

Quantum Chemistry in the Age of Quantum Computing

Yudong Cao,^{†,‡} Jonathan Romero,^{†,‡} Jonathan P. Olson,^{†,‡} Matthias Degroote,^{†,¶,§} Peter D. Johnson,^{†,‡} Mária Kieferová,^{||,⊥,‡} Ian D. Kivlichan,^{#,†} Tim Menke,^{#,Ⓜ,△} Borja Peropadre,[‡] Nicolas P. D. Sawaya,[▽] Sukin Sim,^{†,‡} Libor Veis,^{††} and Alán Aspuru-Guzik^{*,†,‡,¶,§,‡‡,¶¶}

[†]*Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA 02138, USA*

[‡]*Zapata Computing Inc., Cambridge, MA 02139, USA*

[¶]*Department of Chemistry, University of Toronto, Toronto, Ontario M5G 1Z8, Canada*

[§]*Department of Computer Science, University of Toronto, Toronto, Ontario M5G 1Z8, Canada*

^{||}*Department of Physics and Astronomy, Macquarie University, Sydney, NSW, 2109, Australia*

[⊥]*Institute for Quantum Computing and Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

[#]*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

[Ⓜ]*Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

[△]*Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

[▽]*Intel Labs, Intel Corporation, Santa Clara, CA 95054 USA*

^{††}*J. Heyrovský Institute of Physical Chemistry, Academy of Sciences of the Czech Republic v.v.i., Dolejškova 3, 18223 Prague 8, Czech Republic*

^{‡‡}*Canadian Institute for Advanced Research, Toronto, Ontario M5G 1Z8, Canada*

^{¶¶}*Vector Institute for Artificial Intelligence, Toronto, Ontario M5S 1M1, Canada*

E-mail: alan@aspuru.com

Abstract

Practical challenges in simulating quantum systems on classical computers have been widely recognized in the quantum physics and quantum chemistry communities over the past century. Although many approximation methods have been introduced, the complexity of quantum mechanics remains hard to appease. The advent of quantum computation brings new pathways to navigate this challenging complexity landscape. By manipulating quantum states of matter and taking advantage of their unique features such as superposition and entanglement, quantum computers promise to efficiently deliver accurate results for many important problems in quantum chemistry such as the electronic structure of molecules.

In the past two decades significant advances have been made in developing algorithms and physical hardware for quantum computing, heralding a revolution in simulation of quantum systems. This article is an overview of the algorithms and results that are relevant for quantum chemistry. The intended audience is both quantum chemists who seek to learn more about quantum computing, and quantum computing researchers who would like to explore applications in quantum chemistry.

Table of Contents

1	Introduction and historical overview	5
2	Quantum chemistry in the age of quantum computing	11
2.1	Basics and challenges of classical quantum chemistry	13
2.2	Classical approximation techniques and their limitations	16
2.2.1	Static properties: wave functions and energy spectra	16
2.2.2	Dynamical properties: time evolution of the wave function	21
2.3	Going beyond classical limitations using quantum computers	26
2.3.1	Statics: phase estimation and variational quantum eigensolvers	29
2.3.2	Dynamics: Hamiltonian simulation and hybrid methods	31

2.4	Chapter 2 Glossary	34
3	Computational complexity	35
3.1	Complexity classes	36
3.2	Complexity theory of the electronic structure problem	39
3.3	Complexity theory in molecular vibronics	41
3.4	Chapter 3 Glossary	45
4	Quantum simulation algorithms for fault-tolerant quantum computers	46
4.1	Quantum algorithms for energy estimation	49
4.1.1	General Hamiltonian simulation	50
4.1.2	Quantum phase estimation	57
4.1.3	Quantum Fourier transform	60
4.2	Determining the ground states energies of quantum chemistry Hamiltonians	62
4.2.1	State preparation	63
4.2.2	Hamiltonian simulation for quantum chemistry	65
4.2.3	Measuring properties of many-body systems	70
4.3	Chapter 4 Glossary	73
5	Quantum algorithms for noisy intermediate-scale quantum devices	74
5.1	Variational quantum eigensolver (VQE) algorithm	76
5.1.1	Ansätze for state preparation	78
5.1.2	Energy estimation in VQE	100
5.1.3	Optimization methods	107
5.1.4	Physical realizations of VQE	112
5.1.5	VQE for excited states	113
5.1.6	Calculation of molecular properties with VQE	116
5.2	Other hybrid quantum-classical algorithms for chemistry	117
5.2.1	Variational quantum simulator	117

5.2.2	Imaginary-time variational quantum simulator	118
5.2.3	Imaginary-time variational quantum simulator for excited states . . .	119
5.3	Non-gate-based methods for quantum chemistry	120
5.3.1	Adiabatic	120
5.3.2	Linear optics	122
5.4	Chapter 5 Glossary	125
6	Summary and outlook	127
	Acknowledgements	131
A	Quantum chemistry basis sets	132
B	Mappings to qubits	137
B.1	Jordan-Wigner mapping	138
B.2	Bravyi-Kitaev mapping	139
C	From quantum chemistry to quantum computation: example of molecular hydrogen	141
C.1	Introduction	141
C.2	Defining the chemistry problem	141
C.3	Mapping the problem	143
C.4	A brief introduction to quantum computation	145
C.5	Variational quantum eigensolver for quantum chemistry	146
C.6	Appendix Glossary	150
	References	153

1 Introduction and historical overview

Since the advent of electronic computers in the last century, computation has played a fundamental role in the development of chemistry. Numerical methods for computing the static and dynamic properties of chemicals have revolutionized chemistry as a discipline. With the recent emergence of quantum computing, there is potential for a similarly disruptive progress. This review provides a comprehensive picture of the current landscape of quantum computation for chemistry to help establish a perspective for future progress.

Where can quantum computers help chemistry? As our understanding of quantum computers continues to mature, so too will the development of new methods and techniques which can benefit chemistry. For now, at least, we are confident that quantum computers can aid those quantum chemistry computations that require an explicit representation of the wave function, either because of a high accuracy requirement of simulated properties or because of a high degree of entanglement in the system. In these cases, the exponential growth of the dimension of the wave function makes manipulation and storage very inefficient on a classical computer. Indeed, for even moderately large systems it is already intractable to explicitly maintain the full wave function.

As both the quantum hardware and software communities continue to make rapid progress, the immediate role of quantum computing for quantum chemistry becomes clearer. While it seems that there will be no shortage in the demand for new classical methods, the replacement of specific classical subroutines with quantum computations can improve the accuracy and tractability of chemical predictions. This also brings up a new challenge: advancing quantum algorithms for quantum chemistry requires the synergy of quantum information theory and classical quantum chemistry techniques. While these two areas are closely related, they have evolved following different disciplines—the former from physics and computer science, and the latter from chemistry. In this review, our aim is to be pedagogical to scientists from each background in order to bridge this gap and stimulate new developments. We will describe how quantum computing methods can be used to replace or augment classical quantum

chemistry methods and survey an array of state-of-the-art quantum computing techniques that may, one day, become as common-place as density functional theory.

Quantum computers are quantum systems which can be initialized, sufficiently controlled, and measured, in order to perform a computational task¹. There are many candidate physical realizations of quantum computers². Leading proposals include ion traps³⁻⁷, superconducting architectures⁸⁻¹², quantum optics¹³⁻¹⁶, nitrogen vacancy centers¹⁷⁻¹⁹, nuclear magnetic resonance²⁰⁻²² and topological qubits^{23,24}. As will be done in this review, it is convenient to abstract the description of the computation away from the particular physical implementation. However, for a concrete example in the case of ion trap quantum computing, one approach is to use trapped Ca^+ ions as qubits²⁵. The quantum computation consists of trapping the ions in an array, altering the valence-electron collective quantum state of the ions with a series of precise laser pulses, and measuring the occupation of two Zeeman states of the s and d orbital manifolds in each ion. Instead of referring to the energy levels of trapped ions, we refer to the parts of an abstract quantum computer; namely, interactions between two-level quantum systems called quantum bits or *qubits*. Instead of laser pulses, the operations on an abstract quantum computer are unitary transformations. After a sequence of unitary transformations, each qubit is measured, returning a binary outcome which is labeled “0” or “1”.

A classical computation process can be broken down to elementary logical operations, such as AND, on a few bits at a time. Similarly, controlled local unitary operations, called *quantum gates*, can be used as elementary operations to be applied on quantum bits. The novelty of quantum computing is derived from processes which “entangle” the qubits. Just as atomic orbitals fail to appropriately describe molecular electronic wave functions, entangled qubits cannot be independently assigned definite quantum states. In designing a sequence of quantum gates (called a *quantum circuit*) to output the solution to a problem, entangling quantum gates can afford shortcuts. A process for constructing a quantum circuit designed to solve a particular problem is called a *quantum algorithm*. Just like a classical algorithm, the

performance of a quantum algorithm is characterized by the number of elementary operations and the run time as a function of the problem instance size (e.g. the number of basis set functions used in the calculation).

A number of quantum algorithms have been invented which stand to outperform their classical counterparts. With regard to chemistry, these include algorithms for estimating the ground state energies of molecular Hamiltonians and computing molecular reaction rates. Existing quantum computers have yet to solve instances of problems which are intractable for classical computers. However, the pace of progress bears promise for achieving this feat^{11,13,26,27} and is certain to drive the discovery of new useful quantum algorithms.

The idea of quantum computing originated in the 1980s when Manin²⁸ and Feynman²⁹ independently described a vision for using quantum mechanical systems to perform computation. Both argued that simulating quantum mechanics on classical computers requires resources growing exponentially in the problem instance size, so that certain problems will remain out of reach, regardless of the ingenuity of the algorithm designers. It is not true that all quantum mechanical systems are difficult to simulate; some of them have easily computable exact solutions and others have very clever computational shortcuts leading to approximate solutions. However, simulation of a general quantum mechanical system has proven to be difficult. A quantum computer, as envisioned by Feynman and Manin, can potentially circumvent the roadblock of exponential cost by being quantum mechanical itself.

During the same decade, there were also developments of abstract models of quantum mechanical computation by Benioff³⁰ and Deutsch³¹, raising the still-open question of “what are the problems for which quantum computers can have a speedup over the best-known classical algorithms?”. Motivated by this question, researchers in the 1990s developed several quantum algorithms of major importance³²⁻³⁸. These algorithms not only solve their respective problems with provable speedup over the best known classical counterparts, but also provide meaningful frameworks for developing subsequent quantum algorithms. We refer the reader to the review article by Childs and van Dam³⁹ for a comprehensive treatment of the

algorithm developments prior to 2011.

This review will focus on the subset of quantum algorithms for solving problems in chemistry. In 1996, Feynman’s vision was manifested in a theoretical proposal by Lloyd³⁷, providing a concrete method for using quantum computers to efficiently simulate the dynamics of other quantum systems. Around the same time, Wiesner⁴⁰ and Zalka⁴¹ also suggested using quantum computers for simulating the dynamics of many-body systems.

The general set of techniques for using a quantum computer to simulate another quantum system falls under the name of *Hamiltonian simulation*. Broadly speaking, for a given Hamiltonian H , Hamiltonian simulation is an efficient and accurate implementation of e^{-iHt} using elementary operations that can be executed on a quantum computer. Since these early works of the 1990s, Hamiltonian simulation has grown to be an important subfield of quantum computing, offering many valuable insights towards the development of further quantum algorithms. As a prominent example, the ability to efficiently perform Hamiltonian simulation is used in conjunction with another technique known as the *quantum phase estimation algorithm (QPEA)*⁴² for efficiently obtaining the eigen-energies and eigenstates of a quantum system³⁸.

The first quantum algorithms suited particularly for quantum chemistry appeared as early as the late 1990s. These include, for instance, simulating fermionic Hamiltonians⁴³ and the quantum algorithm for calculating thermal rate constant efficiently⁴⁴. The opening decade of the 21st century witnessed the first quantum algorithms for quantum chemistry built on the insights of QPEA. For example, the quantum algorithms developed for computing molecular spectra with exponential speedup over classical computers^{45,46}. The basic idea of this work is to use Hamiltonian simulation techniques to efficiently simulate the dynamics of a quantum molecular Hamiltonian and apply QPEA to extract the eigen-energies. Since these initial contributions, a flurry of results have appeared in the literature which address quantum chemistry problems of various forms. The quantum computational cost of these algorithms has continued to be reduced⁴⁷⁻⁵⁰. However, practical implementation of these al-

gorithms is widely believed to be decades away because they require scalable, error-corrected quantum computers. This perspective has driven researchers to ask “What are the problems we can address with near-term quantum computers of moderate size and without error correction?”. Such devices are increasingly referred to as noisy intermediate-scale quantum (NISQ) devices⁵¹.

This question has led to the new paradigm of *variational quantum algorithms*. The approach has attracted substantial attention from both the theoretical and experimental communities in quantum computing. In short, variational quantum algorithms utilize a hybrid approach involving a quantum computer and a classical computer working in tandem. Unlike the algorithms based on QPEA, which require the quantum computer to perform a long sequence of operations, in the variational quantum algorithm framework the quantum computer only needs to perform a short sequence of operations. In this way, the shortcomings of present-day hardware are partly circumvented.

With variational quantum algorithms, each operation has parameters which can be set by the classical computer. Hence the quantum computer can be regarded as a machine which produces states $|\psi(\vec{\theta})\rangle$ lying on a manifold determined by the classical parameters $\vec{\theta}$. The machine then measures these states in the computational basis. This is useful, for instance, for the estimation of the ground state energy of some Hamiltonian H ⁵². For each state, we measure the expectation $\langle\psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle$ with respect to the Hamiltonian H . We then use the classical computer to tune $\vec{\theta}$ to minimize this energy expectation. By off-loading a portion of the computation to a classical computer, a variational quantum algorithm is a far more practical alternative to its QPEA-based counterparts. A diverse set of experimental groups worldwide have already implemented various variational quantum algorithms on their physical systems^{52–59}.

The quantum algorithms discussed so far fall within the standard *gate model* of quantum computing, meaning they can always be described by a sequence of quantum gates which manipulate the quantum state of a register of qubits. In parallel to the developments in

the gate model, there are other models of quantum computation that are also relevant for chemistry applications. Two important models are *adiabatic quantum computing* (AQC) and *Boson Sampling model*. The basic principle underlying AQC is the adiabatic theorem of quantum mechanics⁶⁰. Suppose we initialize a quantum system as the ground state of some Hamiltonian H_0 , and slowly evolve the Hamiltonian to some other Hamiltonian H_1 via a linear interpolation $H(s) = (1 - s)H_0 + sH_1$, $s \in [0, 1]$. Assuming there is always a gap between the ground state of $H(s)$ and the first excited state, the end of the adiabatic evolution of the state of the system should be close to the ground state of H_1 . The approach of AQC is to initialize a quantum system according to some Hamiltonian H_0 whose ground state is easy to prepare and verify, and choose H_1 such that preparing its ground state is equivalent to solving a hard computational problem. Extensive efforts have been made in developing AQC and progress in the field has also been thoroughly reviewed in the literature⁶¹.

The AQC model has been shown to be computationally equivalent to the gate model⁶², meaning that algorithms written in either model can be translated to ones in the other model without incurring prohibitive computational overhead. In contrast, the Boson Sampling model of computation does not have this property⁶³. However, it does give rise to a class of sampling problems⁶⁴ which are hard to solve on a classical computer. This model is discussed in more detail in Section 3.3.

Apart from quantum computation, another related intersection between computer science and quantum physics, which has become relevant to quantum chemistry, is quantum computational complexity theory. Well-known problems in quantum chemistry have been characterized in terms of their quantum computational complexity by taking advantage of the rigorous machinery developed in theoretical computer science for determining the worst-case hardness of a computational problem. Examples include finding the universal functional in density functional theory⁶⁵, finding the globally optimized Hartree-Fock wave function⁶⁵, the N -representability problem⁶⁶, and calculating the density of states of quantum systems⁶⁷. These characterizations, discussed in Chapter 3, provide further motivation for the develop-

ment of quantum algorithms targeted at solving them. A comprehensive review⁶⁸ of relevant results has been presented in literature, which can be a useful resource for interested readers.

The remainder of the review is organized as follows. We have provided a dependency graph of the sections and subsection in Figure 1. In Section 2 we outline several important computational roadblocks in quantum chemistry, describe state of the art classical methods which grapple with these issues, and then introduce quantum algorithms which may resolve them. In Section 3 we assess the computational hardness of classical and quantum algorithms for chemistry from a perspective of computational complexity. Section 4 provides a more detailed description of the state of the art quantum algorithms for quantum chemistry for fault-tolerant quantum computation. Section 5 describes the approaches that are suitable for calculations on present-day quantum hardware. We conclude with a summary and outlook in Section 6. In Appendix C, we give a hands-on demonstration of how to calculate a dissociation curve of H₂ with a quantum algorithm. Due to a vast amount of notations, symbols and nomenclatures introduced throughout the article, at the end of each chapter (including the Appendix) we provide a glossary for terminologies used inside the respective chapter.

2 Quantum chemistry in the age of quantum computing

A fundamental goal of quantum chemistry is to solve the time-independent, non-relativistic Schrödinger equation for molecular systems,

$$\hat{H}_{mol}(\vec{r})\psi(\vec{r}, t) = E\psi(\vec{r}, t), \tag{1}$$

where \hat{H}_{mol} is the molecular Hamiltonian, $\psi(\vec{r}, t)$ is the multi-particle wave function of the system, and E is the energy eigenvalue. In atomic units, the molecular Hamiltonian is given




Figure 1: Graphical illustration of the dependency between the subsections of this article. Each dashed box corresponds to a section and each node on the graph corresponds to a subsection. In general the materials are organized in a tree structure, with more details being covered in later sections. We broadly divide the subjects into computational complexity (Section 3) and algorithmic techniques (Section 4 and 5). In particular, the sections on algorithms are further divided into methods suited for fault-tolerant quantum computers (FTQC), as covered in Section 4, and noisy intermediate-scale quantum (NISQ) devices⁵¹ covered in Section 5.

by

$$\hat{H}_{mol} = -\sum_i \frac{\nabla_{\vec{R}_i}^2}{2M_i} - \sum_i \frac{\nabla_{\vec{r}_i}^2}{2} - \sum_{i,j} \frac{Z_i}{|\vec{R}_i - \vec{r}_j|} + \sum_{i,j>i} \frac{Z_i Z_j}{|\vec{R}_i - \vec{R}_j|} + \sum_{i,j>i} \frac{1}{|\vec{r}_i - \vec{r}_j|}, \quad (2)$$

where \vec{R}_i , M_i and Z_i indicate the spatial coordinates, masses and charges of the nuclei in the molecule and \vec{r}_i are the electronic coordinates. However, an exact solution quickly becomes intractable due to the exponential growth of the size of the wave function with the particle number⁶⁹. This has inspired formulations of various physically-motivated wave function approximations as well as quantum chemistry algorithms on classical computers. This section starts from challenges as such and discusses how classical computers may fall short in addressing some of the challenges while quantum computers are able to circumvent the limitations. The dependency among the subsections is shown in Figure 1. In order to facilitate the discussion, we partition the problem space into *static* and *dynamic* problems. Static problems computing the spectrum of the Hamiltonian, most notably ground state.

2.1 Basics and challenges of classical quantum chemistry

As an approach to solve the Schrödinger equation, which is comprised of a coupled set of differential equations, the most common strategy is to introduce a complete set of independent functions, a *basis set*, to describe the physical wave function of the system. Although this does not resolve the exponential increase of parameters with system size, it allows the balancing of computational resources with accuracy.

The full molecular Hamiltonian consists of terms associated with electrons and nuclei as well as the interactions between them. Because the mass of a nucleus is three orders of magnitude greater than that of an electron, a common simplifying assumption made in quantum chemistry calculations, called the *Born-Oppenheimer approximation* (BOA), treats the nuclei as stationary point charges. The problem is then transformed such that it only involves electrons moving in a stationary potential due to the nuclei. The molecular

Hamiltonian is split up into a nuclear part and an electronic part

$$\hat{H}_{mol} = \hat{H}_{nucl}(\vec{R}) + \hat{H}_{elec}(\vec{R}, \vec{r}). \quad (3)$$

The wave function equally can be separated into a nuclear and electronic part

$$\psi(\vec{R}, \vec{r}) = \phi_{nucl}(\vec{R})\chi_{elec}(\vec{R}, \vec{r}), \quad (4)$$

such that for every nuclear configuration \vec{R} , a separate electronic eigenvalue problem

$$\hat{H}_{elec}\chi_{elec}(\vec{R}, \vec{r}) = E_{elec}(\vec{R})\chi_{elec}(\vec{R}, \vec{r}) \quad (5)$$

needs to be solved. When the BOA is valid, the Hamiltonian of interest to quantum chemistry is

$$\hat{H}_{elec} = -\sum_i \frac{\nabla_{\vec{r}_i}^2}{2} - \sum_{i,j} \frac{Z_i}{|\vec{R}_i - \vec{r}_j|} + \sum_{i,j>i} \frac{1}{|\vec{r}_i - \vec{r}_j|}. \quad (6)$$

The electronic degrees of freedom influence the nuclear motion through the appearance of the energy in the nuclear eigenvalue equation

$$(\hat{H}_{nucl} + E_{elec}(\vec{R}))\phi_{nucl}(\vec{R}) = E\phi_{nucl}(\vec{R}), \quad (7)$$

where E is now the full molecular energy as in Equation 1.

Although in recent times there has been significant effort to treat chemistry calculations without the BOA⁷⁰⁻⁷³, in this review we focus on treating quantum chemistry problems within this regime. However, we note that though most of the applications are for cases within the BOA, some of the previous formalisms have been combined with phase estimation to treat non-BOA cases⁷⁴. Quantum simulation in first quantization can be readily applied to non-BOA instances as well^{75,76}.

For systems of electrons, one common methodology to avoid treating the size of the full

many-electron wave function is density functional theory (DFT)^{77,78}. All degrees of freedom of the electronic system are integrated out, except for the density. However, the quality of DFT depends on the chosen exchange-correlation functional. There have been several efforts to improve DFT functionals systematically (e.g. GGA, meta-GGA, hybrid approaches) but there is no uniform answer on which functional to use for which chemical systems. At the same time, these more involved approaches negate some of the more attractive characteristics of regular DFT. Most functionals tend to give reasonable results around equilibrium geometries but behave unpredictably in the regime of strong correlations or in the presence of dispersive electron interactions, e.g. bond breaking or solvation chemistry.

Another way to avoid an exponential parameter scaling is by reducing the electron Hamiltonian to its self-consistent field (SCF) part. This approximates the system as being comprised of electrons experiencing a mean field potential, which has to be determined self-consistently. As a first step, an initial guess is made for the single-electron orbitals, most frequently the atomic orbitals. After this, the influence of all electrons is averaged into the mean-field potential and new orbitals are generated from it until convergence is reached. The ground-state wave function is described using a single Slater determinant or configuration. Correlation effects are thereby completely neglected, meaning computational approximations for various physical quantities are usually of low quality. Nevertheless, the SCF wave function is often used as a starting point to construct more sophisticated wave function ansätze.

The exact solution to Equation 1 for an electronic system described with a fixed set of basis set functions consists of a variationally determined superposition of all determinants in the N -particle Fock space, also known as a full configuration interaction (FCI) expansion. Such wave function cannot be efficiently manipulated and stored on a classical computer. Correspondingly, quantum chemistry methods based on a wave function approach comprise a hierarchy of models ranging from the single determinant SCF description to FCI. Each model assumes a specific type of wave function parametrization in addition to other approximations,

offering different compromises between accuracy and computational cost. Some of the most promising wave functions methods employed in quantum chemistry are described in the next section.

2.2 Classical approximation techniques and their limitations

Quantum computation proposes methods which may overcome the challenges of quantum chemistry described in the previous subsection. These algorithms are reviewed in Section 2.3. However, we emphasize that there are various methods and techniques in quantum chemistry applied using classical computers that address said challenges. In this section, we review some representative techniques and motivate the development of quantum algorithms.

2.2.1 Static properties: wave functions and energy spectra

Coupled cluster theory. Coupled cluster⁷⁹ is formulated as a compact wave function parametrization in terms of an exponential functional of a cluster operator. The algorithm leverages several approximations to be tractable on a classical computer. Most codes utilize the similarity transformation formulation with the cluster operator as a pure excitation operator, taking advantage of the termination of the Baker-Campbell-Hausdorff series. This makes the problem non-Hermitian and requires a projection onto a small subspace to solve the cluster amplitudes, introducing a further approximation. Although this formulation of coupled cluster is still exact in the limit of an N -body excitation operator, in practice, excitation operators are truncated. Such approximations are not guaranteed to behave well. The breakdown of coupled cluster with only single and double excitations for molecules in dissociation⁸⁰ is well-documented. This behavior can be generalized to any level of truncation of the excitation operator when strong correlation is exhibited by the hamiltonian⁸¹. There are several approaches^{82,83} that improve upon this framework by approximating higher order cluster operators or including static correlation effects, but development of a generally applicable method is still an active field of research.

While coupled cluster theory is a promising method for quantum chemistry, its non-Hermitian formulation makes the obtained energy non-variational. Variationally optimized coupled cluster⁸⁴ overcomes this problem but can only be applied to small systems⁸⁵⁻⁸⁷ or models with a reduced parameter space⁸⁸ on a classical computer. The introduction of a complex conjugate excitation operator implies contractions between the operators in the bra and the ket. The transformed hamiltonian is a full N -body operator with the number of terms scaling factorially as $N!$. The problem becomes more challenging when the operators in the exponential are not pure excitation operators. In this case the expectation values that need to be calculated feature contraction among the cluster operators and there are an infinite number of terms. Unitary coupled cluster⁸⁹ (in which the excitation and de-excitation operators are combined into a unitary cluster operator) and generalized coupled cluster⁹⁰ (in which the cluster operator lacks a specific substructure) fall in this category. We note that on a quantum computer, physical operations on qubits are often realized in terms of unitary operators. If the cluster operator can be represented as a low-depth quantum circuit, it is efficient to prepare a unitary coupled cluster wave function. This method of state preparation together with a classical optimizer (see Section 5.1) enables unitary coupled cluster for general hamiltonians.

Quantum Monte Carlo. The idea of efficiently sampling a distribution of states has a counterpart on classical computers in Quantum Monte Carlo (QMC). It has been very successful in performing large-scale calculations for extended systems in quantum chemistry. QMC relies on the stochastic estimation of the energy of a trial wave function

$$E_T = \frac{\int d\vec{r}^{3N} \Psi_T(\vec{r}) H(\vec{r}) \Psi_T^*(\vec{r})}{\int d\vec{r}^{3N} \Psi_T(\vec{r}) \Psi_T^*(\vec{r})}. \quad (8)$$

In the real-space variational Monte Carlo formulation, the position vectors are randomly generated according to the Metropolis algorithm from the norm of the trial wave function.

After sufficient convergence in the energy with respect to the number of Monte Carlo samples, the wave function parameters can be updated. This process is repeated until a desired accuracy has been met.

More accurate QMC methods have been developed, in which more elaborate distributions are sampled. Diffusion Monte Carlo (DMC) uses a population of walkers to stochastically sample the position space. The composition of the population is algorithmically controlled from generation to generation. This process depends on the knowledge of the nodes of the exact wave function. A popular implementation of DMC fixes the nodes of the trial wave function (fixed-node DMC) which biases the population but is shown to be more robust and provides accurate results. Another possibility is to sample the manifold of electronic configurations. The Auxiliary Field QMC (AFQMC) is performed as an imaginary-time evolution of the transformed Hamiltonian in terms of the auxiliary fields. A sign problem occurs in both DMC and AFQMC, but is mitigated by defining phase constraints on the auxiliary fields from a trial wave function.

There is a natural connection between the underlying principles of Monte Carlo algorithms and the representation of a state on a quantum computer. Measurement outcomes from the output of a quantum circuit are probabilistic and several algorithms have been proposed for Monte Carlo integration on a quantum computer that outperform their classical counterparts⁹¹. There are certain drawbacks to the approximations in classical QMC. The most notorious one is the sign problem that plagues most fermionic implementations. The nature of fermions makes accurately estimating quantities using the wave function a difficult task. It has been shown that quantum computers avoid the dynamical sign problem⁴³ providing a more effective alternative to classical methods for quantum simulation.

Exact diagonalization. The exact diagonalization (ED) methods^{79,92,93} provide the exact answer for the wave function and energy within a certain basis set (i.e. the Full Configuration Interaction (FCI) answer). While QMC arrives at this answer through extrapolation,

ED achieves it in a single calculation by exact diagonalization of the Hamiltonian matrix without stochastic error bars. This comes at the cost of storing the coefficients of all (relevant) determinants, which becomes prohibitive even for medium-sized molecules^{94,95}. The rapid increase in computing power and the development of libraries that take advantage of distributed computing clusters have meant a steady increase in the feasible number of determinants⁹⁶⁻⁹⁸. Despite all these advances, FCI is still the most useful as a benchmarking method for less-costly quantum chemical methods. Exact simulation of quantum chemistry systems is widely regarded as one of the problems that would benefit enormously from quantum hardware²⁹. The Quantum Phase Estimation (QPE) algorithm^{38,99} is the natural translation of the FCI procedure to quantum computers⁴⁵.

Tensor network methods. Compared to the classical methods discussed so far, entanglement-based tensor product methods¹⁰⁰, which have been applied to study strongly correlated many-body physics and quantum chemistry, are closer to the formulation of gate-model quantum computation. In the case of strongly correlated (multireference) systems where the concept of a single dominating configuration breaks down, traditional single reference electronic structure methods are usually not applicable. On the contrary, tensor product methods¹⁰⁰ have proved to be efficient and accurate as multireference computational methods. The most commonly used entanglement-based tensor product method so far is the quantum chemistry implementation of density matrix renormalization group method (DMRG)¹⁰⁰⁻¹⁰⁹. Specifically, DMRG variationally optimizes a wave function in the form of a matrix product state (MPS)¹¹⁰:

$$|\Psi_{\text{MPS}}\rangle = \sum_{\{\alpha\}} \mathbf{A}^{\alpha_1} \mathbf{A}^{\alpha_2} \dots \mathbf{A}^{\alpha_n} |\alpha_1 \alpha_2 \dots \alpha_n\rangle, \quad (9)$$

which is a one-dimensional chain of tensors. In the quantum chemistry setting, α_i represents the local Hilbert space of a single (i -th) molecular orbital (MO), i.e. $\alpha_i \in \{|0\rangle, |\downarrow\rangle, |\uparrow\rangle, |\downarrow\uparrow\rangle\}$. The DMRG computational complexity is governed by the size of the MPS matrices (\mathbf{A}^{α_i}),

formally characterizing the entanglement between different MOs. Matrix product states allow for a compact representation of one-dimensional systems and also satisfy an area law¹¹¹ in one dimension. This has led to the application of DMRG to describe ground states of one-dimensional gapped local Hamiltonians. However, the long range Coulomb interaction in molecules causes the entanglement among the MOs to be multidimensional as opposed to one-dimensional. As a result, larger dimension MPS matrices (or “high bond dimension” MPS) have to be employed in case of generic (non-one-dimensional) molecules in order to properly capture correlations among MOs. These problems can be alleviated to some extent by optimizing the orbital ordering along a one-dimensional lattice^{112,113} as well as by basis optimization¹¹⁴. However, these problems pose a practical limit on the applicability of tensor network states (TNS) for quantum chemistry.

DMRG applied to quantum chemistry can generally handle large active spaces (~ 50 MOs¹¹⁵). It handles much larger active spaces than conventional active-space methods such as complete active space self consistent field¹¹⁶. In certain cases it can reach the exact FCI limit^{115,117}. In case of linear (or quasi-linear) molecular systems, active spaces of a hundred of MOs are reachable¹¹⁸. However, DMRG wave functions at a low bond dimension miss a sizable amount of the dynamic correlation that is important in quantum chemistry. In order to account for the missing dynamic correlation not included in the active space, several post-DMRG methods have been developed^{119–127}.

TNS¹²⁸ represent a generalization of MPS aimed at an improved representation of multidimensional entanglement. TNS include e.g. the projected entangled pair states¹²⁹, tree tensor network states (TTNS)^{130–134}, multiscale entanglement renormalization ansatz¹³⁵ or complete-graph TNS^{136,137}. A new TTNS variant was recently presented, the three-legged TTNS, which combines the tree-like network topology and branching tensors without physical indices¹³⁴. It is especially appealing due to its low computational cost and network topology capable of taking into consideration the underlying entanglement of generic molecules more efficiently.

The FCI wave function expansion coefficients can be converted into the MPS form by successive applications of the singular value decomposition^{110,138}. In the case of the TTNS ansatz, which is the most general TNS without any loops, higher-order singular value decomposition is employed. Factorization of the FCI wave function into the MPS form itself does not bring any computational savings as it results in bond dimensions of the size of the original problem. In order to achieve polynomial scaling, bounded bond dimensions have to be applied. Consequently the entanglement between two subsystems of the bipartite splitting governs the quantum information loss¹³⁹ and accuracy of the ansatz. It must be noted that low bond dimension MPS are not candidates for quantum speedup as slightly entangled quantum computations may be *efficiently* simulated on a classical computer using the MPS formalism¹³⁸.

2.2.2 Dynamical properties: time evolution of the wave function

Zero-temperature ground-state quantum chemistry covers only a restricted set of the chemistry in nature. Many processes occur at finite temperature where the system is propagated in time and described by the time-dependent Schrödinger equation

$$i \frac{\partial}{\partial t} \Psi(\vec{r}, t) = H \Psi(\vec{r}, t). \quad (10)$$

When the Hamiltonian is time-independent, this equation is formally solved by

$$\Psi(\vec{r}, t) = \exp(-iHt) \Psi(\vec{r}, 0). \quad (11)$$

The expansion of this equation in terms of eigenvectors $\psi_j(\vec{r})$ of the time-independent Hamiltonian

$$\Psi(\vec{r}, t) = \sum_j c_j \exp(-iE_j t) \psi_j(\vec{r}) \quad (12)$$

reveals the difficulty of the problem: exact time evolution requires knowledge of the full spectrum of the Hamiltonian and the expansion of the initial state in terms of the eigenvectors. This is only feasible for very small systems.

Several levels of approximations are often applied to approach the problem. Here we review two broad categories of methods. Molecular dynamics avoids the explicit wave function propagation by treating the nuclei as classical charged particles and only solving for the electrostatic field from electron density with fixed nuclei positions. Other methods avoid maintaining the full wave function with similar strategies as those described in Section 2.2.1, with an additional consideration of time dependence.

Molecular dynamics (MD). When studying the dynamics of large chemical systems (approximately hundreds of atoms), one often is interested in properties for which no quantum effects are necessary. It is sufficient to use classical MD methods which simulate the problem using an effective force field and Newtonian mechanics¹⁴⁰. MD is often applied to problems such as diffusion rates, reaction rates, conformational changes, infrared spectra, and binding energies, to study processes such as protein folding, gas absorption in porous materials, stress-induced material deformation, and electrolyte solvation, among many others. Because these quantities can often be calculated to desired accuracy by treating each atom as a classical Newtonian particle, MD avoids the construction of a Hilbert space under time evolution.

In order to advance one time step in a Newtonian simulation, the forces on the particles need to be approximated. This is achieved in MD by introducing a force field that acts on the atoms without explicitly considering electronic degrees of freedom. A force field comprises of a set of terms that are parametrized by electronic structure calculations or experiment. The dynamics of the system are propagated using algorithms such as Euler or Verlet integration, after which one can obtain an intuitive qualitative understanding of the dynamics of a process (e.g. protein-ligand docking), or calculate quantitative results from

correlation functions (such as spectra or diffusion constants). If the primary interest is in finding low-energy states (which is the case in finding the lowest-energy molecular crystal conformation or in protein-ligand docking), then the dynamics are less relevant, and one of many optimization methods is used to find the ground state of the system. We note that one way to improve accuracy is to use *ab initio* electronic structure methods in conjunction with MD methods as a way to consider dynamics for the whole system. Such methods include Born-Oppenheimer MD¹⁴¹ and Carr-Parinello MD¹⁴².

There are many chemical systems for which nuclear quantum effects must be considered, even if the electronic degrees of freedom can be safely absorbed into the force field. Such effects include nuclear tunneling and zero-point vibrational energy, which are not captured by Newtonian mechanics. Nuclear quantum effects are needed for simulating condensed matter systems containing low-temperature helium and hydrogen¹⁴³, as well as for obtaining high accuracy simulations of hydrogen-containing room-temperature systems such as liquid water¹⁴⁴. A variety of flavors of path integral methods¹⁴⁵ have been developed to model nuclear quantum effects, including path-integral MD¹⁴⁶, centroid MD¹⁴⁷, and ring-polymer MD¹⁴⁸. Conceptually, these methods operate by discretizing an atom into many “beads” that represent a state of the system in the path-integral. While these methods include some quantum effects listed above, they avoid using a Hilbert space, which limits their computational cost to that of classical MD methods.

It is important to note that, even though MD has opened a whole field of research that was previously inaccessible, the method has limitations that are hard to overcome without fundamentally altering the procedure. The most severe problem is the limited amount of time propagation that MD currently supports. Some processes in biochemistry such as conformational transition in hemoglobin occur on a time scale that is many times the current limit of MD calculations¹⁴⁹.

Quantum time evolution. Most of the methods in section 2.2.1 can be adapted to perform

time evolution over which the quantum character is preserved at all steps. One of the strategies for time propagation is directly propagating the solution of the Schrödinger equation, i.e. updating an explicit wave function or density matrix at each time step. This quickly becomes memory intensive and requires a method to approximate the matrix exponential of the Hamiltonian¹⁵⁰.

Moler and Loan¹⁵¹ categorized the many methods for approximating the operation of a matrix exponential. The Taylor expansion, Padé approximant, and Chebyshev approximant are series methods based on powers of the Hamiltonian^{150,151} to approximate the evolution operator. They are accurate for small matrices but prohibitively expensive for larger systems. Standard ordinary differential equation (ODE) solvers such as Runge-Kutta are a common workhorse for time propagation of classical equations. They can be adapted to perform quantum propagation^{150,152}, although they are not the most computationally efficient choice for a given accuracy. Polynomial methods form another algorithmic category. They rely on the exact expression of high-order matrix powers as lower-order polynomials of the Hamiltonian¹⁵³. The bottleneck for those methods is calculating the characteristic polynomial, which is expensive for large matrices. Note that in practice these methods are often combined—for instance, after performing the Lanczos decomposition one may use the Padé approximant on the resulting smaller approximate matrix¹⁵⁴. For all of these methods, the propagation or matrix exponentiation algorithm is chosen based on desired accuracy, the size of the system, ease of implementation, and properties of the matrix such as its condition number.

Matrix decomposition methods try to approximate the operator exponential of the Hamiltonian directly. The prototypical example of this is the decomposition of the exponential as $e^A = Se^V S^{-1}$, where $A = SVS^{-1}$ and V is diagonal. This approach requires far fewer matrix-matrix operations than the naive implementation. Improved stability and efficiency are often observed by rewriting the exponential in terms of triangular (e.g. Schur) or block-diagonal matrix decompositions¹⁵⁵. There are also *splitting methods* which apply an approximation

that can be used if the Hamiltonian is expressible as a sum of matrices, and where each term of the sum can be simulated by a known efficient method. This is the same procedure used in the Trotter-Suzuki method, as well as the split-operator method for applying the kinetic and potential energy terms separately¹⁵⁶. Finally, the Krylov subspace methods (e.g. the Lanczos method for Hermitian matrices) make it possible to approximate $e^{-iH}v$ for a given vector v , by solely performing matrix-vector multiplications^{154,157}. This is efficient for large matrices.

One of the most cost effective ways to obtain time-dependent properties without explicitly doing time propagation is the first-order perturbation approach or linear response theory. Even though the original formulation of Time-Dependent DFT¹⁵⁸ does not require the time-dependent part of the Hamiltonian to be small, it has mostly been applied using this assumption^{159–161}. In the same spirit, a linear-response approximation to DMRG^{162,163} and coupled cluster^{164–166} allow for the calculation of dynamical properties with high-level quantum chemistry methods. Linear response is the first-order expansion of a time-dependent property and perturbation theory falls short of describing the full time dependence. Ideally we would like to be able to truly propagate quantum wave functions for an arbitrary amount of time.

The time-dependent variational principle (TDVP)^{167–169} has enabled time-dependent variational Monte Carlo calculations for bosons¹⁷⁰ and strongly correlated fermions¹⁷¹. Although the TDVP has also found its application in DMRG¹⁷², the most popular algorithm for time evolution with a DMRG-style ansatz is called time-evolving block decimation¹⁷³. This method relies on a Trotterization of the time-evolution and the DMRG truncation of the Hilbert space after every time-step. A more thorough review of time evolution algorithms related to DMRG can be found in Schollwöck¹⁷⁴. Despite these impressive efforts to formulate analogous time-dependent methods, the shortcomings of the tensor-based methods they are based on are equally present and even aggravated. After longer time evolution, errors due to the inherent approximations accumulate and lead to loss of accuracy and a rapidly

growing bond dimension.

Simulating dynamical behavior around conical intersections is an important topic in theoretical chemistry, relevant for studying reaction mechanisms and in photochemistry, including photocatalysis and the interpretation of spectroscopic results. Conical intersections occur when two electronic states, e.g. the ground state and first excited state of a molecule, intersect at a point in nuclear coordinate space. Close to these points, the Born-Oppenheimer approximation breaks down, resulting in entanglement between electronic and nuclear degrees of freedom. This means that the full quantum state is no longer well-approximated with a product state, necessitating methods that allow for multireference character in the wave function, of which we mention three popular methods here. The multiconfiguration time-dependent Hartree (MCTDH) algorithm¹⁷⁵ models all relevant quantum degrees of freedom, propagating superpositions of product states. Evolution of a wave packet over longer time periods in MCTDH requires considerable computer resources and therefore larger systems are out of the scope of current capabilities. The matching-pursuit/split-operator Fourier-transform¹⁷⁶ and the more classical multiple spawning method¹⁷⁷ use careful metrics to continually update the basis set in which the system is propagated to ensure efficient simulation. Recent reviews have been published on these and other algorithms for studying nonadiabatic dynamics and conical intersections^{178,179}.

2.3 Going beyond classical limitations using quantum computers

All of the classical techniques introduced in the previous section are designed to avoid two features in a computation. Namely, these are explicitly maintaining the full many-body wave function $|\psi\rangle$ and propagating the wave function in time by general matrix-vector multiplication $e^{-iHt}|\psi\rangle$. However, as we will discuss in this section, quantum computers allow for efficient implementation of these two features. For the former, i.e. state representation, there are quantum states for which no known efficient classical representation exists, but that can be prepared on a quantum computer. The quantum computer thus has access

to a richer state space for the ground-state search for certain systems. For the latter feature, time evolution, a long line of inquiry in quantum computing that spans more than two decades has led to a refined set of quantum algorithms for efficiently simulating time evolution under both time-independent and time-dependent Hamiltonians. These two features of a quantum computer allow for algorithm design techniques that are quite different from the classical algorithms for simulating quantum systems. Both state preparation and Hamiltonian evolution are instrumental to the techniques we review in this section.

