec{\theta})\rangle = U(\vec{\theta})|\psi_0\rangle = \prod_{i=N_U}^{1} e^{\theta_i A_i}|\psi_0\rangle. \tag{C.3}$$

To calculate the statevector $|\psi(\vec{\theta})\rangle$, we need to calculate the N_U exponentials $\{e^{\theta_i A_i}\}$ in Eq. (C.3), and then multiply them sequentially to the statevector $|\psi_0\rangle$. For N_{so} qubits, each A_i is represented by an $2^{N_{so}} \times 2^{N_{so}}$-dimensional matrix. Hence, the complexity of directly calculating each exponential $e^{\theta_i A_i}$ is $O(2^{3N_{so}})$ [178]. However, we can make use of relation

(C.2) and write each exponential in Eq. (C.3) as

$$e^{\theta_i A_i} = \sum_{k=0}^{\infty} \frac{\theta_i^k A_i^k}{k!} = I + \sum_{k_1=0}^{\infty} \frac{(-1)^{k_1} \theta_i^{2k_1+1}}{(2k_1+1)!} A_i + \sum_{k_2=1}^{\infty} \frac{(-1)^{k_2} \theta_i^{2k_2}}{(2k_2)!} A_i^2 = \qquad (C.4)$$

$$I + \sin\theta A_i + (1 - \cos\theta) A_i^2. \qquad (C.5)$$

The operators $\{A_i\}$ are fixed throughout a simulation. Therefore, if we compute in advance and store the matrix representations of each A_i and A_i^2, we can evaluate the expression in Eq. (C.4) by performing matrix addition only, which has a complexity of $O(2^{2N_{so}})$. Hence, the calculation of $|\psi(\vec{\theta})\rangle$, requires N_U matrix-to-vector multiplications and N_U matrix additions, which have a total complexity of $O(N_U 2^{2N_{MO}})$.

Using this method, the statevector simulator of the CQCC is approximately 20 faster than the Qiskit statevector simulator for 14 qubits (BeH$_2$). However, the method has a drawback in that it needs to store the matrices for all A_i and A_i^2 operators. For example, the most memory demanding simulation in this work, running the qubit-ADAPT-VQE for BeH$_2$, required around 2GB of RAM to store the matrices for all Pauli string operators, and their squares, that define the ansatz element pool of the qubit-ADAPT-VQE (see Sec. 4.2.2).

C.2   Energy-gradient vector of an ansatz

When using a gradient-descent minimizer, e.g. the BFGS [157], we have the option to supply a method that returns the gradient vector of the minimized function. If the minimized function is differentiable, supplying a gradient vector guarantees a faster optimization of the variational parameters.

Let us consider minimizing the Hamiltonian expectation value $E(\vec{\theta}) = \langle \psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle$, where the statevector is

$$|\psi(\vec{\theta})\rangle = U(\vec{\theta})|\psi_0\rangle = \prod_{i=N_U}^{1} e^{\theta_i A_i}|\psi_0\rangle. \qquad (C.6)$$

Similarly to above, $\{A_i\}$ are skew-hermitian operators that correspond to qubit/fermionic excitation operators or Pauli strings. The i^{th} component of the energy gradient vector, $\nabla E(\vec{\theta})$,

can be expressed as

$$\mathbf{\nabla}_i E(\vec{\theta}) = \frac{\partial}{\partial \theta_i} \langle \psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle =$$

$$\frac{\partial}{\partial \theta_i} \langle \psi_0| \prod_1^{k_1=N_U} e^{\theta_{k_1} A_{k_1}^\dagger} H \prod_{k_2=N_U}^1 e^{\theta_{k_2} A_{k_2}} |\psi_0\rangle =$$