In addition to distinguishing between static and dynamic problems as in the previous section, we also make a distinction between NISQ devices and fault-tolerant quantum computing (FTQC) devices. The first are available at present or are using technologies that will be available in the near-term, while the latter require much more research and are a much longer-term prospect. The main difference between NISQ and FTQC devices concerns decoherence due to the environmental noise and whether it sets an upper limit on the time duration of a quantum computation. In this sense, NISQ devices are sensitive to sources of noise and error (hence the term “noisy” in NISQ) while FTQC devices can, in principle, carry out coherent quantum computation for an arbitrary amount of time.

Many well-known quantum algorithms that have become standard examples of quantum speedups over classical algorithms, such as Shor’s quantum factorization algorithm, quantum search algorithm by Grover and quantum phase estimation (which we will review here), in fact assume the availability of FTQC devices in order to deliver the promised quantum advantage. This does not *a priori* exclude the possibility of quantum advantage on NISQ devices. In fact, there are promising ideas which may bring tangible results for quantum chemistry on near-term devices.

Before discussing any quantum simulation algorithms, we would like to start with how wave functions of quantum systems can be represented on a quantum computer. The basic building blocks of quantum computation are controllable two-level systems called *qubits*. If we denote the two levels of a qubit as $|0\rangle$ and $|1\rangle$, the wave function of an n -qubit quantum

computer can be expressed as a superposition of 2^n *computational basis states*:

$$|\psi\rangle = \sum_{i_1, i_2, \dots, i_n \in \{0,1\}} a_{i_1 i_2 \dots i_n} |i_1\rangle \otimes |i_2\rangle \otimes \dots \otimes |i_n\rangle. \quad (13)$$

For describing a computational basis state the Kronecker product symbol \otimes is commonly dropped for simplicity in notation. For example, we will use $|110\rangle$ for representing the 3-qubit state $|1\rangle \otimes |1\rangle \otimes |0\rangle$. Early proposals^{40,41} for representing wave functions of quantum systems consider a real-space representation where the spatial wave function $\psi(x)$ is discretized into grid points $\sum_i \psi_i |i\rangle$ with each $|i\rangle$ being a computational basis state. Similar representations have been improved in later studies applying them to simulation of quantum dynamics in real space^{75,76}. An alternative representation considers the second-quantized picture where each computational basis state $|i_1 i_2 \dots i_n\rangle$ corresponds to a Slater determinant in the Fock space. Each $i_j = 1$ if the spin-orbital j is occupied and 0 otherwise. This representation underlies the majority of quantum algorithms for quantum chemistry on both NISQ and FTQC devices. One of the appealing features of this approach is that it allows one to transform the molecular Hamiltonian into a standard second-quantized form

$$H_{2q} = \sum_{p,q} h_{pq} a_p^\dagger a_q + \sum_{p,q,r,s} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s \quad (14)$$

where a_i^\dagger and a_i are raising and lowering operators acting on the i -th basis function. Here the coefficients h_{pq} and h_{pqrs} are one-electron and two-electron integrals which can be efficiently computed classically for many choices of basis functions. The fermionic operators obey anti-commutation relations $\{a_i, a_j^\dagger\} = a_i a_j^\dagger + a_j^\dagger a_i = \delta_{ij}$ and $\{a_i, a_j\} = \{a_i^\dagger, a_j^\dagger\} = 0$. One can transform the Hamiltonian expressed in terms of fermionic operators into a more natural form for quantum simulation of electronic structure on a quantum computer. The basic idea is to replace each fermionic operator with a tensor product of Pauli matrices $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. While there are several methods for accomplishing this

transform^{180,181}, all result in a Hamiltonian that is a linear combination of tensor products of matrices $\{I, X, Y, Z\}$ with $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ in such a way that preserves the fermionic anti-commutation relations. In general we can write such Hamiltonian as a *k-local Hamiltonian*

$$H_{k\text{-local}} = \sum_i c_i \sigma_{i,1} \sigma_{i,2} \cdots \sigma_{i,k}. \quad (15)$$

Here “*k*-local” means that each term in the Hamiltonian acts non-trivially on at most *k* qubits. In Equation 15, the notation $\sigma_{i,j}$ means the *j*-th operator in the *i*-th term of the Hamiltonian. Each $\sigma_{i,j}$ is a Pauli operator acting on one qubit. For a detailed illustration of mapping from a molecular Hamiltonian to a *k*-local Hamiltonian representation in the Fock space, the reader is encouraged to refer to Appendix C for an example using a hydrogen molecule.

The organization of this section is similar to the previous section where the problem space is partitioned into static and dynamic problems. In addition, for each category we discuss quantum techniques devised for both NISQ and FTQC devices. The goal is to introduce some representative ideas in the space, leaving more detailed discussions to Sections 4, 5 and Appendix C.

2.3.1 Statics: phase estimation and variational quantum eigensolvers

The idea of quantum phase estimation can be traced back to early works on quantum mechanics by von Neumann¹⁸² inquiring the nature of quantum measurements. Contemporary discussions on phase estimation would commonly place its origin at the work of Kitaev et al.⁴². However, textbooks¹ on quantum computing typically use another version of phase estimation which differs from Kitaev’s version. The basic idea stems from the fact that for a given unitary operator U , its eigenvalues take the form of a phase $\lambda_j = e^{i\varphi_j}$. We assume that U has an efficient implementation on a quantum computer. That is, if U acts on n qubits, there is a polynomial-length sequence of elementary operations that one can perform on a

quantum computer which equals or well approximates U . For an eigenstate $|\psi_j\rangle$, the phase estimation¹ algorithm transforms the state on a quantum computer into

$$|\psi_j\rangle|0\rangle \mapsto |\psi_j\rangle|\tilde{\varphi}_j\rangle, \quad (16)$$

where $|\tilde{\varphi}_j\rangle$ is a state that encodes an approximation of the phase $\varphi_j/(2\pi)$. The approximation error comes from the attempt to capture the value of φ_j , which is a continuous variable, with a finite qubit register and hence a finite-dimensional quantum system. By measuring the state $|\tilde{\varphi}_j\rangle$ one could extract the phase φ_i up to the precision allowed by the qubit register.

Phase estimation is an algorithmic framework that encapsulates a broad set of quantum algorithms. For example, Shor’s algorithm for integer factorization³² and amplitude estimation¹⁸³ both can be cast as a form of phase estimation for specific construction of unitary operators U . Most relevant to this review, is the use of phase estimation for extracting the eigenspectrum of a Hamiltonian H . This corresponds to the special case where $U = e^{-iHt}$. Efficient implementation of e^{-iHt} is the subject of *Hamiltonian simulation* which merits a separate discussion in Section 2.3.2. It should be noted that in the case where the objective is a specific eigenstate of the Hamiltonian such as the ground state $|\psi_0\rangle$, an appropriate initial state $|\phi_0\rangle$, reasonably close to $|\psi_0\rangle$ is needed. To see this, consider a general state $|\phi\rangle = \sum_j \beta_j |\psi_j\rangle$. Applying phase estimation on the initial state $|\phi\rangle|0\rangle$ yields $\sum_j \beta_j |\psi_j\rangle|\tilde{\varphi}_j\rangle$. Upon measurement of the second qubit register, one would like to maximize the chance of obtaining $|\tilde{\varphi}_0\rangle$, which is dictated by $|\beta_0|^2$. In order to obtain the ground state energy in a time duration that scales polynomially in the number of qubits n acted on by the Hamiltonian H , we require that $|\beta_0|^2$ be at least as large as an inverse polynomial in n . Several methods of state preparation on a quantum computer have been discussed in the literature and we review them in detail in Section 4.2.1.

For quantum chemistry applications, there are well motivated methods for preparing initial wave functions based on efficiently computable ansatzes. In addition, the Hamiltonians

for molecular electronic structure problems are also well studied in the context of implementing time evolution $U = e^{-iHt}$. Both ingredients give hope that quantum computers can take advantage of quantum phase estimation to efficiently compute the spectrum of quantum systems to accuracy comparable to FCI. However, an important technical point is that the sequences of operations (or quantum circuits, see Section 4.1.3) yielded from phase estimation are often too deep to be feasible on today’s NISQ devices⁵¹. Instead they require fault-tolerant quantum computers which can in principle perform quantum computation indefinitely. However, there are significant technical challenges which need to be met before fault-tolerance can be achieved. For NISQ devices, the recently proposed paradigm of *hybrid quantum-classical algorithm* is a more practical framework for developing heuristics to solve the same problems of eigenstates and energy levels.

A salient feature of hybrid quantum-classical algorithms is that much of the computational burden is offloaded to a classical computer. The quantum computer is only used for its ability to prepare entangled states that are otherwise hard to represent classically and to make measurements with respect to the states. Specifically, the setting is that the quantum computer can prepare a set of states $|\psi(\vec{\theta})\rangle$ parametrized by classical parameters $\vec{\theta}$. One then makes measurement with respect to the state and the classical computer updates the parameter setting to $\vec{\theta}'$ and feeds it back into the state preparation on the quantum device. By iterating between classical and quantum processors, the hybrid algorithm eventually trains the quantum computer to prepare the state that is desired for the problem. An important class of hybrid quantum-classical algorithms that are useful for solving static problems in quantum chemistry is the variational quantum eigensolver (VQE), which will be discussed in greater detail in Section 5.1.

2.3.2 Dynamics: Hamiltonian simulation and hybrid methods

Many quantum chemistry problems concern the time evolution of a quantum state under a Hamiltonian H . Propagating an initial wave function $|\psi_0\rangle$ in time by calculating $|\psi(t)\rangle =$

$e^{-iHt}|\psi_0\rangle$ is in general hard on a classical computer. As we have alluded to in Section 2.2, classically tractable cases include restricted settings where

1. the system is small enough for the Hamiltonian to be treatable with explicit methods;
2. the time step t is small, for which $e^{-iHt} \approx I - iHt$. This yields a linear response theory;
3. efficient approximation of the wave function is possible using, for instance, tensor networks, as in t-DMRG, or more recently neural networks¹⁸⁴;
4. the dynamical sign problem can be effectively suppressed¹⁸⁵.

Quantum computation circumvents the main issues encountered in classical methods by maintaining a highly entangled quantum state which may not admit efficient classical description, enabling a fundamentally different approach to realizing the unitary evolution and in doing so avoids the dynamical sign problem⁴³.

On a (fault-tolerant) quantum computer, there are efficient methods for implementing an approximation of e^{-iHt} for arbitrary t and a broad class of Hamiltonians H . Early works on Hamiltonian simulation^{37,45,186} assume that H can be written as a sum of local terms $H = \sum_j H_j$ where each H_j acts on a subsystem. By application of the Trotter-Suzuki formula, the problem becomes how to implement each $e^{-iH_j t}$ individually on a quantum computer. For molecular electronic structure problems the recipes for such implementations have been explicitly specified¹⁸⁶. Using such recipes, implementing e^{-iHt} for an n -qubit Hamiltonian takes time polynomial in n on a quantum computer, while a classical computer takes time at least $\sim 2^n$. In recent years there has been an extensive line of inquiry⁴⁷⁻⁵⁰ in estimating and improving Hamiltonian simulation techniques in the context of phase estimation for quantum chemistry applications.

An alternative setting of Hamiltonian simulation^{187,188} considers an ‘‘oracle model’’ where the Hamiltonian is assumed to be a sparse matrix and its elements are provided by a black box (or *oracle*) such that when queried with the row number i and the index j (namely 1st,

2nd etc.), the black box returns the j^{th} non-zero element of the matrix. It is also assumed that one can make queries to the oracle in superposition and obtain answers in superposition. The total cost of the algorithm is then taken to be either the number of total queries, or the sum of total queries and the total number of elementary operations performed on a quantum computer. Some of the recently developed paradigms for Hamiltonian simulation^{187–190} achieve exponential improvements in precision compared with their predecessors^{187,188}. They have also been applied to quantum chemistry^{191,192} yielding a similar exponential improvement. We will expand on the details in Section 4.1.

Naturally, any discussion of realizing the operation e^{-iHt} is restricted to simulating time evolution of *closed, time-independent* systems. More generally, algorithms for simulating open quantum systems have also been explored^{193–195}. These algorithms focus on the Hamiltonian and Lindblad operators being sparse and given by oracles (black-boxes). For time-dependent Hamiltonians, early works¹⁹⁶ focus on variants of Trotter-Suzuki formula, which are subsequently improved with black-box models¹⁸⁸. Recently, simulation techniques in the interaction picture¹⁹⁷ have been proposed. A similar technique has also been applied in quantum chemistry¹⁹⁸ to produce the first quantum simulation algorithm of cost scaling *sublinearly* in the number of basis functions. In parallel with previous methods for simulating time-independent Hamiltonians¹⁸⁸ yielding near-optimal cost scaling, techniques for simulating time-dependent Hamiltonians using Dyson series¹⁹⁹ have also been recently proposed, with scaling matching the time-independent case¹⁸⁸.

All of the quantum algorithms mentioned in this subsection rely on the possibility of fault-tolerant quantum computers in order to outperform any of their classical counterparts. However, the prospect of physical fault-tolerant devices will most likely require decades of experimental progress. This prospect has led researchers to investigate what quantum dynamics one can simulate with the quantum devices that we have today. Indeed, there are already proposals²⁰⁰ for NISQ algorithms based on variational principles^{201,202}. In these methods, a quantum state $|\psi(\vec{\theta})\rangle$ is generated by a low depth circuit with classical parame-

ters $\vec{\theta}$. The time evolution of an initial state $|\psi(\vec{\theta}_0)\rangle$ then translates to updating the classical parameters iteratively as the simulation progresses. More details on using variational principles for both real and imaginary time propagation will be discussed in Section 5.2.

2.4 Chapter 2 Glossary

H, H_i	Hamiltonian and terms in the Hamiltonian, respectively
c_n, β_j	Probability amplitudes
t	Time
$ \psi\rangle, \psi, \phi$	Quantum states
A, S, V	Arbitrary matrix, its similarity transform and its diagonalized form
$\vec{\theta}, \vec{\theta}_0$	Variational parameters
a, a^\dagger	Fermion annihilation and creation operators
\vec{r}	Electron coordinate position vector
\vec{R}	Nuclear coordinate position vector
h_{pq}, h_{pqrs}	One and two body Hamiltonian elements
X, Y, Z, I	Pauli matrices and identity
E	Energy
$\sigma_{i,j}$	i -th Pauli operator acting on the j -th qubit
N	Particle number
U	Unitary operator
A	Tensor
n	Number of qubits
E_j, ψ_j	Eigenvalue and eigenvector of Hamiltonian
λ_j, φ_j	Complex eigenvalue of a unitary and its phase
QPEA	Quantum phase estimation algorithm
VQA	Variational quantum algorithm
AQC	Adiabatic quantum computing

DFT	Density functional theory
GGA, meta-GGA	Generalized gradient approximation
SCF	Self-consistent field
FCI	Full configuration interaction
QMC	Quantum Monte Carlo
DMC	Diffusion Monte Carlo
AFQMC	Auxiliary field quantum Monte Carlo
ED	Exact diagonalization
QPE	Quantum phase estimation
DMRG	Density matrix renormalization group
QC-DMRG	Quantum chemistry density matrix renormalization group
MPS	Matrix product states
MO	Molecular orbital
TNS	Tensor network state
TTNS	Tree tensor network states
MD	Molecular dynamics
ODE	Ordinary differential equation
TDVP	Time-dependent variational principle
MCTDH	Multiconfiguration time-dependent hartree
NISQ	Noisy intermediate-scale quantum
FTQC	Fault-tolerant quantum computing
VQE	Variational quantum eigensolver

3 Computational complexity

There are many examples of elusive problems in quantum physics or quantum chemistry that seem to defy efficient solutions despite decades of effort. These include the 3D Ising

problem, Hubbard models, and finding the universal functional in density functional theory. In the face of these hard open problems, one may ask the question of why these problems remain unanswered. Is it a lack of human ingenuity and understanding, or a fundamental difficulty that is inherent to these problems? Being able to answer this question not only provides theoretical insights on the problems themselves, but could also provide guidance for research efforts to choose the most fruitful direction toward a solution. To make progress, one may turn to the theory of *computational complexity*, which seeks to rigorously characterize the computational hardness of various problems⁶⁸. In order to help capture the inherent hardness of solving various families of problems, the theoretical computer science community has developed the concept of *complexity classes*.

It is important to note that the computational complexity of a given problem refers to its worst-case hardness, and does not always inform us of how difficult a problem is to solve in practice. A pertinent example is that of determining the optimal Hartree-Fock approximation of a molecular eigenstate. This problem is known to have the same complexity (NP-completeness) as many infamously hard problems. Yet, there are heuristic methods which regularly solve the Hartree-Fock problem in practice. Complexity classes account for the worst-case instances of a problem, even if these instances may rarely occur in practice. The utility of defining complexity classes, then, is that they provide a reality check when developing new methods for solving previously-classified problems.

The remainder of the section is organized as follows (see also Figure 1): we start with Section 3.1, where we establish the basic notions in the theory of computational complexity. Building on these definitions, Sections 3.2 and 3.3 discuss results related to problems of electronic structure and molecular vibronic structure respectively.

3.1 Complexity classes

This section intends to provide an intuitive introduction into the notion of complexity classes. Consider the example of finding the ground states of a family of Hamiltonians of various

system sizes. Formally, this is an example of a *problem*, while finding the ground state of a particular Hamiltonian within this family is an example of a *problem instance*. Problems and problem instances are two important primitives in the study of computational complexity classes.

From the computational viewpoint, an algorithm is *efficient* for a problem if for any problem instance of size n , it can solve the instance in a time that is less than some polynomial function of n . Here the “size” is a parameter of the problem that characterizes how much information is needed to specify the problem completely (e.g. basis set size in the ground state search of a molecular Hamiltonian, number of grid points in a real space calculation). This allows us to rigorously characterize how hard it is to solve a computational problem.

The simplest category of complexity classes involve *decision problems*, which are problems for which the answer is binary (“yes” or “no”). If a decision problem has an efficient algorithm then it is in complexity class **P**, meaning *deterministic polynomial time*. If a decision problem is such that, when provided a candidate solution, one can efficiently check whether the candidate is correct, then the problem is in the complexity class **NP**, which stands for *non-deterministic polynomial time*. It is unknown whether or not a problem being efficient to check implies that it is efficient to solve. This question is captured by the famous **P** vs **NP** conjecture²⁰³ that is still unsolved. Another important complexity class is known as **NP-hard**. The qualification “hard” is used to indicate that a problem is at least as hard as any problem in that class. Formally, a problem is said to be **NP-hard** if an algorithm for solving such problems can be efficiently translated into an algorithm for solving any problem in **NP**. If a problem is both **NP-hard** and in **NP**, it is said to be **NP-complete**. Roughly speaking, a problem being **NP-complete** is strong evidence that there does not exist a provably efficient algorithm to solve it, be it of classical or quantum nature²⁰⁴.

In addition to decision problems, another important category of problem are *counting problems*. While a decision problem can be formulated as “Does a solution exist?”, a counting problem can be formulated as “How many solutions exist?”. From the class of problems

in \mathbf{NP} derives a complexity class of counting problems known as $\#\mathbf{P}$ (“Sharp-P”). The complexity class $\#\mathbf{P}$ is roughly defined as the set of problems which ask to count the number of solutions to problem instances, where candidate solutions can be efficiently checked. There is a sense in which problems in $\#\mathbf{P}$ are at least as hard as the corresponding problems in \mathbf{NP} . If one had a way to compute the number of solutions to a problem, then one could determine whether or not a solution exists.




Figure 2: Known relationship between the complexity classes discussed in this section. Note that some classes reside entirely inside other classes, which naturally defines a difficulty hierarchy. Note that $\mathbf{P}^{\#\mathbf{P}}$ and the polynomial hierarchy \mathbf{PH} are powerful complexity classes that relate to some of the major developments in quantum complexity theory.

The complexity classes \mathbf{P} , \mathbf{NP} , and $\#\mathbf{P}$ assume that the underlying computation is classical. In other words, \mathbf{P} is the class of problems efficiently solvable on a classical computer, \mathbf{NP} is the class of problems whose solutions are efficiently checkable on a classical computer, and $\#\mathbf{P}$ is the class of problems which count the number of solutions for problems whose solutions are efficiently checkable on a classical computer. The advent of quantum computing has also led to research on the quantum generalizations of these complexity classes. In particular the complexity class \mathbf{BQP} , short for *Bounded-error Quantum Polynomial-time*, refers to the set of problems that can be efficiently solved on a quantum computer. \mathbf{QMA} ,

short for *Quantum Merlin-Arthur*, refers to the set of problems such that when presented a candidate solution as a quantum state, one can use a quantum computer to efficiently verify the solution. Drawing a parallel from the discussion above, if a problem is **QMA**-complete then it is unlikely that there is an efficient solution even on a quantum computer. While the quantum analog of **#P**, **#BQP**, has been defined and has been shown to contain important physical problems, it turns out to be of equal complexity to **#P**⁶⁷. Although the exact relationships between the complexity classes remain largely open, certain containment relationships have already been proven (Figure 2).

At the intersection between theoretical computer science and condensed matter physics has emerged a research program that has recently come under the name *Quantum Hamiltonian Complexity*^{205,206}. Quantum Hamiltonian complexity seeks to gain insights on some of the most challenging open problems in quantum physics and chemistry. Most notably, it has been shown that finding the ground state energy of many simple locally interacting systems is **QMA**-complete^{207,208}. The three example open problems we mentioned in the opening of this section all have their own computational complexity characterizations: the 3D Ising model is **NP**-complete²⁰⁹, finding the ground state of Bose-Hubbard model is **QMA**-complete²¹⁰ and finding the universal functional in density functional theory is also **QMA**-complete⁶⁵. Finally we note that a comprehensive list²¹¹ of **QMA**-complete problems have been compiled and the reader may use it for further exploration.

3.2 Complexity theory of the electronic structure problem

The exponential complexity in handling the exact wave function is well known to be a major obstacle in solving many important problems in quantum chemistry such as electronic structure calculations. In Section 2.1 we elaborated on classical methods for circumventing the exponential complexity. Here we review rigorous results for characterizing computational hardness of various problems in electronic structure. Naturally, the **QMA**-completeness results for the ground state of locally-interacting systems mentioned before leads to specu-

lations about potential consequences for electronic structure²¹². Indeed, other than finding the universal functional for density functional theory⁶⁵, the problem of checking whether a given set of two-electron reduced density matrices (2RDM) is consistent with a single many-electron wave function, known as the *N-representability problem*, has also been proven **QMA**-complete⁶⁶. Solving the Hartree-Fock (HF) method, which intends to find a single Slater determinant that best approximates the true wave function, has been shown to be **NP**-complete⁶⁵. The **NP**-completeness has also been proven to persist even if one restricts to translationally invariant systems²¹³.

Although it is unlikely that one can find general and provably effective (quantum or classical) algorithms for **QMA**-complete or **NP**-complete problems, in practice, there are heuristic methods that have gained some empirical success. For example in the case of Hartree-Fock, the self-consistent field iteration serves as a local search algorithm which can converge to a local optimum. By chance, this can also be the global optimum, but proving that it is, or systematically finding the global minimum, is **NP**-complete. Although the *N*-representability problem is **QMA**-complete, heuristic methods building on constrained optimization of 2RDMs can nonetheless produce useful results in many cases^{214,215}. On quantum computers, the strategies for navigating the computational hardness of finding the eigenstates and eigenenergies revolve around exploiting the ability to maintain full wave functions encoded as multi-qubit entangled states. For example, the variational quantum eigensolver may start by preparing a physically motivated wave function ansatz that is otherwise hard to prepare classically, such as unitary coupled cluster²¹⁶. Another cause for optimism in light of the **QMA**-completeness of the ground state problem for general local Hamiltonians is that the Hamiltonians occurring in quantum chemistry problems are more restricted than the general setting considered in the computational complexity results⁶⁸.

Aside from the problems that are provably hard in the worst case, there are problems that *can* be efficiently solved on a quantum computer. Many such problems are related to the time evolution of a wave function under a Hamiltonian. In Section 2.3.2 we mentioned

some representative techniques for realizing unitary time evolution on a quantum computer. Each polynomial-time quantum algorithm is by itself a proof that the corresponding problem addressed by the algorithm is in **BQP**. In contrast to the **QMA**-completeness of evaluating the universal functional in density functional theory⁶⁵, the problem of evaluating the time-dependent effective Kohn-Sham potential can be efficiently addressed on a quantum computer²¹⁷ and is therefore in **BQP**. We categorize the problems and their respective computational complexity classes in Table 2.

Table 2: Representative problems in quantum physics and quantum chemistry and their computational complexity classes.

BQP	NP -complete	QMA -complete
Simulating unitary time evolution ⁷⁵	Ground state of 3D Ising models ²⁰⁹	N -representability ⁶⁶
Approximating Kohn-Sham potential in TDDFT ²¹⁷	Hartree-Fock on general quantum systems ⁶⁵	Universal functional in DFT ⁶⁵
	Hartree-Fock on translationally invariant systems ²¹³	Bose-Hubbard model ²¹⁰

3.3 Complexity theory in molecular vibronics

The motivation for solving chemical problems on a quantum computer often stems from the ability of the computer to *natively* simulate some or all of the dynamics of the chemical system directly on a quantum device. An appealing consequence of such native simulation is that the computing device often does not need to use a great deal of resources encoding the appropriate quantum states onto the hardware, thereby reducing the resource overhead significantly. While there are as yet no quantum computer systems whose underlying physical platform is naturally fermionic, the field of linear optical quantum computing has seen substantial progress toward performing quantum information processing with photons, which are bosonic systems. Given the resource-limited nascent state of gate-model quantum computing, algorithms that look toward solving bosonic chemical problems may be particularly valuable for showing early demonstrations of post-classical computing and to making an impact to chemistry in the immediate future.

In particular, the optical *boson sampling* model is one such proposal for a near-term application. The model consists of a series of optical modes, or pathways, where light is allowed to propagate. Along these modes are passive linear-optics devices, such as beam-splitters and phase shifters, which allow the exchange of photons to take place between two modes. This generates so-called number-mode entanglement in a bosonic Fock space, much like the fermionic Fock space corresponding to the occupation number basis, which conserves the total number of photons (note that a major difference between bosonic and fermionic Fock spaces is the ability for bosonic modes to be occupied by any number of photons, as opposed to spin-orbitals which have at most a single fermion). It is natural, then, to consider chemical systems which are described by bosonic processes as candidates for a quantum algorithm on this platform.

One particular process is molecular vibronic spectroscopy, which considers the vibrational transitions between two electronic states of a molecule. This process is fully characterized by the Franck-Condon profile (FCP)²¹⁸, which describes the spectra of vibronic modes in the harmonic regime. Recall that vibrational modes can be described by excitations of phonons, which are virtual bosonic particles. There are a number of important molecular properties that depend on the vibrational modes of a molecule, particularly those related to its interaction with light. Emission, absorption, and Raman scattering all depend on the electronic states of the molecule which in turn are influenced by the vibrational frequencies. For instance, in a spectroscopic setting, these properties could be used to extract information about unknown molecules. For molecular engineering of solar cells or photosynthetic materials, it may be necessary to have molecules meeting particular constraints on these properties. However, it may be prohibitively expensive in time or resources to experimentally test every candidate molecule; simulation of vibronic spectra then becomes a valuable tool. A quantum algorithm to solve this problem is detailed later in Section 5.3.2; here, we will motivate the development of such algorithm by discussing the complexity-theoretic limitations of solving the problem with only classical methods.

To describe molecular vibronic processes, it was shown by Duschinsky²¹⁹ that a linear relationship between the initial and final normal coordinates (\vec{q} and \vec{q}' , respectively) could be given by

$$\vec{q}' = U_{Dus}\vec{q} + \vec{d}, \quad (17)$$

where U_{Dus} is a real-valued matrix called the *Duschinsky* matrix, and \vec{d} is a displacement vector. Doktorov et al.²²⁰ show that this corresponds to a unitary rotation \hat{U}_{Dok} of the bosonic ladder operators

$$a'^{\dagger} = \hat{U}_{Dok}^{\dagger} a^{\dagger} \hat{U}_{Dok}, \quad (18)$$

$$\hat{U}_{Dok} = \hat{S}_{\Omega'}^{\dagger} \hat{R}_U \hat{S}_{\Omega} \hat{D}_{\delta}, \quad (19)$$

where \hat{D}_{δ} corresponds to a displacement operator dependent on \vec{d} , $\hat{S}_{\Omega'}$ (\hat{S}_{Ω}) a squeezing operator whose parameters depend on the harmonic angular frequencies of the transition, and \hat{R}_U a rotation corresponding to U_{Dus} ²²¹. The distribution of non-trivial transition amplitudes ω at 0 K gives rise to the FCP,

$$FCP(\omega) = \sum_{\mathbf{m}} |\langle \mathbf{m} | U_{Dok} | 0 \rangle|^2 \delta \left(\omega_{\text{vib}} - \sum_k^N \omega'_k m_k \right) \quad (20)$$

where $|0\rangle$ is the ground state of the initial vibronic potential, $|\mathbf{m}\rangle$ is an eigenmode of the final vibronic potential of eigenenergies $\mathbf{m} = (m_1, \dots, m_M)$, where $M = 3N - 6$ is the number of degrees of freedom (vibronic modes) in an N -atom molecule. The set $\{m_k\}$ consists of all possible configurations of phonons leading to allowed transitions of energy ω_{vib} . The inner product $\langle \mathbf{m} | U_{Dok} | 0 \rangle$, known as Franck-Condon integral, is thus a product-sum over the all phonon configurations contributing to ω . This product-sum corresponds to the computation of a *matrix permanent*. The permanent of an $n \times n$ matrix M is defined by

$$\text{perm}(M) = \sum_{\sigma \in S_n} \prod_{i=1}^n m_{i,\sigma(i)}, \quad (21)$$

which has a similar definition to the *determinant* of M , sans an additional product of each term with $\text{sgn}(\sigma)$. Colloquially, the permanent is often described as the “determinant with all + signs.”

Despite the permanent’s similarity to the determinant, the two functions in general exhibit dramatically different complexities. While the determinant can be efficiently computed (in $O(n^3)$ steps) by reduction to row-echelon form via elementary row operations (called *Gaussian elimination*), the permanent of a matrix is not invariant under row operations. Valiant²²² realized that the permanent was an archetypal problem for a complexity class of counting problems, which he coined $\#\mathbf{P}$. Recall that \mathbf{NP} problems can be thought of as asking “Does at least one satisfying assignment exist for some criteria?” Problems in $\#\mathbf{P}$, on the other hand, are characterized by finding *how many* satisfying assignments exist for the same criteria, a much more difficult task. It then followed, Valiant showed, that an exact and efficient algorithm for permanent computation, even for a matrix of only binary valued entries, would imply that $\mathbf{P}=\mathbf{NP}$. The expectation is, then, that $\#\mathbf{P}$ problems will be hard to compute. Indeed, the best known general algorithm today for exactly solving matrix permanents is due to Ryser²²³, requiring at least $O(n2^n)$ operations.

It is important to realize that, although matrix permanents may be hard to compute or even estimate, determining the FCP is more precisely a sampling problem. That is, although the hardness of the permanent is a necessary condition for the FCP sampling problem to be hard, it is not known to be a sufficient condition. There is, however, compelling evidence in the literature to suggest that the sampling problem as such is indeed hard as well. Particular instances of the FCP reduce to an instance of the recently-developed linear optics architecture, boson sampling^{64,224–227}. This prompted Huh and Yung²²⁸ to devise an algorithm for sampling from the FCP using an optical quantum computer, which is summarized in Section 5.3.2.

It remains unclear whether or not the family of chemically relevant problems is generally in the set of instances for which an optical quantum computer could significantly

speed up the calculation. Complexity-theoretic results hold for average-case Haar-random or Gaussian-random unitary matrices, but chemical problems do not correspond to choosing randomly from these sets. This is consistent with classical algorithmic approaches whose techniques leverage approximations that take into account the physical symmetries of the system to generate solutions²²⁹. However, these techniques will only give efficient and accurate solutions in some problem instances where the approximations are valid. It remains to be seen whether or not an efficient and fully generalizable classical approach exists. While the search for this approach continues, there is room for quantum computers to tackle the hardest of these chemical problems.

3.4 Chapter 3 Glossary

\vec{q} (\vec{q}')	Initial (final) normal coordinates
U_{Dus}	Duschinsky rotation
\hat{R}_U	Rotation corresponding to U_{Dus}
U_{Dok}	Doktorov rotation
\vec{d}	Displacement vector
a, a^\dagger	Bosonic annihilation and creation operators
\hat{D}_δ	Displacement operator
\hat{S}_Ω	Squeezing operator
\mathbf{m}	Eigenmode corresponding to photon transitions
m_k	Phonon configuration
ω	Distribution of vibronic transition amplitudes
ω_{vib}	Allowed energy transitions
N	Number of atoms in a molecule
M	Number of vibronic modes in an N -atom molecule
$FCP(\omega)$	Franck-Condon profile distribution
$\text{perm}(\cdot)$	Matrix permanent

$\delta(\cdot)$	Dirac delta function
S_n	Symmetric group of n elements
σ	Permutation of S_n
$sgn(\cdot)$	Sign, or parity, of a permutation (equal to $+1$ if even, and -1 otherwise)
M	An arbitrary matrix
n	Dimension of M
$m_{i,j}$	The i, j -th entry of the matrix M
NP	Non-deterministic polynomial time (complexity class)
P	Deterministic polynomial time (complexity class)
#P	Sharp-P (complexity class)
BQP	Bounded-probability quantum polynomial (complexity class)
QMA	Quantum Merlin-Arthur (complexity class)
#BQP	Sharp-BQP (complexity class)
2RDM	Two-electron reduced density matrices
HF	Hartree-Fock
TDDFT	Time-dependent density functional theory
FCP	Franck-Condon profile

4 Quantum simulation algorithms for fault-tolerant quantum computers

Quantum computation was born out of the idea that a controllable quantum system could be used to simulate the dynamics of a different quantum system. As described in the previous section, it has been shown that a quantum computer consisting of quantum bits and elementary unitary operations can solve the computational problem of simulating unitary evolution for a physically realistic system efficiently (c.f. Table 2). The purpose of this section is to provide a motivating narrative, guiding the reader through the significant advances in the

field of quantum simulation and its applications to the electronic structure problem. These advances have carried the idea of quantum simulation from a theoretical proposal to the premier application of near-term quantum computers.

In its simplest form, quantum simulation entails two steps. First, the quantum system of interest must be mapped onto the qubits of the quantum computer. Second, the unitary evolution of the simulated system must be translated into a sequence of elementary operations. The original proposals of quantum simulation in the early 1980s, by Manin²⁸ and Feynman²³⁰, were motivated by the inefficiency of simulating quantum systems using classical processors. Their revolutionary idea was that, while the scaling of the number of classical bits needed to store a wave function of N quantum systems grows as $O(\exp(N))$, the scaling of *quantum bits* needed to store this wave function is¹ $O(N)$. The caveat is that the amplitudes of a wave function stored in quantum bits cannot be efficiently accessed. Obtaining a classical description of the quantum state would require repeated quantum simulation and tomography which essentially eliminates all the savings. Instead, one can use quantum simulation to compute expectation values of observables³⁸, sample from a given distribution⁶³, or use the simulation as a subroutine in more complicated quantum algorithms^{231,232}.

In this setting, quantum computers can provide an exponential improvement in memory resources for quantum simulation compared to a naive approach using a wave function. Of course, there are classical methods for computing expectation values that do not require storing the wave function, but the ability to manipulate a discretization of the wave function directly opened a floodgate of new results in quantum algorithms for quantum chemistry^{37,40,41,75,233}.

It is important that the wave function can be efficiently evolved in time for many “physically realistic” Hamiltonians³⁷. That is, with respect to a desired accuracy and duration of the time evolution, the number of quantum gates (roughly, computation time) grows polyno-

¹Here the notation $O(N)$ indicates an asymptotic upper bound on the scaling of the algorithm that is linear in N . A tilde on top of the bound notation, e.g. $\tilde{O}(N)$, indicates suppression of polylogarithmic factors. In contrast to formally rigorous bounds, a tilde inside of a bound, e.g. $O(\sim N)$, indicates the bound is obtained empirically.

mially in system size, time, and inverse-precision. In practice, we are interested not only in the scaling but also in the actual amount of resources needed to execute an algorithm. Over the years, requirements for quantum simulation of a real-world system have been steadily lowered. The state-of-art quantum algorithms achieve scaling roughly logarithmic in inverse precision and linear in time. We review the progress in quantum simulations in Section 4.1.

In quantum chemistry, it is relatively uncommon to directly simulate the dynamics of a quantum wave function. Instead, the standard task is to determine the ground state energy of a quantum system. Solving such an eigenvalue problem exactly is a much more challenging problem than the Hamiltonian simulation. In fact, this problem known to be **QMA**-complete, for a general Hamiltonian, as discussed in Section 3. However, there are several assumptions that make this problem easier for realistic systems, as discussed in Section 3.

The quantum algorithms that we discuss in this section were designed for quantum computers that function without error. In practice, quantum computers are error-prone due to both imprecision of processes and unwanted interaction with the environment. These defects limit the length of the programs that can be accurately executed. The algorithms described in this section are typically beyond this limit. Thus, only proof-of-principle demonstrations for simple problems have been executed thus far^{55,234–236}.

Given the difficulty to engineer precise gates and stable qubits, one might be skeptical about the prospect of quantum algorithms requiring coherence. However, it is possible to correct the errors and the system of detecting and correcting for errors is called an *error correcting code*²³⁷. In brief, the principle behind quantum error correction is to use more physical qubits to establish a set of fewer, but higher-fidelity, equivalent qubits called *logical qubits*²³⁸. A rough, conservative estimate for the ratio of physical to logical qubits needed for reliable quantum computation is 10,000-to-1²³⁷ but depends on the realization of physical qubits, type of noise and the properties of the error correcting code. The ability to gain a digit in precision for a quantum computation per fixed cost of quantum resources (e.g. the

number of qubits and operations) is known as *fault-tolerance*²³⁹. Given enough time, a fault-tolerant quantum computer is able to execute an arbitrarily long circuit without sacrificing accuracy.

The emphasis on fault-tolerance is important for this section because here we focus on quantum algorithms which assume the underlying quantum computer is fault tolerant. In Section 5 we will discuss quantum algorithms for quantum chemistry which, *a priori*, do not require quantum error correction because of more shallow circuits. The remainder of this section is organized as follows (see also Figure 1). In Section 4.1 we discuss techniques for simulating general quantum systems on a quantum computer, which contains simulating time evolution under a given Hamiltonian (Section 4.1.1) as well as other components needed for extracting the spectrum of the Hamiltonian by phase estimation (Sections 4.1.2 and 4.1.3). In Section 4.2 we discuss applications of these techniques to quantum chemistry problems. Section 4.2 is further split into three subsections corresponding to high-level components of a quantum algorithm for quantum chemistry. In Section 4.2.1 we discuss methods for preparing the initial states for an algorithm. Next, we consider simulation methods in Section 4.2.2 and we conclude with methods for efficiently extracting useful information out of an algorithm in Section 4.2.3.

4.1 Quantum algorithms for energy estimation

The observation that quantum systems appear to be difficult to simulate was one of the forces driving early quantum computing research. Together with the algorithm for factoring³², Hamiltonian simulation³⁷ was one of the landmark achievements of quantum algorithms in the 1990s. Quantum simulations are closely connected to other quantum computing schemes such as linear systems²⁴⁰, thermal state generation²⁴¹ and quantum machine learning^{242–244}. In addition, Hamiltonian simulation together with phase estimation can be used for estimating eigenvalues.

4.1.1 General Hamiltonian simulation

We begin our discussion of Hamiltonian simulation by describing the breakthrough paper of Lloyd³⁷, which laid the theoretical foundation for much of the later work in this field. As described above, while the work of Manin²⁸ and Feynman²³⁰ proposed the idea of quantum computers for solving the memory storage issue with quantum simulation, the work of Lloyd gave a rigorous upper bound on the computational time for quantum simulation. We emphasize that Lloyd’s result applies to a particular approach to quantum simulation, known as gate-based or *digital quantum simulation*. This approach is distinct from an *analog quantum simulation*, for which there have been several recent ground-breaking experiments^{245–249}.

In an analog quantum simulation, the physical Hamiltonian of the controllable quantum system is engineered to correspond directly to the Hamiltonian of the system of interest. In a digital Hamiltonian simulation, the dynamics of the targeted system are approximated by a sequence of elementary quantum gates. We give a brief overview of quantum gates (see Appendix C.4) but recommend Nielsen and Chuang¹ for a reader unfamiliar with this formalism. The advantage of the digital simulation approach is its aim for *universality*: any feasible Hamiltonian can be digitally simulated. Unless otherwise stated, by “Hamiltonian simulation”, we mean digital Hamiltonian simulation of a closed quantum system under a time-independent Hamiltonian.

Informally, the task of Hamiltonian simulation is to construct a sequence of quantum gates which approximate the Hamiltonian evolution of an input state under the action of the operator e^{-iHt} . For an arbitrary Hamiltonian H , the number of elementary gates needed to construct $U(t) = e^{-iHt}$ grows exponentially with the number of qubits¹. Accordingly, Lloyd’s result, or any other efficient Hamiltonian simulation algorithm, requires the simulated Hamiltonian to have a special structure. The structure assumed in Lloyd’s analysis is that the Hamiltonian describes only *local interactions*. That is, the Hamiltonian of interest is

mapped to an N -qubit Hamiltonian which takes the form

$$H = \sum_{j=1}^{\ell} H_j, \quad (22)$$

where each H_j acts on at most k qubits and each $e^{-iH_j\Delta t}$ is easy to implement for a short time segment Δt .

The local interaction assumption leads to a quantum algorithm for approximating the time dynamics. The key insight is to use the *Trotter decomposition* to approximate the exact evolution as $U(t) \approx (e^{-iH_1 t/n} \dots e^{-iH_\ell t/n})^n$. From the local interactions assumption, each of the factors $e^{-iH_j t/n}$ is a unitary which acts on a constant number of qubits (i.e. independent of N). These local unitary transformations can, in principle, be decomposed into a number of elementary gates that is independent of N .

The formal result from Lloyd's paper is that, for a Hamiltonian of ℓ terms describing k -body interactions, $U(t)$ can be approximated within error ϵ by a sequence of $O(\ell\tau^2/\epsilon)$ elementary quantum gates²⁵⁰, where $\tau = \|H\|t$. Assuming that the number of terms ℓ in the Hamiltonian scales polynomially with the system size N (such as the number of particles or the number of qubits), then the scaling of the number of gates is polynomial in N , τ , and $1/\epsilon$. Thus, the Trotter decomposition for local-interaction Hamiltonian simulation is efficient. It is worth noting that Lloyd also gives an algorithm for simulating open-system dynamics. Following this first algorithm for digital quantum simulation, several questions drove the field of quantum simulation:

- Does there exist an algorithm for Hamiltonians which are not necessarily local?
- Is it possible to improve the scaling in terms of t , $1/\epsilon$, and the norm of the Hamiltonian $\|H\|$?
- Can we move beyond Hamiltonian evolution and implement evolution of open systems or evolution under a time-dependent Hamiltonian?

- Can we exploit structure in quantum chemistry Hamiltonians to improve algorithmic performance?

This last question, in particular, has led to major advances in quantum simulation for quantum chemistry and will be the focus of Section 4.2.2. But first, we will briefly review the major advances in the field of general Hamiltonian simulation.

Dodd et al.²⁵¹ and Nielsen et al.²⁵² built on the work of Lloyd³⁷ and gave algorithms for simulating Hamiltonians that are a sum of tensor products of Pauli matrices. Independently, Aharonov and Ta-Shma²⁵³ introduced the *sparse Hamiltonian* paradigm and thus addressed the first two questions above. This approach provides a means for efficiently simulating (i.e. with gate count $\text{poly}(N, t, 1/\epsilon, d)$) Hamiltonians which have fewer than d entries in each column (d -sparse) and includes local Hamiltonians as a special case. A sparse Hamiltonian is a Hamiltonian for which $d \in \text{poly}(n)$. It is natural to ask how to represent such a Hamiltonian since storing (an exponential number of) the coefficients in memory would require an exponential overhead. Indeed, one needs to have a way of efficiently computing these coefficients in superposition. Instead of constraining ourselves to one implementation, it is customary to assume access to the Hamiltonian through oracles

$$O_{loc} |k, r, z\rangle = |k, r, z \oplus l\rangle, \quad (23)$$

$$O_{val} |l, r, z\rangle = |l, r, z \oplus H_{l,r}\rangle. \quad (24)$$

The oracle O_{loc} locates the column l of the k -th non-zero element in the row r . The oracle O_{val} then gives the (non-zero) value of the matrix element $H_{r,l}$. Both oracles are by construction involutory and therefore unitary.

If such oracles can be constructed from a polynomial number of elementary gates, we say that the Hamiltonian is row-computable. An important subclass of d -sparse Hamiltonians are *d-local* Hamiltonians. A Hamiltonian is d -local if it can be written as a sum of polynomially (in the number of qubits) many terms where each term acts non-trivially on at most d

qubits. Depending on the context, the qubit may or may not be spatially close. As such many physically-relevant Hamiltonians are sparse²⁵⁴ or even local^{186,255}.

Several other oracles are used in Hamiltonian simulation literature. Childs and Wiebe²⁵⁶ introduced the unitary access oracle. In this model, the Hamiltonian is represented as a linear combination of unitaries $H = \sum_{j=0}^{L-1} \alpha_j V_j$. For an efficient simulation, $L \in \text{poly}(N)$ and therefore the coefficients α_j can be represented classically. The unitaries V_j are given through an oracle

$$\text{select}(V) |j\rangle |\psi\rangle = |j\rangle V_j |\psi\rangle \quad (25)$$

that acts directly on the data. Berry et al.¹⁸⁷ showed that one query to $\text{select}(V)$ can be implemented with d^2 calls to O_{loc} and O_{val} . Tensor products of Pauli matrices are a special case of the unitaries V_j .

Later, Low and Chuang¹⁸⁹ introduced the signal oracle U_{signal} such that in a subspace flagged by an ancilla in a signal state $|G\rangle$, $H = \langle G| U_{signal} |G\rangle$. This case arises from a linear combination of unitaries. Other Hamiltonian access models include efficiently *index-computable* Hamiltonians²⁵⁷, special data structures²⁵⁸ and Hamiltonian encoded into a density matrix as a *quantum software state*^{243,259}. Since a density matrix is Hermitian, it is possible to exponentiate it and use for generating time evolution. While not commonly used in quantum chemistry yet, density matrix exponentiation has been used in multiple quantum machine learning algorithms²⁴³.

We estimate the cost of the Hamiltonian simulation in terms of the number of calls or *queries* to these oracles as well as the number of gates. Since each query requires a construction of the Hamiltonian, query complexity often dominates the cost of an algorithm. The concept of an oracle provides a useful abstraction for designing algorithms for a wide spectrum of Hamiltonians. In practice, however, it is not always necessary to explicitly construct an oracle directly if the Hamiltonian has a structure that can be exploited directly. An example of such a structure is a sum of Pauli terms^{188,196,260}. Along the line of oracular versus non-oracular approaches, we review the results in Hamiltonian simulations following

two different avenues.

In the first one, the Hamiltonian is first decomposed into a sum of terms $H = \sum_j H_j$, where each $e^{iH_j t}$ can be implemented directly, and use a Trotter decomposition to approximate the evolution under the sum of Hamiltonians²⁵³. These approaches are typically referred to as *product formula* algorithms. The original Aharonov and Ta-Shma work established a query complexity of $O(\text{poly}(N, d)(\tau^2/\epsilon))$, which matches the gate complexity of Lloyd’s algorithm, but applies to a broader class of d -sparse Hamiltonians²⁵⁰. This set the stage for a wave of improvements in the query complexity. First, by more sophisticated Trotter-Suzuki decompositions²⁶¹, the approximation error can be made to depend on progressively-higher orders in τ ^{196,260,261}. Intuitively, one can achieve a more favorable error scaling by using more precise approximations for the exponential of a sum. In the simplest case, the symmetric Trotter formula $e^{(A+B)\Delta t} \approx e^{A\Delta t/2} e^{B\Delta t} e^{A\Delta t/2}$ suppresses $O(\Delta t^2)$ errors. Second, by considering more favorable schemes for Hamiltonian decomposition^{256,257}, most modern-day simulation algorithms scale in terms of $\tau_{max} = t\|H\|_{max}$ instead of τ , where $\|H\|_{max}$ is the absolute value of the largest entry of H as a matrix in the computational basis.

The second approach comes from the equivalence between continuous- and discrete-time quantum walks²⁶². This work improved the dependence on time and sparsity but did not match the error scaling of contemporaneous algorithms. Unlike the previous results, it did not rely on Hamiltonian decomposition and can be used for certain non-sparse Hamiltonians. Childs²⁵⁷ showed that given oracle access to a Hamiltonian with eigenvalues λ , one can construct a unitary quantum walk operator with eigenvalues $e^{\pm i \arcsin \lambda}$. The computationally difficult part is to un-compute the arcsin in the phase. Childs does this by estimating the phase $\arcsin \lambda$ using phase estimation and coherently applying the sin function onto the phase. The use of phase estimation is responsible for a relatively poor scaling in the error ϵ .

While these algorithms reduced the asymptotic complexity of Hamiltonian simulation, this came at the cost of increased sophistication of the methods. A breakthrough in quantum

simulation methods came in 2014 when Berry et al. presented an algorithm that exponentially improved the dependence on precision¹⁸⁷. A series of improvements resulted in the development of a relatively straight-forward technique, which considers a decomposition of the Hamiltonian into a linear combination of unitary operators (LCU) $H = \sum_j \alpha_j U_j$ ¹⁸⁸. The Hamiltonian evolution is first divided into shorter segments. Then, each segment is approximated by a truncated Taylor series

$$e^{-iH\Delta t} \approx \sum_{k=0}^K \frac{(-i\Delta t)^k}{k!} \left(\sum_j \alpha_j U_j \right)^k, \quad (26)$$

which, itself, is a linear combination of unitaries. A technique called *oblivious amplitude amplification*¹⁸⁷ is then used to turn black-box implementations of the individual unitaries into a desired linear combination. Oblivious amplitude amplification allowed for the implementation of algorithms based on existing probabilistic schemes²⁵⁶ with near-perfect success. LCU has a query complexity of $O(d^2\tau \frac{\log(d^2\tau_{\max}/\epsilon)}{\log \log(d^2\tau_{\max}/\epsilon)})$, which is optimal in inverse-precision. This approach is particularly appealing for quantum chemistry application where each U_j is a tensor product of Pauli matrices. These Hamiltonians can be implemented without the use of oracles requiring $O(T \frac{L(n+\log L)\log(T/\epsilon)}{\log \log(T/\epsilon)})$ gates, where n is the number of the qubits in the simulated system, L is the number number of terms in the Hamiltonian and $T = \sum_{j=0}^{L-1} \alpha_j$.

The major open question was whether or not this query complexity scaling could be reduced to being additive in its τ and ϵ dependence, rather than multiplicative. Such a dependence was known to be a lower bound on the asymptotic query complexity¹⁸⁷. In 2016, Low and Chuang developed algorithms based on *quantum signal processing*¹⁹⁰ and *qubitization*¹⁸⁹ which achieve this provably optimal asymptotic scaling in query complexity of $O(d\tau_{\max} + \frac{\log(1/\epsilon)}{\log \log(1/\epsilon)})$, achieving the sought-after “additive” approximation error.

Quantum signal processing refers to the ability to transform a unitary $W = \sum_{\lambda} e^{i\theta_{\lambda}} |u_{\lambda}\rangle\langle u_{\lambda}|$ as

$$W \rightarrow V_{\text{ideal}} = \sum_{\lambda} e^{ih(\theta_{\lambda})} |u_{\lambda}\rangle\langle u_{\lambda}| \quad (27)$$

for any real function h . Using the phase-kickback trick²⁶³, one can modify the spectrum by manipulating an ancilla qubit. The transformation is then implemented through single qubit rotations on the ancilla and repeated applications of control- W . Low et al.²⁶⁴ fully characterize the transformations that can be implemented this way. In short, the transformations are first approximated by a tuple of polynomials whose degree determines the query complexity of the algorithm. If the polynomials form an *achievable tuple*, they can be turned into single qubit rotations and implemented through quantum signal processing.

Since the walk operator given by Childs²⁵⁷ has eigenvalues $e^{\pm i \arcsin \lambda}$ where λ is an eigenvalue of the Hamiltonian, $h(\theta) = -\tau \sin(\theta)$ gives time evolution. The transformation is approximated by the Jacobi-Anger expansion that satisfies the criteria of achievability and its quick convergence results in the optimal scaling of the algorithm. Qubitization²⁶⁵ takes these ideas one step further and explains how to combine LCU with quantum signal processing. The core of this idea is to modify the signal oracle to obtain 2-dimensional invariant subspaces analogous to the construction in Grover's algorithm³⁵.

The quantum signal processing approach was later simplified and generalized into singular value transformations by Gilyén et al.²⁶⁶ and can be applied to a wide array of applications beyond quantum simulation. More recent work has focused on the possibility of improvements in the scaling with different norms of the Hamiltonian and spectral amplification^{265,267}.

Low and Wiebe¹⁹⁷ introduced Hamiltonian simulation in the interaction picture. A Hamiltonians with components $H = A + B$ can be transformed into an interaction picture Hamiltonian as $H_I(t) = e^{iAt} B e^{-iAt}$. The time evolution in the interaction picture then corresponds to the evolution under a time-dependent Hamiltonian $H_I(t)$. The dominant scaling comes from simulating the interaction terms B while A is chosen to be a part of the Hamiltonian that is easy to simulate. If the norm of A is much larger than the norm of B , say for a diagonally dominant matrix, interaction picture simulations can lead to an exponential improvement over existing techniques. Recently, Low²⁶⁵ combined multiple methods^{189,190,197} to further improve scaling with respect to the norm of the Hamiltonian. This approach

improves the dependence on sparsity by utilizing information about the spectral norm.

Quantum simulation algorithms have been developed in other important dynamical regimes. The work of Wang and Wossnig²⁶⁸ and Kimmel et al.²⁵⁹ considered simulation of non-sparse Hamiltonians. Advances paralleling the time-independent results have been carried out for time-dependent Hamiltonians, using Trotter-Suzuki decompositions^{269,270}. Very recently, a truncated Dyson series algorithm was developed¹⁹⁹ which achieves the logarithmic scaling in inverse-precision, analogous to the truncated Taylor series method for time-independent simulation.

Beyond unitary dynamics, a small line of inquiry has investigated general algorithms for open quantum dynamics, beginning with the case of local Markovian (or local Lindblad) dynamics²⁷¹. Algorithms with improved scaling were recently developed, drawing on techniques from unitary simulation such as improved product decompositions, sparse Hamiltonian methods, and the linear-combination of unitaries approach^{194,195}.

It is quite remarkable that algorithms with optimal asymptotic scaling have been discovered. In practice, however, we must consider the exact gate counts to determine which algorithm is the best in a particular instance. A recent study²⁷² showed that different algorithms fare better in different regimes. In particular, product-formula based algorithms required fewer resources than quantum signal processing in a numerical simulation for systems between 10 and 100 qubits. In Section 4.2, we will turn to such considerations in the context of Hamiltonian simulation for quantum chemistry.

4.1.2 Quantum phase estimation

Hamiltonian simulation is rarely used as a standalone algorithm. For the purposes of quantum chemistry, its main application is as a subroutine in the quantum phase estimation algorithm (QPEA)^{231,263}, which is sometimes referred to as the “von Neumann trick”²⁷³. It gives an exponential advantage to many of the algorithms reviewed in the next section. Given a unitary operator $e^{i\Phi}$ and an approximate eigenstate of that operator prepared on the

quantum computer, the quantum phase estimation algorithm is used to efficiently output a binary representation of the corresponding phase Φ . In our case, the operator corresponds to a unitary evolution e^{iHt} and phase estimation is used to extract information about the spectrum of the Hamiltonian.




Figure 3: Circuit performing quantum phase estimation with three ancilla qubits. First, the Hadamard gate (labeled “Had”) is applied on each ancilla qubit to create a uniform superposition. Next, a series of controlled operations U that corresponds to Hamiltonian evolution e^{-iHt} are applied. Finally, the inverse QFT is used, see 4.1.3. The input is given by one register (consisting of the top three qubits in this specific diagram) initialized to zero and a second register containing the desired eigenstate of H . (The general case where the second register contains a superposition of eigenstates is discussed in the main text.) After execution, the first register stores an approximation of the eigenenergy of $|\psi_m\rangle$. Please see AppendixC.4 or Nielsen and Chuang¹ for an introduction to quantum circuits.

A circuit representation of the algorithm is shown in Figure 3. The QPEA uses two separate registers of qubits.

The first register of T ancilla qubits is used to read out the binary representation of λ_m . For simplicity, assume that the state $|\psi_m\rangle$ in the second register is an eigenstate of the unitary U for which we want to compute the eigenenergy λ_m .

After each ancilla qubit is initialized in the state $|0\rangle$, a Hadamard gate is applied to each to prepare this register in an equal superposition of all computational basis states $\frac{1}{\sqrt{2^T}} \sum_x |x\rangle$. In the first step of reading out λ_m , the phase $(e^{-i\lambda_m t})^{2^k}$ is imprinted on the k th ancilla qubit as

$$(|0\rangle + |1\rangle) |\psi_m\rangle \rightarrow \left(|0\rangle + e^{-2^k i \lambda_m t} |1\rangle \right) |\psi_m\rangle. \quad (28)$$

This is achieved via “phase-kickback”^{34,263} by applying a controlled- U^{2^k} between this qubit and the state preparation register in $|\psi_m\rangle$. Finally, the inverse quantum Fourier transform (see Section 4.1.3) is applied to the ancilla register to convert the state to a computational basis state $|x_1x_2\dots x_T\rangle$ that expresses the binary representation of $\lambda_m t/2\pi \approx 0.x_1x_2\dots x_T$. After execution of the algorithm, measurement of the first register will yield a binary approximation of the eigenenergy λ_m with a high probability. It can be shown that with probability at least $1 - \epsilon$, the estimation of λ_m is accurate to $(T - \lceil \log(2 + \frac{1}{2\epsilon}) \rceil)$ bits¹. Roughly speaking, each ancilla adds another digit of accuracy.

Let us now examine the case when the second register is not an eigenstate but rather a superposition $\sum_m a_m |\psi\rangle$. One can perform the above analysis and convince oneself that phase estimation will yield $\sum_m a_m |\tilde{\lambda}_m t/2\pi\rangle |\psi_m\rangle$ where $\tilde{\lambda}_m$ is a fixed-point approximation of the eigenenergy λ_m . After measuring of the first register, the second register will collapse in the corresponding eigenstate. The probability of obtaining the eigenstate λ_m is $|a_m|^2$, i.e. the squared overlap between the eigenstate with the prepared superposition. Therefore, it is possible to obtain multiple eigenstates and eigenenergies of H by running QPEA repeatedly with an appropriately chosen initial state $|\phi\rangle$, highlighting the importance of the state preparation step, as discussed in Section 4.2.1.

QPEA is closely linked to Hamiltonian simulation algorithms because of the controlled- U^{2^k} transformations. Each of these transformations can be decomposed into a sequence of gates by first decomposing U^{2^k} into a sequence of gates using Hamiltonian simulation methods from Section 4.1, then augmenting each gate into a controlled gate with a certain ancilla qubit. Accordingly, advances in quantum algorithms for Hamiltonian simulation lead to improvements in the performance of QPEA. We note, however, that certain methods, such as quantum random access memory (quantum RAM)²⁷⁴, could provide an alternative means of implementing the powers of controlled unitaries.

In the standard QPEA, the accuracy of the estimated eigenvalue is determined by the number of ancilla qubits. As detailed by Aspuru-Guzik et al.⁴⁵ and Dobšiček et al.²⁷⁵, the

number of ancilla qubits can be greatly reduced, while maintaining the same precision by using the *iterative quantum phase estimation algorithm* (IPEA). Feeding back on the rotation angle of the quantum gates, the phase estimation is improved in each step.

Further improvements to the quantum phase estimation algorithm have been made by introducing Bayesian inference techniques^{276,277}. In this way, the maximum amount of information is extracted from previous measurements in order to inform the rotation angles for the next evaluation of the algorithm. Wiebe and Granade²⁷⁷ have shown that the eigenenergy and its uncertainty can thus be inferred directly rather than iteratively. This Bayesian estimation approach to quantum phase estimation has recently been shown experimentally²⁷⁸, demonstrating its robustness to noise and decoherence.

The use of the time evolution operator in phase estimation is not necessary for the quantum phase estimation algorithm. Recent work^{279–281} has investigated using unitaries which encode the spectrum of the Hamiltonian, but are easier to implement than e^{-iHt} . As long as the alternative unitary is determined by a known invertible function f of the Hamiltonian, as $e^{-if(H)t}$, then the measured eigenphase of this unitary can be used to infer the corresponding eigenvalue of H . The unitary $e^{i\arccos(H/\lambda)}$ can be implemented using a quantum walk operator¹⁸⁹ which requires fewer gates than time evolution.

4.1.3 Quantum Fourier transform

Many important quantum algorithms, including Shor’s factoring algorithm³², rely on the quantum Fourier transform (QFT)²⁸² as a subroutine. The QFT underlies the problem of Forrelation²⁸³, which is said to “capture the maximal power of quantum computation”. In the context of the quantum phase estimation algorithm, the QFT is the final step before measuring the ancilla qubits that carry the binary representation of the energy estimate. After the various powers of the eigenphases have been applied to the ancilla qubits (c.f. Figure 3), the QFT is used to convert these powers of phases into a binary representation of the eigenenergy on these qubits.




Figure 4: Quantum Fourier transform circuit for four qubits, reproduced from Nielsen and Chuang¹. SWAP gates that reverse the order of the states at the end are not shown. The Hadamard gates are labeled “Had”, gates R_k indicate a rotation over the angle $e^{2\pi i/2^k}$, $0.j_1 j_2 \dots j_n$ denotes a binary fraction.

The quantum Fourier transform operation can be understood as a unitary change of basis from the computational basis to the Fourier basis. Mathematically, it closely resembles its classical counterpart, the discrete Fourier transform. Given an orthonormal basis $|0\rangle, |1\rangle, \dots, |N-1\rangle$, the action of the quantum Fourier transform (QFT) on a basis state $|j\rangle$ is given by

$$\text{QFT } |j\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{2\pi i j k / N} |k\rangle. \quad (29)$$

We now change to a binary representation $j = j_1 2^{n-1} + j_2 2^{n-2} + \dots + j_n 2^0 =: j_1 j_2 \dots j_n$ with n bits, where $N = 2^n$. This allows us to write the state $|j\rangle$ as a multi-qubit state $|j_1, j_2, \dots, j_n\rangle$. The Fourier transform now assumes the following form:

$$\text{QFT } |j_1, j_2, \dots, j_n\rangle = \frac{(|0\rangle + e^{2\pi i 0.j_n} |1\rangle) (|0\rangle + e^{2\pi i 0.j_{n-1} j_n} |1\rangle) \dots (|0\rangle + e^{2\pi i 0.j_1 j_2 \dots j_n} |1\rangle)}{2^{n/2}} \quad (30)$$

Here, the terms $0.j_1 j_2 \dots j_n$ denote a binary fraction.

Computing the quantum Fourier transform requires a sequence of Hadamard gates and controlled rotations. In Figure 4, we show this circuit for four qubits. The following operations are applied to each qubit, successively from top to bottom qubit: a Hadamard gate followed by rotations R_k that are controlled by the qubits below. The rotation angle of the gate R_k is given by $e^{2\pi i/2^k}$. As indicated in the figure, the final state of the register is equal to

Eq. 30 up to a reversal of the qubit state order. Although this circuit is complex and dense, it often serves merely as a subroutine within other algorithms. It will usually be written as “QFT” or “FT” in circuit diagrams. The inverse Fourier transform, given by the Hermitian conjugate QFT^\dagger , is at least as ubiquitous.

The quantum Fourier transform has also found application in simulating chemical dynamics. The quantum “split-operator” method^{75,273} is analogous to classical exact wave function simulation techniques such as matching pursuit/split-operator Fourier transform¹⁷⁶. During simulation, the kinetic and interaction/potential energy terms can both be implemented diagonally by applying the quantum Fourier transform to change between the momentum and position basis. It has been shown that this technique leads to quantum speedups both for first- and second-quantized problem formulations^{40,75,254,284}.

4.2 Determining the ground states energies of quantum chemistry Hamiltonians

Beneath the massive computational resources devoted to determining properties of molecules and materials lies the electronic structure problem. The importance of this problem derives from the fact that the quantum mechanical eigenenergies and wave functions of the electrons in a molecular system are essential for understanding and predicting industry-relevant quantities, including reaction rates, binding energies, and molecular pathways. In particular, the ground state energy as a function of the nuclear coordinates (the ground state energy manifold) of a molecular system is sufficient for gleaning many of these properties.

In the late 1990s, quantum computing scientists began to develop approaches for estimating static properties of molecular systems using techniques for simulating time dynamics. The standard approach requires first preparing the quantum computer in the desired quantum state (e.g. an approximation of the ground state of a molecular system) and then extracting the properties of interest (e.g. the ground state energy) using QPEA. We will review the development of necessary techniques in this section.

4.2.1 State preparation

The first step of any quantum algorithm is the preparation of an initial state. The success of an algorithm for determining ground state energy depends on the quality of the state preparation. The ansatz for the preparation comes from a classical approximation for a ground state of a given Hamiltonian. Such a state must have two properties. First, the ansatz needs to have a significant overlap with the ground state. Second, the state has to be efficiently preparable on a quantum computer. Examples of state preparation include the Hartree-Fock (HF) state, states from coupled-cluster methods, or states obtained using the method of adiabatic quantum evolution. Quantum computers are particularly useful when classical methods cannot provide answers with high-accuracy such as for strongly correlated systems.

In his pioneering work, Zalka⁴¹ considered discretizing the wave function and initializing simulation through a series of controlled rotations. Grover and Rudolph²⁸⁵ gave a similar algorithm to prepare a state corresponding to an efficiently integrable probability distribution. However, this would be a costly procedure in practice.

Abrams and Lloyd³⁸ focused on state preparation in more detail and described concrete initialization procedures. They pointed out that state preparation is conceptually easier in second quantization, where a Fock state is represented by a simple product state of qubits in $\{|0\rangle, |1\rangle\}$ states. Their algorithm begins with all qubits reset to $|0\dots 0\rangle$. Fermi statistics are automatically accounted for in the creation and annihilation operators. In first quantization, one must initialize a completely antisymmetric state. Antisymmetrization can be enforced by reversing a suitable sorting algorithm and uncomputing ancillas to make the procedure reversible²⁸⁶. Recently, Berry et al.²⁷⁹ improved the antisymmetrization procedure by using a sorting network and optimizing components of the algorithm. This algorithm achieves a gate count scaling of $O(\eta \log \eta \log N)$ and a circuit depth of $O(\log \eta \log \log N)$ where η stands for the number of particles and $N \geq \eta$ for the number of single-particle basis functions. These are, respectively, a polynomial and an exponential improvement over the previous algorithm.

The choice of the initial state for ground state energy estimation was considered by Aspuru-Guzik et al.⁴⁵. HF is the usual starting point for state preparation, however, it might not always have a sufficient overlap with the ground state. Tubman et al.²⁸⁷ recently investigated the support of HF on the ground state. For small molecules, HF provides a good approximation for the ground state, for example, 0.77 for FeMoCo²⁸⁷. The applicability of other methods was investigated in cases when HF fails. Wang et al.⁴⁶ investigated preparing states which are generated with the multi-configurational self-consistent field (MCSCF) approach. These states have non-zero overlap with a polynomial number of computational basis states, for which there are efficient methods of preparation²⁸⁸. In addition to having better overlap with the ground state, MCSCF states express the electron correlation needed to represent low-level excited states.

Kassal et al.⁷⁵ combined state preparation with a proposal for preparing the nuclear ground-state wave function to investigate the simulation of a chemical reaction without the Born-Oppenheimer approximation. They conclude that simulating the complete nuclear and electronic wave function in a chemical reaction of four or more particles is more efficient than using the Born-Oppenheimer approximation.

Recently, Tubman et al.^{287,289} introduced the adaptive sampling configuration interaction method (ASCI). ASCI allows for the generation of a small number of determinants while accounting for at least 90% of the wave function. Sugisaki et al.²³³ extended the HF state preparation by proposing an algorithm for efficiently preparing an exponential number of Slater determinants. Their proposal is particularly appealing for preparing configuration state functions of open-shell molecules.

Adiabatic state preparation (ASP), proposed by Aspuru-Guzik et al.⁴⁵, is a method for transforming the initially-prepared HF state into an approximation of the FCI ground-state wave function. The idea draws on the method of adiabatic quantum computing²⁹⁰ (c.f. Section 5.3.1). ASP initializes the register in the HF state. For a second-quantized Hamiltonian, the HF state may be prepared as a product state on the qubits and is, therefore,

easy to prepare⁴⁵. The state is then evolved with respect to a circuit that approximates an adiabatic change in the Hamiltonian starting from the HF Hamiltonian and ending with the full Hamiltonian. The adiabatic evolution is then digitally simulated either with a product formula decomposition²⁷⁰ or truncated Dyson series^{197,199}. If the Hamiltonian and its first derivative are upper-bounded by a constant, a good approximation of the ground state can be reached for evolution time $T = O\left(\frac{1}{\gamma^2}\right)$ where γ is the minimum gap (except for pathological cases²⁹¹). This is a so-called adiabatic regime where the evolved state has support at least $1 - O\left(\frac{1}{T}\right)$ on the ground state²⁹². A number of strategies can be used to increase the support, see Wiebe and Babcock²⁹³, Lidar et al.²⁹⁴, Kieferová and Wiebe²⁹⁵. Roland and Cerf²⁹⁶ proposed a method for reaching the adiabatic regime faster by modifying the speed of Hamiltonian change throughout the evolution. Note that it is not always necessary to reach the adiabatic regime to achieve a significant overlap with the ground state. Symmetries in the spectrum can be exploited even when the gap is exponentially small²⁹⁷. Crosson et al.²⁹⁸ demonstrated that this effect is common even for instances that do not exhibit noticeable symmetries. Standard quantum chemistry methods can be used to lower-bound this gap⁴⁵ even though exactly computing it can be demanding.

In Section 5 we will discuss several other state preparation methods for second quantization which have been introduced for the purpose of the variational quantum eigensolver algorithm, but which also apply to the quantum phase estimation algorithm.

Another possibility is to initialize the registers to approximate a thermal state⁴⁵. While creating a thermal state is a computationally difficult problem by itself, an approximation may suffice. Several quantum algorithms can be used for accomplishing this task^{241,266,299}.

4.2.2 Hamiltonian simulation for quantum chemistry

Independently of Lloyd’s quantum simulation algorithm³⁷, Zalka²⁷³ introduced a quantum algorithm for simulation of single- and many-body quantum systems. Zalka suggested that a wave function can be discretized into an l -bit quantum register. The Green’s function for

a quantum particle is then approximated by a product of exponentials with potential and kinetic energy in the exponents. The term corresponding to potential energy is diagonal and thus relatively straightforward to simulate. Kinetic energy terms can be diagonalized using the quantum Fourier transform. Zalka proposed to take similar steps for simulating many-body systems and field theories. Similar ideas also appear in Wiesner⁴⁰.

At the same time, Abrams and Lloyd⁹⁹ discussed the subject of Hamiltonian simulation in first and second quantization in detail using Trotter decompositions and block diagonalization. They pointed out that the Hamiltonian can be expressed more easily in first quantization but the basis set necessary for expressing the quantum state is smaller in second quantization.

Aspuru-Guzik et al.⁴⁵ first realized that quantum computers can be used for determining ground-state energies. This work lays the grounds for further research in quantum simulations for quantum chemistry, discussing choices of chemical basis, Hamiltonian mapping, simulation, and energy estimation. To map a many-body system onto a discretized Hamiltonian, one needs to make several choices.

First, one must choose how to map the physical Hamiltonian onto the qubits. The upshot is that the additional structure in these Hamiltonians can be exploited to improve the performance over general algorithms. Such improvements should also take constant overheads into account, which are necessary for determining the feasibility of running a useful instance of a problem.

Second, in traditional quantum chemistry, it is standard to make several simplifications to the Hamiltonian which yield a finite-dimensional operator that can be handled more easily. The Born-Oppenheimer approximation allows one to freeze out the slow-moving ionic degrees of freedom, leaving only the important electronic degrees of freedom. Up to a few exceptions, most quantum algorithms work with this approximation.

Third, one needs to make a choice between first and the second quantization. Classical methods for simulating evolution in the first quantization are relatively rare because of the

difficulty of storing the wave function.

Last, one needs to find a suitable basis. The most common choice of basis considered so far are the Gaussian basis sets of Pople quantum chemistry³⁰⁰, for example the STO-nG basis set¹⁹². The plane wave basis was long seen as unsuitable because of the high number of orbitals required for approximating a wave function but provides a much simpler representation of the Hamiltonian. However, the recent improvement in the scaling in the number of basis terms may make it a front-runner for quantum chemistry quantum algorithms^{198,281,301,302}. Babbush et al.³⁰¹ showed that the electronic structure Hamiltonian requires only $O(N^2)$ terms in a plane wave basis while $O(N^4)$ terms are necessary for Gaussians. Kivlichan et al. further refined this to show that single Trotter steps of the electronic structure Hamiltonian can be simulated in $O(N)$ depth with only nearest-neighbor connectivity on a line³⁰². White³⁰³ proposed a so-called Gausslet set that combines Gaussian with wavelet features.

The number of choices led to several avenues of research. In second quantization, Whitfield et al.¹⁸⁶ used the Jordan-Wigner transform to create a local Hamiltonian and gave explicit formulas to compute the matrix elements. Their algorithm is then applied to H_2 , decomposing the operations into elementary gates and simulating their quantum algorithm. A similarly detailed study for LiH with quantum error correction was performed by Jones et al.³⁰⁴.

Simulations in real space have also been proposed⁷⁵. Building on the work of Zalka⁴¹ and Wiesner⁴⁰, they use the QFT to simulate the kinetic term in the Hamiltonian but implement the gate sequence and give complexity estimates for their algorithm.

Toloui and Love²⁵⁴ chose a different approach and used sparse matrix simulation techniques to evolve a state under the configuration interaction matrix. The representation of the Hamiltonian in the basis of Slater determinants (the configuration interaction matrix) is sparse and therefore techniques from Section 4.1 can be used for simulation. This technique is also more space efficient compared to second quantized techniques using Fock states.

While the above algorithms are efficient in terms of asymptotic scaling, their practicality

for modest-sized quantum computers remained unclear. The advent of quantum simulators brought a number of landmark *resource estimate* studies^{48,186,305}. Wecker et al.⁴⁷ investigated the quantum computational resources needed to simulate a molecular Hamiltonian with twice the number of orbitals that can be handled with classical methods. They paint an ambivalent picture for the future of Hamiltonian simulation in chemistry. The pessimistic conclusion was that, while the required spatial resources (i.e. the number of qubits) were just an order of magnitude more than current devices, the required time resources demanded a quantum computer whose coherence time was many orders of magnitude beyond the capabilities of current devices. The algorithm used in their analysis was a “naive Trotter decomposition”, similar to the proposal of Aspuru-Guzik et al.⁴⁵, which gives a gate count scaling of $O(N^{11})$, and an empirical scaling of $O(N^9)$. The optimistic conclusion of this work, however, was that, through improving quantum algorithms for Hamiltonian simulation, these required time resources could be drastically reduced.

The following years saw a sequence of developments^{48,49,306} in analyzing and improving the asymptotic scaling of resources needed for performing quantum chemical simulations on a quantum computer, in terms of the numbers of electrons and spin-orbitals. This work drastically reduced the requirements originally estimated by Wecker et al.⁴⁷.

A series of papers showed how to merge sequential Jordan-Wigner strings to reduce their cost from linear to constant as well how to parallelize evolution under different terms in the Hamiltonian⁴⁸. Poulin et al.⁴⁹ empirically studied the cost of simulation for real-world molecules, finding that the cost was closer to $O(\sim N^6)$. Later, it was also demonstrated that the Trotter errors depend on the maximum nuclear charge rather than the number of spin-orbitals⁵⁰. Subsequently, Reiher et al.³⁰⁵ carried out a detailed study of the computational cost, including the number of costly gates for quantum error correction, for FeMoCo, a model of the nitrogenase enzyme that suggests that it is indeed feasible to employ future error-corrected architectures for the simulation of realistic chemical systems of scientific and industrial interest using Trotter-based approaches. The gate counts were further

lowered thanks to advancements in circuit synthesis^{76,281,302,307,308} and Hamiltonian simulations^{188,189,197,266,270}.

Application of LCU techniques¹⁸⁸ in second quantization were studied by Babbush et al.¹⁹² and first quantization by Babbush et al.¹⁹¹. These techniques exponentially increased the precision Hamiltonian simulation. In second quantization, the number of gates required is $\tilde{O}(N^5)$. In first quantization, it is $\tilde{O}(\eta^2 N^3)$, where η is the number of particles in the system. Importantly, these early LCU-based algorithms have been shown to scale better asymptotically as a function of molecule size than prior Trotter algorithms for quantum simulation. The same scaling holds for the qubitization paradigm with better constant factors¹⁸⁹.

Concepts such as wave function locality³⁰⁶ can be introduced to further reduce the cost in terms of quantum gates for molecular simulation. The combination of these ideas with sparse algorithms and an intelligent choice of basis functions has been shown to reduce the cost of quantum simulation of chemical systems, but further improvements and generalizations may yet be possible. There are several ideas from the domain of classical molecular electronic structure that can be applied in the field of quantum simulation for further reduction of quantum computational cost.

Recently work has studied the chemistry simulation problem in bases where the Coulomb operator is diagonal^{301,302}. This allows a representation of the chemistry Hamiltonian with a number of terms scaling quadratically in the number of spin-orbitals, a significant reduction on the number of terms in the Hamiltonian when using molecular orbitals. This representation was first used to construct Trotter steps of the chemistry Hamiltonian requiring only a grid of qubits with nearest-neighbor connectivity³⁰¹. Following this, it was shown that Trotter steps exactly N layers of gates deep (with N the number of spin-orbitals) could be performed even with the restriction of the qubits being on a line³⁰².

This has led to particularly efficient simulation algorithms. For real state simulations, Kivlichan et al.⁷⁶ simulated interacting particles and achieved super-polynomially higher accuracy than previous algorithms. Kivlichan et al.³⁰² gave an algorithm linear in depth and

quadratic in the number of two-qubit gates in term of spin-orbitals. Soon after, Babbush et al.²⁸¹ gave an algorithm which has been fully priced to the level of fault-tolerant gates. This algorithm uses quantum signal processing and qubitization to achieve T complexity $O(N + \log 1/\epsilon)$ for Hamiltonian simulation and $O(N^3/\epsilon + N^2 \log 1/\epsilon)$ for energy estimation.

The most recent series of improvements in asymptotic scaling was by Low and Wiebe¹⁹⁷ for second and Babbush et al.¹⁹⁸ for first quantization. Both results use interaction picture simulations¹⁹⁷ and plane wave basis and achieve $O(N^{11/3})$ and $O(\eta \log N)$ complexities.

Researchers have also considered extending the use of quantum algorithms in exploring non-traditional regimes in quantum chemistry such as relativistic dynamics^{309,310} and quantum dynamics beyond the Born-Oppenheimer approximation^{75,76,311}. As reliable quantum computers begin to come online, there will be a continued demand for improvement of these quantum algorithms.

We give a brief overview of simulation techniques used in quantum chemistry with corresponding Hamiltonian simulation algorithms in Figure 5. A summary of query complexity and T-counts for a majority of these algorithms can be found in Babbush et al.²⁸¹.

4.2.3 Measuring properties of many-body systems

The last step of quantum chemistry algorithms is extracting information about the system from the wave function. Such information includes the ground state energy, scattering amplitudes, electronic charge density or k -particle correlations³⁸. In principle, it is possible to estimate any physical quantity or observable that can be expressed through a low-depth quantum circuit and single-qubit measurements.

The most sought-for information in these algorithms is the ground state energy. Zalka⁴¹ sketched a method for preparing a known quantum state and introduced the “von Neumann trick” for extracting properties from the prepared wave function. Lidar and Wang⁴⁴ then applied these techniques to develop an algorithm for determining the thermal rate constant of a chemical reaction. Soon after, Abrams and Lloyd³⁸ further developed these techniques




Figure 5: Chronological overview of quantum chemistry simulation algorithms. On the left-hand side we list quantum simulation algorithm grouped by the techniques they use. The right-hand side outlines the improvement for time evolution in quantum chemistry. We indicate the underlying simulation technique with an arrow from the left column to the right. Furthermore, we color-code the Hamiltonian representation – yellow (lighter color) for first quantization and red (darker color) for second quantization.

to specifically apply them towards calculating static properties of a quantum system. In 2005, Aspuru-Guzik et al.⁴⁵ adapted these techniques for the electronic structure problem. This allowed the number of qubits in the QPEA ancilla register to be reduced from 20 to 4, enabling the study of quantum algorithms for electronic structure problems on a classical computer⁴⁵. This work shows that even with modest quantum computers of 30-100 error-corrected qubits, ground state energy calculations of H₂O and LiH could be carried out to an accuracy beyond that of classical computers. The essential algorithm underlying this approach is *quantum phase estimation* and its modifications (see Section 4.1.2).

An iterative version of QPEA for ground state estimation was recently introduced by Berry et al.²⁷⁹. The authors assume that there is an upper bound on the ground state energy, say from a classical variational method, that is guaranteed to be lower than the energy of the first excited state. Given this assumption, one can perform QPEA gradually, measuring the ancilla qubits sequentially instead of postponing the measurement to the end of the circuit. If the outcome of the QPEA is likely to be a state above this threshold, i.e. not the ground state, it is possible to abort phase estimation and restart the algorithm.

QPEA can be also used for estimating the energies of excited states. In the simplest setting, one can use a state that is not an energy eigenstate and use it to sample multiple energies from the spectrum of the Hamiltonian. Santagati et al.³¹⁶ introduced a more sophisticated technique for approximating excited state energies.

A quantum algorithm for ground state energy estimation (as discussed here) requires techniques from state preparation (Section 4.2.1) and as a subroutine simulation (Section 4.2.2). Combining the recent development all these areas, Babbush et al.²⁸¹ give detailed resource estimates, or “pricing”, for classically intractable ground state energy calculations of diamond, graphite, silicon, metallic lithium, and crystalline lithium hydride. Incorporating state-of-the-art error correction methods, they show that the estimates for the number of gates (which is dominated by the number of T -gates) required to estimate the ground state energy of FeMoCo is millions of times smaller than the number needed in the methods used

earlier³⁰⁵.

Without access to fault-tolerant quantum computers, only proof-of-principle quantum chemistry calculations have been demonstrated. In particular, few-qubit simulations, often without error correction, have been carried out in most major architectures used for quantum information. In 2010, Lanyon et al.²³⁴ demonstrated the use of the IPEA to measure the energy of molecular wave functions. In this case, the wave function of molecular hydrogen (H_2) in a minimal basis set was encoded in a one-qubit state and the IPEA was realized using a two-qubit photonic chip, calculating the molecular hydrogen spectrum to 20 bits of precision. A similar procedure was applied to H_2 using nuclear magnetic resonance Lanyon et al.²³⁴, Du et al.²³⁵, Dolde et al.²³⁶ and to helium hydride (HeH^+) using nitrogen vacancies in diamond²³⁶.

Although these proof-of-principle experiments are groundbreaking, it is not clear how to scale them because of their reliance on Hamiltonians simplifications and tomography. The first scalable demonstration of the IPEA (and the variational quantum eigensolver algorithm, as discussed in Section 5) employed three superconducting qubits for simulating H_2 in a minimal basis and was carried out by a Google Research and Harvard collaboration⁵⁵. We will explain hybrid quantum-classical algorithms better suitable for these devices in the next section.

4.3 Chapter 4 Glossary

N	Number of quantum systems, number of qubits, number of orbitals
$O(\cdot)$	Big-O notation
H	Hamiltonian
e^{-iHt}	Hamiltonian evolution
$U(t)$	Unitary dynamics
ℓ	Number of k-body interactions
k	Support size of largest Hamiltonian term

τ	$\ H\ t$
ϵ	Error
$\text{poly}(\cdot)$	Polynomial function of order \cdot
d	Sparsity of Hamiltonian
η	Number of particles in system
λ_m	Eigenenergy
T	Number of ancillary qubits
$ \psi_m\rangle$	Eigenstate
$ \phi\rangle$	State prepared for phase estimation
f	(Sect. 4.0) Oracular function
f	(Sect. 4.1.2) invertible function used in generalization of PEA
R_k	Rotation gate
n	Number of qubits
QPEA	Quantum phase estimation algorithm
RAM	Random access memory
IPEA	Iterative phase estimation algorithm
ASP	Adiabatic state preparation
MCSCF	Multi-configurational self-consistent field
QFT	Quantum Fourier transform

5 Quantum algorithms for noisy intermediate-scale quantum devices

Despite recent improvements in the resource estimates for fully-quantum chemistry algorithms such as the quantum phase estimation, the number of gates and circuit depth required for their implementation surpasses the capabilities of existing and near-term quantum devices. Furthermore, these early devices, also called “pre-threshold” or NISQ⁵¹ devices, can-

not support quantum error correction schemes, which require additional overhead. Such limitations have motivated the development of quantum algorithms that assume the use of an imperfect, noisy quantum computer that can still outperform classical algorithms at providing approximate solutions to relevant problems.