$$\langle \psi(\vec{\theta})|H \prod_{k_1=N_u}^{i+1} e^{\theta_{k_1} A_{k_1}} A_i \prod_{k_1=i}^1 e^{\theta_{k_2} A_{k_2}} |\psi_0\rangle + \langle \psi_0| \prod_{k_1=1}^i e^{\theta_{k_1} A_{k_1}^\dagger} A_i^\dagger \prod_{k_1=i+1}^{N_U} e^{\theta_{k_2} A_{k_2}^\dagger} |H|\psi(\vec{\theta})\rangle =$$

$$2\langle \alpha_i(\vec{\theta})|A_i|\beta_i(\vec{\theta})\rangle, \tag{C.7}$$

where

$$|\beta_i(\vec{\theta})\rangle = \prod_{k=i}^1 e^{\theta_k A_k} |\psi_0\rangle \text{ and} \tag{C.8}$$

$$|\alpha_i(\vec{\theta})\rangle = \prod_{k=i+1}^{N_u} e^{\theta_k A_k^\dagger} H |\psi(\vec{\theta})\rangle. \tag{C.9}$$

The N_U components of $\mathbf{\nabla}E(\boldsymbol{\theta})$ can be calculated with minimum number of matrix multiplications by updating $|\beta_i(\vec{\theta})\rangle$ and $|\alpha_i(\vec{\theta})\rangle$ in the following way:

1. For $i = N_U$, initiate

$$|\alpha_{N_U}(\vec{\theta})\rangle = H|\psi(\vec{\theta})\rangle \text{ and} \tag{C.10}$$

$$|\beta_{N_U}(\vec{\theta})\rangle = |\psi(\vec{\theta})\rangle \tag{C.11}$$

2. For $1 < i < N_U$, update

$$|\alpha_{i-1}(\vec{\theta})\rangle = e^{\theta_i A_i^\dagger} |\alpha_i(\vec{\theta})\rangle \text{ and} \tag{C.12}$$

$$|\beta_{i-1}(\vec{\theta})\rangle = (e^{\theta_i A_i})^{-1}|\beta_i(\vec{\theta})\rangle = e^{\theta_i A_i^\dagger}|\beta_i(\vec{\theta})\rangle, \tag{C.13}$$

where in Eq. (C.13), we use that A_i is skew-Hermitian ($A_i^\dagger = -A_i$).

Let us assume that we have already computed and stored the matrices of the exponentials $\{e^{\theta_i A_i}\}$, when calculating the statevector $|\psi(\vec{\theta})\rangle$. Hence, to calculate each component of $\mathbf{\nabla}E(\vec{\theta})$ we need to perform 3 matrix-to-vector multiplications. Thus, overall to calculate $\mathbf{\nabla}E(\vec{\theta})$ we need to perform $3N_U$ matrix-to-vector multiplications, resulting in a total cost of $O(3N_U 2^{2N_{MO}})$ operations.

The cost of calculating $\mathbf{\nabla}E(\vec{\theta})$ is about 3 times the cost of calculating $|\psi(\vec{\theta})\rangle$. However, using the energy-gradient vector in the optimization subroutine of the VQE, reduces the

number of VQE iterations by at least an order of magnitude, which justifies the use of the energy-gradient vector.

The energy-gradient vector of an ansatz can also be calculated using a quantum computer. Each component of $\boldsymbol{\nabla} E(\vec{\theta})$ can be obtained by explicitly measuring the expectation value in Eq. (C.7). The most efficient way to do this, in terms of circuit-depth, is to measure the overlap of $\left(\prod_{k=i+1}^{N_u} e^{\theta_k A_k^\dagger} \right)^\dagger A_i \prod_{k=i}^{1} e^{\theta_k A_k} |\psi_0\rangle$ and $H|\psi(\vec{\theta})\rangle$, using the SWAP test [171]. This would require a circuit with maximum depth roughly equal to the depth of the ansatz circuit generating $|\psi(\vec{\theta})\rangle$, and $2N_{so}$ qubits.