An algorithm for a NISQ device requires a low circuit depth that allows for execution within the limited coherence time of the device. Additionally, the quantity of interest should be easy to extract from direct measurements of the quantum states prepared in the device. To make NISQ algorithms as efficient as possible, one can allocate computational tasks between quantum and classical devices based on the inherent advantages of each device. These observations led to the rise of hybrid quantum-classical (HQC) algorithms, which leverage strengths of quantum and classical computation, utilizing each where appropriate.




Figure 6: Outline of the hybrid quantum-classical (HQC) framework. The first part of the algorithm is performed as a quantum circuit on quantum hardware going from an input state to an output state. After measurement the data is passed on to a classical computer where the parameter optimization of the quantum circuit happens. This process allows one to take advantage of the strengths of present-day hardware and to avoid some of the weaknesses.

The general layout of HQC algorithms, as illustrated in Figure 6, comprises of three steps. The first is state preparation, achieved by application of a sequence of parametrized quantum gates on an initial state. The second step is measurement, in which the value of the objective function is estimated from measurements performed on the prepared quantum state. The

third step involves feedback or optimization on the classical computer to determine new parameters to improve the quality of the state.

While the algorithms developed within this framework can be implemented on fault-tolerant quantum computers, their original intent is to be implemented on NISQ devices. One particular challenge of this implementation is the presence of errors in the machine, which impacts the quality of the observables measured on the device. While some HQC algorithms have shown robustness against certain types of errors, specific techniques to mitigate noise on NISQ devices have been proposed and started to be incorporated into the implementations of HQC algorithms^{317–319}.

In this chapter, we will discuss some of the HQC algorithms developed for the simulation of quantum chemistry. See Figure 1 for its connection with previous sections. Specifically, Section 5.1 describes the variational quantum eigensolver (VQE) approach^{52,320}, the first example of an HQC algorithm to be proposed. VQE applies the time-independent variational principle to optimize the parameters of a quantum circuit implementing an ansatz on the quantum computer. This method provides approximate solutions to the time-independent Schrödinger equation. Section 5.2 describes the variational quantum simulation (VQS) algorithm, which is analogous to VQE for finding approximate solutions to the time-dynamics of the Schrödinger equation by applying one formulation of the time-dependent variational principle and other variations of these algorithms for the static problem. Finally, Section 5.3 covers algorithms for quantum chemistry simulation developed for non-gate-based near-term quantum devices, which includes devices for adiabatic quantum computing and linear optics.

5.1 Variational quantum eigensolver (VQE) algorithm

The VQE algorithm is an application of the time-independent variational principle, where the wave function ansatz preparation and energy estimation are both implemented on a quantum device. This paradigm allows for much flexibility in practice because the ansatz can be chosen to take into account the specifications of the quantum hardware. Furthermore,




Figure 7: Illustration of the VQE algorithm. Traditionally, the quantum computer is used to prepare a set of parametrized quantum states followed by applications of rotations $R_i \in \{I, R_X(-\pi/2), R_Y(\pi/2)\}$ depending on the Pauli term of the Hamiltonian to be measured. The classical computer then takes the individual estimates of the Pauli term expectation values $\langle P_i(\vec{\theta}) \rangle$ and averages them to compute a single value $\bar{E}(\vec{\theta})$. This cost function value is fed into an optimization routine, which produces an updated set of parameters $\vec{\theta}'$ as input for the quantum circuit in the next optimization loop. This procedure is repeated until the energy converges. We note that recent and future efforts are to improve each VQE component as well as replace certain “classical” subroutines with quantum counterparts to further leverage the capabilities of quantum computers.

the energy estimation for VQE is performed by a technique called *Hamiltonian averaging*, described below, that exchanges coherence time with the ability to sample the quantum computer multiple times. These two aspects make VQE a good candidate to simulate the ground states of quantum systems using current and near-term quantum devices. The core of the algorithm can be broken down into the following steps:

1. **State Preparation:** A parametrized quantum state $|\Psi(\vec{\theta})\rangle$ is prepared on the quantum device. This is achieved by applying a *parametrized unitary* to an easy-to-prepare initial state $|\Psi_0\rangle$ (e.g. a computational basis state): $U(\vec{\theta})|\Psi_0\rangle = |\Psi(\vec{\theta})\rangle$. The parametrized unitary is defined by the choice of ansatz, which should correspond to a family of states that cannot be efficiently represented and manipulated on a classical computer;
2. **Energy Estimation:** The expectation value of the energy $\langle H \rangle(\vec{\theta})$ is estimated using a Hamiltonian averaging procedure, which involves taking measurements of tensor products of Pauli terms corresponding to the qubit representation of the target Hamiltonian²¹⁶;
3. **Classical Feedback:** The parameters $\vec{\theta}$ of the quantum state are updated using a classical non-linear optimization routine;
4. Steps 2 and 3 are repeated until convergence criteria (e.g. energy) are satisfied.

The basic framework of VQE is modular in design such that various types of extensions and improvements are possible. In the following subsections, we describe each VQE step in greater detail, outlining improvements as well as open questions pertaining to the particular step.

5.1.1 Ansatzes for state preparation

The flexibility of the VQE algorithm arises from the ability to choose the parametrized trial state used to approximate the eigenstates of the target Hamiltonian. Consequently, the

performance of the algorithm depends largely on the quality and structure of this ansatz. In general, the construction of the ansatz involves preparing an initial state and building upon it using a parametrized quantum circuit, as shown in Figure 7.

Intuitively, a well-chosen ansatz can dramatically improve the capabilities of the classical optimization in the VQE algorithm. One aspect of this intuition was recently formalized by McClean et al.³²¹. The authors show that, for sufficiently random circuits, the variational parameter landscape is plagued by the overabundance of “barren plateaus”, large regions in which the cost function gradient is nearly zero. This implies that using arbitrary or unstructured ansätze will lead to poor convergence of the optimization routine. This issue can be circumvented by using physically motivated ansätze that have measurable gradients and thus better guide the optimizer. Another important quality of a circuit ansatz, particularly important for NISQ devices, is the ability to implement it with a high-fidelity, low-depth circuit.

The two broad strategies in circuit ansatz design account for these considerations to different extents. They are

1. *physically-motivated ansätze* (PMA), which are methods based on or inspired by numerical techniques that systematically approximate the exact electronic wave function;
2. *hardware heuristic ansätze* (HHA), which correspond to parametrized circuits comprising single-qubit rotations and entangling blocks generally chosen to take advantage of specific quantum hardware capabilities.

The prototypical example of the PMA category is the Unitary Coupled Cluster (UCC) ansatz, which consists of a systematic expansion that approximates the exact many-electron wave function as the excitation rank of the operators in the cluster operator increases. While the UCC ansatz is intractable to express with a classical computer^{322–324}, a Trotter-Suzuki decomposition approximation (c.f. Section 2.3) of this transformation can be implemented efficiently on a quantum computer^{52,320}. The relatively large number of parameters and

long circuit depth required to implement UCC motivated the development of alternative approaches. In particular, Wecker et al. proposed a *Hamiltonian variational ansatz* approach³²⁵, which consists of a parametrized version of the circuit implementing Hamiltonian simulation for the target Hamiltonian. This approach allows for a reduction in the number of variational parameters, however, the depth of the circuit depends on the complexity of the target Hamiltonian. When combined with the simplified representations of the Hamiltonian described in the previous section³⁰¹, the Hamiltonian variational method can be used to construct variational circuits with depth equal to the number of spin-orbitals for a linear qubit architecture, using the fermionic swap network³⁰². Independently, Dallaire-Demers et al.³²⁶ proposed a related PMA ansatz with the same scaling in depth based on a circuit employed to prepare fermionic Gaussian states (FGSs) on a quantum register, called the *low-depth circuit ansatz* (LDCA). The LDCA comprises parallel layers of nearest-neighbor matchgates augmented by ZZ interactions.

On the other hand, HHA approaches have been mainly motivated by limitations of existing quantum hardware. The general structure of the quantum circuits employed as HHA comprises two interleaved blocks: the first one comprising single-qubit rotations and the second one composed of fixed entangling operations that can be easily executed when taking into account particular constraints of the quantum hardware employed⁵⁶. This concept was applied in the experimental demonstration of the VQE protocol for small molecules⁵⁶, and it has been recently refined using sets of gates that conserve particle number³²⁷. Although these approaches can be implemented more easily than existing PMA on NISQ devices, to date, there lacks theoretical evidence one way or the other regarding their ability to efficiently approximate many-electron wave functions.

Evidently, significant progress has been made in the development and design of ansatze. However the existence of a general formulation to construct meaningful ansatze for particular target Hamiltonians remains an open question. In what follows we describe in more detail some of the ansatzes outlined above.

Unitary Coupled Cluster. Traditional coupled-cluster (CC) methods have long been considered the “gold standard” of quantum chemistry because they offer chemically accurate results for a variety of systems, including molecules near equilibrium, and because they are tractable on classical computers.

Recall that the ansatz for the traditional coupled-cluster method can be written as

$$|\Psi_{CC}\rangle = e^T |\Phi_0\rangle, \quad (31)$$

where $|\Phi_0\rangle$ is a reference state, usually the solution to the Hartree-Fock equations, and T corresponds to excitation operators of different excitation rank defined in terms of fermionic annihilation and creation operators as

$$T = \sum_{i=1}^{\eta} T_i \quad (32)$$

$$T_1 = \sum_{\substack{i \in \text{occ} \\ a \in \text{virt}}} t_a^i a_a^\dagger a_i \quad (33)$$

$$T_2 = \sum_{\substack{i > j \in \text{occ} \\ a > b \in \text{virt}}} t_{ab}^{ij} a_a^\dagger a_b^\dagger a_i a_j \quad (34)$$

...

where the *occ* and *virt* spaces are defined as the occupied and unoccupied sites in the reference state and η is the number of electrons in the system. For practical implementations on classical computers, T is usually truncated to lower excitation ranks e.g. single (T_1) and double (T_2) excitations. This allows one to find optimal CC amplitudes by projecting the Schrödinger equation in the form

$$e^{-T} H e^T |\Phi_0\rangle = E_{CC} |\Phi_0\rangle \quad (35)$$

against a set of configurations $\{|\mu\rangle\}$. This set spans the space of all the states that can be reached by applying the truncated cluster operator T linearly to the reference state. The similarity-transformed Hamiltonian, $e^{-T}He^T$ can be expanded using the Baker-Campbell-Hausdorff (BCH) formula. The expansion applied to Equation 35 has a finite truncation leading to a finite set of non-linear equations for the CC energy and amplitudes. These equations are then solved using standard numerical methods such as Newton’s approach⁷⁹. A method such as CCSD(T) (where the effect of triple excitations is included using perturbation theory) suffices to achieve chemical accuracy for energies of molecules near equilibrium. Unfortunately the scaling of this approach, $O(N^7)$, limits its applicability to small molecules.

Unitary coupled-cluster (UCC) is an alternative CC formulation constructed as

$$|\Psi_{UCC}\rangle = e^{T-T^\dagger} |\Phi_0\rangle, \tag{36}$$

where the T operator is replaced by the anti-Hermitian operator $T - T^\dagger$. Consequently the corresponding CC operator is unitary and the formulation is variational, which offers an advantage compared to standard CC. However, UCC is not a practical ansatz in classical quantum chemistry because the BCH expansion of the similarity transformed Hamiltonian, $e^{-(T-T^\dagger)}He^{T-T^\dagger}$, does not terminate³²⁴. Fortunately, unitary operations such as e^{T-T^\dagger} are natural operations on quantum computers. They can be readily implemented by mapping fermionic excitation operators to qubit operators and applying the Trotter-Suzuki decomposition to the cluster operator. The different Trotter approximations to the UCC operator constitute a family of CC-inspired ansätze that can be implemented on a quantum computer using standard circuit compilation techniques for exponentiating Pauli matrices¹. Given the inherent advantage in implementing the ansatz as well as the high accuracies of CC methods, the application of the UCC ansatz truncated up to single and double excitations (UCCSD) has been investigated both theoretically and experimentally^{52,53,59,325–330}.

In the case of a Trotterized implementation of UCCSD, the ansatz parameters correspond

to the cluster amplitudes, whose number scales as $O(N^2\eta^2) < O(N^4)$. Correspondingly, a serial implementation of these terms would result in a scaling of the number of gates of $O(fN^4)$, where f indicates the number of gates required for implementing each term. The scaling of f with N depends on the fermion-to-qubit mapping employed: in the case of the Bravyi-Kitaev transformation $f \in O(\log(N))$ while in the case of Jordan-Wigner $f \in O(N)$ (see Appendix B). Similarly, the circuit depth can be upper-bounded to $O(fN^4)$, which corresponds to a serial execution, however this can be improved depending on the connectivity constraints of the hardware and the ability to parallelize operations in the quantum registers. For architectures with arbitrary connectivity that allow parallelization, some strategies developed for Trotterized Hamiltonian evolution⁴⁹ could save up to a linear factor in depth using a constant number of ancilla qubits when a Jordan-Wigner mapping is used.

The scaling in the number of parameters of the UCCSD ansatz poses a challenge for typical optimization strategies, as the number of amplitudes can easily increase to the order of thousands for systems with tens of spin-orbitals and half filling. To reduce this cost, practical applications of UCCSD must consider typical strategies employed in classical quantum chemistry to prioritize those excitation operators that are more relevant for achieving accuracy. Some of these strategies include: limiting the calculation to active spaces, freezing core orbitals, applying spin-symmetry constraints and discarding operators with low amplitudes based on classical efficient approximations e.g. Møller-Plesset perturbation theory^{216,327,329}. The last strategy also has the potential to significantly reduce the number of gates required and provides approximate amplitudes for initializing the optimization, which has been pointed out as critical for improving convergence³²⁹.

Most numerical demonstrations of UCCSD applied to VQE have explored the impact of different optimization approaches, including both gradient-free and gradient-based methods, as described in more detail in Section 5.1.3. In practice, gradient-based methods can be implemented using either numerical or analytical gradients. It has been pointed out that

analytical gradients require orders-of-magnitude fewer measurements than numerical gradients³²⁹. The effect of the number of Trotter steps has been also investigated, indicating that the accuracy of the exact implementation of the UCCSD exponential operator is virtually the same as the one obtained with Trotterized UCCSD for different Trotter numbers. This can be understood considering that different Trotterized versions of UCCSD correspond to different ansatze and the variational optimization tends to compensate the differences between the non-Trotterized and the Trotterized version, which would have different optimal amplitudes. UCCSD has been implemented in combination with different basis set representations of the chemistry Hamiltonian, including the particle/hole basis³²⁷ and the Bogoliubov basis³²⁶. The latter allows for the application of UCC to Hamiltonians with pairing fields, extending the range of applicability of VQE to problems in condensed matter and nuclear physics. Finally, UCCSD has been demonstrated experimentally for HeH⁺ using a single trapped ion⁵³, for H₂ on a superconducting quantum computer⁵⁵, for H₂ and LiH on an ion-trapped quantum computer⁵⁹ and for the wave function of a deuterium nucleus using cloud quantum computing³³⁰.

Hamiltonian variational ansatze. The Hamiltonian variational approach, proposed by Wecker et al.⁴⁷, defines a strategy for building variational ansatze for arbitrary Hamiltonians. The general idea is described as follows: consider a target Hamiltonian that can be written as a sum of local terms, $H = \sum_i h_i O_i$, where h_i is a scalar and O_i is a local operator. Furthermore, the terms in this Hamiltonian can be grouped into arbitrary subsets, $H = \sum_j H_j$. The criteria for grouping can differ depending on the target Hamiltonian as shown below. The variational ansatz is defined as

$$|\Psi(\theta)\rangle = \left(\prod_{b=1}^S [U_{H_j}(\theta_j^b) \dots U_{H_1}(\theta_1^b)] \right) |\Psi_0\rangle, \quad (37)$$

where $|\Psi_0\rangle$ can be chosen to match the desired symmetry properties of the target ground

state, $U_{H_k}(\theta_k^b) = \exp(\theta_k^b H_j)$, θ is a vector grouping the scalar variational parameters to be minimized $[\theta_1^1, \dots, \theta_k^b, \dots, \theta_j^S]$, and S is the number of repetitions of the circuit block and can be varied to control the approximation accuracy. Notice that in the implementation, each term $U_{H_k}(\theta_k^b)$ has to be further broken down into a product of exponentials of individual terms since H_j may be a sum of terms. If all the terms within a subset were to commute, this could be implemented exactly. The definition of the circuit in the product of Equation 37 resembles a first order Trotterized evolution under the Hamiltonian, where now the terms have been multiplied by arbitrary parameters. Other versions of Equation 37 can be obtained by mimicking higher-order Trotter-Suzuki formulas.

The Hamiltonian variational approach is inspired by both adiabatic state preparation (ASP) and the quantum approximate optimization algorithm (QAOA)^{331,332}. In ASP, the ground state of a target Hamiltonian, H_1 , is obtained by adiabatically evolving the ground state of a Hamiltonian H_0 , $|\Psi_0\rangle$, that is easy to prepare. This process assumes that the transition between H_0 and H_1 is continuous and that the ground state energy of $\lambda H_0 + (1 - \lambda)H_1$ remains sufficiently separated from the first excited state energy during the preparation (see Section 5.3.1 for discussion of the importance of the “energy gap” remaining large). In practice, ASP is implemented by rotating the initial state by a time-dependent Hamiltonian of the form $H(t) = A(t)H_0 + B(t)H_1$, where $A(t)$ and $B(t)$ are continuous functions defining the relative weights between the initial Hamiltonian and the target one, which constitutes the annealing schedule. This evolution is performed during some annealing time t and is discretized into S steps of size $\Delta t = \frac{t}{S}$, such that the annealing becomes a sequence of S unitary rotations by Hamiltonians interpolating between H_0 and H_1 . If the initial and target Hamiltonian are both expressed as a sum over the same Hamiltonian terms, $H_i = \sum_j h_j^i O_j$, then the evolution in Equation 37 can be seen as analogous to a discretized adiabatic state preparation, where the angles of the rotations are obtained by a variational optimization instead of being fixed by a predefined annealing schedule. A related idea is the concept of adiabatically parametrized states, introduced in McClean et al.²¹⁶, where

the functions defining the annealing schedule are parametrized and subsequently optimized using variational minimization.

Wecker et al.³²⁵ also proposed two different optimization strategies for minimizing the energy of the ansatz defined in Equation 37. The first approach, described as *global variational*, corresponds to the simultaneous optimization of all the parameters in the ansatz, given a fixed value of S . The second approach, called *annealed variational*, is also inspired by discretized adiabatic state preparation. It consists of performing a block-by-block optimization that starts by optimizing the ansatz with $S = 1$. The state resulting from this optimization, $|\Psi^{(1)}\rangle$, is used as the initial state for an optimization with another variational circuit with $S = 1$, obtaining a new state, $|\Psi^{(2)}\rangle$. This process is repeated until S steps have been completed. The final parameters could be further refined by a global optimization, or at each step k in the sequential optimization a global optimization could be applied. The block-by-block optimization procedure improves convergence and is analogous to the segment-by-segment pre-training strategy employed in machine learning to avoid barren plateaus in the optimization of neural networks^{333,334}.

The performance of the Hamiltonian variational approach applied to the Fermi-Hubbard model in a 2D lattice and the electronic structure Hamiltonian is analyzed in Wecker et al.³²⁵. In the first case, the target Hamiltonian has the form

$$H = -t \sum_{\langle i,j \rangle, \sigma} a_{i,\sigma}^\dagger a_{j,\sigma} + U \sum_i a_{i,\uparrow}^\dagger a_{i,\uparrow} a_{i,\downarrow}^\dagger a_{i,\downarrow}, \quad (38)$$

where $a_{i,\sigma}^\dagger$ and $a_{i,\sigma}$ respectively create and annihilate an electron at site i with spin $\sigma \in \{\uparrow, \downarrow\}$. The summation in the first term runs over nearest neighbors, denoted as $\langle i, j \rangle$. In this case, the Hubbard Hamiltonian is divided as $H = h_h + h_v + h_U$, where h_h and h_v are the sums of hopping terms in the horizontal and vertical directions, respectively, and h_U is the repulsion

term. The authors use a second-order Trotter formula for the Hamiltonian variational ansatz

$$|\Psi(\theta)\rangle = \prod_{b=1}^S \left[U_U^\dagger \left(-\frac{\sigma_U^b}{2} \right) U_h(\theta_h^b) U_v(\theta_v^b) U_U \left(\frac{\theta_U^b}{2} \right) \right] |\Psi_0\rangle, \quad (39)$$

where $U_X(\theta_X)$ denotes a Trotter approximation to $\exp(i\theta_X h_X)$ where $X \in \{U, h, v\}$. The initial state $|\Psi_0\rangle$ is chosen to be in the correct spin sector e.g. a computational state with equal number of spin-up and spin-down qubits for singlet ground states.

For the application to the electronic structure Hamiltonian, the terms are separated into three groups corresponding to diagonal terms, hopping terms and exchange terms as follows

$$H_{diag} = \sum_p \sum_p \epsilon_p a_p^\dagger a_p + \sum_{p,q} h_{pqqp} a_p^\dagger a_p a_q^\dagger a_q, \quad (40)$$

$$H_{hop} = \sum_{p,q} h_{pq} a_p^\dagger a_q + \sum_{p,q,r} h_{prrq} a_p^\dagger a_q a_r^\dagger a_r, \quad (41)$$

$$H_{ex} = \sum_{p,q,r,s} h_{pqrs} a_p^\dagger a_q^\dagger a_s a_r, \quad (42)$$

where the sums are taken with p, q, r, s all distinct from one another. In this case the ansatz is built as:

$$|\Psi(\theta)\rangle = \prod_{b=1}^S \left[U_{ex}^\dagger \left(-\frac{\sigma_{ex}^b}{2} \right) U_{hop}^\dagger \left(-\frac{\theta_{hop}^b}{2} \right) U_{diag}(\theta_{diag}^b) U_{hop} \left(\frac{\theta_{hop}^b}{2} \right) U_{ex} \left(\frac{\theta_{ex}^b}{2} \right) \right] |\Psi_0\rangle, \quad (43)$$

where $U_X(\theta_X)$ denotes a Trotter approximation to $\exp(i\theta_X h_X)$ where $X \in \{ex, hop, diag\}$. The initial state is chosen to be the ground state of H_{diag} and the basis is a Hartree-Fock basis such that $H_{hop} |\Psi_0\rangle = 0$. Both Equations 39 and 40 were numerically tested on lattices of size $2 \times N$ with $N = 4, 6, 8, 10, 12$ and for small molecules such as HeH^+ , H_2O , BeH_2 , and hydrogen chains (H_x with $x \in \{6, 8, 10\}$). The results show that this strategy can achieve good overlaps with the ground state for values of S varying between 3 and 20 (see Wecker et al.³²⁵). However, it remains to investigate how S would scale with the size of the system and the accuracy required for generic Hamiltonians.

The Hamiltonian variational approach greatly reduces the number of variational parameters, which eases the optimization. In the case of the Hubbard model, it also provides an ansatz that scales only linearly with the size of the system and therefore is a good candidate for experimental demonstration in near-term devices. However, in the case of quantum chemistry Hamiltonians, the depth of the corresponding ansatz is directly related to the number of terms in the Hamiltonian, which scales formally as $O(N^4)$. This scaling might vary depending on the analytical properties of the basis sets employed to represent the Hamiltonian. In the case of Gaussian-type orbitals (GTOs), which are the most popular basis sets in quantum chemistry, the number of non-negligible terms scales nearly quadratically for systems of sufficient size, as explained in more detail in Section 5.1.2. The sizes at which this scaling is observed, however, depends on the nature of the system and the basis set and might correspond to tens to hundreds of angstroms^{79,306}. This fact has motivated the exploration of alternative basis sets that can show more favorable scalings and therefore facilitate the computation.

In particular, Babbush et al. have proposed to use a dual form of the plane wave basis³⁰¹ which diagonalizes the potential operator, leading to the following representation of the electronic structure Hamiltonian that has only $O(N^2)$ terms:

$$\begin{aligned}
H &= \frac{1}{2N} \sum_{\nu, p', q', \sigma} k_\nu^2 \cos[k_\nu \cdot r_{q'-p'}] a_{p', \sigma}^\dagger a_{q', \sigma} - \frac{4\pi}{\Omega} \sum_{\substack{p', \sigma \\ j, v \neq 0}} \frac{\zeta_j \cos[k_\nu \cdot (R_j - r_{p'})]}{k_\nu^2} n_{p', \sigma} \\
&+ \frac{2\pi}{\Omega} \sum_{\substack{(p', q') \neq (q', \sigma') \\ \nu \neq 0}} \frac{\cos[k_\nu \cdot r_{p'-q'}]}{k_\nu^2} n_{p', \sigma} n_{q', \sigma'}, \\
&= \sum_{pq} T_{pq} a_p^\dagger a_q + \sum_p U_p n_p + \sum_{p \neq q} V_{pq} n_p n_q
\end{aligned} \tag{44}$$

where, n_p is the number operator, ζ_j are nuclei charges and k_ν is a vector of the plane wave frequencies at the harmonics of the computational cell in three dimensions, with volume Ω . Here, σ is the spin degree of freedom and ν is a three-dimensional vector of integers

with elements in $[-N^{1/3}, N^{1/3}]$. Letter indices with an apostrophe denote orbitals while regular letters indicate spin-orbitals. The plain wave dual basis set can be viewed as a discrete variable representation (DVR) and is particularly suited for periodic systems. When mapped to qubit operators, the Hamiltonian in Equation 44 is equivalent to the following local Hamiltonian in the Jordan-Wigner representation:

$$\begin{aligned}
H = \sum_p \left(\frac{T_{pq} + U_p}{2} + \sum_q \frac{V_{pq}}{2} \right) Z_p + \sum_{p \neq q} V_{pq} Z_p Z_q + \\
\sum_{p \neq q} \frac{T_{pq}}{2} (X_p Z_{p+1} \dots Z_{q-1} X_q + Y_p Z_{p+1} \dots Z_{q-1} Y_q). \tag{45}
\end{aligned}$$

Notice that this Hamiltonian also encompasses the Fermi-Hubbard Hamiltonian and can be readily combined with the Hamiltonian variational strategy to produce a variational ansatz. In Babbush et al.³⁰¹, this idea is developed to propose a VQE simulation for Jellium, that requires depth $O(N \cdot S)$ and optimizes only over $O(N)$ parameters by exploiting the translational invariance of this system.

Finally, Kivlichan et al.³⁰² demonstrated that a single Trotter step of the Hamiltonian in Equation 44 can be implemented with a circuit of depth N using a *Fermionic Swap Network*, which has a constant improvement with respect to the proposal in Babbush et al.³⁰¹ and is conjectured to be optimal. In the particular case of the Fermi-Hubbard Hamiltonian, the scaling is further reduced to $O(\sqrt{N})$. The main insight of this proposal is that when the Jordan-Wigner transformation acts on the terms $(a_p^\dagger a_q + a_q^\dagger a_p)$, it produces terms of locality $|p - q| + 1$, which has significant overhead when mapped to a circuit. To get rid of this problem, the authors propose to execute the Trotter step by executing operations only on nearest neighbor qubits, and then swapping the spin-orbital indexes, until all the $N(N - 1)/2$ interactions between spin-orbitals are implemented. The swapping is performed using the fermionic-SWAP (f-SWAP) gate, defined as

$$f_{SWAP}^{p,q} = 1 + a_p^\dagger a_q + a_q^\dagger a_p - a_p^\dagger a_q - a_q^\dagger a_p. \tag{46}$$

Furthermore, rotations by the interaction and repulsion terms between orbitals with indexes p and q could be implemented together with the f-SWAP gates, forming the so-called *fermionic simulation gate*,

$$\begin{aligned}
F(p, q) &= \exp(-iV_{pq}n_p n_q t) \exp(-iT_{pq}(a_p^\dagger a_q + a_q^\dagger a_p)t) f_{swap}^{p,q} \\
&\equiv \exp\left(-\frac{i}{4}V_{pq}(1 - Z_p - Z_q - Z_p Z_q)t\right) \exp\left(-\frac{i}{2}T_{pq}(X_p X_q + Y_p Y_q)t\right) \\
&\quad \times \frac{1}{2}(X_p X_q + Y_p Y_q + Z_p + Z_q)
\end{aligned} \tag{47}$$

$$\equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -i \sin(T_{pq}t) & \cos(T_{pq}t) & 0 \\ 0 & \cos(T_{pq}t) & -i \sin(T_{pq}t) & 0 \\ 0 & 0 & 0 & -e^{-iV_{pq}t} \end{bmatrix} \tag{48}$$

The fermionic swap network circuit could be readily employed in combination with the Hamiltonian variational strategy to create a variational ansatz for preparing ground states of the Hamiltonian of Equation 44 using VQE. For example, one could adopt a parametrization with distinct variables for the interaction and the hopping terms, in analogy with Equation 39. Alternatively, one could free the T and V terms as variational parameters, which scales as $O(N^2)$ compared to $O(N^4)$ of the UCCSD ansatz.

Low depth circuit ansatz. Another ansatz for state preparation in VQE is the so-called low depth circuit ansatz (LDCA) proposed by Dallaire-Demers et al.³²⁶. This ansatz is inspired by a circuit for preparing fermionic Gaussian states (FGSs) on quantum computers. In what follows, we briefly describe FGSs and some of the techniques for preparing these states in quantum registers before describing the LDCA.

FGSs are states of which the density operators can be expressed as an exponential of a quadratic function of fermionic creation and annihilation operators. Typical single determinant ansätze such as the Hartree-Fock (HF) and the Bardeen-Cooper-Schreiffer (BCS) states

for superconductivity, are FGSs. The advantage of simulating this family of states is that they are completely characterized by a quadratic number of parameters and therefore, can be manipulated efficiently on a classical computer. A particularly convenient representation of an FGS is via its covariance matrix (Γ), which corresponds to a $2N \times 2N$ matrix

$$\Gamma_{kl} = \frac{i}{2} \text{tr}(\rho [\gamma_k, \gamma_l]). \quad (49)$$

where $\gamma_j = a_j^\dagger + a_j$ and $\gamma_j^\dagger = -i(a_j^\dagger - a_j)$ correspond to Majorana operators and ρ is the corresponding FGSs. For pure states $i\Gamma - \mathbf{1} \leq 0$ and $\Gamma^2 = -\mathbf{1}$. Importantly, any expectation value of an operator can be computed from the covariance matrix as:

$$i^p \text{tr}(\rho \gamma_{j_1} \dots \gamma_{j_{2p}}) = \text{Pf}(\Gamma|_{j_1 \dots j_{2p}}), \quad (50)$$

where $1 \leq j_1 < \dots < j_{2p} \leq 2N$, $\Gamma|_{j_1 \dots j_{2p}}$ is the corresponding submatrix of Γ and

$$\begin{aligned} \text{Pf}(\Gamma) &= \frac{1}{2^N N!} \sum_{s \in S_{2N}} \text{sgn}(s) \prod_{j=1}^N \Gamma_{s(2j-1), s(2j)} \\ &= \sqrt{\det(\Gamma)} \end{aligned} \quad (51)$$

is the Pfaffian of a $2N \times 2N$ matrix defined from the symmetric group S_{2N} and $\text{sgn}(s)$ is the signature of the permutation s . Similarly, the time evolution of a state can also be described via the covariance matrix. FGSs have special properties, including being ground states of quadratic Hamiltonians of the form $H = \sum_{pq} h_{pq} \gamma_p \gamma_q$. In addition, pure FGSs can be brought into the form $|\psi\rangle = \prod_{k=1}^{2N} (u_k + v_k a_k^\dagger a_k^\dagger) |0\rangle$ where $|v_k|^2 + |u_k|^2 = 1$, using an appropriate basis set transformation. Finally, we point out that FGSs also include thermal states of quadratic fermionic Hamiltonians and in the case where the number of particles is well-defined, FGSs correspond to Slater determinants.

Preparing FGSs on a quantum computer is often the first step in the quantum simulation of fermionic Hamiltonians, whether quantum phase estimation or variational schemes such

as VQE are employed. Methods for preparing Slater determinants and general FGSs employ a quadratic number of gates^{302,326,335,336}. In the case of Slater determinants, the preparation is often described in terms of a rotation \mathcal{U} applied to a single determinant,

$$|\psi_s\rangle = \mathcal{U}c_1^\dagger \cdots c_N^\dagger |vac\rangle, \quad (52)$$

where c_k^\dagger and c_k represent fermionic operators in an easy to prepare basis (e.g. the computational basis). The transformation \mathcal{U} is a $2^N \times 2^N$ matrix, however it is equivalent to rotating the basis,

$$\phi_p = \sum_q \psi_q u_{pq}; \quad a_p^\dagger = \sum_q c_q u_{pq}^*; \quad a_p = \sum_q c_q^\dagger u_{pq}, \quad (53)$$

where ϕ_p , a_p^\dagger , a_p are spin-orbitals and fermionic operators in the rotated basis and u is an $N \times N$ unitary matrix. The operator \mathcal{U} is related to u via

$$\mathcal{U} = \exp \left(\sum_{pq} u_{pq} (a_p^\dagger a_q - a_p a_q^\dagger) \right). \quad (54)$$

\mathcal{U} can be decomposed into a sequence of Givens rotations by applying a QR decomposition scheme to u , generating a circuit composed of two-qubit gates corresponding to these Givens rotations applied on neighboring qubits^{302,335}.

In the case of FGSs, there exist two related schemes. Jiang et al.³³⁵ proposed a method that specifies the FGSs as the corresponding quadratic fermionic Hamiltonian for which the target FGSs are ground states. The Hamiltonian is

$$H = \sum_{j,k=1}^N \Sigma_{jk} c_j^\dagger c_k + \frac{1}{2} \sum_{j,k=1}^N \left(\Delta_{jk} c_j^\dagger c_k^\dagger + h.c. \right), \quad (55)$$

where $\Sigma = \Sigma^\dagger$ and $\Delta = -\Delta^\dagger$ are complex matrices. This Hamiltonian is brought into its

diagonal form using standard methods, yielding

$$H = \sum_j \epsilon_j b_j^\dagger b_j + C, \quad (56)$$

where C is a constant and b_j and b_j^\dagger are new fermionic operators. The transformation that maps the original fermionic operators to the new ones is described by the matrix $2N \times 2N$ matrix W , defined by

$$\begin{bmatrix} \mathbf{b}^\dagger \\ \mathbf{b} \end{bmatrix} = W \begin{bmatrix} \mathbf{c}^\dagger \\ \mathbf{c} \end{bmatrix}. \quad (57)$$

Using matrix manipulation on W , the corresponding circuit implementing this transformation, \mathcal{W} , can be decomposed into a sequence of $(N - 1)N/2$ Givens rotations applied on neighboring qubits and N particle-hole transformations implemented as single qubit rotations. This circuit can be parallelized in depth $O(N)$ ³³⁵.

Independently, Dallaire-Demers et al.³²⁶ proposed a method that prepares FGSs specified by their corresponding covariance matrix. The prepared state has the form

$$|\Phi_0\rangle = C \prod_{k=1}^M b_k |\text{vac}\rangle, \quad (58)$$

where $|\text{vac}\rangle$ is the Fock vacuum and the operators b_k and their Hermitian conjugates, b_k^\dagger , correspond to Bogoliubov operators which are defined as in Equation 57 and c_k^\dagger and c_k being the original fermionic operators with respect to the Fock vacuum. From the covariance matrix it is possible to obtain the corresponding W matrix, which is later decomposed into $2N(N - 1)$ matchgates acting on neighboring qubits and N local rotations, using the Hoffmann algorithm³³⁷. The final circuit can be parallelized such that its depth scales only as $O(N)$. Notice that both methods^{326,335} achieve the same scaling but differ in the types of gates and the numerical methods employed to compute the angles of the preparation

circuits. While the other linear-depth methods have better constant factors, their circuits are less general^{302,326,335}.

The circuit requires a single layer of single qubit rotations followed by blocks of neighboring matchgates, $G_{ij}^{(k)}$, acting on qubits i and j . The matchgates can be decomposed into a product of two-qubit evolutions under XX , XY , YX , and YY , with the evolution times obtainable from the covariance matrix. Each parallel cycle of the LDCA circuit interleaves gates between even and odd neighboring qubits,

$$U_{\text{MG}}^{(k)} = \prod_{i \in \text{odd}} G_{i,i+1}^{(k)} \prod_{i \in \text{even}} G_{i,i+1}^{(k)}. \quad (59)$$

There are $\lceil \frac{N}{2} \rceil$ cycles in total, i.e.

$$U_{\text{MG}}^{\text{NN}} = \prod_{k=1}^{\lceil \frac{M}{2} \rceil} U_{\text{MG}}^{(k)}, \quad (60)$$

such that the transformation W can be composed as

$$W = U_{\text{MG}}^{\text{NN}} \prod_{i=1}^M R_i^Z. \quad (61)$$

Dallaire-Demers et al.³²⁶ suggest that a similar circuit template could be used to create a variational ansatz by complementing the list of matchgates with neighboring ZZ interactions, allowing for the preparation of non-Gaussian fermionic states. A similar idea has been used to design ansätze for preparing non-Gaussian states efficiently on classical computers, with high success³³⁸. The so-called low depth circuit ansatz (LDCA), acts on the quasiparticle vacuum state for the Bogoliubov picture, such that the initial states correspond to the generalized Hartree-Fock optimal state for the VQE target Hamiltonian. The circuit comprises L cycles of variational blocks composed of the simulation gates K_{ij} , where each K_{ij} is a product of evolutions under YX , XY , ZZ , YY , and XX with the evolution times again given by the covariance matrix.

Each layer k applies variational rotations in parallel, first on the even pairs and then on the odd pairs such that

$$U_{\text{VarMG}}^{(k,l)}(\Theta^{(k,l)}) = \prod_{i \in \text{odd}} K_{i,i+1}^{(k,l)}(\Theta_{i,i+1}^{(k,l)}) \times \prod_{i \in \text{even}} K_{i,i+1}^{(k,l)}(\Theta_{i,i+1}^{(k,l)}). \quad (62)$$

A cycle l is composed of $\lceil \frac{N}{2} \rceil$ layers, in analogy with the FGSs preparation:

$$U_{\text{VarMG}}^{\text{NN}(l)}(\Theta^{(l)}) = \prod_{k=1}^{\lceil \frac{N}{2} \rceil} U_{\text{VarMG}}^{(k,l)}(\Theta^{(k,l)}). \quad (63)$$

Finally, the L cycles are assembled sequentially to form the complete variational ansatz

$$U_{\text{VarMG}}(\Theta) = \prod_{l=1}^L U_{\text{VarMG}}^{\text{NN}(l)}(\Theta^{(l)}) \prod_{i=1}^M R_i^Z(\theta_i^Z), \quad (64)$$

with only one round of variational phase rotations

$$R_i^Z(\theta_i^Z) = e^{i\theta_i^Z \sigma_z^i}. \quad (65)$$

The variational state therefore has the form

$$|\Psi(\Theta)\rangle = U_{\text{Bog}}^\dagger U_{\text{VarMG}}(\Theta) \prod_{i=1}^N X_i |0\rangle^{\otimes M}, \quad (66)$$

where U_{Bog} is the \mathcal{W} transformation corresponding to the optimal Bogoliubov transformation for the VQE target Hamiltonian, which can be included if measuring the Hamiltonian is more efficient in the original basis. There are 5 variational angles per $K_{ij}^{(k,l)}$ and $N - 1$ of those terms per layer. Since each cycle has $\lceil \frac{N}{2} \rceil$ layers, an L -cycle circuit has $5L(N - 1)\lceil \frac{N}{2} \rceil + N$ variational angles, the extra term arising from the round of Z rotations. Considering parallelization, the final circuit depth is $(10L + 8)\lceil \frac{N}{2} \rceil + 4$ when we account for U_{Bog}^\dagger and the initial round of single-qubit X gates.

The linear scaling of the depth of LDCA makes it a viable option for implementation on near-term NISQ devices. The same global variational and annealed variational optimization approaches employed in the optimization of Hamiltonian variational approaches can be applied in this case. The LDCA also has the advantage of being able to treat systems where the number of particles is not conserved, such as superconductors, and therefore expands the range of applications of VQE to problems in condensed matter physics. For the same reason, the optimization of the energy must be performed by constraining the number of particles in those cases where the number of particles is conserved, e.g. in molecular systems. Numerical studies of LDCA have shown great promise. In the original paper³²⁶, the authors applied LDCA to a system with 8 qubits, corresponding to a 2×2 Fermi-Hubbard lattice and an extended Fermi-Hubbard model describing the isomerization of cyclobutadiene. The results demonstrated the ability of LDCA to describe the ground states of these systems in both weakly-correlated and strongly-correlated regimes. More recently, the LDCA has been applied to lithium hydride³³⁹, showing superior accuracy compared to hardware efficient ansätze. Still, further numerical benchmarks and theoretical analysis are required to better understand the scaling of the number of layers with the accuracy of the ansatz.

Finally, we point out that approaches analogous to LDCA could be developed by interleaving the blocks of Givens rotations obtained in the procedure proposed by Jiang et al.³³⁵, with ZZ interactions and optimizing the parameters in the Givens rotations variationally. In the same vein we note some similarities between the LDCA circuit and the fermionic swap network circuit proposed in the context of Hamiltonian evolution for Equation 44. The circuits for the fermionic swap network comprise fermionic swap gates applied on neighboring qubits that implement evolution under ZZ , Z , XX and YY . This is very similar to the construction of the LDCA circuit with the difference that LDCA also implements XY and YX interactions, which do not conserve the number of particles. This suggests that the fermionic swap network circuit could be employed as a template for variational ansätze with potentially similar results as those observed for LDCA.

Hardware efficient ansatz (HEA). This is a specific kind of hardware heuristic ansatz (HHA) mentioned previously. These ansatz comprise a series of parametrized quantum circuits originally motivated by the limitations of existing NISQ hardware⁵⁶. The general construction of these circuits comprises L entangling blocks, each block containing parametrized gates chosen to facilitate implementation according to the limitations in fidelity and connectivity of the hardware employed for the simulation. In general, these circuits can be expressed as:

$$|\Psi(\theta)\rangle = U^{(L)}(\theta^L) \dots U^{(1)}(\theta^1) U^{(0)}(\theta^0) |\Psi_0\rangle \quad (67)$$

where the variational parameters correspond to $\theta = [\theta^0, \dots, \theta^{L-1}, \theta^L]$, $U^{(k)}$ denotes the k -th entangling block and $|\Psi_0\rangle$ is a suitable initial state. In the first example of this ansatz⁵⁶, each entangling block comprised a single layer of parametrized single-qubit rotations followed by a fixed set of entangling two-qubit gates, specially chosen for a six-qubit quantum processor employed in the experiment. The circuit has the following form:

$$U^{(k)}(\theta^k) = \left(\prod_i^N R_i^Z(\theta_{i,1}^k) R_i^Z(\theta_{i,2}^k) R_i^Z(\theta_{i,3}^k) \right) E_{ent}, \quad (68)$$

where E_{ent} denotes a fixed set of two-qubit gates. In this case, the total number of parameters scales as $3N(L+2)$. For the experiment, the qubit register was initialized with the $|0\rangle^{\otimes N}$ state.

More recently, Barkoutsos et al.³²⁷ proposed a new set of entangling blocks. Two of these new blocks are designed to implement number-conserving operators, more suitable for simulations of molecular systems, and can be implemented natively on superconducting processor with transmon qubits³⁴⁰⁻³⁴². The third type of entangling block introduced by Barkoutsos et al.³²⁷ is similar to the original proposal (Equation 68), with E_{ent} being a circuit with all-to-all CNOT gates and with single-qubit rotations corresponding to R_z .

For these new ansatzes, the Hartree-Fock state was employed as initial state in the VQE calculation.

The accuracy of HEA have been tested both experimentally⁵⁶ and numerically³²⁷ for small molecular systems including H₂, LiH, BeH₂ and H₂O. These implementations have incorporated the use of active spaces and methodologies for reducing the number of qubits in the Hamiltonian³⁴³. Since not all the HEA ansatzes conserve the number of particles, a constrained minimization method must be employed in cases where the number of particles is conserved, using the number of particles as a constraint. In general, in order to approach the accuracy of methods like UCCSD, the number of layers has to grow rapidly, even to an extent that the number of parameters approaches or surpasses the size of the Hilbert space³²⁷. As in the case of other variational algorithms, a better understanding of the scaling of accuracy with the number of layers for different versions of this family of ansatzes is required.

Alternative approaches to state preparation. In gate-model quantum computers, it is important to first initialize the computer into a known state, and to do so with high accuracy. Typically, the form of the initial state is a product state such as $|0 \dots 0\rangle$ which is simple to prepare because each qubit can be addressed individually. Once an initial state is prepared, the evolution of the quantum state through the computer is determined by the sequence of gates that the algorithm calls for. While improving circuit ansatzes help reduce the overhead in the depth of a quantum circuit after the quantum computer has been initialized, there can be an advantage in preparing more complicated initial states that are closer to the final state. For example, VQE relies on ultimately preparing (and then measuring) an entangled state of many qubits. Thus, an initial state that is already highly entangled could be advantageous to preparing a desired ground state.

Despite rapid advances in both theory and experiment for quantum computing, significant efforts must be dedicated toward maximizing the use of NISQ devices while accounting for their limitations. As we have reviewed earlier in this section, the choice/design of the ansatz




Figure 8: a.) Illustration of the quantum autoencoder (QAE) algorithm. After preparing the training set of quantum states, a variational circuit comprised of encoding and decoding components is trained using measurements from either the trash state or the final decoded state. Once successfully trained, the encoding circuit of the overall QAE circuit can compress a set of n -qubit states to m -qubit states, where $m \leq n$. b.) A classical 4-2-4 autoencoder.

is a crucial component for algorithmic performance. One particular strategy for lowering the resource requirements for expressing quantum data (of quantum states) is through the use of the quantum autoencoder that combines ideas from both quantum simulation and quantum machine learning. Specifically, the quantum autoencoder³⁴⁴ is an HQC algorithm that encodes a set of quantum states using a fewer number of quantum bits, as shown in Figure 8. For instance, the wave function of molecular hydrogen (in the minimal basis) at a particular bond length requires four qubits to represent the occupancy of the ground state. However, the set of these ground states lives on only a one-dimensional manifold in the four-qubit space. This means that the quantum information corresponding to this four-qubit state can be represented by a single qubit. The task of a quantum autoencoder is to find a unitary mapping that coherently maps this four-qubit state to a single qubit. The payoff for finding such a unitary mapping is that it enables a type of generative modeling approach called *Compressed Unsupervised State Preparation* (CUSP) that explores the latent space to find a shorter pathway to a desired ground state³⁴⁵. CUSP uses an existing VQE circuit ansatz, together with a quantum autoencoder, to generate a new ansatz that is shorter in gate depth and hence less vulnerable to noise on a NISQ computer, resulting in more accurate VQE energies.

5.1.2 Energy estimation in VQE

After preparing the parametrized trial state $|\Psi(\vec{\theta})\rangle$, the next stage for the VQE algorithm is to estimate the expectation value of the target Hamiltonian, $E(\vec{\theta}) = \langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle$. In practice, this implies estimating the value of $E(\vec{\theta})$ up to a given precision, ϵ . Arguably, the most common application of *ab initio* quantum chemistry is the calculation of thermochemical properties, such as reaction rates and binding energies. Reaction rates, for example, are exponentially sensitive to changes in energy: a change of only 1.4 kcal/mol in the free energy (which is proportional to the energy) generates a change of one order of magnitude in the estimate of a reaction rate. This sensitivity motivates the concept of chemical accuracy,

which is a standard for the desired energy accuracy of 1 kcal/mol (1.59×10^{-3} Hartrees, 43.3 meV) or less⁷⁹. This requirement sets an upper bound to the precision, ϵ , for energy estimation in VQE and in general to any quantum simulation algorithm for chemistry.

One possible approach to estimating the energy is the PEA approach, which provides an estimate of the energy in a single state preparation with cost $O(\frac{1}{\epsilon})$ for error tolerance ϵ . Unfortunately, to guarantee the success of PEA estimation, all the operations must be executed fully coherently on the quantum computer, which is only possible using quantum error correction. This requirement is unattainable on NISQ devices because the operations available on these devices are too error-prone to achieve useful quantum error correction. To avoid the overhead of PEA, Peruzzo et al.⁵² suggested an alternative scheme called *Hamiltonian averaging* for computing the expectation value of local Hamiltonians.

Hamiltonian averaging. The basic idea behind this approach is that any Hermitian operator, such as the electronic structure Hamiltonian, can be decomposed in the basis of Pauli operators. Therefore, any Hermitian operator can be written as:

$$H = \sum_{i_1 \alpha_1} h_{\alpha_1}^{i_1} \sigma_{\alpha_1}^{i_1} + \sum_{i_1 \alpha_1 i_2 \alpha_2} h_{\alpha_1 \alpha_2}^{i_1 i_2} \sigma_{\alpha_1 \alpha_2}^{i_1 i_2} + \dots, \quad (69)$$

and by linearity the expectation value can be expressed as:

$$\langle H \rangle = \sum_{i_1 \alpha_1} h_{\alpha_1}^{i_1} \langle \sigma_{\alpha_1}^{i_1} \rangle + \sum_{i_1 \alpha_1 i_2 \alpha_2} h_{\alpha_1 \alpha_2}^{i_1 i_2} \langle \sigma_{\alpha_1 \alpha_2}^{i_1 i_2} \rangle + \dots, \quad (70)$$

where $\sigma_{\alpha_1 \alpha_2 \dots}^{i_1 i_2 \dots}$ represent a product of Pauli matrices, where i_j indicates the type of Pauli matrix and α_j qubit indices. In the case of approximation based on second-quantized approaches, the Hamiltonian is directly obtained as a sum of Pauli operators after a mapping from fermions to qubits is applied (see Appendix B). Unlike the PEA approach, where each k -local Pauli operator requires $O(k)$ circuit depth for the implementation, Hamiltonian averaging only adds a constant factor to the state preparation procedure, assuming parallel

rotation and orbital readout is available. However, not all the Hamiltonian terms can be measured simultaneously, which increases the number of measurements needed to achieve a given accuracy. In general, the number of state preparations and measurement cycles required to converge the energy to a precision ϵ will scale roughly as M^2/ϵ^2 , where M is the number of terms in the Pauli decomposition of the Hamiltonian. A more detailed analysis of the number of measurements takes into consideration the structure of the Hamiltonian terms. In general, any Hamiltonian for simulation can be written as $H = \sum_i^M h_i O_i$, where $\{h_i\}$ represents the set of Hamiltonian coefficients and $\{O_i\}$ correspond to a set of Pauli operators. Assuming that all the terms are measured independently i.e $\text{Cov}(\langle O_i \rangle, \langle O_j \rangle) = 0 \forall i \neq j$), the total variance of the estimator of $\langle H \rangle$, will be given by the sum of the variances of the individual Hamiltonian terms

$$\epsilon^2 = \sum_i^M |h_i|^2 \frac{\text{Var}(\langle O_i \rangle)}{\epsilon_i^2}, \quad (71)$$

where $\text{Var}(\langle O_i \rangle)$ denotes the variance of the estimator of $\langle O_i \rangle$. In practice, we need to define the precision at which each term is going to be measured. A simple approach would be to measure all the terms to the same accuracy i.e. taking $\epsilon_i^2 = \frac{\epsilon^2}{M}$. Considering that $\text{Var}(O_i) = 1 - \langle O_i \rangle^2 = \sigma_i$, where σ_i is a Pauli term, we can estimate the required number of measurements, m , as:

$$m = \frac{1}{\epsilon^2} \sum_i^M |h_i|^2 \sigma_i \leq \frac{1}{\epsilon^2} \sum_i^M |h_i|^2, \quad (72)$$

A more practical approach is to choose ϵ_i^2 to be proportional to $|h_i|$. This reduces the estimate in the number of measurements to:

$$m = \left(\frac{1}{\epsilon} \sum_i^M |h_i| \sigma_i \right)^2 \leq \left(\frac{1}{\epsilon} \sum_i^M |h_i| \right)^2, \quad (73)$$

which can be proved to be optimal by application of the Lagrange conditions³⁴⁶. Therefore, the effective number of measurements will be dependent on the magnitude of the Hamiltonian

coefficients, which is the case of electronic structure correspond to linear combinations of 1-electron and 2-electron integrals, which in turn depend on the choice of basis set for the calculation (see Appendix A for an introduction to basis sets).

As noted before, the electronic structure Hamiltonians by definition have $O(N^4)$ terms, however, in a local basis the number of non-negligible terms scale as $O(N^2)$ and advanced techniques can evaluate these contributions in $O(N)$ time. This is the basis of linear scaling methods employed in quantum chemistry^{79,347}. In the particular case of Gaussian-type orbitals (GTOs), this scaling can be verified by considering the cutoffs due to exponentially vanishing overlaps between GTOs and bounds on the largest integral to estimate the number of non-vanishing terms in the Hamiltonian^{79,306}. Although this scaling is guaranteed asymptotically, in practice, the molecular size at which this scaling is observed depends on the type of atoms in the system and the family of basis set used³⁰⁶. This analysis applied to Hamiltonian averaging showed that the number of measurements required for achieving precision ϵ in the energy is expected to scale as $O(N^6/\epsilon^2)$ for molecules of sufficient size, compared to the formal $O(N^8/\epsilon^2)$ estimated based on the number of terms in the Hamiltonian³⁰⁶. Moreover, this analysis justifies the intuitive strategy of neglecting Hamiltonian terms that are under a certain threshold δ , such as the maximum $\delta \leq \frac{\epsilon}{N_r}$, with N_r being the number of removed terms. In practice, it has been observed that $\delta = 10^{-10}$ is more than sufficient to retain quantitatively accurate results in molecular calculation with GTOs^{79,306,347}.

Alternatively, for Hamiltonians in the plane wave dual basis (Equation 44), the number of measurements required to estimate the expectation values can be bounded as³⁰¹

$$O\left(\frac{N^{14/3}}{\epsilon^2\Omega^{2/3}}\left(1 + \frac{1}{N^{4/3}\Omega^{2/3}}\right)\right). \quad (74)$$

This bound is based on upper bounds to the norm of the plane wave dual basis Hamiltonian applied to Equation 73. In this case, it is possible to obtain further savings in the number of measurements by exploiting the fact of the interaction (V) and potential terms

(U) of the Hamiltonian are diagonal and commute with each other. Therefore they allow the use of a separate, unbiased estimation for the mean of the sum of these terms³¹⁶. The kinetic operator is not diagonal in the plane wave dual basis, but it can be made diagonal by the application of a fermionic fast Fourier transform (FFFT). This trick can be used to measure the kinetic operator. Correspondingly, the number of measurements can be bounded as³⁰¹:

$$m \in O\left(\frac{\langle V^2 \rangle + \langle T^2 \rangle}{\epsilon^2}\right) \subseteq O\left(\frac{\eta^{10/3} N^{2/3}}{\epsilon^2} + \frac{\eta^{2/3} N^{4/3}}{\epsilon^2}\right), \quad (75)$$

where η is the number of electrons in the system, and T represents the kinetic energy terms of the Hamiltonian in Equation 44. This bound is obtained by bounding the norms of the different terms in the Hamiltonian assuming a neutral system and a fixed density. Since in finite molecules and bulk materials, $\eta \in O(N)$, this shows that the number of measurements does not scale worse than $O(N^4/\epsilon^2)$ for this representation of the Hamiltonian. To avoid the extra cost of implementing the FFFT, one could also measure $U + V$ at once and measure T by typical Hamiltonian averaging, which would scale as $O(N^4/\epsilon^2)$ ³⁰¹.

Strategies for reducing the number of measurements in Hamiltonian averaging.

The computational cost of operator averaging is one of the factors contributing to the overall cost of VQE prediction, the others being the depth of the circuit for state preparation and the number of function calls required by the optimization routine. Different strategies have been proposed to reduce the cost of operator averaging. The first of these approaches is term truncation, briefly described in the previous section. Implementation of this approach^{216,306} requires ordering the Hamiltonian terms according to the norm of their coefficients, which correspond to the maximum expected contribution to the total expectation value. After this, we can select sequences of particular sums $\delta_k = \sum_{i=1}^k |h_i|$, that correspond to the maximal bias introduced by neglecting the k smallest terms. Subsequently, one could choose a constant $C \in [0, 1)$ and remove the k^* lowest terms such that $\delta_{k^*} < C\epsilon^2$, where ϵ is the

intended accuracy. Correspondingly, the remaining terms must be computed to accuracy $(1 - C)^2 \epsilon^2$, and the number of measurements required would be $\frac{1}{(1-C^2)\epsilon^2} \left(\sum_{i=k^*+1}^M |h_i| \sigma_i \right)^2$, where the constant C can be adjusted according to the problem instance and implementation constraints to minimize the number of measurements.

A second strategy is to group commuting terms in the Hamiltonian to reduce the number of state preparations required for a single estimate of the energy, which was first described in McClean et al.²¹⁶. Commuting terms can be measured in sequence without biasing the final expectation. This can be achieved, for example, by using ancilla qubits to measure the Pauli operators, which only adds a small overhead in gate depth. When using this technique, terms within the same measurement group will be correlated, and therefore the covariance between their expectation values will be different from zero. This could increase or decrease the number of measurements required for the set of terms depending on whether the covariances are positive or negative, an effect similar to the one observed in correlated sampling in classical Monte Carlo simulations³⁴⁸. In practice, implementing this approach requires classical pre-processing to group the terms into commuting groups using sorting algorithms³¹², with a classical computational cost of $O(M^2)$ as a result of checking commutativity between pairs of terms. Furthermore, one must check whether a specific grouping lowers the number of measurements, which requires access to efficient approximations of the state of interest, since covariances are state dependent. Therefore, efficient classical algorithms for grouping Hamiltonian terms could benefit from efficient classical approximation to the electronic structure problem²¹⁶. A stronger condition that can be imposed on groups of commuting terms is to have all the terms being diagonal in the same tensor product basis, as proposed in Kandala et al.⁵⁶ (for example, the set of terms $\{ZIXZ, ZIXI, ZZXZ, ZZXI\}$). This guarantees that all the expectation values within the set can be obtained from a single state preparation followed by a specific set of single qubit rotations to measure in the appropriate tensor product basis.

More recently, Rubin et al.³⁴⁶ pointed out that N-representability conditions on the

fermionic reduced density matrices can be exploited to reduce the number of measurements in VQE calculations. N-representability conditions define a set of mathematical constraints which must be satisfied by the reduced density matrices (RDM). The p -electron RDMs (p -RDMs) are defined as the set of expectation values $\langle a_{i_1}^\dagger a_{j_1} \dots a_{i_p}^\dagger a_{j_p} \rangle$, for fermionic or bosonic systems. In the specific case of fermionic RDMs, the main constraints are the following³⁴⁹:

1. RDMs are Hermitian;
2. the 2-RDM is antisymmetric;
3. the $(p - 1)$ -RDM can be obtained by contracting the p -RDM;
4. the trace of each RDM is fixed by the number of particles in the system e.g. the trace of the 1-RDM is equivalent to the number of particles;
5. RDMs correspond to positive semidefinite density operators.

In addition, constraints related to the total spin and spin quantum numbers (S^2 and S_z) can be formulated for each marginal³⁵⁰. The equality N-representability conditions can be expressed in a qubit basis by mapping the RDM elements using a standard transformation such as Jordan-Wigner or Bravyi-Kitaev. The k -th condition can be written as: $C_k = \sum_i^M c_{k,i} \langle O_i \rangle = 0$, with O_i being the same terms appearing in the target Hamiltonian for VQE. These constraints provide extra information about the relations between expectation values that can be exploited to reduce the number of measurements. The basic idea of the procedure proposed in Rubin et al.³⁴⁶ is to add the constraints to the original Hamiltonian to reduce the sum in Equation 73. Specifically, this is equivalent to running VQE for an alternative Hamiltonian of the form:

$$\hat{H} = H + \sum_{k=1}^K \beta_k C_k = \sum_{i=1}^M \left(h_i + \sum_k^K \beta_k c_{k,i} \right) O_i, \quad (76)$$

where β is a set of real parameters. Notice that $\langle H \rangle = \langle \hat{H} \rangle$ due to the definition of the

N-representability conditions. Correspondingly, based on Equation 73, the number of measurements can be minimized by performing the following optimization:

$$\min_{\beta} \left(\sum_{i=1}^M \left| h_i + \sum_k^K \beta_k c_{k,i} \right| \right) \quad \text{or} \quad \min_{\beta} \left(\sum_{i=1}^M \left| h_i + \sum_k^K \beta_k c_{k,i} \right| \text{Var}(O_i) \right), \quad (77)$$

with the second optimization depending on the availability of a meaningful prior on the expectation values of the Hamiltonian terms. This optimization problem can be recast as a convex L_1 minimization method by writing the constraints and the original Hamiltonian as a vector, where each position of the vector represents a fermionic term³⁴⁶. The optimal β can be found by applying standard optimization methods. The application of this strategy to example molecular Hamiltonians showed reductions of around one order of magnitude in the number of measurements required³⁴⁶. Finally, another direction that deserves more exploration is the use of Bayesian inference to estimate the expectation values of Hamiltonian terms, described in McClean et al.²¹⁶. The use of sensible choices of classical approximations could provide priors for Bayesian estimation potentially leading to reductions in the number of measurements.

5.1.3 Optimization methods

Several numerical studies comparing the efficacy of different optimization methods have been published since VQE was first introduced; the original paper⁵² used Nelder-Mead, a simple derivative-free optimization algorithm. McClean et al.²¹⁶ studied VQE for optimizing H_2 using the software package TOMLAB³⁵¹ to compare four algorithms: Nelder-Mead, MULT-MIN, GLCCLUSTER, and LGO, finding that LGO was usually superior. Romero et al.³²⁹ simulated VQE for H_4 (a system of four hydrogen atoms), using both gradient and derivative-free methods for the classical optimization step. They used algorithms known as L-BFGS-B, COBYLA, Powell, and Nelder-Mead, observing that the derivative-free methods Powell and Nelder-Mead often had difficulties converging to the correct answer in the large parameter

space required by VQE, with COBYLA and L-BFGS-B showing a better performance.

Romero et al.³²⁹ and Guerreschi and Smelyanskiy³⁵² proposed quantum circuits to compute analytical gradients for HQC algorithms, the former in the context of VQE and the latter in the context of the QAOA. Both studies compared the cost of numerical gradients relative to analytical ones, which depends on the cost function employed. In the case of QAOA applied to the Max-Cut problem³⁵² the authors found that for quasi-Newton methods, numerical gradients required fewer overall function calls than analytical gradients, and both out-performed Nelder-Mead. This is in sharp contrast with VQE applied to electronic structure Hamiltonians³²⁹, where the conclusion is that analytical gradients offer a practical advantage over numerical gradients. The numerical gradients require a few orders of magnitude more measurements to achieve the same accuracy. Both these studies also point out the dependence of the achieved accuracy and convergence times on the hyper-parameters chosen for the classical optimizer^{329,352}.

In experimental implementations of VQE on NISQ devices, the accuracy of the optimization can be significantly lowered by the fluctuations on the measured properties caused by noise and finite measurement statistics. This is a major hurdle to the implementation of VQE for larger systems, where the number of parameters can be significant even for linear scaling ansätze. Correspondingly, most of the experimental implementations so far have employed methods that better tolerate these fluctuations, such as Nelder-Mead and Particle-Swarm Optimization (PSO)³⁵³, which are gradient free. More recently the Simultaneous Perturbation Stochastic Approximation (SPSA) has been applied^{354,355}. This method relies on numerical gradients computed in random directions. References to these experiments are provided in Table 5.

For small instances such as H_2 in a minimal basis set that require a single variational parameter, it is possible to sample the energy landscape. In this case, the expectation values of the Hamiltonian terms are computed within a certain range of the variational parameters and subsequently fitted using methods such as trigonometric interpolations or

Gaussian process regression^{55,59}. This procedure provides classical functions (which model the values of the expectation values of the Hamiltonian terms) as a function of the variational parameters: $\langle O_i(\theta) \rangle \approx f_i(\theta)$. Subsequently, the optimization can be performed on the fitted models to simplify classical post-processing. These functions are convenient when applying VQE to a family of Hamiltonians (e.g. ground state calculations along a path in the potential energy surface), since the same expectation values are used to compute all the ground states, saving measurements.

Methods such as Nelder-Mead are limited to problems with few parameters³⁵⁶ and more likely their application is not scalable to large molecules. In contrast, SPSA is better for treating more parameters and is also expected to be robust to statistical fluctuations³⁵⁵. However, being a method based on numerical gradients, the accuracy and precision of SPSA is also limited by the step size chosen for the gradient discretization. Choosing a smaller step size implies lower energy differences and therefore more precision required for the energy estimation. In fact, the number of measurements required for the estimation of the gradient at a fixed precision increases quadratically with the inverse of the step size, as shown in Romero et al.³²⁹. Consequently, methods employing numerical gradients require careful tuning of the hyper-parameters throughout the optimization to minimize the number of measurements employed. Some heuristic have been proposed for dynamically updating the hyperparameters of SPSA in VQE calculations (see appendix of Kandala et al.⁵⁶). Nonetheless, more efforts are needed to design efficient strategies to estimate energies and gradients to sufficient precision in VQE.

A different and complimentary strategy to improve optimization, and the quality of VQE results in general, is error mitigation. This term groups a series of recent proposals for improving the quality of the expectation values of observables measured on NISQ devices. Since estimation of expectation values is at the heart of the VQE protocol, these approaches have a direct benefit in the accuracy of the energies and gradients measured in VQE, which in turn could improve the overall performance of the optimization. Some of the proposals for error

mitigation assume that the first-order contributions of the noise to the expectation values can be removed by introducing a controllable source of noise in the circuit of interest^{317,318}. The expectation values are estimated at different error levels and an extrapolation to zero noise is performed using different mathematical techniques. Some of these methods have been successfully applied in experimental demonstrations of VQE for small molecules³¹⁹. However, it has been pointed out that the successful application of some of these approaches to bigger systems might require lower error rates than those existing today³⁵⁷.

Other proposals for mitigating errors are the VQE subspace expansion^{58,320} and the updating of expectation values using marginal constraints³⁴⁶. In the former case, described in more detail in Section 5.1.5, the target Hamiltonian is expanded in a basis built from the ground state obtained from a regular VQE procedure at the cost of an additional number of measurements. This new expansion allows for the calculation of excited state energies and a refined ground state energy which is found to be closer to the exact result. The second approach exploits the physical constraints of fermionic reduced density matrices, described in Section 5.1.2, to update the reduced density matrices measured experimentally. Although the enforcement of these conditions can increase the energy, the procedure can improve the overall shape of the predicted potential energy surfaces and derived properties such as forces. Finally, McArdle et al.³⁵⁷ have also proposed the use of easy-to-implement stabilizer checks for molecular properties (e.g. electron number parity) to detect errors and discard faulty state preparations, which improves the overall VQE performance. The so-called stabilizer-VQE, can be implemented in combination with other error mitigation strategies. For a more detailed description of error mitigation techniques we refer the reader to McArdle et al.³⁵⁸.

The recent observation³²¹ that the variational parameter landscape is often plagued by the overabundance of “barren plateaus” implies that care must be taken when designing a circuit ansatz for the VQE algorithm. By Levy’s lemma, one can prove that expectation values only substantially deviate from their mean over an exponentially small fraction of the Hilbert space. This implies that the use of arbitrary or unstructured ansatze may be uninformative

and expensive, as the optimization algorithm would have to search through most of the space before measuring a non-zero gradient. However, this issue can be circumvented by using physically motivated ansätze that have measurable gradients and thus could better guide the optimizer. This is the case of UCC, where classical quantum chemistry methods such as Møller-Plesset second order perturbation theory and traditional CC or CI can provide approximate amplitudes to initialize the optimization. As described in Section 5.1.1, the barren plateau phenomenon presents a challenge to unstructured energy minimization in VQE. Besides selecting a well-motivated circuit ansatz, several other strategies may be taken to address this issue. One of these strategies is the “annealed variational” optimization⁴⁷ described above. This method can be naturally applied to Hamiltonian variational ansätze, HEA and LDCA, where the structure of the variational ansatz is naturally divided into layers or circuit blocks that can be optimized following a block-by-block fashion. We could also adapt this method to ansätze such as UCC, by separating the excitation operators into subsets that are implemented and optimized sequentially.

Finally, an alternative strategy to break the optimization into easier, smaller parts is the *Adiabatically Assisted Variational Quantum Eigensolver* (AAVQE) method³⁵⁹. This method is inspired by the adiabatic evolution of a ground state in quantum annealing. As in adiabatic quantum computing, the method employs a parametrized Hamiltonian of the form:

$$H(s) = (1 - s)H_0 + sH_P \tag{78}$$

where H_0 is the initial Hamiltonian, for which the ground state preparation is easily implemented (e.g. a local Hamiltonian may be used), H_P is the problem Hamiltonian, and s interpolates between the two Hamiltonians as it is incremented from 0 to 1. However, rather than dynamically tuning the interpolation parameter (see Adiabatic Quantum Computing in Section 5.3.1), the interpolation parameter is used to adjust the Hamiltonian from one VQE run to the next. The state preparation parameters at each step are initialized by the optimized parameters of the previous step. As long as the gap of $H(s)$ does not become too

small (see Section 5.3.1), the ground state of $H(s)$ will lie close to the ground state of the Hamiltonian $H(s + \Delta s)$ in the subsequent step. Thus, AAVQE, which runs VQE while incrementally adjusting the optimization landscape, appears to be another promising approach for circumventing the barren plateau issue. Notice that this method is not equivalent to implementing a Hamiltonian variational approach, since it is the Hamiltonian and not the variational ansatz that has an adiabatic parametrization. Consequently, AAVQE could be combined with any choice of variational ansatz for its execution.

5.1.4 Physical realizations of VQE

To demonstrate the capabilities of existing quantum devices, few-qubit simulations, often without error correction, have been carried out in most major architectures used for quantum information, as shown in Table 5. The first experimental demonstration of VQE was performed in 2014⁵². This experiment employed a two-qubit photonic chip to variationally minimize the energy of HeH^+ . A VQE experiment of the same molecule was achieved using a system comprising a single trapped ion⁵³. These experiments served as proofs-of-principle for the VQE approaches applied to chemistry but involved non-scalable simplifications of the Hamiltonians or non-scalable experimental procedures such as full tomography. The first scalable demonstration of both the IPEA and VQE algorithms employed three superconducting qubits for simulating H_2 in a minimal basis and was carried out by a Google Research and Harvard collaboration⁵⁵. These demonstrations were followed by the demonstration of VQE with a hardware-based ansatz for H_2 , lithium hydride (LiH) and beryllium hydride (BeH_2) in a system with six qubits⁵⁶ by the IBM quantum computing team. Similarly, the first scalable simulation of VQE in ion trap quantum computers was demonstrated for H_2 and LiH using an approximate implementation of the UCC ansatz with Molmer-Sorensen gates. In this demonstration, LiH was simulated using an active space approach that exploited the Bravyi-Kitaev mapping to generate an effective Hamiltonian for this molecule in three qubits⁵⁹.

In recent years, various full-stack platforms have been introduced by IBM, Rigetti Computing, and Microsoft, supporting quantum computations via cloud computing. Leveraging this availability of quantum resources, a recent study computed the deuteron binding energy using the VQE framework on quantum devices from IBM and Rigetti Computing³³⁰.

5.1.5 VQE for excited states

Since the development of the variational quantum eigensolver (VQE) algorithm by Peruzzo et al.⁵², numerous studies and demonstrations of VQE focused on approximating ground states of physical systems. In principle, the VQE algorithm was designed as a modular framework, treating each component (i.e. state preparation, energy estimation, classical optimization) as a black box that could be easily improved and/or extended. Recent studies have leveraged this flexibility, specifically applying different formulations of objective functions to compute excited states, which are fundamental to understanding photochemical properties and reactivities of molecules. In the following subsections, we highlight several methods extending the original VQE algorithm for approximating excited states for molecular systems.

Folded Spectrum and Lagrangian-Based Approaches. The first and perhaps the simplest extension consists of the application of the folded spectrum method, which utilizes a variational method to converge to the eigenvector closest to a shift parameter λ .

This is achieved by variationally minimizing the operator $H_\lambda = (H - \lambda I)^2$ according to Peruzzo et al.⁵². This methodology, though relatively straightforward to implement, requires a quadratic increase in the number of terms of the effective Hamiltonian. This translates to a significant increase in the number of measurements needed, especially in the case of quantum chemistry Hamiltonians³¹⁶.

A related approach consists of adding a set of constraints to the original Hamiltonian in

the VQE calculation to construct the Lagrangian²¹⁶:

$$L = H + \sum_i \lambda_i (S_i - s_i I)^2 \tag{79}$$

where λ_i are energy multipliers, S_i are sets of operators associated with the desired symmetries, and s_i are the desired expectation values for these sets of operators. The set of operators S_i , e.g. spin numbers, accounts for symmetries whereby the energies are minimized by the appropriate excited states (with respect to the original Hamiltonian). We note that S_i^2 and S_i must be efficiently measurable on the quantum computer to ensure the efficiency of the method. By solving VQE for the Lagrangian instead of the original Hamiltonian, it is possible to converge to an approximation of the excited state.

More recently, two new Lagrangian-based approaches were introduced for calculating excited state energies without the measurement overhead³⁶⁰. The first doubles the circuit depth to measure overlaps between the ground state and the prepared state to put a penalty on the overlap with the ground state. Choosing this penalty to be large enough ensures that the first excited state becomes the minimizer of the new cost function. The second method works by a similar principle, except the size of the quantum register is doubled instead of the circuit depth. A SWAP test, a circuit construction to compare two quantum states (or quantum registers) by computing the overlap, is then applied to incorporate a penalty for the prepared state having overlap with the ground state.

Linear Response: Quantum Subspace Expansion (QSE). More recently, a methodology based on linear response has been developed³²⁰ and demonstrated on existing hardware⁵⁸. In summary, this framework, called the Quantum Subspace Expansion (QSE), extends the VQE algorithm and requires additional measurements to estimate the excited state energies. That is, after obtaining the the ground state $|\psi\rangle$ of a molecule using VQE, an approximate subspace of low-energy excited states is found by taking the linear combinations

of states of the form $O_i |\psi\rangle$, where the O_i are physically motivated quantum operators. For example, in fermionic systems, these operators could correspond to fermionic excitation operators. In the algorithm, the matrix elements $\langle \psi | O_i H O_j | \psi \rangle$ are computed on the quantum device. The classical computation then diagonalizes the matrix to find the excited state energies. While the QSE method benefits from the low coherence time requirements of VQE, the quality of the excited states obtained is subject not only to the quality of the ansatz employed in VQE but also to the errors induced by the linear-response expansion.

Witness-Assisted Variational Eigenspectra Solver (WAVES). An alternative protocol that also utilizes VQE as a subroutine to compute the ground state is the witness-assisted variational eigenspectra solver (WAVES)³¹⁶. The objective function in WAVES is augmented to include the energy (E) as well as an approximation for the entropy (a purity term $\text{Tr}[\rho_C^2]$):

$$\mathcal{F}_{obj}(\mathcal{P}, E) = E - T \cdot \text{Tr}[\rho_C^2]. \quad (80)$$

In this setup, a control ancilla qubit is considered along with the trial state. Here, the control qubit behaves as an “eigenstate witness” where its entropy measurement nears zero if the optimized trial state is arbitrarily close to an eigenstate of the Hamiltonian.

A tunable parameter T (that can be pre-optimized) is used to bias towards excited states. In the first iteration of WAVES, T is set to 0 (i.e. implementing regular VQE) to compute the ground state. Then, T is tuned such that when the objective function is optimized, the resulting trial states correspond to approximate excited states. These states are fed into the iterative phase estimation algorithm (IPEA) to extract the corresponding excited state energies. For near-term devices, the IPEA procedure could be replaced by a Hamiltonian averaging approach³²⁰.

5.1.6 Calculation of molecular properties with VQE

Obtaining the ground state for a molecular system is usually just the first step in a quantum chemistry study. In most cases, chemists are interested in a variety of molecular properties that provide insights into the chemistry of the target system. Some of these properties, such as the dipole moment and the charge density, can be computed by measuring the corresponding operators represented in the basis set employed for the VQE. For example, the expectation values of 1-electron (E^1) and 2-electron (E^2) operators can be calculated using the same 1-RDM and 2-RDM measured for the optimal state obtained from VQE:

$$\langle E^1 \rangle = \sum_{pq} o_{pq} \langle a_p^\dagger a_q \rangle, \quad (81)$$

$$\langle E^2 \rangle = \sum_{pqrs} o_{pqrs} \langle a_p^\dagger a_q^\dagger a_r a_s \rangle. \quad (82)$$

where the coefficients o represent expansion coefficients. Other properties might require the estimation of k -RDMs, with $k > 2$, which requires a larger number of measurements.

The information of the RDMs measured in the VQE procedure can be used to connect the VQE algorithm with complementary classical algorithms for the calculation of other properties or energy refinement. One possibility is to combine VQE calculations on active spaces with multi-reference perturbation theory techniques implemented on a classical computer to add dynamical correlation to the calculations, as suggested in³⁴⁶. For example, n -valence electron perturbation theory³⁶¹ requires measurement of the 3-RDM and 4-RDM in order to compute the corresponding energy correction. To avoid the measurement cost of estimating these RDMs, a cumulant approach that approximates this correction using only the 1-RDM and 2-RDM could be employed³⁶². Other techniques that could be implemented using this strategy are canonical transformation theory¹¹⁹ and perturbative explicit correlated corrections^{363,364}. Similarly, the VQE algorithm can be used as a subroutine for energy calculation and interfaced with classical routines for geometry optimization and calculation

of thermodynamic properties. In all these calculations, an important factor to consider is the impact of measurement noise, especially when iterative procedures are involved. Investigating the extension of VQE to the calculation of molecular properties common in classical quantum chemistry is an important direction of research for HQC techniques.

5.2 Other hybrid quantum-classical algorithms for chemistry

In this section, we highlight several other HQC algorithms that may be useful in the context of quantum chemistry. These recent algorithms highlight the opportunity for continued innovation in the development of HQC algorithms.

5.2.1 Variational quantum simulator

A common explanation for the efficiency of Hamiltonian simulation is the following: as a controllable quantum system, the quantum computer can be made to evolve in a way which mimics the to-be-simulated quantum system. Surprisingly, this canonical approach to quantum simulation, that was detailed in Section 4, is not the only method for accomplishing this task. Motivated by the hopeful prospects of near-term quantum algorithms, such as VQE, Li and Benjamin²⁰⁰ introduced a variational quantum algorithm which implements Hamiltonian simulation, known as the variational quantum simulator algorithm.

Instead of applying quantum gates which simulate Hamiltonian dynamics, the variational quantum simulator algorithm applies a series of tunable gates to prepare a tunable output state. The evolution is discretized (t_1, t_2, \dots) , such that the task of the algorithm is to determine a series of circuit parameter settings $(\vec{\theta}_1, \vec{\theta}_2, \dots)$ for which the subsequent output states $(U(\vec{\theta}_1)|\psi_0\rangle, U(\vec{\theta}_2)|\psi_0\rangle, \dots)$ approximate the time-evolved states $(|\psi(t_1)\rangle, |\psi(t_2)\rangle, \dots)$ according to the Schrödinger equation. The algorithm achieves this by invoking the variational principle: the optimal “path” in parameter space for expressing the dynamics is the one which is stationary with respect to variation of the action $\int_{t_i}^{t_f} dt \langle \psi(\vec{\theta}(t)) | (i \frac{d}{dt} - H) | \psi(\vec{\theta}(t)) \rangle$. The Euler-Lagrange equation derived from this variation-minimization problem determines a

differential equation for the evolution of the parameters $M\dot{\vec{\theta}} = \vec{V}$. The crux of the algorithm is that the entries of the matrix M and vector \vec{V} can be estimated from measurements made on the quantum computer. It is important to note that these measurements are facilitated by using an ancilla qubit which is coupled to the system register through controlled-unitary gates. For details of the measurement scheme used to determine M and \vec{V} , refer to the original publication²⁰⁰.

Each subsequent setting for the parameters is then determined by the update rule

$$\vec{\theta}_{i+1} = M^{-1}\vec{V}(t_{i+1} - t_i) + \vec{\theta}_i. \quad (83)$$

This matrix inversion and multiplication is carried out by a classical processor. The role of the quantum computer is to explore families of output states, $U(\vec{\theta})|\psi_0\rangle$, which are unlikely to be representable on a classical computer, and to use these states to efficiently estimate the entries of M and \vec{V} .

5.2.2 Imaginary-time variational quantum simulator

The variational quantum simulator technique has been developed further for problems beyond quantum simulation. In McArdle et al.³⁶⁵ a variant of the variational quantum simulator is used for preparing certain useful quantum states. The key insight of this paper is that the time evolution in the variational quantum simulator algorithm can be replaced by imaginary-time evolution $e^{-i(-i\tau)H} = e^{-\tau H}$.

The functioning of the algorithm is similar to the variational quantum simulator algorithm, apart from replacing t with $-i\tau$. Imaginary-time evolution is useful for preparing a thermal state at temperature T with respect to the Hamiltonian: $\rho(T) = e^{-H/T}/\text{tr}(e^{-H/T})$. If the system is initialized in the completely mixed state \mathbb{I}/d , then the imaginary-time evolution of $\tau = 1/(2T)$ will yield a thermal state $e^{-\tau H} \cdot \frac{\mathbb{I}}{d} \cdot e^{-\tau H} \propto e^{-H/T}$. Preparation of a thermal state can be used for preparing approximations to the ground state

of the Hamiltonian by taking the limit $T \rightarrow 0$, since, for a non-degenerate Hamiltonian $|\psi_{\text{gs}}\rangle\langle\psi_{\text{gs}}| = \lim_{T \rightarrow 0} e^{-H/T} / \text{tr}(e^{-H/T})$.

With approximate preparation of a Hamiltonian’s ground state, this algorithm can be used as an alternative to the variational quantum eigensolver algorithm. Initial simulations show that this method is competitive with the traditional gradient descent approach of the variational quantum eigensolver. Hardware implementations which directly compare the two methods are still needed.

5.2.3 Imaginary-time variational quantum simulator for excited states

The imaginary-time variational quantum simulator can be used as a substitute for VQE. Just as variants of VQE have been developed for determining excited state energies of Hamiltonians^{320,360}, a variant of the imaginary-time variational simulator has been proposed to do the same³³⁹.

Similar to the excited-state energy methods described in Section 5.1.5, the VQS version of determining excited state energies iteratively builds up the spectrum starting with the ground state energy. Furthermore, in both approaches, each step introduces a penalty term to the Hamiltonian that is proportional to the projector onto the previously-found eigenstate. As proposed before³⁶⁰, a short-depth swap test circuit is added³³⁹ to measure the contributions to the Hamiltonian from the additional projectors. The difference in the VQS version for determining excited states is that these penalty terms added to the Hamiltonian alter the imaginary-time dynamics, rather than the cost-function that is optimized.

The value of these algorithms is that they provide alternative approaches to several established quantum algorithms. An important caveat regarding these alternative algorithms is that, as with VQE, they are heuristic. There are no guarantees on their performance, and their proper functioning can only be established on a case-by-case basis. Furthermore, the true test of such heuristic algorithms will be when they are implemented on hardware. So, while these algorithms, along with VQE, are promising candidates for making use of near-

term quantum devices, more substantial experimental testing is needed to determine their potential for scalability.

5.3 Non-gate-based methods for quantum chemistry

5.3.1 Adiabatic

The results discussed so far are based on gate model quantum computers, which can be broadly categorized as the “digital” approach to quantum simulation. The overall unitary quantum evolution being simulated is eventually decomposed into a sequence of elementary operations called quantum gates in a way similar to how logic gates underlie classical computers. However, there is an alternative “analog” approach to quantum computing based on adiabatic quantum evolution⁶⁰ that has also gained much attention in the quantum information community. Rapid experimental progress has been underway to scale up devices for realizing adiabatic quantum computing architectures (D-Wave Systems Inc., 2018).

In the context of quantum chemistry, research has focused on the analog simulation of molecules by mapping the electronic structure problem to a Hamiltonian of spins with additional coupling terms³⁶⁶, although proposals exist for simulating chemical reactions^{367–369} and protein folding simulations^{370,371}. We will illustrate the analog simulation of electronic structure in the case of hydrogen in a minimal basis set, as shown in Babbush et al.³⁶⁶. The Hamiltonian of molecular hydrogen in the minimal basis mapped to a four-qubit system, using the Bravyi-Kitaev transformation, takes the following form^{181,366}:

$$H_{elec} = \mu_1 \mathbb{I} + \mu_2 Z_1 + \mu_3 Z_2 + \mu_4 Z_3 + \mu_5 Z_1 Z_2 \tag{84}$$

$$+ \mu_6 Z_1 Z_3 + \mu_7 Z_2 Z_4 + \mu_8 X_1 Z_2 X_3 + \mu_9 Y_1 Z_2 Y_3 \tag{85}$$

$$+ \mu_{10} Z_1 Z_2 Z_3 + \mu_{11} Z_1 Z_3 Z_4 + \mu_{12} Z_2 Z_3 Z_4 \tag{86}$$

$$+ \mu_{13} X_1 Z_2 X_3 Z_4 + \mu_{14} Y_1 Z_2 Y_3 Z_4 + \mu_{15} Z_1 Z_2 Z_3 Z_4, \tag{87}$$

We refer to Appendix C.3 for a derivation and description of this formula. A suitable

adiabatic quantum computer would be able to effectively encode such a spin Hamiltonian. It would then evolve the system from an “easy” Hamiltonian (such as the transverse field $H_0 = \sum_i X_i$) with known ground state to the spin Hamiltonian of interest, therefore encoding the ground state of the molecule in its final state.