Appendix D

Additional results and simulation data

D.1   Energy convergence data

	iteration	E	element	var_parameters
0	1	-7.877119623435884	d_q_exc_[2, 3]_[10, 11]	-0.11031109472387213
1	2	-7.878915001650162	d_q_exc_[2, 3]_[5, 10]	0.05700525215310389
2	3	-7.880581566203068	d_q_exc_[2, 3]_[4, 11]	-0.05737310731341327
3	4	-7.880668708374632	d_q_exc_[0, 1]_[4, 5]	-0.003928578568135717
4	5	-7.8812514499854585	d_q_exc_[2, 3]_[6, 7]	-0.028015551814893413
5	6	-7.881806943859643	d_q_exc_[2, 3]_[8, 9]	-0.028026661663886527
6	7	-7.882217915759775	d_q_exc_[2, 3]_[4, 5]	-0.03138833150495772
7	8	-7.882233950748895	d_q_exc_[0, 1]_[8, 9]	-0.001959697647985118
8	9	-7.882249868462785	d_q_exc_[0, 1]_[6, 7]	-0.001959583309577129
9	10	-7.882270611904488	d_q_exc_[1, 2]_[6, 7]	-0.0035499803675991416
10	11	-7.882291346534323	d_q_exc_[0, 3]_[8, 9]	0.0034230829074340914
11	12	-7.882311595621427	d_q_exc_[0, 3]_[6, 7]	0.003423409616028284
12	13	-7.882331839546545	d_q_exc_[1, 2]_[8, 9]	-0.003550533798326349
13	14	-7.88233763339666	d_q_exc_[0, 1]_[10, 11]	-0.0011497317831317589
14	15	-7.882541770888251	s_q_exc_2_4	-0.03850398164110951
15	16	-7.882724374672856	s_q_exc_3_5	0.03883460917281597
16	17	-7.882739715836327	s_q_exc_2_10	0.005798213133724907
17	18	-7.8827420181984005	d_q_exc_[1, 2]_[4, 11]	0.0009570729120199658
18	19	-7.8827443222474765	d_q_exc_[0, 3]_[5, 10]	0.000956659940102136
19	20	-7.882746119842321	d_q_exc_[1, 3]_[5, 11]	0.013899535004236947
20	21	-7.882747912380555	d_q_exc_[0, 2]_[4, 10]	0.0029821381122755036
21	22	-7.882748761684853	s_q_exc_1_11	-0.0005236492575660854
22	23	-7.882749474720335	s_q_exc_0_10	0.0005144177367431057
23	24	-7.882749790026824	d_q_exc_[1, 2]_[10, 11]	-0.0002927907900392912
24	25	-7.88275005649819	d_q_exc_[0, 3]_[10, 11]	0.0002970333290100582
25	26	-7.882751665773464	s_q_exc_3_11	-0.002114870567031991
26	27	-7.882751834270224	s_q_exc_1_5	-0.0002912486350685454
27	28	-7.882751886532499	d_q_exc_[0, 1]_[4, 11]	-4.6102121494920244e-05
28	29	-7.882752009391789	d_q_exc_[1, 2]_[4, 5]	-0.008651612377156632
29	30	-7.882752058562675	d_q_exc_[0, 1]_[5, 10]	0.0003097995650207747
30	31	-7.8827521621056755	d_q_exc_[0, 3]_[4, 5]	0.011379025673876449
31	32	-7.88275226100295	s_q_exc_0_4	-0.009221817747341363

Table D.1 Simulation data for an energy convergence plot for LiH. The simulation parameters are summarised in the table below. Pairs of rows of the same colour indicate spin-complement pairs of qubit-evolutions that are appended at either adjacent or close positions in the ansatz.