Two major challenges in successfully implementing this adiabatic evolution are keeping the energy scales well below the size of energy gaps in the Hamiltonian throughout the evolution and realizing the many-body interactions. For the first issue, the energy scales that must be considered are temperature and speed with which the Hamiltonian is changed. Thermal fluctuations which drive the system out of the ground state can be mitigated by keeping the system at extremely low temperatures. Additionally, the Hamiltonian must be changed slowly enough relative to the transition energy of the first excited state (the energy gap) so that non-adiabatic transitions do not drive population out of the ground state. For the second issue, certain many-body terms in the Hamiltonian, such as three- and four-body terms $X_1Z_2X_3$ and $Y_1Z_2Y_3Z_4$, are difficult to implement experimentally. Techniques exist to reduce them to two-body terms by adding ancilla qubits and encoding the correct couplings in the lowest eigenstate of a different Hamiltonian, which is commonly called *gadgetization*^{42,208,372–374}. Although gadgetization schemes have greatly improved over the last decade, they provide limited tunability over the couplings that are required to construct the problem Hamiltonian. This poses a challenge as molecular Hamiltonians and the adiabatic evolution require the spread of coupling term strengths over multiple orders of magnitude. Therefore, experimental progress in the field of adiabatic quantum computing for chemistry applications hinges on the development of non-gadgetized, highly tunable multi-spin couplers that either directly encode the problem Hamiltonian or reduce the mapping overhead of gadgetization.

5.3.2 Linear optics

As we have seen earlier, much of the work in quantum simulation is related to the structure of atoms and molecules, such as electronic structure, which are typically fermionic in nature. However, bosonic dynamics govern many other aspects of chemical systems; for instance, the vibrational modes of molecules can be described by phonon transitions, which obey bosonic statistics. While one could imagine devising a gate-model quantum algorithm to simulate this system, optical quantum systems are native to bosonic statistics and represent a platform for carrying out this simulation more efficiently than encoding the problem into a qubit formalism. The first realization of this protocol was introduced by Huh and Yung²²⁸ and targeted the simulation of the Franck-Condon profile. This work and its successors^{375,376} are the only examples of native simulation of bosonic chemical systems on a quantum device to date.

In Section 3, we discussed some of the complexity issues related to computing distributions whose probability amplitudes are derived from matrix permanents. Here we give a description of how the algorithm works. The goal is to compute the FCP by sampling from the output distribution of a *modified* boson sampling device, as shown in Figure 9.

This modification of the optical apparatus stems from the unitary transformation (Equation 19) that encodes the molecular problem²²⁰. On top of the M -dimensional rotation of the bosonic modes given by the Duschinsky rotation \hat{R}_U , one has to include the displacement \hat{D}_δ caused by the different minima of the nuclei in the electronic ground and excited states, together with the squeezing $\hat{S}_{\Omega'}^\dagger$ that photons undergo for transitioning to different frequency states ($\Omega' \rightarrow \Omega$). These effects are encoded in the initial state preparation, which can be written as a squeezed coherent state $|\psi\rangle = \hat{S}_\Omega \hat{D}_\delta |\text{vac}\rangle$. The unitary rotation \hat{R}_U is implemented in a linear optics setup, and precedes the last squeezing operation $\hat{S}_{\Omega'}^\dagger$. The second squeezing operation might be challenging to implement in optical setups, for which Huh et al. proposed a more efficient setup, compressing the two squeezing operations into the initial state, so that the FCP simulation is isomorphic to a Gaussian Boson Sampling




Figure 9: Left: Sketch of the photoelectric spectrum of the ionization of a molecule. The energy level diagram shows the electronic ground and excited state of the molecule, with its corresponding vibrational energy levels. Vibronic transitions are allowed between the two distinct electronic states, with the intensity of the peaks given by the wave function overlap integral $\langle v'_k m | v_0 \rangle$. According to the Franck-Condon principle, due to the slower motion of nuclei the most intense peak in the vibronic spectrum corresponds to the vertical transition in the energy level diagram (blue arrow). The whole spectrum is known as the Franck-Condon profile (FCP). Right: Modified boson sampling device to simulate the FCP. An initially prepared squeezed-coherent state of light (dashed blob represents the resulting photonic entangled state) is sent through a photonic circuit, where a judicious choice of beam splitters and phase shifters implement the Doktorov rotation U . Sampling from the output distribution of this apparatus yields the FCP, up to a certain error ϵ , given by the number of measurements.

problem³⁷⁷.

Given this setup, the following hybrid classical-quantum sampling algorithm can efficiently generate the FCP profile at $T = 0$ K:

1. With the information of the potential energy surfaces obtained from electronic structure calculations, calculate the parameters that define the Doktorov transformation. The electronic structure method could be a classical approach such as time-dependent density functional theory (TD-DFT), but eventually also a quantum computation.
2. Build a Gaussian boson sampler as described above, that implements this Doktorov transformation and samples on the phonon number basis, and collect the photon numbers $\{m_k\}$ in each mode.
3. For every detection event estimate the corresponding associated energy $E = \hbar\omega_k m_k$ where m_k is the number of detected photons and increment its corresponding bin by one unit.
4. Repeat the sampling N_{samp} times until the estimated statistical error is below a given threshold ϵ_{FCP} .
5. Output the normalized profile of the counting statistics.

This result has recently motivated proof-of-principle experimental quantum simulations of molecular spectroscopic systems with linear optics^{375,376}. The paper by Clements et al.³⁷⁵ is the first instance of an experimental FCP of a small molecule, tropolone. The improved algorithm accounts for the unavoidable effects of experimental imperfections. As such it is an excellent example of a scalable simulation in the absence of error-correcting techniques. In Sparrow et al.³⁷⁶, the authors simulate the vibrational quantum dynamics of a variety of molecules (H_2CS , SO_3 , HNCO , HFHF , N_4 and P_4), both in the harmonic and anharmonic regime, using a reprogrammable photonic chip.

Linear optics is the most natural platform for the simulation of vibrational quantum dynamics, as photons have the same boson statistics as the vibrational excitations (phonons). Nevertheless, other physical platforms such as trapped ions or superconducting circuits have proven useful in simulating molecular dynamics. In particular, Shen et al.³⁷⁸ reproduce the molecular vibronic signal of SO_2 , using solely the vibrational excitations in a trapped ion device. Superconducting circuits have no experimental realizations as of yet, however promising architectures have been theoretically proposed. In particular the proposal by Olivares et al.³⁷⁹ consists of superconducting resonators interacting through tunable couplers. In their approach the authors perform a quantum quench in the superconducting simulator, initially prepared in the molecular ground state. The abrupt change takes the system out of equilibrium, in a way that the system relaxes to the ground state populating the resonators with a distribution that resembles that of the FCP.

5.4 Chapter 5 Glossary

$\text{Var}(X)$	Variance of X
$\text{Cov}(X, Y)$	Covariance between X and Y
N	Number of spin orbitals, number of qubits
η	Number of electrons
$ \Psi(\vec{\theta})\rangle$	Variational state
$U(\vec{\theta})$	Variational unitary
H	Hamiltonian
$ \Psi_0\rangle$	Reference state
T	Cluster operator
O	Observable
S	Number of steps in Hamiltonian variational ansatz
t	Time
a_p^\dagger, a_p	Fermionic second quantization operators in canonical basis

c_p^\dagger, c_p	Fermionic second quantization operators in alternative basis
b_p^\dagger, b_p	Bogoliubov second quantization operators
$\gamma_p^\dagger, \gamma_p$	Majorana second quantization operators
M	Number of terms in the Hamiltonian
Γ	Covariance matrix for fermionic gaussian state
\mathcal{U}	Rotation matrix
C	Constant
R_j	Rotation matrix
U_{bog}	Bogoliubov transformation
L	Number of layers in LDCA ansatz
T_{ij}	Coefficient of kinetic term in Hamiltonian
U_{ij}	Coefficient for potential term in Hamiltonian
V_{ij}	Coefficient for interaction term in Hamiltonian
X, Y, Z, I	Pauli Operators
$\sigma^x, \sigma^y, \sigma^z, I$	Pauli Operators
$O(\cdot)$	Big-O notation
O	Observable
NISQ	Noisy intermediate-scale quantum
HQC	Hybrid quantum-classical
VQE	Variational quantum eigensolver
VQC	Variational quantum-classical
CUSP	Compressed unsupervised state preparation
PEA	Phase estimation algorithm
PMA	Physically motivated ansatz
HHA	Hardware heuristic ansatz
HEA	Hardware efficient ansatz (a specific kind of HHA)
CC	Coupled cluster

UCC	Unitary coupled cluster
UCCSD	Unitary coupled cluster singles and doubles
LDCA	Low-depth circuit ansatz
FGS	Fermionic gaussian state
BCH	Baker-Campbell-Hausdorff
ASP	Adiabatic state preparation
DVR	Discrete variable representation
FSN	Fermionic swap network
BCS	Bardeen-Cooper-Schrieffer
HF	Hartree-Fock

6 Summary and outlook

Quantum chemists have come to embrace a set of computational tools which have driven much innovation in chemistry over the last half century. As quantum computers become available, many of these computational tools will be supplemented with, or even replaced by, quantum computations. Furthermore, quantum computation will likely inspire new theoretical approaches and computational methods for solving problems in chemistry. As a precedence, consider the way that early computational tools gave rise to the invention of Monte Carlo methods³⁸⁰. Today's quantum chemists face an impending disruption of their field. As with the previous computational revolution in quantum chemistry, the practitioners and early adopters will drive much of the innovation. This article has aimed to provide the perspective and basic tools helpful in beginning the practice of using quantum computation for chemistry.

In the near term, quantum computation will likely first be used for *ab initio* electronic structure calculations in quantum chemistry. Quantum algorithms for chemistry, then, should be viewed as alternatives to the state-of-the-art post-Hartree-Fock methods such

as coupled-cluster, Møller-Plesset perturbation theory, and density-matrix renormalization group techniques. Such calculations often suffer from inaccuracies due to approximations of the quantum electronic wave function (c.f. Section 2.2). Quantum computers provide a solution to this problem by naturally handling wave functions that span the full Hilbert space to estimate energies, electric polarization, magnetic dipoles, reduced density matrices, etc. The canonical problem that is expected to be solved is the computation of ground and excited state energies of small molecules. Such calculations serve as the starting point for computing many useful quantities, such as reaction pathways, binding energies, and rates of chemical reactions.

There are two distinct approaches to estimating the electronic ground state energies. The first method, proposed by Aspuru-Guzik et al.⁴⁵, uses the quantum phase estimation algorithm to make quantum measurements of the energy of a trial wave function, as described in Section 4.1.2. Although promising, this method requires the use of quantum error correction to properly function, and is therefore not feasible using near-term quantum devices.

In contrast, the second method for estimating ground state energies was specifically developed to be deployable on currently available quantum hardware, without a dependence on quantum error correction for proper functioning. This algorithm, known as the variational quantum eigensolver, as described in Section 5.1 has become the focus of the state-of-the-art quantum experiments^{55,56}. It is possible that the first commercial use of quantum computers will involve applying the variational quantum eigensolver algorithm to predict the electronic structure properties of small molecules. In the few years since its invention, there has been significant algorithmic development surrounding the variational quantum eigensolver algorithm^{320,326,327,360}. Yet, there remains room for innovation surrounding this and other variational quantum algorithms. It is possible that these innovations, coupled with advances in hardware, will enable commercial utility of quantum devices far sooner than anticipated. This crossover point will likely be arrived at through the synergistic collaboration between quantum chemists, quantum information scientists, and quantum device engineers.

Already, insights from quantum chemistry have led to substantial improvements in quantum algorithms. Variational ansätze borrowed from classical quantum chemistry such as the unitary coupled cluster have informed state preparation techniques³²⁹ in several quantum algorithms. The quantum subspace expansion technique, developed in McClean et al.³²⁰, was motivated by the concept of linear response of the ground state. Finally, the N-representability problems and the 2-RDM constraints inspired the methods in Rubin et al.³⁴⁶ for improving expectation value estimates in VQE. Yet, many avenues remain to be explored. For example, a recent paper³⁸¹ initiated the use of low-rank tensor decompositions for improving the performance of certain quantum algorithms for chemistry simulation. Such techniques should find a broader range of applicability in quantum algorithms. Furthermore, as techniques in quantum algorithms mature, it is likely that methods from perturbation theory will be useful in pushing the capabilities of quantum computers, especially in regard to active space methods.

It is difficult to predict the course of quantum computing in the long term. In particular, it is difficult to predict when or if algorithms on near-term noisy intermediate-scale quantum devices will outperform classical computers for useful tasks. But it is likely that, at some point, the achievement of large-scale quantum error correction will enable the deployment of a host of so-called error-corrected quantum algorithms, which in many cases have theoretical guarantees on their performance. These error-corrected algorithms include the celebrated Grover search algorithm, Shor’s factoring algorithm, and the HHL linear system of equations algorithm. Regarding quantum chemistry, quantum error correction would enable the use of the various Hamiltonian simulation algorithms, as described in Section 4.1. These Hamiltonian simulation algorithms can be used to simulate the dynamics of quantum systems. But, more importantly, Hamiltonian simulation serves as a subroutine for other algorithms, such as the quantum phase estimation algorithm. Thus, with quantum error correction, the quantum phase estimation algorithm could become a viable alternative option to the variational quantum eigensolver for ground state energy estimation. That said, many of the

techniques that have, and will be developed for the variational quantum eigensolver, will likely be employed for preparing the initial state in the phase estimation algorithm.

The use of quantum error correction also stands to improve the performance of variational quantum algorithms such as VQE. For instance, error-correction could lead to an improvement in the preparation of good ansatz states for the variational quantum eigensolver, as described in Section 5.1.1. However, a possible future scenario is that, by the time useful, large-scale quantum error correction is achieved, the dichotomy between error-corrected algorithms and non-error-corrected algorithms will have become blurred. Already, several methods have been introduced which aim to bridge the gap between these two classes of quantum algorithms³⁸²⁻³⁸⁴.

The landscape of quantum algorithm development has recently been changing quite dramatically. With the arrival of usable quantum devices^{385,386}, we have now entered a *prototype era* in the field of quantum algorithms. Use of these devices has and will continue to spur the development of many new quantum algorithms at a substantially increased rate. In particular, the ability to test performance on real quantum devices facilitates the development of heuristic quantum algorithms^{52,200,320,326,331,344,382}. In the past few years, several papers have proposed new quantum algorithms along with experimental demonstrations on quantum devices^{387,388}. This points to the increasing value in having quantum algorithm developers working closely with quantum machines. Fortunately, this synergy has been amplified by the development of software platforms such as Rigetti’s pyQuil³⁸⁹, IBM’s QISKit³⁹⁰, ETH Zürich’s ProjectQ³⁹¹, Google’s Cirq³⁹², and Xanadu’s Strawberry Fields³⁹³.

Any reasonable forecast on the timing of the quantum utility crossover point is given with a wide margin of error. The more conservative estimates predict that this crossover will occur in 15 to 20 years³⁹⁴. However, the most optimistic estimates predict that quantum computers will solve useful problems in 2 to 5 years³⁹⁵. The estimated dates for reaching this horizon are strongly determined by progress in hardware development and are extrapolated based on recent advances in quantum device technology^{56,385,396,397}. However, quantum algorithm

development also serves to influence these estimates. Quantum simulation for quantum chemistry provides a rather striking example. Over the last five years, as outlined in Section 4.2.2, the asymptotic scaling of quantum simulation algorithms for quantum chemistry have been dramatically improved from a high-degree polynomial to sublinear in the number of orbitals. Furthermore, continued development in quantum error correction^{237,398} and error-mitigation techniques^{200,317} will also improve prospects. So, while progress in quantum hardware development carries us toward the utility horizon, progress in quantum algorithm development moves this landmark itself closer into our reach.

With this perspective, then, there is ample opportunity for quantum chemists and quantum theorists to make valuable algorithmic contributions toward the quest for useful quantum computation. So far, many of the quantum approaches take inspiration from the standard classical techniques. For example, the variational quantum eigensolver algorithm can be viewed as a quantum version of Ritz’s variational method. With novel computational means, there is an opportunity for developing more quantum chemistry methods which truly have no classical analogs. Quantum computing for quantum chemistry is likely to develop into a rich subfield of quantum chemistry. Now is an opportune time to enter into this emerging research field. Through cross-disciplinary engagements, early practitioners of these novel computational tools will usher in a renaissance for computational methods in chemistry.

Acknowledgements

A.A.-G acknowledges support from the Army Research Office under Award No. W911NF-15-1-0256 and the Vannevar Bush Faculty Fellowship program sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering (Award number ONR 00014-16-1-2008). A.A.-G. also acknowledges generous support from Anders G. Froseth and from the Canada 150 Research Chair Program. L.V. acknowledges support by the Czech Science Foundation (Grant No. 18-18940Y). S.S. is supported by the DOE Computational

Science Graduate Fellowship under grant number DE-FG02-97ER25308. TM was supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via U.S. Army Research Office Contract No. W911NF-17-C-0050. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

A Quantum chemistry basis sets

Wave function ansatzes for the electronic structure problem where the total wave function is approximated as a product of one-electron functions, or methods based in second quantization techniques where the system is represented in the Fock basis, require the introduction of a set of system-dependent one-particle functions $\phi(x_i)$, where x_i represents the spatial and spin coordinates of a single electron. In practice, the set of one-particle functions is finite, introducing an approximation error associated with a truncated representation of the Hilbert space, or basis-set incompleteness. Furthermore, these one-electron functions, also called spin-orbitals, are usually expanded as linear combinations of a set of standard system-independent functions, which are called the *basis set functions*. Choosing basis set functions with the right mathematical properties can significantly facilitate the evaluation of the Hamiltonian elements and expectation values.

Ideally, basis set functions should be designed to capture the physics of the problem, such that a good representation can be achieved using as few functions as possible. Correspondingly, the performance of the basis set is usually measured as the difference with the undiscretized problem, which would correspond with a hypothetical infinite basis set. A good basis set should allow for a systematic and quickly converging extrapolation to this

basis set limit. Furthermore, basis sets should have a mathematical form that facilitates the evaluation of molecular integrals and should be able to accurately describe not only total energies but also other properties. The existing basis sets in computational chemistry offer different levels of compromise among these desired qualities.

The first and most important aspect to consider in the choice or design of a basis set is the nature of the problem. Molecular systems can be roughly categorized into two types: periodic systems and isolated molecules. The first case describes extended systems modeled as unit cells with periodic boundary conditions, such as crystals. In periodic systems, the periodic density does not vanish exponentially far away from the nuclei, unlike the case of isolated molecules. Molecular crystals and large biomolecules can be considered intermediate cases between these two extremes. Traditionally, periodic systems have been described using plane-wave basis sets whereas the basis sets for isolated molecules have been dominated by basis sets with atom-centered functions.

In isolated molecules, the electronic density is highly peaked around the nuclei and vanishes exponentially away from them, which justifies the chemical view of molecules as atoms interacting mostly through their outermost regions. This intuition applied to basis sets inspired the linear combination of atomic orbitals (LCAO) method, which conceives molecular spin-orbitals as linear combinations of a fixed set of atomic basis set functions, also called atomic orbitals (AOs). This framework has the advantage of providing a systematic way of constructing basis sets for molecules by keeping a dataset of AOs for most elements in the periodic table.

The functional form of atomic basis sets is inspired by the solutions of the Schrödinger equation for hydrogen-like atoms (one electron and one nucleus), which have the following general form:

$$\Phi(\vec{r}) = R_{nl}(r)Y_{l,m}(\theta, \phi) \quad (88)$$

where n is the natural quantum number, l and m are the angular momentum and magnetic quantum numbers, $Y_{lm}(\theta, \phi)$ is a spherical harmonic function, $R_{nl}(r)$ is a product of Laguerre

polynomials and a term decaying exponentially with r , and (r, θ, ϕ) are spherical coordinates.

Slater-type orbitals (STOs) have the same structure of the orbitals of hydrogenic atoms, with the radial function taking the form:

$$R_n^{STO}(r) = Nr^{n-1}e^{-\zeta r} \quad (89)$$

where N is a normalization constant and ζ is called the orbital exponent. The exponent controls the degree of “diffuseness” of the orbital, meaning how fast the density vanishes as a function of the nuclear distance. Consequently, within a basis set comprised of several STOs, the minimum and maximum orbital exponents determines how far and how close the resulting wave function can be represented.

AO basis sets are then constructed by choosing a particular combination of individual basis set functions for each of the orbitals in a given atom. For example, one could construct a basis set for any element in the second period of the periodic table by using five STOs with different exponents and appropriate spherical harmonics to represent the $1s, 2s, 2p_x, 2p_y, 2p_z$ orbitals of the electronic configuration of these elements. One could further improve the flexibility of the basis set by expressing each orbital as a linear combination of n STOs with different exponents, instead of using a single one. This strategy is regarded as the n-zeta representation.

The main advantage of STOs is that they capture the appropriate behavior of electron density near and far from the nuclei, with appropriate cusps for s orbitals at the nuclei and exponential decaying tails. Unfortunately, the calculation of the molecular integrals using STOs has to be carried out numerically because of the lack of analytical solutions, restricting the use of STOs to small molecules. This difficulty motivated the adoption of Gaussian type orbitals (GTOs) as basis sets. Spherical GTOs have the same angular form of STOs, but differ in the form of the radial function which adopts a gaussian form instead of the original

exponential:

$$R_n^{GTO}(r) = Nr^{n-1}e^{-\zeta r^2}. \quad (90)$$

Furthermore GTOs can be also cartesian, where the spherical harmonic factor is replaced by a cartesian product of the form $x^i y^j z^k$, where x , y and z are the components of the radial distance between the electron and the center of the gaussian and the sum $i + j + k$ determines the angular momentum. GTOs have convenient analytical properties that allow for efficient schemes for the computational evaluation of molecular integrals. However, unlike STOs, they do not describe the cusp and exponential tails of the electronic density correctly. A middle ground between accuracy in the representation and computational ease is the contracted GTOs scheme, where a linear combination of GTOs, also referred as *primitives* is used to emulate a single STO. The set of primitive STOs is called a *contraction*. Contracted GTOs have been adopted as the main basis set functions for electronic structure calculations in isolated molecules³⁹⁹, with different families of basis sets created to fit different purposes.

Although different families of Gaussian basis sets differ in the specific parameters and optimization strategy employed, they share similar structures and nomenclatures. Most families are created by augmenting the number of basis sets using the N-Zeta strategy. A minimal basis set corresponds to a 1-zeta, denoted as SZ (single zeta), comprising a number of orbitals corresponding to the orbitals in the electronic configuration of the period of the corresponding element. For example, elements in the second row of the periodic table would have five orbitals corresponding to $(1s, 2s, 2p_x, 2p_y, 2p_z)$ for an SZ basis set. The flexibility of the basis set can be increased by multiplying the number of contractions employed to represent valence orbitals, which play a more prominent role in the molecule energetics. For example, a DZ (double zeta) basis set would have twice the number of valence orbitals of a SZ basis set $((1s, 2s, 2s', 2p_x, 2p'_x, 2p_y, 2p'_y, 2p_z, 2p'_z))$. The number of orbitals can be also augmented by adding polarization functions. Correspondingly, a DZP (double zeta plus

polarization) would include orbitals of angular momenta 2 for elements in the second row of the periodic table. In addition to polarization functions, diffuse functions (GTOs with small exponents) can be included to improve the description of tails of the density, especially in system such as anions. Similarly, tight functions (GTOs with large orbitals exponents) can be included when the description of core electrons plays an important role in the accuracy.

Apart from the strategies for augmentation of the basis set size, there are two different contraction schemes: segmented and general contracted basis sets. In segmented contractions, different sets of primitive GTOs are used to represent different orbitals. An example of segmented contracted basis is the 3-21G basis that belongs to the split-valence family (also called Pople family)³⁰⁰. In a 3-21G basis, core electrons are described with a single contraction of 3 GTOs while valence electrons are made of contractions of 2 and 1 GTOs. In contrast, the general contractions approach employs the same set of primitive GTOs expands all the orbitals, with only the combination coefficients differing between different orbitals. Most basis sets in use share elements of both segmented and general contraction. Some Gaussian basis sets have been also developed to improve the calculation of magnetic and electronic properties and to include relativistic effects. For a more detailed description of atomic orbital basis sets we refer the reader to references^{400,401}.

Atomic basis sets are well suited to describe electron densities in isolated molecules, where the densities resemble those of individual atoms. However, energy bands in periodic systems are different from atomic orbitals. For example, in metals, valence electrons have a behavior more similar to free electrons. Correspondingly, solutions to the particle-in-a-box problem can offer a better description of periodic systems, giving rise to plane-waves (PW) basis sets, which have a complex exponential form. In a three-dimension cell, a PW can be expressed as:

$$\phi_\nu(r) = \sqrt{\frac{1}{V}} \exp\left(\frac{2\pi\nu r}{L}\right) \quad (91)$$

for a wave with wavevector corresponding to the ν -th harmonic of a box with length L and volume V . The size of the PW basis set is determined by the chosen value of the maximum energy and the volume V of the unit cell, and in contrast with atomic basis sets, does not depend on the number of atoms within the cell. Correspondingly, the description of core electrons might require a large number of PW basis sets due the large energies, and thus, PW are usually employed in combination with pseudopotentials to describe core electrons. Compared to GTOs, PW basis sets generally require an order of magnitude more functions to achieve a similar accuracy.

Recently, the dual plane-wave basis set (dual PW), corresponding to the Fourier transform of PW, has been proposed as an alternative for electronic structure calculations on quantum computers³⁰¹. PWs and their dual diagonalize the kinetic and potential operators, respectively. The dual PW basis set has the advantage of providing a more compact representation of the molecular Hamiltonian in second quantization. These functions have a form that resembles the Sinc function, which is expressed as:

$$\text{Sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{92}$$

The Sinc function is oscillatory in nature just like a plane wave, but its amplitude decays to zero away from the expansion point x . The dual PW and Sinc functions can be used as basis sets by placing functions at a number of real-space grid points. The quality of the basis set can be improved by decreasing the separation in the grid.

B Mappings to qubits

In this section we introduce the two main approaches to mapping second-quantized quantum chemistry Hamiltonians to qubit Hamiltonians. We define the mappings and discuss their benefits, drawbacks, and variants. For a more-detailed introduction to these various mappings, refer to McArdle et al.³⁵⁸.

B.1 Jordan-Wigner mapping

In the Jordan-Wigner mapping, each electronic orbital, described by creation operator a_j^\dagger is associated to a qubit j via

$$a_j^\dagger \leftrightarrow (\sigma_z)_1 \otimes \dots \otimes (\sigma_z)_{j-1} \otimes (\sigma^+)_j \otimes \mathbb{I}, \quad (93)$$

where $\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$ and $\sigma^+ = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$. This induces a correspondence between each second-quantized basis state $(a_K^\dagger)^{i_N} \dots (a_1^\dagger)^{i_1} |vac\rangle$ and a corresponding computational basis state $|i_1\rangle \otimes |i_2\rangle \otimes \dots \otimes |i_N\rangle$, where each i_1, \dots, i_N is 0 or 1. Notice that the creation operators a_i^\dagger act non-trivially on i qubits. This is a consequence of the creation operator carrying out two actions: the action of changing the occupation is carried out locally by σ_i^+ , while the action of applying a phase according to the *parity* (even or odd) of the occupations for orbital labels less than i is achieved by a string of σ_z (see Equation 93). The high weight of the parity-counting phase can be costly for certain quantum simulation algorithms. So, while each spin-orbital occupation operator $a_i^\dagger a_i$ is local, acting only on the i th qubit, other one-body operators $a_i^\dagger a_j$ can be very non-local due to the high weight of the parity.

A dual version of the Jordan-Wigner mapping is the so-called parity mapping¹⁸¹. Here, the parity operators are low-weight, while the occupation operators become high-weight. In the parity mapping, the creation and annihilation operators are

$$a_j^\dagger \leftrightarrow \mathbb{I} \otimes (\sigma_z)_{j-1} \otimes (\sigma^+)_j \otimes (\sigma_x)_{j+1} \otimes \dots \otimes (\sigma_x)_N \quad (94)$$

This definition induces a transformation from Jordan-Wigner product states to parity product states,

$$|i_1\rangle \otimes \dots \otimes |i_N\rangle \rightarrow |i_1\rangle \otimes |i_1 \oplus i_2\rangle \otimes \dots \otimes |i_1 \oplus \dots \oplus i_N\rangle, \quad (95)$$

where \oplus indicates addition modulo-two. While the Jordan-Wigner mapping stores the occupation of each spin-orbital in each qubit, the parity mapping stores the parity in each

qubit.

While the parity mapping is a valid transformation from fermionic operators to qubits, so far, it has not been considered as a useful contender for use in quantum simulation. Rather, it serves as a useful pedagogical tool for motivating the Bravyi-Kitaev transformation introduced in the following subsection.

B.2 Bravyi-Kitaev mapping

The Bravyi-Kitaev mapping¹⁸⁰ combines the advantages from the Jordan-Wigner and parity mappings to yield creation and annihilation operators which act non-trivially on $O(\log N)$ qubits (i.e. have weight $O(\log N)$). This is important for quantum simulation, as larger-weight terms in the Hamiltonian require longer circuits for their simulation. The original Bravyi-Kitaev encoding is defined for the case that $N = 2^n$. In this case the maximum weight of any encoded a_j^\dagger is exactly $\log_2 N$ according to Havlíček et al.⁴⁰². The creation and annihilation operators in the Bravyi-Kitaev mapping have a more involved description than in the Jordan-Wigner or parity mapping. For a full description of these operators, see Seeley et al.¹⁸¹. We limit ourselves to presenting the transformation from Jordan-Wigner product states to Bravyi-Kitaev product states:

$$|i_1\rangle \otimes \dots \otimes |i_N\rangle \rightarrow |b_1\rangle \otimes \dots \otimes |b_N\rangle, \quad (96)$$

where $b_k = \sum_{l=1}^k [\beta_r]_{kl} i_l \pmod 2$ and the matrix β_r is defined recursively as

$$\beta_1 = 1, \tag{97}$$

$$\beta_{r+1} = \begin{bmatrix} & 1 & \dots & 1 \\ \beta_r & 0 & \dots & 0 \\ & \vdots & & \vdots \\ & 0 & \dots & 0 \\ \mathbf{0} & & & \beta_r \end{bmatrix}. \tag{98}$$

As an example, in the case of $N = 2^3$, the matrix defining the Bravyi-Kitaev transformation is

$$\beta_3 = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{99}$$

The Bravyi-Kitaev transformation, as presented here, applies only for systems with a spin-orbital number equal to a power of two. A closely related variant of the Bravyi-Kitaev transformation, known as the Bravyi-Kitaev tree method, also achieves a mapping with $O(\log N)$ -weight creation and annihilation operators⁴⁰². The algorithm for generating this mapping uses the data structure known as Fenwick trees, which were originally used for arithmetic coding compression algorithms⁴⁰³. Although the scaling of operator weight is the same as that of the standard Bravyi-Kitaev mapping, in practice the Bravyi-Kitaev tree method produces higher-weight creation and annihilation operators³⁵⁸. However, it has been noted³⁵⁸ that, in contrast to the standard Bravyi-Kitaev mapping, the Bravyi-Kitaev tree

method enables the use of qubit-reduction techniques when the spin-orbital number is not a power of two.

C From quantum chemistry to quantum computation: example of molecular hydrogen

C.1 Introduction

The purpose of this appendix is to present a detailed description of the workflow underlying a simple quantum computation (VQE calculation) of a quantum chemistry problem. We will focus on a concrete problem and progress step-by-step through the various assumptions, simplifications, and calculations taken to convert the chemistry problem into results on a quantum computer. This section is intended to speak to both quantum chemists and quantum information scientists and is structured as follows:

- (a) define the chemistry problem,
- (b) map the problem onto the quantum computer,
- (c) provide a brief introduction to (circuit-model) quantum computation,
- (d) apply the variational quantum eigensolver algorithm to treat the problem.

C.2 Defining the chemistry problem

We address the problem of determining the electronic ground state energy of molecular hydrogen as a function of the distance between the nuclei. This relationship is otherwise known as the ground state energy dissociation curve (Figure 10). An accurate description of this energy surface is a key challenge in quantum chemistry that can provide insight on a range of chemical phenomena, e.g. bond breaking and reaction dynamics. To simplify the

problem, we apply the Born-Oppenheimer approximation, in which we treat the nuclei as stationary classical particles. This is justified as the ratio of electronic to nuclear mass is roughly 1:1000, leading to a separation in the timescale of their dynamics. The resulting quantum Hamiltonian describing the electronic system, in atomic units, can be written as

$$H_{elec} = - \sum_i \frac{\nabla_{\mathbf{r}_i}^2}{2} - \sum_{i,j} \frac{Z_i}{|\mathbf{R}_i - \mathbf{r}_j|} + \sum_{i,j>i} \frac{Z_i Z_j}{|\mathbf{R}_i - \mathbf{R}_j|} + \sum_{i,j>i} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}, \quad (100)$$

where \mathbf{r}_i are the position coordinates of the electrons, which parametrically depend on the fixed position coordinates of the nuclei \mathbf{R}_i , and Z_i and M_i denoting the nuclear charges and masses, respectively. The electronic system, as written in the first-quantized picture in Equation 100, assumes an infinite-dimensional Hilbert space. When applying quantum computation for chemistry, the Hamiltonian is instead often considered in the second-quantized formulation, in which the system can be described approximately using a finite basis. For our example of molecular hydrogen, we consider the minimal basis (STO-6G). For a deeper review of basis sets, the reader should refer to Appendix A. Within this framework, states are labeled by the occupation of the orbitals, and the exchange symmetry of the particles is naturally considered through the use of fermionic creation and annihilation operators. The electronic Hamiltonian can then be expressed in terms of these second-quantized operators as

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s, \quad (101)$$

where a_p is an electron annihilation operator that removes an electron from an orbital with label p . The weights of the operators are given by the molecular integrals

$$h_{pq} = \int dx \phi_p^*(x) \left(\frac{\nabla_r^2}{2} - \sum_i \frac{Z_i}{|\mathbf{R}_i - \mathbf{r}|} \right) \phi_q(x) \quad (102)$$

and

$$h_{pqrs} = \int dx_1 dx_2 \frac{\phi_p^*(x_1)\phi_q^*(x_2)\phi_s(x_1)\phi_r(x_2)}{|\mathbf{r}_1 - \mathbf{r}_2|}, \quad (103)$$

where we use x_i to denote the spatial and spin coordinates, i.e. $x_i = (\mathbf{r}_i, \sigma_i)$. In practice, several electronic structure packages and codes have been developed and optimized for computing these integrals. To prepare our quantum computation for molecular hydrogen, we use such classically pre-computed integrals to prepare the second-quantized Hamiltonian.




Figure 10: Dissociation profile of molecular hydrogen in the minimal basis (STO-6G). The ground state energies computed using the full configuration interaction (FCI) method are shown. The ideal goal in applying the VQE algorithm for chemistry is to compute energies that approximately reproduce the FCI results.

C.3 Mapping the problem

In classical computation, the chemistry problem is treated by solving the necessary equations using a method that implements some level of approximations. In quantum computation, an

extra “translation” step is necessary to encode the second-quantized Hamiltonian in a system of quantum bits, or qubits, prior to treating the problem. This is achieved by applying a mapping from the set of creation and annihilation operators in the Hamiltonian to the set of operators on qubits, in which the mapping preserves the algebraic relationships, i.e the fermionic canonical commutation relations. We note that this step is a general prerequisite for implementing Hamiltonian simulations, the quantum phase estimation algorithm, or the variational quantum eigensolver. In Appendix B, we review the two most well-known mappings, the Jordan Wigner and the Bravyi-Kitaev transformations.

For our molecular hydrogen example, we employ the Bravyi-Kitaev transformation, mapping the minimal-basis second-quantized Hamiltonian shown in Equation 101 to a four-qubit Hamiltonian

$$H_{elec} = \mu_0\mathbb{I} + \mu_1Z_1 + \mu_2Z_2 + \mu_3Z_3 + \mu_4Z_1Z_2 \quad (104)$$

$$+ \mu_5Z_1Z_3 + \mu_6Z_2Z_4 + \mu_7X_1Z_2X_3 + \mu_8Y_1Z_2Y_3 \quad (105)$$

$$+ \mu_9Z_1Z_2Z_3 + \mu_{10}Z_1Z_3Z_4 + \mu_{11}Z_2Z_3Z_4 \quad (106)$$

$$+ \mu_{12}X_1Z_2X_3Z_4 + \mu_{13}Y_1Z_2Y_3Z_4 + \mu_{14}Z_1Z_2Z_3Z_4. \quad (107)$$

Here, X_j , Y_j , and Z_j are Pauli operators acting on the j th qubit, such as $X_1 = \sigma_x \otimes \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I}$ and the μ_k are determined by the integrals from Equations 102 and 103. As described in O’Malley et al.⁵⁵, symmetries in this Hamiltonian can be exploited to construct a two-qubit Hamiltonian representing the symmetry sector of the original Hamiltonian that contains the ground state. The resulting two-qubit Hamiltonian is

$$\tilde{H}_{elec} = \nu_0\mathbb{I} + \nu_1Z_1 + \nu_2Z_2 + \nu_3Z_1Z_2 + \nu_4X_1X_2 + \nu_5Y_1Y_2, \quad (108)$$

where the ν_j are linear combinations of the μ_k . This is the final form of the Hamiltonian we will use in our quantum computation to determine the ground state energy as a function of interatomic spacing.

C.4 A brief introduction to quantum computation

After encoding the computational problem or task onto the quantum computer, we can approach the said task (or parts of the overall task) on the quantum computer by manipulating its qubits with relevant quantum operations. Before we describe the application of the variational quantum eigensolver (VQE), an algorithm that involves both classical and quantum computations, we briefly introduce the workings of a quantum computation. The following paragraphs describe the widely used “circuit-model” of quantum computation. However, as described in the main text, other models such as adiabatic quantum computation exist.

Just as bits are the elementary units of a (classical) computation, quantum bits or qubits are the elementary units of a quantum computation. Qubits are controllable two-level quantum systems. Analogous to logic gates, which comprise, *en masse*, a computation, quantum gates are simple actions or operations performed on qubits, which, in sequence, comprise a quantum computation. Specifically, a quantum gate is a unitary transformation which, typically, manipulates just a few qubits at a time. As an example, a common single-qubit gate is a “Z-gate”, or “phase-gate”, which maps the quantum state $|0\rangle$ to itself, and the state $|1\rangle$ to $-|1\rangle$. While this phase-gate only changes the (unobservable) phase of these two states, it affects a non-trivial action on superpositions of these states. For example, the quantum state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ is transformed to the orthogonal state vector $Z(|0\rangle + |1\rangle)/\sqrt{2} = (|0\rangle - |1\rangle)/\sqrt{2} = |-\rangle$. This transformation may be physically realized by subjecting the two level system to a driving term that is diagonal in the qubit basis, i.e. proportional to σ_z . A qubit flip between the 0 and 1 state, in turn, may be realized by a drive along σ_x . Similarly, a two-qubit gate may be implemented by inducing a direct or indirect coupling between two qubits. A two-qubit gate is fully characterized by how it transforms the four two-qubit basis states $|00\rangle, |01\rangle, |10\rangle$, and $|11\rangle$.

In its simplest form, a quantum computation implements the following three steps, as illustrated using a circuit diagram in Fig. 11:

- 1.) initialize the qubits in the state $|0\rangle \otimes \dots \otimes |0\rangle$,

- 2.) apply a sequence of quantum gates $U |0 \dots 0\rangle$, and
- 3.) measure to obtain either 0 or 1 for each qubit.

These steps can be combined and manipulated at a higher level to achieve more sophisticated tasks. Namely, the variational quantum eigensolver (VQE) algorithm complements “quantum” routines comprising of executions of quantum circuits with classical routines to estimate the ground state energy of molecular systems. In the following subsection, we will describe the VQE algorithm in the context of finding the ground state energy of molecular hydrogen.




Figure 11: An illustration of a quantum computation using the circuit model. Various quantum and/or quantum-classical algorithms leverage the capabilities of quantum circuits to achieve computational tasks or subtasks.

C.5 Variational quantum eigensolver for quantum chemistry

As detailed in Section 5.1, the variational quantum eigensolver (VQE) is a hybrid quantum-classical algorithm that estimates molecular properties, often the ground state energies, of

quantum systems using the variational principle. Consequently, a promising application of VQE is quantum chemistry. At a high level, VQE allocates subtasks between quantum and classical resources based on the inherent strengths and capabilities of each device. In this framework, the role of the quantum computer is to *prepare* the parametrized trial quantum state $|\psi(\vec{\theta})\rangle$ (also known as the *ansatz*) and *estimate* the energy with respect to the Hamiltonian. The ansatz is constructed by applying a *variational circuit*, that is a parametrized quantum circuit $U(\vec{\theta})$ with classical parameters $\vec{\theta}$ to an initial or reference state $|\phi_0\rangle$. The role of the classical processor is then to orchestrate the minimization of the energy expectation through feedback to the parameters $\vec{\theta}$. Procedurally, the VQE algorithm can be summarized in the following steps:

- 1.) prepare the parametrized trial quantum state $|\psi(\vec{\theta})\rangle = U(\vec{\theta})|\phi_0\rangle$ on the quantum computer,²
- 2.) estimate the expectation value of energy $\langle\psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle$ using measurements of terms in the Hamiltonian,
- 3.) update the parameter(s) $\vec{\theta}$ of the quantum state using a classical optimization routine,
- 4.) repeat the previous steps until convergence criteria (e.g. in energy and/or iteration number) are satisfied.

Often the challenge in VQE is the choice and/or design of the ansatz, which largely influences the performance of the algorithm³²¹. This has motivated numerous studies and designs of ansätze, several of which are reviewed in Section 5.1.1. For our case of simulating molecular hydrogen, we selected an ansatz based on the unitary coupled cluster (UCC) method, as shown in Figure 12. To construct the ansatz, the Hartree-Fock reference state (i.e. $|01\rangle$) is first prepared, followed by quantum operations corresponding to the application of the UCC operators. For more detail on the UCC method, the reader should refer to

²Normalization is assumed.

Section 5.1.1. We note that for the H_2 example, this level of theory is equivalent with the exact solution.

Once the ansatz is selected, its variational circuit implementation is executed on the quantum computer to compute the objective function value, which, in the case of VQE, is the energy expectation. We note that initialization of the variational circuit parameters should ideally be informative. For instance, in the case of unitary coupled-cluster ansatz, the classically computed MP2 amplitudes can be used to initialize the VQE parameters (i.e. UCC amplitudes). The energy expectation can then be estimated using the *Hamiltonian averaging* procedure. Given that the Hamiltonian is written as a sum of Pauli terms acting on subsets of qubits, we can compute the energy expectation by averaging over the expectation values of the individual Pauli terms, as shown below:

$$\langle H \rangle = \sum_i h_i \langle O_i \rangle, \quad (109)$$

where O_i is a Pauli term, a tensor product of Pauli operators (i.e. X , Y , Z , or I) acting on some subset of qubits, and h_i is the corresponding weight.

In the case of molecular hydrogen, the energy expectation expression becomes

$$\langle H \rangle = \nu_0 \mathbb{I} + \nu_1 \langle Z_1 \rangle + \nu_2 \langle Z_2 \rangle + \nu_3 \langle Z_1 Z_2 \rangle + \nu_4 \langle X_1 X_2 \rangle + \nu_5 \langle Y_1 Y_2 \rangle. \quad (110)$$

We note that when measuring each Pauli expectation, post-rotations may need to be applied to make measurements in the Z basis. These measurements are then collected and processed to approximate the total energy. In practice, we can only obtain a finite number of measurements, leading to errors in the energy estimation. For a deeper analysis of the sampling error, the reader should refer to Section 5.1.2 or McClean et al.²¹⁶. While VQE is a near-term alternative to the quantum phase estimation algorithm due to its low coherence time requirements, the trade-off or cost of the algorithm is the large number of measurements needed to approximate the ground state energy with high precision.




Figure 12: The quantum circuit used in VQE to estimate the ground state energy for molecular hydrogen in the minimal basis. After preparing the parametrized quantum state, using the Hartree-Fock reference state followed by the application of the unitary coupled cluster-inspired variational circuit. To estimate the energy expectation, necessary post-rotations ($R_t \in \{R_X(-\pi/2), R_Y(\pi/2), I\}$) are applied before measuring the qubits in the Z basis.

After computing the energy expectation with respect to some values assigned to the parameters, VQE employs a classical optimization routine to update the parameters to ideally reach a quantum state that better approximates the ground state. Section 5.1.3 reviews a number of optimization routines used and benchmarked for VQE in previous studies. Provided that the ansatz can well describe the ground state and the classical optimizer is robust against noise in the cost function landscape, VQE can provide a high-quality estimation for the ground state energy. This is observed for our small example of molecular hydrogen, shown in Figure 13a, in which a simulation of the VQE algorithm was able to compute ground state energies along the energy surface that were numerically equal to the corresponding FCI energy values. Note that we also show the sampling error at a particular geometry in Figure 13b.

Since first presented in 2014, VQE has been widely studied and improved from both theoretical and experimental standpoints. In particular, VQE was experimentally implemented for molecular systems beyond hydrogen as highlighted in Section 5.1, demonstrating the utility and potential of the algorithm for applications in quantum chemistry even on early quantum computers.




Figure 13: VQE simulation results for molecular hydrogen in the minimal basis (STO-6G). a) Dissociation profile computed using the VQE algorithm. At each bond length, the energy computed using VQE is numerically equal to that of the FCI method. The L-BFGS-B method was used for parameter optimization. b) Energy expectation plotted over a range of parameter values for bond length of 1.2 Angstroms. A parameter scan using a finite number of samples is overlaid with that generated using the wave function simulator. These simulations were implemented using OpenFermion⁴⁰⁴ and Forest³⁸⁹.

C.6 Appendix Glossary

l	angular momentum quantum number
m	magnetic quantum number
$R_{lm}(r)$	a product of Laguerre polynomials and a term decaying exponentially with r
H_{elec}	Electronic Hamiltonian
\mathbf{r}_i	Position coordinates of electrons
\mathbf{R}_i	Fixed position coordinates of nuclei
Ψ_i	Basis wave functions
σ_i	Spin degree of freedom for electrons
$\phi_{i,j}$	Molecular orbitals
$Y_{l,m}$	Spherical harmonic with orbital angular momentum quantum numbers l, m
a_i, a_i^\dagger	Fermionic annihilation and creation operator on orbital ϕ_i
$ \text{vac}\rangle$	Fermionic vacuum state

Z_i	(Eqn. 106-108) i -th nuclear charge
σ_z	Pauli Z matrix
σ^+	$(\sigma_x + i\sigma_y)/2$
n	A positive integer
\mathbb{I}	Identity operator
X_j, Y_j, Z_j	(Eqn. 110-116) Pauli X, Y, and Z operator acting on the j -th qubit
μ_i	Coefficients of electronic Hamiltonian terms
\tilde{H}_{elec}	2-qubit electronic Hamiltonian (symmetry sector of the original Hamiltonian)
ν_i	Coefficients of 2-qubit electronic Hamiltonian terms; linear combinations of μ_i
U	A unitary operator
$\Psi'(\vec{t})$	Ansatz state
U	Unitary operation
$ \psi(\vec{t})\rangle$	Ansatz state
$U(\vec{t})$	Variational quantum circuit
$ \phi_0\rangle$	Initial/reference quantum state
O_i	A tensor product of Pauli operators
h_i	Weight of Pauli operator corresponding to O_i
I	Identity operator
VQE	Variational quantum eigensolver
HF	Hartree-Fock
UCC	Unitary coupled cluster
SPSA	Simultaneous perturbation stochastic approximation
PSO	Particle swarm optimization
STO	Slater-type orbitals
GTO	Gaussian type orbitals
PW	Plane-wave

Table 5: Representative experimental demonstrations of the VQE algorithm using various quantum computer architectures. Here SPSA and PSO stand for simultaneous perturbation stochastic approximation and particle swarm optimization, respectively.

Architecture/ Platform	System- of-interest	Number of physical qubits	Ansatz	Optimization Routine/ Strategy	Computed properties	Reference
Photonic chip	HeH ⁺	2	Hardware-specific parametrized ansatz	Nelder-Mead	Ground state energy	52
Single trapped ion	HeH ⁺		UCC	Nelder-Mead	Ground and excited state energies	53
Superconducting processor (transmon qubits)	H ₂	2	UCC	Grid scan and locally optimize	Ground state energy	55
Superconducting processor (transmon qubits)	H ₂	2	“Hardware-efficient” ansatz	SPSA	Ground state energy	56
	LiH	4	“Hardware-efficient” ansatz	SPSA	Ground state energy	56
	BeH ₂	6	“Hardware-efficient” ansatz	SPSA	Ground state energy	56
Ion trap processor (Ca ⁺ ions)	H ₂	2	UCC	Grid scan and locally optimize	Ground state energy	59
	LiH	3	Approximate UCC	Grid scan and locally optimize	Ground state energy	59
Superconducting processor (transmon qubits)	H ₂	2	Hardware-specific parametrized ansatz	PSO	Ground and excited state energies	58
Silicon photonic chip	Two chlorophyll units in 18-mer ring of LHII complex	2	“Parametrized Hamiltonian” ansatz with truncation scheme	PSO	Ground and excited state energies	316
Superconducting processor (transmon qubits) via Cloud	Deuteron	2-3	UCC	Grid scan	Ground state energy	330

References

- (1) Nielsen, M. A.; Chuang, I. L. *Quantum Computation and Quantum Information*; Cambridge University Press, 2010.
- (2) Divincenzo, D. P. *Scalable Quantum Computers*; Wiley-VCH Verlag GmbH & Co. KGaA, 2005; pp 1–13.
- (3) Brown, K. R.; Kim, J.; Monroe, C. Co-designing a scalable quantum computer with trapped atomic ions. *npj Quantum Inf.* **2016**, *2*.
- (4) Monroe, C.; Kim, J. Scaling the Ion Trap Quantum Processor. *Science* **2013**, *339*, 1164–1169.
- (5) Blatt, R.; Wineland, D. Entangled states of trapped atomic ions. *Nature* **2008**, *453*, 1008–1015.
- (6) Wineland, D.; Monroe, C.; Itano, W.; King, B.; Leibfried, D.; Meekhof, D.; Myatt, C.; Wood, C. *Quantum Computing*; Wiley-VCH Verlag GmbH & Co. KGaA, 2004; pp 57–84.
- (7) Steane, A. The ion trap quantum information processor. *Appl. Phys. B: Lasers Opt.* **1997**, *64*, 623–643.
- (8) Devoret, M. H.; Martinis, J. M. *Experimental Aspects of Quantum Computing*; Springer US, 2005; pp 163–203.
- (9) Devoret, M. H.; Wallraff, A.; Martinis, J. M. Superconducting Qubits: A Short Review. **2004**, cond-mat/0411174.
- (10) Wendin, G. Quantum information processing with superconducting circuits: a review. *Rep. Prog. Phys.* **2017**, *80*, 106001.

- (11) Gambetta, J. M.; Chow, J. M.; Steffen, M. Building logical qubits in a superconducting quantum computing system. *npj Quantum Inf.* **2017**, *3*.
- (12) Devoret, M. H.; Schoelkopf, R. J. Superconducting Circuits for Quantum Information: An Outlook. *Science* **2013**, *339*, 1169–1174.
- (13) Kok, P.; Munro, W. J.; Nemoto, K.; Ralph, T. C.; Dowling, J. P.; Milburn, G. J. Linear optical quantum computing with photonic qubits. *Rev. Mod. Phys.* **2007**, *79*, 135–174.
- (14) Adami, C.; Cerf, N. J. In *Quantum Computing and Quantum Communications*; Williams, C. P., Ed.; Springer Berlin Heidelberg: Berlin, Heidelberg, 1999; pp 391–401.
- (15) Krovi, H. Models of optical quantum computing. *Nanophotonics* **2017**, *6*.
- (16) O'Brien, J. L. Optical Quantum Computing. *Science* **2007**, *318*, 1567–1570.
- (17) Doherty, M. W.; Manson, N. B.; Delaney, P.; Jelezko, F.; Wrachtrup, J.; Hollenberg, L. C. The nitrogen-vacancy colour centre in diamond. *Phys. Rep.* **2013**, *528*, 1–45.
- (18) Childress, L.; Hanson, R. Diamond NV centers for quantum computing and quantum networks. *MRS Bulletin* **2013**, *38*, 134–138.
- (19) Gordon, L.; Weber, J. R.; Varley, J. B.; Janotti, A.; Awschalom, D. D.; de Walle, C. G. V. Quantum computing with defects. *MRS Bulletin* **2013**, *38*, 802–807.
- (20) Jones, J. A. In *Quantum Computing and Communications*; Brooks, M., Ed.; Springer London: London, 1999; pp 71–78.
- (21) Oliveira, I. S.; Bonagamba, T. J.; Sarthour, R. S.; Freitas, J. C.; deAzevedo, E. R. In *NMR Quantum Information Processing*; Oliveira, I. S., Bonagamba, T. J.,

- Sarthour, R. S., Freitas, J. C., deAzevedo, E. R., Eds.; Elsevier Science: Amsterdam, 2007; pp 137–181.
- (22) Ramanathan, C.; Boulant, N.; Chen, Z.; Cory, D. G.; Chuang, I.; Steffen, M. In *Experimental Aspects of Quantum Computing*; Everitt, H. O., Ed.; Springer US: Boston, MA, 2005; pp 15–44.
- (23) Nayak, C.; Simon, S. H.; Stern, A.; Freedman, M.; Sarma, S. D. Non-Abelian anyons and topological quantum computation. *Rev. Mod. Phys.* **2008**, *80*, 1083–1159.
- (24) Lahtinen, V.; Pachos, J. A Short Introduction to Topological Quantum Computation. *SciPost Phys.* **2017**, *3*.
- (25) Blatt, R.; Häffner, H.; Roos, C. F.; Becher, C.; Schmidt-Kaler, F. *Experimental Aspects of Quantum Computing*; Springer US, 2015; pp 61–73.
- (26) Monz, T.; Schindler, P.; Barreiro, J. T.; Chwalla, M.; Nigg, D.; Coish, W. A.; Harlander, M.; Hänsel, W.; Hennrich, M.; Blatt, R. 14-Qubit Entanglement: Creation and Coherence. *Phys. Rev. Lett.* **2011**, *106*.
- (27) Mohseni, M.; Read, P.; Neven, H.; Boixo, S.; Denchev, V.; Babbush, R.; Fowler, A.; Smelyanskiy, V.; Martinis, J. Commercialize quantum technologies in five years. *Nature* **2017**, *543*, 171–174.
- (28) Manin, Y. *Computable and Uncomputable*; Sovetskoye Radio: Moscow, 1980.
- (29) Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **1982**, *21*, 467–488.
- (30) Benioff, P. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *J. Stat. Phys.* **1980**, *22*, 563–591.

- (31) Deutsch, D. Quantum Theory the Church-Turing Principle and the Universal Quantum Computer. *Proc. Royal Soc. Lond. A* **1985**, *400*, 97–117.
- (32) Shor, P. W. Polynomial time algorithms for discrete logarithms and factoring on a quantum computer. *Algorithmic Number Theory*. Berlin, Heidelberg, 1994; pp 289–289.
- (33) Bernstein, E.; Vazirani, U. Quantum Complexity Theory. Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing. New York, NY, USA, 1993; pp 11–20.
- (34) Deutsch, D.; Jozsa, R. Rapid Solution of Problems by Quantum Computation. *Proc. Royal Soc. Lond. A* **1992**, *439*, 553–558.
- (35) Grover, L. K. A Fast Quantum Mechanical Algorithm for Database Search. Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing. New York, NY, USA, 1996; pp 212–219.
- (36) Simon, D. On the power of quantum computation. Proceedings 35th Annual Symposium on Foundations of Computer Science. 1994.
- (37) Lloyd, S. Universal Quantum Simulators. *Science* **1996**, *273*, 1073–1078.
- (38) Abrams, D. S.; Lloyd, S. Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors. *Phys. Rev. Lett.* **1999**, *83*, 5162–5165.
- (39) Childs, A. M.; van Dam, W. Quantum algorithms for algebraic problems. *Rev. Mod. Phys.* **2010**, *82*, 1–52.
- (40) Wiesner, S. Simulations of Many-Body Quantum Systems by a Quantum Computer. **1996**, quant-ph/9603028.
- (41) Zalka, C. Efficient Simulation of Quantum Systems by Quantum Computers. *Fortschr. Phys.* **1998**, *46*, 877–879.

- (42) Kitaev, A.; Shen, A.; Vyalyi, M. *Classical and Quantum Computation*; American Mathematical Society: Providence, Rhode Island, 2002.
- (43) Ortiz, G.; Gubernatis, J. E.; Knill, E.; Laflamme, R. Quantum algorithms for fermionic simulations. *Phys. Rev. A* **2001**, *64*.
- (44) Lidar, D. A.; Wang, H. Calculating the thermal rate constant with exponential speedup on a quantum computer. *Phys. Rev. E* **1999**, *59*, 2429–2438.
- (45) Aspuru-Guzik, A.; Dutoi, A.; Love, P. J.; Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **2005**, *309*, 1704–1707.
- (46) Wang, H.; Kais, S.; Aspuru-Guzik, A.; Hoffmann, M. R. Quantum algorithm for obtaining the energy spectrum of molecular systems. *Phys. Chem. Chem. Phys.* **2008**, *10*, 5388.
- (47) Wecker, D.; Bauer, B.; Clark, B. K.; Hastings, M. B.; Troyer, M. Gate-count estimates for performing quantum chemistry on small quantum computers. *Phys. Rev. A* **2014**, *90*.
- (48) Hastings, M. B.; Wecker, D.; Bauer, B.; Troyer, M. Improving Quantum Algorithms for Quantum Chemistry. *Quantum Info. Comput.* **2015**, *15*, 1–21.
- (49) Poulin, D.; Hastings, M. B.; Wecker, D.; Wiebe, N.; Doberty, A. C.; Troyer, M. The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry. *Quantum Info. Comput.* **2015**, *15*, 361–384.
- (50) Babbush, R.; McClean, J.; Wecker, D.; Aspuru-Guzik, A.; Wiebe, N. Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation. *Phys. Rev. A* **2015**, *91*.
- (51) Preskill, J. Quantum Computing in the NISQ era and beyond. **2018**, arXiv:1801.00862 [quant-ph].

- (52) Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'Brien, J. L. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **2014**, *5*.
- (53) Shen, Y.; Zhang, X.; Zhang, S.; Zhang, J.-N.; Yung, M.-H.; Kim, K. Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys. Rev. A* **2017**, *95*.
- (54) Wang, Y.; Dolde, F.; Biamonte, J.; Babbush, R.; Bergholm, V.; Yang, S.; Jakobi, I.; Neumann, P.; Aspuru-Guzik, A.; Whitfield, J. D.; Wrachtrup, J. Quantum Simulation of Helium Hydride Cation in a Solid-State Spin Register. *ACS Nano* **2015**, *9*, 7769–7774.
- (55) O'Malley, P. J. J. et al. Scalable Quantum Simulation of Molecular Energies. *Phys. Rev. X* **2016**, *6*, 031007.
- (56) Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; Gambetta, J. M. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **2017**, *549*, 242–246.
- (57) Colless, J.; Ramasesh, V.; Dahlen, D.; Blok, M.; McClean, J.; Carter, J.; de Jong, W. A.; Siddiqi, I. Implementing a Variational Quantum Eigensolver using Superconducting Qubits. Quantum Information and Measurement (QIM) 2017. 2017; p QF6A.2.
- (58) Colless, J. I.; Ramasesh, V. V.; Dahlen, D.; Blok, M. S.; Kimchi-Schwartz, M. E.; McClean, J. R.; Carter, J.; de Jong, W. A.; Siddiqi, I. Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm. *Phys. Rev. X* **2018**, *8*, 011021.
- (59) Hempel, C.; Maier, C.; Romero, J.; McClean, J.; Monz, T.; Shen, H.; Jurcevic, P.; Lanyon, B. P.; Love, P.; Babbush, R.; Aspuru-Guzik, A.; Blatt, R.; Roos, C. F.

- Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator. *Phys. Rev. X* **2018**, *8*, 031022.
- (60) Farhi, E.; Goldstone, J.; Gutmann, S.; Sipser, M. Quantum Computation by Adiabatic Evolution. **2000**, quant-ph/0001106.
- (61) Albash, T.; Lidar, D. A. Adiabatic quantum computation. *Rev. Mod. Phys.* **2018**, *90*.
- (62) Aharonov, D.; van Dam, W.; Kempe, J.; Landau, Z.; Lloyd, S.; Regev, O. Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation. 45th Annual IEEE Symposium on Foundations of Computer Science. 2004.
- (63) Aaronson, S.; Arkhipov, A. The Computational Complexity of Linear Optics. Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing. New York, New York, 2011; pp 333–342.
- (64) Arkhipov, A.; Aaronson, S. The Computational Complexity of Linear Optics. Quantum Information and Measurement 2014. Messe Berlin, Berlin, 2014.
- (65) Schuch, N.; Verstraete, F. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nat. Phys.* **2009**, *5*, 732–735.
- (66) Liu, Y.-K.; Christandl, M.; Verstraete, F. Quantum Computational Complexity of the N-Representability Problem: QMA Complete. *Phys. Rev. Lett.* **2007**, *98*.
- (67) Brown, B.; Flammia, S. T.; Schuch, N. Computational Difficulty of Computing the Density of States. *Phys. Rev. Lett.* **2011**, *107*.
- (68) Whitfield, J. D.; Love, P. J.; Aspuru-Guzik, A. Computational complexity in electronic structure. *Phys. Chem. Chem. Phys.* **2013**, *15*, 397–411.
- (69) Kohn, W. Nobel Lecture: Electronic structure of matter—wave functions and density functionals. *Rev. Mod. Phys.* **1999**, *71*, 1253–1266.

- (70) Markland, T. E.; Ceriotti, M. Nuclear quantum effects enter the mainstream. *Nat. Rev. Chem.* **2018**, *2*, 0109.
- (71) Bubin, S.; Pavanello, M.; Tung, W.-C.; Sharkey, K.; Adamowicz, L. Born–Oppenheimer and Non-Born–Oppenheimer, Atomic and Molecular Calculations with Explicitly Correlated Gaussians. *Chem. Rev.* **2013**, *113*, 36–79.
- (72) Nakai, H. Nuclear orbital plus molecular orbital theory: Simultaneous determination of nuclear and electronic wave functions without Born–Oppenheimer approximation. *Int. J. Quantum Chem.* **2007**, *107*.
- (73) Ishimoto, T.; Tachikawa, M.; Nagashima, U. Review of multicomponent molecular orbital method for direct treatment of nuclear quantum effect. *Int. J. Quantum Chem.* **2009**, *109*.
- (74) Veis, L.; Višňák, J.; Nishizawa, H.; Nakai, H.; Pittner, J. Quantum chemistry beyond Born–Oppenheimer approximation on a quantum computer: A simulated phase estimation study. *Int. J. Quantum Chem.* *116*, 1328–1336.
- (75) Kassal, I.; Jordan, S. P.; Love, P. J.; Mohseni, M.; Aspuru-Guzik, A. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proc. Natl. Acad. Sci.* **2008**, *105*, 18681–18686.
- (76) Kivlichan, I. D.; Wiebe, N.; Babbush, R.; Aspuru-Guzik, A. Bounding the costs of quantum simulation of many-body physics in real space. *J. Phys. A* **2017**, *50*, 305301.
- (77) Hohenberg, P.; Kohn, W. Inhomogeneous Electron Gas. *Phys. Rev.* **1964**, *136*, B864–B871.
- (78) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133–A1138.

- (79) Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular Electronic-Structure Theory*; John Wiley & Sons Ltd: Hoboken, New Jersey, 2000.
- (80) Fan, P.-D.; Piecuch, P. In *The Usefulness of Exponential Wave Function Expansions Employing One- and Two-Body Cluster Operators in Electronic Structure Theory: The Extended and Generalized Coupled-Cluster Methods*; Sabin, J., Brändas, E., Eds.; Advances in Quantum Chemistry; Academic Press, 2006; Vol. 51; pp 1 – 57.
- (81) Bulik, I. W.; Henderson, T. M.; Scuseria, G. E. Can Single-Reference Coupled Cluster Theory Describe Static Correlation? *J. Chem. Theory Comput.* **2015**, *11*, 3171–3179, PMID: 26575754.
- (82) Kowalski, K.; Piecuch, P. The method of moments of coupled-cluster equations and the renormalized CCSD[T], CCSD(T), CCSD(TQ), and CCSDT(Q) approaches. *J. Chem. Phys.* **2000**, *113*, 18–35.
- (83) Degroote, M.; Henderson, T. M.; Zhao, J.; Dukelsky, J.; Scuseria, G. E. Polynomial similarity transformation theory: A smooth interpolation between coupled cluster doubles and projected BCS applied to the reduced BCS Hamiltonian. *Phys. Rev. B* **2016**, *93*, 125124.
- (84) Pal, S.; Prasad, M. D.; Mukherjee, D. On certain correspondences among various coupled-cluster theories for closed-shell systems. *Pramana* **1982**, *18*, 261–270.
- (85) Cooper, B.; Knowles, P. J. Benchmark studies of variational unitary and extended coupled cluster methods. *J. Chem. Phys.* **2010**, *133*, 234102.
- (86) Voorhis, T. V.; Head-Gordon, M. Benchmark variational coupled cluster doubles results. *J. Chem. Phys.* **2000**, *113*, 8873–8879.
- (87) Evangelista, F. A. Alternative single-reference coupled cluster approaches for multireference problems: The simpler the better. *J. Chem. Phys.* **2011**, *134*, 224102.

- (88) Harsha, G.; Shiozaki, T.; Scuseria, G. E. On the difference between variational and unitary coupled cluster theories. *J. Chem. Phys.* **2018**, *148*, 044107.
- (89) Taube, A. G.; Bartlett, R. J. New perspectives on unitary coupled-cluster theory. *Int. J. Quantum Chem.* **2006**, *106*, 3393–3401.
- (90) Nooijen, M. Can the Eigenstates of a Many-Body Hamiltonian Be Represented Exactly Using a General Two-Body Cluster Expansion? *Phys. Rev. Lett.* **2000**, *84*, 2108–2111.
- (91) Heinrich, S. From Monte Carlo to quantum computation. *Math Comput Simul.* **2003**, *62*, 219–230.
- (92) Sherrill, C. D.; Schaefer, H. F. *Adv. Quantum Chem.*; Elsevier, 1999; pp 143–269.
- (93) Knowles, P. J.; Handy, N. C. Unlimited full configuration interaction calculations. *J. Chem. Phys.* **1989**, *91*, 2396–2398.
- (94) Thøgersen, L.; Olsen, J. A coupled cluster and full configuration interaction study of CN and CN-. *Chem. Phys. Lett.* **2004**, *393*, 36–43.
- (95) Ansaloni, R.; Bendazzoli, G. L.; Evangelisti, S.; Rossi, E. A parallel Full-CI algorithm. *Comput. Phys. Commun.* **2000**, *128*, 496–515.
- (96) Rolik, Z.; Szabados, Á.; Surján, P. R. A sparse matrix based full-configuration interaction algorithm. *J. Chem. Phys.* **2008**, *128*, 144101.
- (97) Rossi, E.; Bendazzoli, G. L.; Evangelisti, S. Full configuration interaction algorithm on a massively parallel architecture: Direct-list implementation. *J. Comput. Chem.* **1998**, *19*, 658–672.
- (98) Harrison, R.; Stahlberg, E. *Massively parallel full configuration interaction. Benchmark electronic structure calculations on the Intel Touchstone Delta*; 1994.

- (99) Abrams, D. S.; Lloyd, S. Simulation of Many-Body Fermi Systems on a Universal Quantum Computer. *Phys. Rev. Lett.* **1997**, *79*, 2586–2589.
- (100) Szalay, S.; Pfeiffer, M.; Murg, V.; Barcza, G.; Verstraete, F.; Schneider, R.; Legeza, O. Tensor product methods and entanglement optimization for ab initio quantum chemistry. *Int. J. Quantum Chem.* **2015**, *115*, 1342–1391.
- (101) White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863–2866.
- (102) White, S. R. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **1993**, *48*, 10345–10356.
- (103) White, S. R.; Martin, R. L. Ab initio quantum chemistry using the density matrix renormalization group. *J. Chem. Phys.* **1999**, *110*, 4127–4130.
- (104) Chan, G. K.-L.; Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *J. Chem. Phys.* **2002**, *116*, 4462–4476.
- (105) Legeza, O.; Röder, J.; Hess, B. A. Controlling the accuracy of the density-matrix renormalization-group method: The dynamical block state selection approach. *Phys. Rev. B* **2003**, *67*.
- (106) Chan, G. K.-L.; Sharma, S. The Density Matrix Renormalization Group in Quantum Chemistry. *Annu. Rev. Phys. Chem.* **2011**, *62*, 465–481.
- (107) Wouters, S.; Neck, D. V. The density matrix renormalization group for ab initio quantum chemistry. *Eur. Phys. J. D* **2014**, *68*.
- (108) Yanai, T.; Kurashige, Y.; Mizukami, W.; Chalupský, J.; Lan, T. N.; Saitow, M. Density matrix renormalization group for ab initio calculations and associated dynamic

- correlation methods: A review of theory and applications. *Int. J. Quantum Chem.* **2014**, *115*, 283–299.
- (109) Marti, K. H.; Reiher, M. The Density Matrix Renormalization Group Algorithm in Quantum Chemistry. *Z. Phys. Chem.* **2010**, *224*, 583–599.
- (110) Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **2011**, *326*, 96–192.
- (111) Eisert, J.; Cramer, M.; Plenio, M. B. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.* **2010**, *82*, 277–306.
- (112) Moritz, G.; Hess, B. A.; Reiher, M. Convergence behavior of the density-matrix renormalization group algorithm for optimized orbital orderings. *J. Chem. Phys.* **2005**, *122*, 024107.
- (113) Barcza, G.; Legeza, O.; Marti, K. H.; Reiher, M. Quantum-information analysis of electronic states of different molecular structures. *Phys. Rev. A* **2011**, *83*.
- (114) Krumnow, C.; Veis, L.; Legeza, O.; Eisert, J. Fermionic Orbital Optimization in Tensor Network States. *Phys. Rev. Lett.* **2016**, *117*.
- (115) Olivares-Amaya, R.; Hu, W.; Nakatani, N.; Sharma, S.; Yang, J.; Chan, G. K.-L. The ab-initio density matrix renormalization group in practice. *J. Chem. Phys.* **2015**, *142*, 034102.
- (116) Olsen, J. The CASSCF method: A perspective and commentary. *Int. J. Quantum Chem.* **2011**, *111*, 3267–3272.
- (117) Chan, G. K.-L.; Kállay, M.; Gauss, J. State-of-the-art density matrix renormalization group and coupled cluster theory studies of the nitrogen binding curve. *J. Chem. Phys.* **2004**, *121*, 6110–6116.

- (118) Hachmann, J.; Cardoen, W.; Chan, G. K.-L. Multireference correlation in long molecules with the quadratic scaling density matrix renormalization group. *J. Chem. Phys.* **2006**, *125*, 144101.
- (119) Neuscamman, E.; Yanai, T.; Chan, G. K.-L. A review of canonical transformation theory. *Int. Rev. Phys. Chem.* **2010**, *29*, 231–271.
- (120) Kurashige, Y.; Yanai, T. Second-order perturbation theory with a density matrix renormalization group self-consistent field reference function: Theory and application to the study of chromium dimer. *J. Chem. Phys.* **2011**, *135*, 094104.
- (121) Saitow, M.; Kurashige, Y.; Yanai, T. Multireference configuration interaction theory using cumulant reconstruction with internal contraction of density matrix renormalization group wave function. *J. Chem. Phys.* **2013**, *139*, 044118.
- (122) Sharma, S.; Chan, G. K.-L. Communication: A flexible multi-reference perturbation theory by minimizing the Hylleraas functional with matrix product states. *J. Chem. Phys.* **2014**, *141*, 111101.
- (123) Jeanmairet, G.; Sharma, S.; Alavi, A. Stochastic multi-reference perturbation theory with application to the linearized coupled cluster method. *J. Chem. Phys.* **2017**, *146*, 044107.
- (124) Yanai, T.; Saitow, M.; Xiong, X.-G.; Chalupský, J.; Kurashige, Y.; Guo, S.; Sharma, S. Multistate Complete-Active-Space Second-Order Perturbation Theory Based on Density Matrix Renormalization Group Reference States. *J. Chem. Theory Comput.* **2017**, *13*, 4829–4840.
- (125) Freitag, L.; Knecht, S.; Angeli, C.; Reiher, M. Multireference Perturbation Theory with Cholesky Decomposition for the Density Matrix Renormalization Group. *J. Chem. Theory Comput.* **2017**, *13*, 451–459.

- (126) Hedegård, E. D.; Knecht, S.; Kielberg, J. S.; Jensen, H. J. A.; Reiher, M. Density matrix renormalization group with efficient dynamical electron correlation through range separation. *J. Chem. Phys.* **2015**, *142*, 224108.
- (127) Veis, L.; Antalík, A.; Brabec, J.; Neese, F.; Legeza, O.; Pittner, J. Coupled Cluster Method with Single and Double Excitations Tailored by Matrix Product State Wave Functions. *J. Phys. Chem. Lett.* **2016**, *7*, 4072–4078.
- (128) Orús, R. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Ann. Phys.* **2014**, *349*, 117–158.
- (129) Verstraete, F.; Cirac, J. I. Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions. **2004**, cond-mat/0407066.
- (130) Shi, Y.-Y.; Duan, L.-M.; Vidal, G. Classical simulation of quantum many-body systems with a tree tensor network. *Phys. Rev. A* **2006**, *74*.
- (131) Murg, V.; Verstraete, F.; Legeza, O.; Noack, R. M. Simulating strongly correlated quantum systems with tree tensor networks. *Phys. Rev. B* **2010**, *82*.
- (132) Nakatani, N.; Chan, G. K.-L. Efficient tree tensor network states (TTNS) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm. *J. Chem. Phys.* **2013**, *138*, 134113.
- (133) Murg, V.; Verstraete, F.; Schneider, R.; Nagy, P. R.; Legeza, O. Tree Tensor Network State with Variable Tensor Order: An Efficient Multireference Method for Strongly Correlated Systems. *J. Chem. Theory Comput.* **2015**, *11*, 1027–1036.
- (134) Gunst, K.; Verstraete, F.; Wouters, S.; Legeza, O.; Neck, D. V. T3NS: three-legged tree tensor network states. *J. Chem Theory Comput.* **2018**, *14*, 2026–2033.
- (135) Evenbly, G.; Vidal, G. Class of Highly Entangled Many-Body States that can be Efficiently Simulated. *Phys. Rev. Lett.* **2014**, *112*.

- (136) Marti, K. H.; Bauer, B.; Reiher, M.; Troyer, M.; Verstraete, F. Complete-graph tensor network states: a new fermionic wave function ansatz for molecules. *New J. Phys.* **2010**, *12*, 103008.
- (137) Kovyrshin, A.; Reiher, M. Self-adaptive tensor network states with multi-site correlators. *J. Chem. Phys.* **2017**, *147*, 214111.
- (138) Vidal, G. Efficient Classical Simulation of Slightly Entangled Quantum Computations. *Phys. Rev. Lett.* **2003**, *91*.
- (139) Legeza, O.; Sólyom, J. Quantum data compression, quantum information generation, and the density-matrix renormalization-group method. *Phys. Rev. B* **2004**, *70*.
- (140) Leach, A. *Molecular Modelling: Principles and Applications (2nd Edition)*; Pearson: Boston, Massachusetts, 2001.
- (141) Barnett, R. N.; Landman, U. Born-Oppenheimer molecular-dynamics simulations of finite systems: Structure and dynamics of (H₂O)₂. *Phys. Rev. B* **1993**, *48*, 2081–2097.
- (142) Car, R.; Parrinello, M. Unified Approach for Molecular Dynamics and Density-Functional Theory. *Phys. Rev. Lett.* **1985**, *55*, 2471–2474.
- (143) Zillich, R. E.; Paesani, F.; Kwon, Y.; Whaley, K. B. Path integral methods for rotating molecules in superfluids. *J. Chem. Phys.* **2005**, *123*, 114301.
- (144) Paesani, F.; Voth, G. A. The Properties of Water: Insights from Quantum Simulations. *J. Phys. Chem. B* **2009**, *113*, 5702–5719.
- (145) Tuckerman, M. E. *Statistical Mechanics: Theory and Molecular Simulation*; Oxford University Press: Oxford, 2010.
- (146) Berne, B. J.; Thirumalai, D. On the Simulation of Quantum Systems: Path Integral Methods. *Annu. Rev. Phys. Chem.* **1986**, *37*, 401–424.

- (147) Cao, J.; Voth, G. A. The formulation of quantum statistical mechanics based on the Feynman path centroid density. II. Dynamical properties. *J. Chem. Phys.* **1994**, *100*, 5106–5117.
- (148) Habershon, S.; Manolopoulos, D. E.; Markland, T. E.; Miller, T. F. Ring-Polymer Molecular Dynamics: Quantum Effects in Chemical Dynamics from Classical Trajectories in an Extended Phase Space. *Annu. Rev. Phys. Chem.* **2013**, *64*, 387–413.
- (149) Petrenko, R.; Meller, J. *Molecular Dynamics*; John Wiley & Sons, Ltd.: Hoboken, New Jersey, 2010.
- (150) Tannor, D. J. *Introduction to Quantum Mechanics: A Time-Dependent Perspective*; University Science Books: Sausalito, California, 2007.
- (151) Moler, C.; Loan, C. V. Nineteen Dubious Ways to Compute the Exponential of a Matrix Twenty-Five Years Later. *SIAM Rev.* **2003**, *45*, 3–49.
- (152) Bardsley, J. N.; Sundaram, B.; Pinnaduwege, L. A.; Bayfield, J. E. Quantum Dynamics for Driven Weakly Bound Electrons near the Threshold for Classical Chaos. *Phys. Rev. Lett.* **1986**, *56*, 1007–1010.
- (153) Bruderer, M.; Contreras-Pulido, L. D.; Thaller, M.; Sironi, L.; Obreschkow, D.; Plenio, M. B. Inverse counting statistics for stochastic and open quantum systems: the characteristic polynomial approach. *New J. Phys.* **2014**, *16*, 033030.
- (154) Sidje, R. B. Expokit: a software package for computing matrix exponentials. *ACM Trans Math Softw.* **1998**, *24*, 130–156.
- (155) Schomerus, H.; Tworzydło, J. Quantum-to-Classical Crossover of Quasibound States in Open Quantum Systems. *Phys. Rev. Lett.* **2004**, *93*, 154102.

- (156) Chen, X.; Batista, V. S. Matching-pursuit/split-operator-Fourier-transform simulations of excited-state nonadiabatic quantum dynamics in pyrazine. *J. Chem. Phys.* **2006**, *125*, 124313.
- (157) Sawaya, N. P. D.; Huh, J.; Fujita, T.; Saikin, S. K.; Aspuru-Guzik, A. Fast Delocalization Leads To Robust Long-Range Excitonic Transfer in a Large Quantum Chlorosome Model. *Nano Lett.* **2015**, *15*, 1722–1729.
- (158) Runge, E.; Gross, E. K. U. Density-Functional Theory for Time-Dependent Systems. *Phys. Rev. Lett.* **1984**, *52*, 997–1000.
- (159) Burke, K.; Werschnik, J.; Gross, E. K. U. Time-dependent density functional theory: Past present, and future. *J. Chem. Phys.* **2005**, *123*, 062206.
- (160) Casida, M. E. Time-dependent density-functional theory for molecules and molecular solids. *J. Mol. Struct.: THEOCHEM* **2009**, *914*, 3–18.
- (161) Casida, M. E. *Recent Advances in Density Functional Methods*; World Scientific, 1995; pp 155–192.
- (162) Nakatani, N.; Wouters, S.; Neck, D. V.; Chan, G. K.-L. Linear response theory for the density matrix renormalization group: Efficient algorithms for strongly correlated excited states. *J. Chem. Phys.* **2014**, *140*, 024108.
- (163) Dorando, J. J.; Hachmann, J.; Chan, G. K.-L. Analytic response theory for the density matrix renormalization group. *J. Chem. Phys.* **2009**, *130*, 184111.
- (164) Monkhorst, H. J. Calculation of properties with the coupled-cluster method. *Int. J. Quantum Chem.* **2009**, *12*, 421–432.
- (165) Dalgaard, E.; Monkhorst, H. J. Some aspects of the time-dependent coupled-cluster approach to dynamic response functions. *Phys. Rev. A* **1983**, *28*, 1217–1222.

- (166) Pedersen, T. B.; Koch, H. Coupled cluster response functions revisited. *J. Chem. Phys.* **1997**, *106*, 8059–8072.
- (167) Dirac, P. A. M. Note on Exchange Phenomena in the Thomas Atom. *Math. Proc. Cambridge Philos. Soc.* **1930**, *26*, 376.
- (168) Haegeman, J.; Osborne, T. J.; Verstraete, F. Post-matrix product state methods: To tangent space and beyond. *Phys. Rev. B* **2013**, *88*.
- (169) McLachlan, A. A variational solution of the time-dependent Schrodinger equation. *Mol. Phys.* **1964**, *8*, 39–44.
- (170) Carleo, G.; Becca, F.; Schiró, M.; Fabrizio, M. Localization and Glassy Dynamics Of Many-Body Quantum Systems. *Sci. Rep.* **2012**, *2*.
- (171) Ido, K.; Ohgoe, T.; Imada, M. Time-dependent many-variable variational Monte Carlo method for nonequilibrium strongly correlated electron systems. *Phys. Rev. B* **2015**, *92*.
- (172) Haegeman, J.; Cirac, J. I.; Osborne, T. J.; Pižorn, I.; Verschelde, H.; Verstraete, F. Time-Dependent Variational Principle for Quantum Lattices. *Phys. Rev. Lett.* **2011**, *107*.
- (173) Vidal, G. Efficient Simulation of One-Dimensional Quantum Many-Body Systems. *Phys. Rev. Lett.* **2004**, *93*.
- (174) Schollwöck, U. Methods for Time Dependence in DMRG. AIP Conference Proceedings. 2006.
- (175) Manthe, U.; Meyer, H.-D.; Cederbaum, L. S. Wave-packet dynamics within the multi-configuration Hartree framework: General aspects and application to NOCl. *J. Chem. Phys.* **1992**, *97*, 3199–3213.

- (176) Wu, Y.; Batista, V. S. Matching-pursuit for simulations of quantum processes. *J. Chem. Phys.* **2003**, *118*, 6720–6724.
- (177) Ben-Nun, M.; Martinez, T. J. Nonadiabatic molecular dynamics: Validation of the multiple spawning method for a multidimensional problem. *J. Chem. Phys.* **1998**, *108*, 7244–7257.
- (178) Crespo-Otero, R.; Barbatti, M. Recent Advances and Perspectives on Nonadiabatic Mixed Quantum–Classical Dynamics. *Chem. Rev.* **2018**, *118*, 7026–7068, PMID: 29767966.
- (179) Matsika, S.; Krause, P. Nonadiabatic Events and Conical Intersections. *Annu. Rev. Phys. Chem.* **2011**, *62*, 621–643.
- (180) Kitaev, S. B. B. A. Y. Fermionic Quantum Computation. *Ann. Phys.* **2002**, *298*.
- (181) Seeley, J. T.; Richard, M. J.; Love, P. J. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *J. Chem. Phys.* **2012**, *137*, 224109.
- (182) von Neumann, J. *Mathematical foundations of quantum mechanics*, translation - from german ed.; Princeton University Press: Princeton, New Jersey, 1949.
- (183) Brassard, G.; Høyer, P.; Mosca, M.; Tapp, A. *Quantum computation and information*; Contemp. Math.; American Mathematical Society: Providence, Rhode Island, 2002; Vol. 305; pp 53–74.
- (184) Freitas, N.; Morigi, G.; Dunjko, V. Neural Network Operations and Susuki-Trotter evolution of Neural Network States. **2018**, arXiv:1803.02118 [quant-ph].
- (185) Cohen, G.; Gull, E.; Reichman, D. R.; Millis, A. J. Taming the Dynamical Sign Problem in Real-Time Evolution of Quantum Many-Body Problems. *Phys. Rev. Lett.* **2015**, *115*.

- (186) Whitfield, J. D.; Biamonte, J.; Aspuru-Guzik, A. Simulation of electronic structure Hamiltonians using quantum computers. *Mol. Phys.* **2011**, *109*, 735–750.
- (187) Berry, D. W.; Childs, A. M.; Cleve, R.; Kothari, R.; Somma, R. D. Exponential Improvement in Precision for Simulating Sparse Hamiltonians. Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing. New York, NY, USA, 2014; pp 283–292.
- (188) Berry, D. W.; Childs, A. M.; Cleve, R.; Kothari, R.; Somma, R. D. Simulating hamiltonian dynamics with a truncated taylor series. *Phys. Rev. Lett.* **2015**, *114*.
- (189) Low, G. H.; Chuang, I. L. Hamiltonian Simulation by Qubitization. **2016**, arXiv:1610.06546 [quant-ph].
- (190) Low, G. H.; Chuang, I. L. Optimal Hamiltonian Simulation by Quantum Signal Processing. *Phys. Rev. Lett.* **2017**, *118*, 010501.
- (191) Babbush, R.; Berry, D. W.; Sanders, Y. R.; Kivlichan, I. D.; Scherer, A.; Wei, A. Y.; Love, P. J.; Aspuru-Guzik, A. Exponentially more precise quantum simulation of fermions in the configuration interaction representation. *Quantum Sci. Technol.* **2018**, *3*, 015006.
- (192) Babbush, R.; Berry, D. W.; Kivlichan, I. D.; Wei, A. Y.; Love, P. J.; Aspuru-Guzik, A. Exponentially more precise quantum simulation of fermions in second quantization. *New J. Phys.* **2016**, *18*, 033032.
- (193) Wang, H.; Ashhab, S.; Nori, F. Quantum algorithm for simulating the dynamics of an open quantum system. *Phys. Rev. A* **2011**, *83*.
- (194) Childs, A. M.; Li, T. Efficient simulation of sparse Markovian quantum dynamics. *Quantum Inf. Comput.* **2017**, *17*, 901–947.