molecule	SO-basis	bond distance	state	method	method params
LiH	STO-3G	$r_{Li-H} = 1.546$	ground	QEB-ADAPT	$n_{qe} = 1$, step 5 NOT performed

	iteration	E	element	var_parameters
0	1	-7.877119623435893	d_q_exc_[2, 3]_[10, 11]	-0.11033939450288231
1	2	-7.880581566203079	d_q_exc_[2, 3]_[5, 10]	0.05698398343753387
2	2	-7.880581566203079	d_q_exc_[3, 2]_[4, 11]	-0.05732520809517464
3	3	-7.880668708374645	d_q_exc_[0, 1]_[4, 5]	-0.003915773378408674
4	4	-7.881251449985466	d_q_exc_[2, 3]_[8, 9]	-0.028010049292039194
5	5	-7.881806943859642	d_q_exc_[2, 3]_[6, 7]	-0.028021169417229225
6	6	-7.882217915759778	d_q_exc_[2, 3]_[4, 5]	-0.03139017108753022
7	7	-7.882233950748927	d_q_exc_[0, 1]_[6, 7]	-0.0021217090274751787
8	8	-7.882249868462803	d_q_exc_[0, 1]_[8, 9]	-0.0021213195161 93217
9	9	-7.8822911540293825	d_q_exc_[0, 3]_[8, 9]	0.003421147814629044
10	9	-7.8822911540293825	d_q_exc_[1, 2]_[9, 8]	-0.0035506035220592117
11	10	-7.882331839546547	d_q_exc_[0, 3]_[6, 7]	0.00342135193510341
12	10	-7.882331839546547	d_q_exc_[1, 2]_[7, 6]	-0.003551072945642547
13	11	-7.882337633396636	d_q_exc_[0, 1]_[10, 11]	-0.0012290844228192265
14	12	-7.882724374672847	s_q_exc_2_4	-0.038528581457990416
15	12	-7.882724374672847	s_q_exc_3_5	0.038979711049612066
16	13	-7.882741306974246	s_q_exc_2_10	0.005755202876815314
17	13	-7.882741306974246	s_q_exc_3_11	-0.001639538099921524
18	14	-7.882745931382006	d_q_exc_[1, 2]_[4, 11]	0.0009570458644833416
19	14	-7.882745931382006	d_q_exc_[0, 3]_[5, 10]	0.00095690669351 67765
20	15	-7.8827495281959985	d_q_exc_[1, 3]_[5, 11]	-0.008857527398217522
21	15	-7.8827495281959985	d_q_exc_[0, 2]_[4, 10]	0.003423331219033033
22	16	-7.8827510818421995	s_q_exc_1_11	-0.0005276571447948891
23	16	-7.8827510818421995	s_q_exc_0_10	0.0005196832728877846
24	17	-7.882751664140968	d_q_exc_[1, 2]_[10, 11]	-0.0002950990563134189
25	17	-7.882751664140968	d_q_exc_[0, 3]_[11, 10]	0.0002984848527454705
26	18	-7.882751942117987	s_q_exc_1_5	-0.016264238078149616
27	18	-7.882751942117987	s_q_exc_0_4	-0.011827197819832056
28	19	-7.882752044478731	d_q_exc_[0, 1]_[4, 11]	-9.446949692716319e-05
29	19	-7.882752044478731	d_q_exc_[1, 0]_[5, 10]	0.00045161967586454315
30	20	-7.882752261018141	d_q_exc_[1, 2]_[4, 5]	-0.008672994688781101
31	20	-7.882752261018141	d_q_exc_[0, 3]_[5, 4]	-0.0002443567211394347
32	21	-7.8827523380554085	d_q_exc_[0, 3]_[4, 11]	-0.000381069865828933
33	21	-7.8827523380554085	d_q_exc_[1, 2]_[5, 10]	5.7151669729966596e-05

Table D.2 Simulation data for an energy convergence plot for LiH. The simulation parameters are summarised in the table below.

molecule	SO-basis	bond distance	state	method	method params
LiH	STO-3G	$r_{Li-H} = 1.546$	ground	QEB-ADAPT	$n_{qe} = 1$, step 5 performed

	iteration	E	element	var_parameters
1	1	-15.566756541415858	spin_d_q_exc_[4, 5]_[10, 11]	-0.0365422971177455
2	2	-15.569262535720876	spin_d_q_exc_[4, 5]_[12, 13]	-0.010788079384711477
3	3	-15.572885374947806	spin_d_q_exc_[2, 3]_[10, 11]	-0.05769017655578784
4	4	-15.575574310943724	spin_d_q_exc_[2, 3]_[8, 9]	-0.029069596900987475
5	5	-15.578106252821492	spin_d_q_exc_[2, 3]_[6, 7]	-0.029122216025795986
6	6	-15.582672418985373	spin_d_q_exc_[3, 4]_[11, 12]	-0.7103254767463593
7	7	-15.583820316263392	spin_d_q_exc_[2, 3]_[12, 13]	-0.0584229021059216
8	8	-15.588907810928106	spin_d_q_exc_[2, 5]_[10, 13]	-0.04614373906725295
9	9	-15.591462440717539	spin_d_q_exc_[2, 5]_[11, 12]	0.05957457991923902
10	10	-15.594260856812149	spin_d_q_exc_[3, 4]_[10, 13]	0.059524516230038925
11	11	-15.594302831067278	spin_d_q_exc_[0, 1]_[10, 11]	-0.0005412304122581799
12	12	-15.594338503692283	spin_d_q_exc_[0, 1]_[12, 13]	-0.00045017692821196724
13	13	-15.594358985691448	spin_d_q_exc_[0, 1]_[6, 7]	-0.0008794528673684362
14	14	-15.594379267137636	spin_d_q_exc_[0, 1]_[8, 9]	-0.0008784845650083714
15	15	-15.594423360314249	spin_d_q_exc_[1, 2]_[8, 9]	0.0038771748286376296
16	16	-15.594466984266408	spin_d_q_exc_[1, 2]_[6, 7]	0.0038357052194273222
17	17	-15.594509469940203	spin_d_q_exc_[0, 3]_[6, 7]	-0.00011680454200448313
18	18	-15.594551506809633	spin_d_q_exc_[0, 3]_[8, 9]	-0.00012006913743716484
19	19	-15.59459978746691	spin_d_q_exc_[3, 5]_[11, 13]	0.9091534952632134
20	20	-15.5946516602168	spin_d_q_exc_[2, 4]_[10, 12]	0.019832660453935363
21	21	-15.594660647475756	spin_d_q_exc_[0, 3]_[12, 13]	0.002256746002918719
22	22	-15.594669557248602	spin_d_q_exc_[1, 2]_[12, 13]	-0.021407917876685743
23	23	-15.594746643416066	spin_s_q_exc_10_2	-0.019205764184440914
24	24	-15.594822475160214	spin_s_q_exc_11_3	-0.0030174286489695183
25	25	-15.59482545054788	spin_d_q_exc_[0, 3]_[10, 11]	-0.026617575562520097
26	26	-15.594828399050765	spin_d_q_exc_[1, 2]_[10, 11]	-0.0173987187396723 03
27	27	-15.594842129303547	spin_s_q_exc_13_5	-0.008649017103510407
28	28	-15.594855756479141	spin_s_q_exc_12_4	0.00633965849679309
29	29	-15.594859484903104	spin_d_q_exc_[4, 5]_[6, 7]	-0.010677556956574902
30	30	-15.594863000847486	spin_d_q_exc_[4, 5]_[8, 9]	-0.010644681862704633
31	31	-15.594864375801164	spin_d_q_exc_[0, 5]_[10, 13]	0.001434854392674299
32	32	-15.594865564156931	spin_d_q_exc_[0, 5]_[11, 12]	0.0032172172894219388

Table D.3 Simulation data for an energy convergence plot for BeH_2. The simulation parameters are summarised in the table below. Pairs of rows of the same colour indicate spin-complement pairs of qubit-evolutions that are appended at either or adjacent or close positions in the ansatz.