- (195) Cleve, R.; Wang, C. Efficient Quantum Algorithms for Simulating Lindblad Evolution. **2016**, arXiv:1612.09512 [quant-ph].
- (196) Wiebe, N.; Berry, D.; Høyer, P.; Sanders, B. C. Higher order decompositions of ordered operator exponentials. *J. Phys. A* **2010**, *43*, 065203.
- (197) Low, G. H.; Wiebe, N. Hamiltonian Simulation in the Interaction Picture. **2018**, arXiv:1805.00675 [quant-ph].
- (198) Babbush, R.; Berry, D. W.; McClean, J. R.; Neven, H. Quantum Simulation of Chemistry with Sublinear Scaling to the Continuum. **2018**, arXiv:1807.09802 [quant-ph].
- (199) Kieferova, M.; Scherer, A.; Berry, D. Simulating the dynamics of time-dependent Hamiltonians with a truncated Dyson series. **2018**, arXiv:1805.00582 [quant-ph].
- (200) Li, Y.; Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **2017**, *7*, 021050.
- (201) Dirac, P. A. M. *The Principles of Quantum Mechanics*; Oxford University Press: Oxford, 1930.
- (202) Frenkel, J. *Wave Mechanics: Advanced General Theory*; Oxford University Press: Oxford, 1934.
- (203) Institute, C. M. P vs NP Problem — Clay Mathematics Institute. <http://www.claymath.org/millennium-problems/p-vs-np-problem>, <http://www.claymath.org/millennium-problems/p-vs-np-problem>, Accessed on Tue, June 26, 2018.
- (204) Aaronson, S. BQP and the Polynomial Hierarchy. Proceedings of the Forty-second ACM Symposium on Theory of Computing. New York, NY, 2010; pp 141–150.

- (205) Osborne, T. J. Hamiltonian complexity. *Rep. Prog. Phys.* **2012**, *75*, 022001.
- (206) Gharibian, S.; Huang, Y.; Landau, Z.; Shin, S. W. Quantum Hamiltonian Complexity. *FnT-TCS* **2015**, *10*, 159–282.
- (207) Oliveira, R.; Terhal, B. M. The Complexity of Quantum Spin Systems on a Two-dimensional Square Lattice. *Quantum Info. Comput.* **2008**, *8*, 900–924.
- (208) Biamonte, J. D.; Love, P. J. Realizable Hamiltonians for universal adiabatic quantum computers. *Phys. Rev. A* **2008**, *78*.
- (209) Barahona, F. On the computational complexity of Ising spin glass models. *J. Phys. A* **1982**, *15*, 3241–3253.
- (210) Childs, A. M.; Gosset, D.; Webb, Z. *Automata Languages, and Programming*; Springer Berlin Heidelberg: Berlin, Heidelberg, 2014; pp 308–319.
- (211) Bookatz, A. D. QMA-complete Problems. *Quantum Info. Comput.* **2014**, *14*, 361–383.
- (212) Rassolov, V. A.; Garashchuk, S. Computational complexity in quantum chemistry. *Chem. Phys. Lett.* **2008**, *464*, 262–264.
- (213) Whitfield, J. D.; Zimborás, Z. On the NP-completeness of the Hartree-Fock method for translationally invariant systems. *J. Chem. Phys.* **2014**, *141*, 234103.
- (214) Mazziotti, D. A. Parametrization of the Two-Electron Reduced Density Matrix for its Direct Calculation without the Many-Electron Wave Function. *Phys. Rev. Lett.* **2008**, *101*.
- (215) Mazziotti, D. A. Parametrization of the two-electron reduced density matrix for its direct calculation without the many-electron wave function: Generalizations and applications. *Phys. Rev. A* **2010**, *81*.

- (216) McClean, J. R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **2016**, *18*, 023023.
- (217) Whitfield, J. D.; Yung, M.-H.; Tempel, D. G.; Boixo, S.; Aspuru-Guzik, A. Computational complexity of time-dependent density functional theory. *New J. Phys.* **2014**, *16*, 083035.
- (218) Franck, J.; Dymond, E. G. Elementary processes of photochemical reactions. *Trans. Faraday Soc.* **1926**, *21*, 536.
- (219) Duschinsky, F. The importance of the electron spectrum in multi atomic molecules. Concerning the Franck-Condon principle. *Acta Physicochim. URSS* **1937**, *7*.
- (220) Doktorov, E.; Malkin, I.; Man'ko, V. Dynamical symmetry of vibronic transitions in polyatomic molecules and the Franck-Condon principle. *J. Mol. Spectrosc.* **1977**, *64*, 302–326.
- (221) Ma, X.; Rhodes, W. Multimode squeeze operators and squeezed states. *Phys. Rev. A* **1990**, *41*, 4625–4631.
- (222) Valiant, L. The complexity of computing the permanent. *Theor. Comput. Sci.* **1979**, *8*, 189–201.
- (223) Ryser, H. J. *Combinatorial Mathematics*; Mathematical Association of America, 1963.
- (224) Gard, B. T.; Motes, K. R.; Olson, J. P.; Rohde, P. P.; Dowling, J. P. *From Atomic to Mesoscale*; Chapter Chapter 8, pp 167–192.
- (225) Seshadreesan, K. P.; Olson, J. P.; Motes, K. R.; Rohde, P. P.; Dowling, J. P. Boson sampling with displaced single-photon Fock states versus single-photon-added coherent states: The quantum-classical divide and computational-complexity transitions in linear optics. *Phys. Rev. A* **2015**, *91*, 022334.

- (226) Olson, J.; Seshadreesan, K.; Motes, K.; Rohde, P.; Dowling, J. Sampling arbitrary photon-added or photon-subtracted squeezed states is in the same complexity class as boson sampling. *Phys. Rev. A* **2015**, *91*, 022317.
- (227) Lund, A. P.; Laing, A.; Rahimi-Keshari, S.; Rudolph, T.; O’Brien, J. L.; Ralph, T. C. Boson Sampling from a Gaussian State. *Phys. Rev. Lett.* **2014**, *113*, 100502.
- (228) Huh, J.; Yung, M.-H. Vibronic Boson Sampling: Generalized Gaussian Boson Sampling for Molecular Vibronic Spectra at Finite Temperature. *Sci. Rep.* **2017**, *7*, 7462.
- (229) Sattasathuchana, T.; Murri, R.; Baldrige, K. An Efficient Analytic Approach for Calculation of Multi-Dimensional Franck–Condon Factors and Associated Photoelectron Spectra. *J. Chem. Theory Comput.* **2017**, *13*, 2147–2158.
- (230) Feynman, R. P. Quantum Mechanical Computers. *Opt. Photonics News* **1985**, *11*, 11–20.
- (231) Kitaev, A. Y. Quantum measurements and the Abelian Stabilizer Problem. **1995**, quant-ph/9511026.
- (232) Clader, B. D.; Jacobs, B. C.; Sprouse, C. R. Preconditioned quantum linear system algorithm. *Phys. Rev. Lett.* **2013**, *110*, 250504.
- (233) Sugisaki, K.; Yamamoto, S.; Nakazawa, S.; Toyota, K.; Sato, K.; Shiomi, D.; Takui, T. Quantum Chemistry on Quantum Computers: A Polynomial-Time Quantum Algorithm for Constructing the Wave Functions of Open-Shell Molecules. *J. Phys. Chem. A* **2016**, *120*, 6459–6466.
- (234) Lanyon, B. P.; Whitfield, J. D.; Gillett, G. G.; Goggin, M. E.; Almeida, M. P.; Kaszal, I.; Biamonte, J. D.; Mohseni, M.; Powell, B. J.; Barbieri, M.; Aspuru-Guzik, A.; White, A. G. Towards quantum chemistry on a quantum computer. *Nat. Chem.* **2010**, *2*, 106–111.

- (235) Du, J.; Xu, N.; Peng, X.; Wang, P.; Wu, S.; Lu, D. NMR Implementation of a Molecular Hydrogen Quantum Simulation with Adiabatic State Preparation. *Phys. Rev. Lett.* **2010**, *104*.
- (236) Dolde, F.; Bergholm, V.; Wang, Y.; Jakobi, I.; Naydenov, B.; Pezzagna, S.; Meijer, J.; Jelezko, F.; Neumann, P.; Schulte-Herbrüggen, T.; Biamonte, J.; Wrachtrup, J. High-fidelity spin entanglement using optimal control. *Nat. Commun.* **2014**, *5*.
- (237) Fowler, A. G.; Mariantoni, M.; Martinis, J. M.; Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* **2012**, *86*, 032324.
- (238) Gottesman, D. An Introduction to Quantum Error Correction and Fault-Tolerant Quantum Computation. **2009**, arXiv:0904.2557 [quant-ph].
- (239) Aharonov, D.; Ben-Or, M. Fault-tolerant Quantum Computation with Constant Error. Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing. New York, NY, USA, 1997; pp 176–188.
- (240) Harrow, A. W.; Hassidim, A.; Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **2009**, *103*, 150502.
- (241) Narayan Chowdhury, A.; Somma, R. D. Quantum algorithms for Gibbs sampling and hitting-time estimation. **2016**, arXiv:1603.02940 [quant-ph].
- (242) Wiebe, N.; Braun, D.; Lloyd, S. Quantum algorithm for data fitting. *Phys. Rev. Lett.* **2012**, *109*, 050505.
- (243) Lloyd, S.; Mohseni, M.; Rebentrost, P. Quantum principal component analysis. *Nat. Phys.* **2014**, *10*, 631.
- (244) Rebentrost, P.; Mohseni, M.; Lloyd, S. Quantum support vector machine for big data classification. *Phys. Rev. Lett.* **2014**, *113*, 130503.

- (245) Islam, R.; Senko, C.; Campbell, W.; Korenblit, S.; Smith, J.; Lee, A.; Edwards, E.; Wang, C.-C.; Freericks, J.; Monroe, C. Emergence and frustration of magnetism with variable-range interactions in a quantum simulator. *Science* **2013**, *340*, 583–587.
- (246) Choi, J.-y.; Hild, S.; Zeiher, J.; Schauß, P.; Rubio-Abadal, A.; Yefsah, T.; Khemani, V.; Huse, D. A.; Bloch, I.; Gross, C. Exploring the many-body localization transition in two dimensions. *Science* **2016**, *352*, 1547–1552.
- (247) Zhang, J.; Pagano, G.; Hess, P. W.; Kyprianidis, A.; Becker, P.; Kaplan, H.; Gorshkov, A. V.; Gong, Z.-X.; Monroe, C. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature* **2017**, *551*, 601.
- (248) Mazurenko, A.; Chiu, C. S.; Ji, G.; Parsons, M. F.; Kanász-Nagy, M.; Schmidt, R.; Grusdt, F.; Demler, E.; Greif, D.; Greiner, M. A cold-atom Fermi–Hubbard antiferromagnet. *Nature* **2017**, *545*, 462.
- (249) Bernien, H.; Schwartz, S.; Keesling, A.; Levine, H.; Omran, A.; Pichler, H.; Choi, S.; Zibrov, A. S.; Endres, M.; Greiner, M.; Vuletić, V.; Luking, M. D. Probing many-body dynamics on a 51-atom quantum simulator. *Nature* **2017**, *551*, 579.
- (250) Childs, A. M. Quantum information processing in continuous time. Ph.D. thesis, Massachusetts Institute of Technology, 2004.
- (251) Dodd, J. L.; Nielsen, M. A.; Bremner, M. J.; Thew, R. T. Universal quantum computation and simulation using any entangling Hamiltonian and local unitaries. *Phys. Rev. A* **2002**, *65*, 040301.
- (252) Nielsen, M. A.; Bremner, M. J.; Dodd, J. L.; Childs, A. M.; Dawson, C. M. Universal simulation of Hamiltonian dynamics for quantum systems with finite-dimensional state spaces. *Phys. Rev. A* **2002**, *66*, 022317.

- (253) Aharonov, D.; Ta-Shma, A. Adiabatic Quantum State Generation and Statistical Zero Knowledge. Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing. New York, New York, 2003; pp 20–29.
- (254) Toloui, B.; Love, P. J. Quantum Algorithms for Quantum Chemistry based on the sparsity of the CI-matrix. **2013**, arXiv:1312.2579 [quant-ph].
- (255) Haah, J.; Hastings, M. B.; Kothari, R.; Low, G. H. Quantum algorithm for simulating real time evolution of lattice Hamiltonians. **2018**, arXiv:1801.03922 [quant-ph].
- (256) Childs, A. M.; Wiebe, N. Hamiltonian Simulation Using Linear Combinations of Unitary Operations. *Quantum Info. Comput.* **2012**, *12*, 901–924.
- (257) Childs, A. M.; Kothari, R. Simulating sparse Hamiltonians with star decompositions. Conference on Quantum Computation, Communication, and Cryptography. 2010; pp 94–103.
- (258) Chunhao, W.; Wossnig, L. A quantum algorithm for simulating non-sparse Hamiltonians. **2018**, arXiv:1803.08273 [quant-ph].
- (259) Kimmel, S.; Lin, C. Y.-Y.; Low, G. H.; Ozols, M.; Yoder, T. J. Hamiltonian simulation with optimal sample complexity. *npj Quantum Inf.* **2017**, *3*, 13.
- (260) Berry, D. W.; Ahokas, G.; Cleve, R.; Sanders, B. C. Efficient quantum algorithms for simulating sparse Hamiltonians. *Communications in Mathematical Physics* **2007**, *270*, 359–371.
- (261) Suzuki, M. Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations. *Phys. Lett. A* **1990**, *146*, 319–323.
- (262) Childs, A. M. On the relationship between continuous-and discrete-time quantum walk. *Commun. Math. Phys.* **2010**, *294*, 581–603.

- (263) Cleve, R.; Ekert, A.; Macchiavello, C.; Mosca, M. Quantum algorithms revisited. *Proc. Royal Soc. Lond. A.* 1998; pp 339–354.
- (264) Low, G. H.; Yoder, T. J.; Chuang, I. L. Methodology of resonant equiangular composite quantum gates. *Phys. Rev. X* **2016**, *6*, 041067.
- (265) Low, G. H. Hamiltonian simulation with nearly optimal dependence on spectral norm. **2018**, arXiv:1807.03967 [quant-ph].
- (266) Gilyén, A.; Su, Y.; Low, G. H.; Wiebe, N. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. **2018**, arXiv:1806.01838 [quant-ph].
- (267) Low, G. H.; Chuang, I. L. Hamiltonian Simulation by Uniform Spectral Amplification. **2017**, arXiv:1707.05391 [quant-ph].
- (268) Wang, C.; Wossnig, L. A quantum algorithm for simulating non-sparse Hamiltonians. **2018**, arXiv:1803.08273 [quant-ph].
- (269) Poulin, D.; Qarry, A.; Somma, R.; Verstraete, F. Quantum simulation of time-dependent Hamiltonians and the convenient illusion of Hilbert space. *Phys. Rev. Lett.* **2011**, *106*, 170501.
- (270) Wiebe, N.; Berry, D. W.; Høyer, P.; Sanders, B. C. Simulating quantum dynamics on a quantum computer. *J. Phys. A* **2011**, *44*, 445308.
- (271) Kliesch, M.; Barthel, T.; Gogolin, C.; Kastoryano, M.; Eisert, J. Dissipative quantum church-turing theorem. *Phys. Rev. Lett.* **2011**, *107*, 120501.
- (272) Childs, A. M.; Maslov, D.; Nam, Y.; Ross, N. J.; Su, Y. Toward the first quantum simulation with quantum speedup. *Proc. Natl. Acad. Sci. U.S.A.* **2018**, *115*, 9456–9461.

- (273) Zalka, C. Simulating quantum systems on a quantum computer. *Proc. Royal Soc. Lond. A: Mathematical, Physical and Engineering Sciences*. 1998; pp 313–322.
- (274) Giovannetti, V.; Lloyd, S.; Maccone, L. Architectures for a quantum random access memory. *Phys. Rev. A* **2008**, *78*, 052310.
- (275) Dobšíček, M.; Johansson, G.; Shumeiko, V.; Wendin, G. Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark. *Phys. Rev. A* **2007**, *76*.
- (276) Svore, K. M.; Hastings, M. B.; Freedman, M. Faster phase estimation. *Quantum Inf. Comput.* **2014**, *14*, 306–328.
- (277) Wiebe, N.; Granade, C. Efficient Bayesian phase estimation. *Phys. Rev. Lett.* **2016**, *117*, 010503.
- (278) Paesani, S.; Gentile, A. A.; Santagati, R.; Wang, J.; Wiebe, N.; Tew, D. P.; O’Brien, J. L.; Thompson, M. G. Experimental Bayesian quantum phase estimation on a silicon photonic chip. *Phys. Rev. Lett.* **2017**, *118*, 100503.
- (279) Berry, D. W.; Kieferová, M.; Scherer, A.; Sanders, Y. R.; Low, G. H.; Wiebe, N.; Gidney, C.; Babbush, R. Improved techniques for preparing eigenstates of fermionic Hamiltonians. *npj Quantum Inf.* **2018**, *4*, 22.
- (280) Poulin, D.; Kitaev, A.; Steiger, D. S.; Hastings, M. B.; Troyer, M. Quantum Algorithm for Spectral Measurement with a Lower Gate Count. *Phys. Rev. Lett.* **2018**, *121*, 010501.
- (281) Babbush, R.; Gidney, C.; Berry, D. W.; Wiebe, N.; McClean, J.; Paler, A.; Fowler, A.; Neven, H. Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity. **2018**, [arXiv:1805.03662](https://arxiv.org/abs/1805.03662) [quant-ph].

- (282) Bernstein, E.; Vazirani, U. Quantum Complexity Theory. *SIAM J. Comput.* **1997**, *26*, 1411–1473.
- (283) Aaronson, S.; Ambainis, A. Forrelation: A problem that optimally separates quantum from classical computing. *SIAM J. Comput.* **2018**, *47*, 982–1038.
- (284) Whitfield, J. D. Unified views of quantum simulation algorithms for chemistry. **2015**, arXiv:1502.03771 [quant-ph].
- (285) Grover, L.; Rudolph, T. Creating superpositions that correspond to efficiently integrable probability distributions. **2002**, quant-ph/0208112.
- (286) Abrams, D. S. Quantum algorithms. Ph.D. thesis, Massachusetts Institute of Technology, 1999.
- (287) Tubman, N. M.; Mejuto-Zaera, C.; Epstein, J. M.; Hait, D.; Levine, D. S.; Huggins, W.; Jiang, Z.; McClean, J. R.; Babbush, R.; Head-Gordon, M.; Whaley, K. B. Postponing the orthogonality catastrophe: efficient state preparation for electronic structure simulations on quantum devices. **2018**, arXiv:1809.05523 [quant-ph].
- (288) Soklakov, A. N.; Schack, R. Efficient state preparation for a register of quantum bits. *Phys. Rev. A* **2006**, *73*, 012307.
- (289) Tubman, N. M.; Lee, J.; Takeshita, T. Y.; Head-Gordon, M.; Whaley, K. B. A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *J. Phys. A* **2016**, *145*, 044112.
- (290) Farhi, E.; Goldstone, J.; Gutmann, S.; Lapan, J.; Lundgren, A.; Preda, D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **2001**, *292*, 472–475.
- (291) Marzlin, K.-P.; Sanders, B. C. Inconsistency in the application of the adiabatic theorem. *Phys. Rev. Lett.* **2004**, *93*, 160408.

- (292) Cheung, D.; Høyer, P.; Wiebe, N. Improved error bounds for the adiabatic approximation. *J. Phys. A* **2011**, *44*, 415302.
- (293) Wiebe, N.; Babcock, N. S. Improved error-scaling for adiabatic quantum evolutions. *New J. Phys.* **2012**, *14*, 013024.
- (294) Lidar, D. A.; Rezakhani, A. T.; Hamma, A. Adiabatic approximation with exponential accuracy for many-body systems and quantum computation. *J. Math. Phys.* **2009**, *50*, 102106.
- (295) Kieferová, M.; Wiebe, N. On the power of coherently controlled quantum adiabatic evolutions. *New J. Phys.* **2014**, *16*, 123034.
- (296) Roland, J.; Cerf, N. J. Quantum search by local adiabatic evolution. *Phy. Rev. A* **2002**, *65*, 042308.
- (297) Somma, R. D.; Nagaj, D.; Kieferová, M. Quantum speedup by quantum annealing. *Phys. Rev. Lett.* **2012**, *109*, 050501.
- (298) Crosson, E.; Farhi, E.; Yen-Yu Lin, C.; Lin, H.-H.; Shor, P. Different Strategies for Optimization Using the Quantum Adiabatic Algorithm. **2014**, arXiv:1401.7320 [quant-ph].
- (299) Yung, M.-H.; Aspuru-Guzik, A. A quantum–quantum Metropolis algorithm. *Proc. Natl. Acad. Sci. U.S.A.* **2012**, *109*, 754–759.
- (300) Ditchfield, R.; Hehre, W. J.; Pople, J. A. Self-Consistent Molecular-Orbital Methods. IX. An Extended Gaussian-Type Basis for Molecular-Orbital Studies of Organic Molecules. *J. Chem. Phys.* **1971**, *54*, 724–728.
- (301) Babbush, R.; Wiebe, N.; McClean, J.; McClain, J.; Neven, H.; Chan, G. K.-L. Low-depth quantum simulation of materials. *Phys. Rev. X* **2018**, *8*, 011044.

- (302) Kivlichan, I. D.; McClean, J.; Wiebe, N.; Gidney, C.; Aspuru-Guzik, A.; Chan, G. K.-L.; Babbush, R. Quantum Simulation of Electronic Structure with Linear Depth and Connectivity. *Phys. Rev. Lett.* **2018**, *120*, 110501.
- (303) White, S. R. Hybrid grid/basis set discretizations of the Schrödinger equation. *J. Chem. Phys.* **2017**, *147*, 244102.
- (304) Jones, N. C.; Whitfield, J. D.; McMahon, P. L.; Yung, M.-H.; Meter, R. V.; Aspuru-Guzik, A.; Yamamoto, Y. Faster quantum chemistry simulation on fault-tolerant quantum computers. *New J. Phys.* **2012**, *14*, 115023.
- (305) Reiher, M.; Wiebe, N.; Svore, K. M.; Wecker, D.; Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc. Natl. Acad. Sci. U.S.A.* **2017**, *114*, 7555–7560.
- (306) McClean, J. R.; Babbush, R.; Love, P. J.; Aspuru-Guzik, A. Exploiting locality in quantum computation for quantum chemistry. *J. Phys. Chem. Lett.* **2014**, *5*, 4368–4380.
- (307) Kliuchnikov, V.; Maslov, D.; Mosca, M. Fast and Efficient Exact Synthesis of Single-qubit Unitaries Generated by Clifford and T Gates. *Quantum Info. Comput.* **2013**, *13*, 607–630.
- (308) Gidney, C. Halving the cost of quantum addition. *Quantum* **2018**, *2*, 74.
- (309) Veis, L.; Višňák, J.; Fleig, T.; Knecht, S.; Saue, T.; Visscher, L.; Pittner, J. Relativistic quantum chemistry on quantum computers. *Phys. Rev. A* **2012**, *85*.
- (310) Veis, L.; Pittner, J. *Advances in Chemical Physics*; John Wiley & Sons Inc., 2014; pp 107–136.
- (311) Welch, J.; Greenbaum, D.; Mostame, S.; Aspuru-Guzik, A. Efficient quantum circuits for diagonal unitaries without ancillas. *New J. Phys.* **2014**, *16*, 033040.

- (312) Raeisi, S.; Wiebe, N.; Sanders, B. C. Quantum-circuit design for efficient simulations of many-body quantum dynamics. *New J. Phys.* **2012**, *14*, 103017.
- (313) Berry, D. W.; Childs, A. M. Black-box Hamiltonian Simulation and Unitary Implementation. *Quantum Info. Comput.* **2012**, *12*, 29–62.
- (314) Berry, D. W.; Childs, A. M.; Kothari, R. Hamiltonian simulation with nearly optimal dependence on all parameters. IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS). 2015; pp 792–809.
- (315) Novo, L.; Berry, D. W. Improved Hamiltonian simulation via a truncated Taylor series and corrections. **2016**, arXiv:1611.10033 [quant-ph].
- (316) Santagati, R.; Wang, J.; Gentile, A. A.; Paesani, S.; Wiebe, N.; McClean, J. R.; Morley-Short, S.; Shadbolt, P. J.; Bonneau, D.; Silverstone, J. W.; Tew, D. P.; Zhou, X.; O’Brien, J. L.; Thompson, M. G. Witnessing eigenstates for quantum simulation of Hamiltonian spectra. *Sci. Adv.* **2018**, *4*, eaap9646.
- (317) Temme, K.; Bravyi, S.; Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **2017**, *119*, 180509.
- (318) Endo, S.; Benjamin, S. C.; Li, Y. Practical Quantum Error Mitigation for Near-Future Applications. *Phys. Rev. X* **2018**, *8*, 031027.
- (319) Kandala, A.; Temme, K.; Corcoles, A. D.; Mezzacapo, A.; Chow, J. M.; Gambetta, J. M. Extending the computational reach of a noisy superconducting quantum processor. **2018**, arXiv:1805.04492 [quant-ph].
- (320) McClean, J. R.; Kimchi-Schwartz, M. E.; Carter, J.; de Jong, W. A. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A* **2017**, *95*.

- (321) McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; Neven, H. Barren plateaus in quantum neural network training landscapes. **2018**, arXiv:1803.11173 [quant-ph].
- (322) Kutzelnigg, W. *In Methods of Electronic Structure Theory; Schaefer, H. F., III, Ed*; Plenum Press: New York, NY, USA, 1977; p 129.
- (323) Hoffmann, M. R.; Simons, J. A unitary multiconfigurational coupled-cluster method: Theory and applications. *J. Chem. Phys.* **1988**, *88*, 993–1002.
- (324) Bartlett, R. J.; Kucharski, S. A.; Noga, J. Alternative coupled-cluster ansätze II. The unitary coupled-cluster method. *Chem. Phys. Lett.* **1989**, *155*, 133–140.
- (325) Wecker, D.; Hastings, M. B.; Troyer, M. Progress towards practical quantum variational algorithms. *Phys. Rev. A* **2015**, *92*.
- (326) Dallaire-Demers, P.-L.; Romero, J.; Veis, L.; Sim, S.; Aspuru-Guzik, A. Low-depth circuit ansatz for preparing correlated fermionic states on a quantum computer. **2018**, arXiv:1801.01053 [quant-ph].
- (327) Barkoutsos, P. K.; Gonthier, J. F.; Sokolov, I.; Moll, N.; Salis, G.; Fuhrer, A.; Ganzhorn, M.; Egger, D. J.; Troyer, M.; Mezzacapo, A.; Filipp, S.; Tavernelli, I. Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions. *Phys. Rev. A* **2018**, *98*, 022322.
- (328) Yung, M.-H.; Casanova, J.; Mezzacapo, A.; McClean, J.; Lamata, L.; Aspuru-Guzik, A.; Solano, E. From transistor to trapped-ion computers for quantum chemistry. *Sci. Rep.* **2014**, *4*.
- (329) Romero, J.; Babbush, R.; McClean, J.; Hempel, C.; Love, P.; Aspuru-Guzik, A. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Sci. Technol.* **2018**,

- (330) Dumitrescu, E. F.; McCaskey, A. J.; Hagen, G.; Jansen, G. R.; Morris, T. D.; Papenbrock, T.; Pooser, R. C.; Dean, D. J.; Lougovski, P. Cloud Quantum Computing of an Atomic Nucleus. *Phys. Rev. Lett.* **2018**,
- (331) Farhi, E.; Goldstone, J.; Gutmann, S. A Quantum Approximate Optimization Algorithm. **2014**, arXiv:1411.4028 [quant-ph].
- (332) Farhi, E.; Harrow, A. W. Quantum Supremacy through the Quantum Approximate Optimization Algorithm. **2016**, arXiv:1602.07674 [quant-ph].
- (333) Bengio, Y.; Lamblin, P.; Popovici, D.; Larochelle, H. Greedy layer-wise training of deep networks. NIPS'06 Proceedings of the 19th International Conference on Neural Information Processing Systems. 2006; pp 153–160.
- (334) He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. Proceedings of the IEEE conference on computer vision and pattern recognition. 2016; pp 770–778.
- (335) Jiang, Z.; Sung, K. J.; Kechedzhi, K.; Smelyanskiy, V. N.; Boixo, S. Quantum algorithms to simulate many-body physics of correlated fermions. *Phys. Rev. Appl.* **2018**, *9*, 044036.
- (336) Wecker, D.; Hastings, M. B.; Wiebe, N.; Clark, B. K.; Nayak, C.; Troyer, M. Solving strongly correlated electron models on a quantum computer. *Phys. Rev. A* **2015**, *92*, 062318.
- (337) Hoffman, D. K.; Raffinetti, R. C.; Ruedenberg, K. Generalization of Euler Angles to N-Dimensional Orthogonal Matrices. *J. Math. Phys.* **1972**, *13*, 528–533.
- (338) Shi, T.; Demler, E.; Cirac, J. I. Variational study of fermionic and bosonic systems with non-Gaussian states: Theory and applications. *Ann. Phys.* **2018**, *390*, 245–302.

- (339) Endo, S.; Jones, T.; McArdle, S.; Yuan, X.; Benjamin, S. Variational quantum algorithms for discovering Hamiltonian spectra. **2018**, arXiv:1806.05707 [quant-ph].
- (340) McKay, D. C.; Filipp, S.; Mezzacapo, A.; Magesan, E.; Chow, J. M.; Gambetta, J. M. Universal gate for fixed-frequency qubits via a tunable bus. *Phys. Rev. Appl.* **2016**, *6*, 064007.
- (341) Roth, M.; Ganzhorn, M.; Moll, N.; Filipp, S.; Salis, G.; Schmidt, S. Analysis of a parametrically driven exchange-type gate and a two-photon excitation gate between superconducting qubits. *Phys. Rev. A* **2017**, *96*, 062323.
- (342) Egger, D. J.; Ganzhorn, M.; Salis, G.; Fuhrer, A.; Mueller, P.; Barkoutsos, P. K.; Moll, N.; Tavernelli, I.; Filipp, S. Entanglement generation in superconducting qubits using holonomic operations. **2018**, arXiv:1804.04900 [quant-ph].
- (343) Bravyi, S.; Gambetta, J. M.; Mezzacapo, A.; Temme, K. Tapering off qubits to simulate fermionic Hamiltonians. **2017**, arXiv:1701.08213 [quant-ph].
- (344) Romero, J.; Olson, J. P.; Aspuru-Guzik, A. Quantum autoencoders for efficient compression of quantum data. *Quantum Sci. Technol.* **2017**, *2*, 045001.
- (345) Olson, J.; Sim, S.; Cao, Y. Implementation of CUSP using Cirq. https://github.com/zapatacomputing/cusp_cirq_demo, 2018; Accessed on Mon, November 5, 2018.
- (346) Rubin, N. C.; Babbush, R.; McClean, J. Application of fermionic marginal constraints to hybrid quantum algorithms. *New J. Phys.* **2018**, *20*, 053020.
- (347) Artacho, E.; Sánchez-Portal, D.; Ordejón, P.; Garcia, A.; Soler, J. M. Linear-scaling ab-initio calculations for large and complex systems. *Phys. Status Solidi B* **1999**, *215*, 809–817.

- (348) Filippi, C.; Umrigar, C. J. Correlated sampling in quantum Monte Carlo: A route to forces. *Phys. Rev. B* **2000**, *61*, R16291–R16294.
- (349) Cioslowski, J. *Many-electron densities and reduced density matrices*; Springer Science & Business Media, 2000.
- (350) Alcoba, D.; Valdemoro, C. Spin structure and properties of the correlation matrices corresponding to pure spin states: Controlling the S-representability of these matrices. *Int. J. Quantum Chem.* **2005**, *102*, 629–644.
- (351) Rios, L. M.; Sahinidis, N. V. Derivative-free optimization: a review of algorithms and comparison of software implementations. *J. Global Optim.* **2012**, *56*, 1247–1293.
- (352) Guerreschi, G. G.; Smelyanskiy, M. Practical optimization for hybrid quantum-classical algorithms. **2017**, arXiv:1701.01450 [quant-ph].
- (353) Kennedy, J. *Encyclopedia of machine learning*; Springer, 2011; pp 760–766.
- (354) Spall, J. C. A one-measurement form of simultaneous perturbation stochastic approximation. *Automatica* **1997**, *33*, 109–112.
- (355) Spall, J. C. Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE Trans. Autom. Control* **2000**, *45*, 1839–1853.
- (356) Han, L.; Neumann, M. Effect of dimensionality on the Nelder–Mead simplex method. *Optim. Methods Softw.* **2006**, *21*, 1–16.
- (357) McArdle, S.; Yuan, X.; Benjamin, S. Error mitigated digital quantum simulation. **2018**, arXiv:1807.02467 [quant-ph].
- (358) McArdle, S.; Endo, S.; Aspuru-Guzik, A.; Benjamin, S.; Yuan, X. Quantum computational chemistry. **2018**, arXiv:1808.10402 [quant-ph].

- (359) Garcia-Saez, A.; Latorre, J. I. Addressing hard classical problems with Adiabatically Assisted Variational Quantum Eigensolvers. **2018**, arXiv:1806.02287 [quant-ph].
- (360) Higgott, O.; Wang, D.; Brierley, S. Variational Quantum Computation of Excited States. **2018**, arXiv:1805.08138 [quant-ph].
- (361) Angeli, C.; Pastore, M.; Cimiraglia, R. New perspectives in multireference perturbation theory: the n-electron valence state approach. *Theor. Chem. Acc.* **2007**, *117*, 743–754.
- (362) Zgid, D.; Ghosh, D.; Neuscamman, E.; Chan, G. K.-L. A study of cumulant approximations to n-electron valence multireference perturbation theory. *J. Chem. Phys.* **2009**, *130*, 194107.
- (363) Kong, L.; Valeev, E. F. Perturbative correction for the basis set incompleteness error of complete-active-space self-consistent field. *J. Chem. Phys.* **2010**, *133*, 174126.
- (364) Torheyden, M.; Valeev, E. F. Universal perturbative explicitly correlated basis set incompleteness correction. *J. Chem. Phys.* **2009**, *131*, 171103.
- (365) McArdle, S.; Jones, T.; Endo, S.; Li, Y.; Benjamin, S.; Yuan, X. Variational quantum simulation of imaginary time evolution. **2018**, arXiv:1804.03023 [quant-ph].
- (366) Babbush, R.; Love, P. J.; Aspuru-Guzik, A. Adiabatic Quantum Simulation of Quantum Chemistry. *Sci. Rep.* **2014**, *4*.
- (367) Smirnov, A. Y.; Savel'ev, S.; Mourokh, L. G.; Nori, F. Modelling chemical reactions using semiconductor quantum dots. *Europhys. Lett.* **2007**, *80*, 67008.
- (368) Torrontegui, E.; Ruschhaupt, A.; Guéry-Odelin, D.; Muga, J. G. Simulation of quantum collinear chemical reactions with ultracold atoms. *J. Phys. B: At., Mol. Opt. Phys.* **2011**, *44*, 195302.

- (369) Georgescu, I. M.; Ashhab, S.; Nori, F. Quantum simulation. *Rev. Mod. Phys.* **2014**, *86*, 153–185.
- (370) Perdomo-Ortiz, A.; Dickson, N.; Drew-Brook, M.; Rose, G.; Aspuru-Guzik, A. Finding low-energy conformations of lattice protein models by quantum annealing. *Sci. Rep.* **2012**, *2*.
- (371) Babbush, R.; Perdomo-Ortiz, A.; O’Gorman, B.; Macready, W.; Aspuru-Guzik, A. *Advances in Chemical Physics*; Wiley-Blackwell, 2014; Vol. 155; Chapter 5, pp 201–244.
- (372) Kempe, J.; Kitaev, A.; Regev, O. The complexity of the local Hamiltonian problem. *SIAM J. Comput.* **2006**, *35*, 1070–1097.
- (373) Cao, Y.; Babbush, R.; Biamonte, J.; Kais, S. Hamiltonian gadgets with reduced resource requirements. *Phys. Rev. A* **2015**, *91*.
- (374) Barkoutsos, P. K.; Moll, N.; Staar, P. W. J.; Mueller, P.; Fuhrer, A.; Filipp, S.; Troyer, M.; Tavernelli, I. Fermionic Hamiltonians for quantum simulations: a general reduction scheme. **2017**, arXiv:1706.03637 [quant-ph].
- (375) Clements, W. R.; Renema, J. J.; Eckstein, A.; Valido, A. A.; Lita, A.; Gerrits, T.; Nam, S. W.; Kolthammer, W. S.; Huh, J.; Walmsley, I. A. Approximating Vibronic Spectroscopy with Imperfect Quantum Optics. **2017**, arXiv:1710.08655 [quant-ph].
- (376) Sparrow, C.; Martín-López, E.; Maraviglia, N.; Neville, A.; Harrold, C.; Carolan, J.; Joglekar, Y. N.; Hashimoto, T.; Matsuda, N.; O’Brien, J. L.; Tew, D. P.; Laing, A. Simulating the vibrational quantum dynamics of molecules using photonics. *Nature* **2018**, *557*, 660–667.

- (377) Hamilton, C.; Kruse, R.; Sansoni, L.; Barkhofen, S.; Silberhorn, C.; Jex, I. Gaussian Boson Sampling. *Phys. Rev. Lett.* **2017**, *119*, 170501.
- (378) Shen, Y.; Lu, Y.; Zhang, K.; Zhang, J.; Zhang, S.; Huh, J.; Kim, K. Quantum optical emulation of molecular vibronic spectroscopy using a trapped-ion device. *Chem. Sci.* **2018**, *9*, 836–840.
- (379) Olivares, D. G.; Peropadre, B.; Huh, J.; García-Ripoll, J. J. Quantum emulation of molecular force fields: A blueprint for a superconducting architecture. *Phys. Rev. Appl.* **2017**, *8*, 064008.
- (380) Eckhardt, R. Stan ulam, john von neumann, and the monte carlo method. *Los Alamos Sci.* **1987**, *15*, 30.
- (381) Motta, M.; Ye, E.; McClean, J. R.; Li, Z.; Minnich, A. J.; Babbush, R.; Kin-Lic Chan, G. Low rank representations for quantum simulation of electronic structure. **2018**, arXiv:1808.02625 [physics.comp-ph].
- (382) Johnson, P. D.; Romero, J.; Olson, J.; Cao, Y.; Aspuru-Guzik, A. QVECTOR: an algorithm for device-tailored quantum error correction. **2017**, arXiv:1711.02249 [quant-ph].
- (383) Wang, D.; Higgott, O.; Brierley, S. A Generalised Variational Quantum Eigensolver. **2018**, arXiv:1802.00171 [quant-ph].
- (384) O’Brien, T. E.; Tarasinski, B.; Terhal, B. M. Quantum phase estimation for noisy, small-scale experiments. **2018**, arXiv:1809.09697 [quant-ph].
- (385) Rigetti 8Q-Agave specification v.2.0.0.dev0. <https://www.rigetti.com/qpu>, 2018; Accessed on Mon, November 5, 2018.

- (386) IBM Q 5 Tenerife backend specification v1.1.0. <https://github.com/QISKit/qiskit-backend-information/tree/master/backends/tenerife/V1>, 2018; Accessed on Mon, November 5, 2018.
- (387) Havlicek, V.; Córcoles, A. D.; Temme, K.; Harrow, A. W.; Kandala, A.; Chow, J. M.; Gambetta, J. M. Supervised learning with quantum enhanced feature spaces. **2018**, [arXiv:1804.11326](https://arxiv.org/abs/1804.11326) [quant-ph].
- (388) Khatri, S.; LaRose, R.; Poremba, A.; Cincio, L.; Sornborger, A. T.; Coles, P. J. Quantum-assisted quantum compiling. **2018**, [arXiv:1807.00800](https://arxiv.org/abs/1807.00800) [quant-ph].
- (389) Smith, R. S.; Curtis, M. J.; Zeng, W. J. A Practical Quantum Instruction Set Architecture. **2016**, [arXiv:1608.03355](https://arxiv.org/abs/1608.03355) [quant-ph].
- (390) QISKit. <https://github.com/QISKit/qiskit-sdk-py>, 2018; Accessed on Mon, November 5, 2018.
- (391) Steiger, D. S.; Häner, T.; Troyer, M. ProjectQ: an open source software framework for quantum computing. *Quantum* **2018**, *2*, 49.
- (392) Cirq: a Python framework for creating, editing, and invoking Noisy Intermediate Scale Quantum (NISQ) circuits. <https://github.com/quantumlib/Cirq>, 2018; Accessed on Mon, November 5, 2018.
- (393) Killoran, N.; Izaac, J.; Quesada, N.; Bergholm, V.; Amy, M.; Weedbrook, C. Strawberry Fields: full-stack Python library for designing, simulating, and optimizing continuous variable quantum optical circuits. <https://github.com/XanaduAI/strawberryfields>, 2018; Accessed on Mon, November 5, 2018.
- (394) Mueck, L. Quantum reform. *Nat. Chem.* **2015**, *7*, 361–363.

- (395) Rigetti, C.; Alidoust, N. The Rigetti Quantum Advantage Prize. <https://medium.com/rigetti/the-rigetti-quantum-advantage-prize-8976492c5c64>, 2018; Accessed on Mon, November 5, 2018.
- (396) Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **2014**, *508*, 500–503.
- (397) Linke, N. M.; Maslov, D.; Roetteler, M.; Debnath, S.; Figgatt, C.; Landsman, K. A.; Wright, K.; Monroe, C. Experimental comparison of two quantum computing architectures. *Proc. Natl. Acad. Sci.* **2017**, *114*, 3305–3310.
- (398) Yoder, T. J.; Kim, I. H. The surface code with a twist. *Quantum* **2017**, *1*, 2.
- (399) Davidson, E. R.; Feller, D. Basis set selection for molecular calculations. *Chem. Rev.* **1986**, *86*, 681–696.
- (400) Jensen, F. Atomic orbital basis sets. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2013**, *3*, 273–295.
- (401) Nagy, B.; Jensen, F. *Reviews in Computational Chemistry*; Wiley Online Library, 2017; Chapter 3.
- (402) Havlíček, V.; Troyer, M.; Whitfield, J. D. Operator locality in the quantum simulation of fermionic models. *Phys. Rev. A* **2017**, *95*.
- (403) Fenwick, P. M. A new data structure for cumulative frequency tables. *Softw. Pract. Exp.* **1994**, *24*, 327–336.
- (404) McClean, J. R. et al. OpenFermion: The Electronic Structure Package for Quantum Computers. **2017**, arXiv:1710.07629 [quant-ph].