molecule	SO-basis	bond distance	state	method	method params
BeH_2	STO-3G	$r_{Li-H} = 1.316$	ground	QEB-ADAPT	$n_{qe} = 1$, step 5 NOT performed

D.2 Energy convergence plots
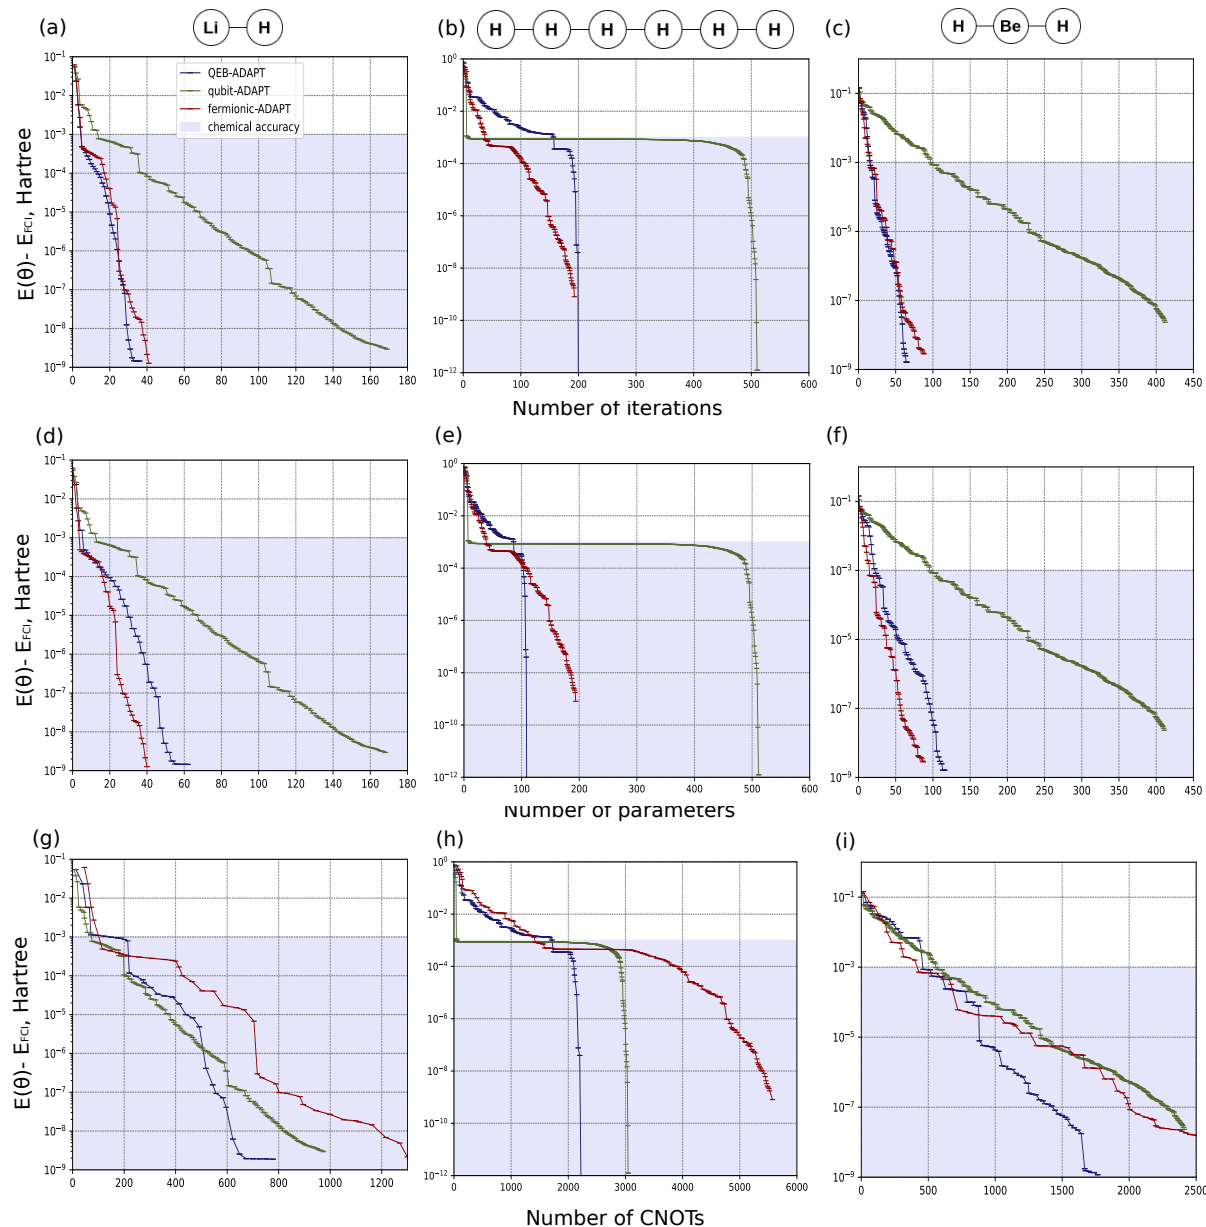

Fig. D.1 Energy convergence plots for the ground states of LiH, H_6 and BeH_2 in the STO-3G orbital basis set, at bond distances $r_{\text{Li-H}} = 3\text{Å}$, $r_{H-H} = 3\text{Å}$ and $r_{\text{Be-H}} = 3\text{Å}$. **Blue** plots are obtained with the QEB-ADAPT-VQE for $n_q e = 1$. **Red** plots are obtained with the fermionic-ADAPT-VQE. **Green** plots are obtained with the qubit-ADAPT-VQE. All convergence plots are terminated for energy decrease threshold of $\varepsilon = 10^{-12}$. **a,b,c** Accuracy as function of ansatz-constructing iterations. **d,e,f** Accuracy as function of ansatz parameters (the number of parameters is the same as the number of ansatz-constructing for the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE). **g,h,i** Accuracy as function of ansatz circuit *CNOT* count. The *CNOT* counts are calculated assuming the use of the circuits derived in Chapter 2.

Appendix E

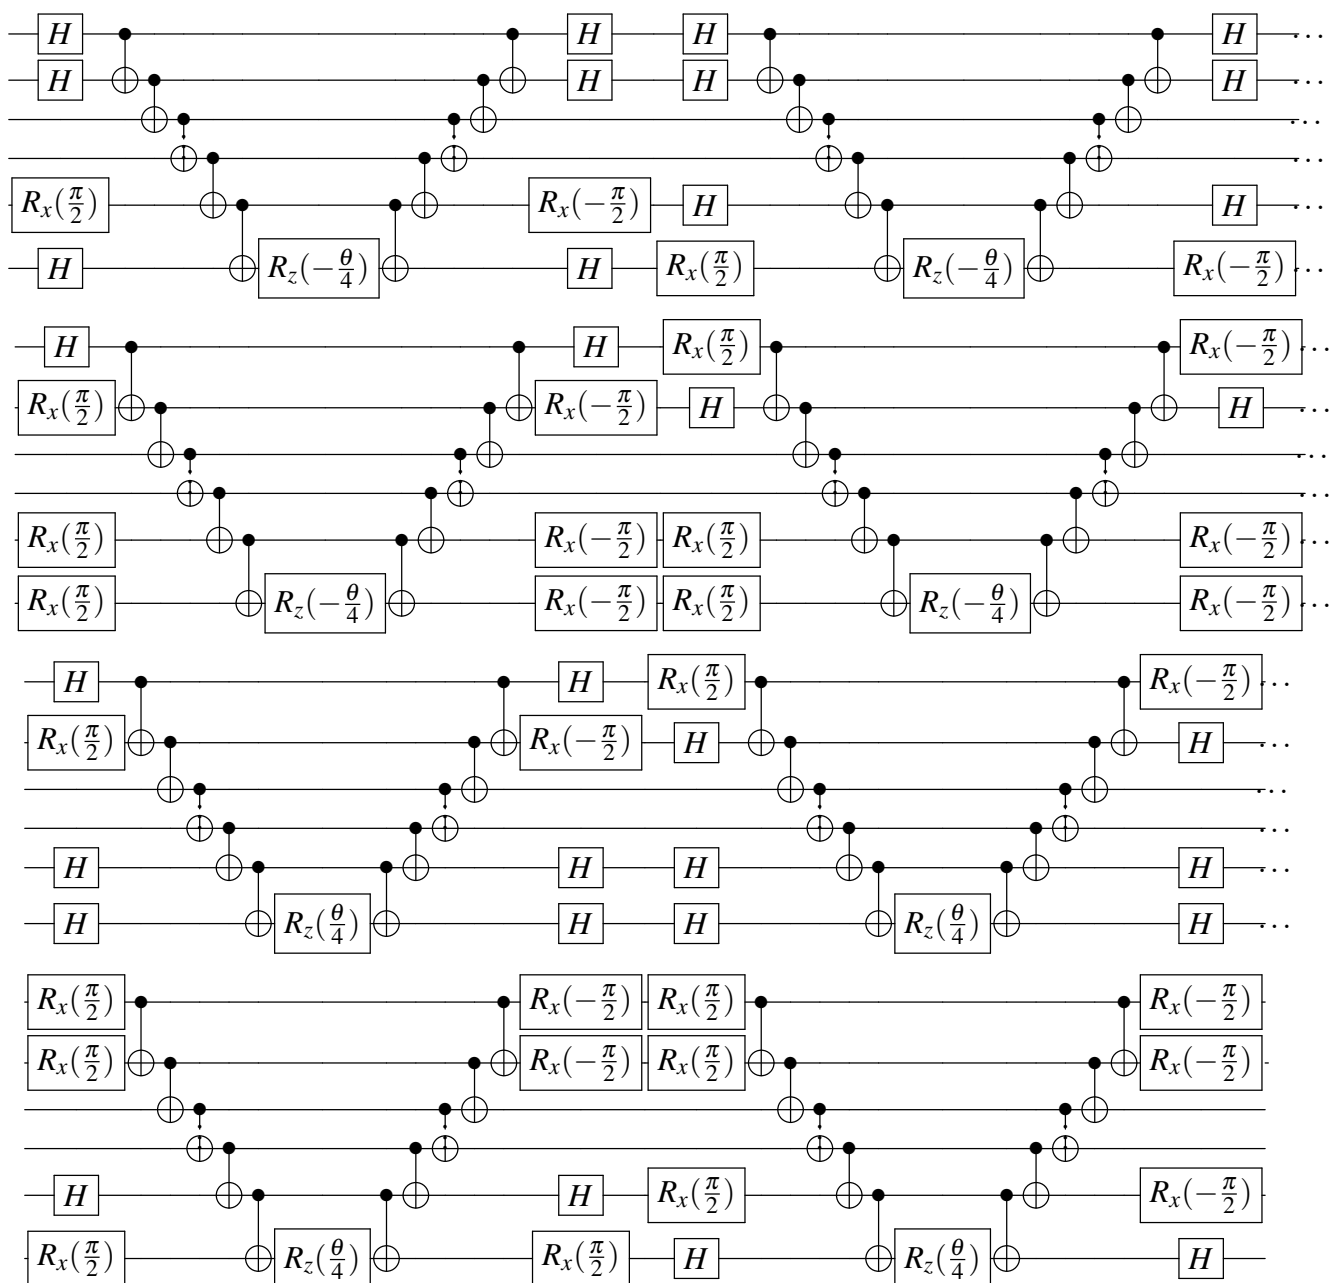# Double-fermionic-evolution circuit

Fig. E.1 A standard circuit to perform a double fermionic evolution [Eq. (2.5)].