





Enabling the quantum leap

*Quantum algorithms
for chemistry and materials*



REPORT ON A NATIONAL SCIENCE FOUNDATION WORKSHOP
Alexandria, Virginia: January 21-24, 2019

Report on the NSF Workshop on Enabling Quantum Leap: Quantum algorithms for quantum chemistry and materials

Bela Bauer¹, Sergey Bravyi², Mario Motta³, and
Garnet Kin-Lic Chan⁴

¹Station Q, Microsoft Corporation, Santa Barbara, California 93106 USA

²IBM T. J. Watson Research Center, Yorktown Heights, USA

³IBM Almaden Research, San Jose, CA 95120, USA

⁴California Institute of Technology, Pasadena, CA 91125, USA

November 12, 2019

Acknowledgements

This report is derived from an NSF workshop held between January 22-24 2019 in Alexandria VA. Funding for the workshop and for the production of the report was provided by the National Science Foundation via award no. CHE 1909531. The workshop participants and website are provided below.

Workshop Participants

Workshop Co-Chairs:

- Garnet Kin-Lic Chan (California Institute of Technology)
- Sergey Bravyi (IBM T. J. Watson Research Center)
- Bela Bauer (Microsoft Station Q)

Participant List:

Chemistry and Molecular Science

- Victor Batista (Yale)
- Tim Berkelbach (UChicago)
- Tucker Carrington (Queens U.)
- Garnet Chan (Caltech)
- Gavin Crooks (Rigetti Computing)
- Francesco Evangelista (Emory)
- Joe Subotnik (UPenn)
- Haobin Wang (CU Denver)
- Dominika Zgid (UMich)

Materials Science, condensed-matter physics

- Bela Bauer (Microsoft)
- Bryan Clark (UIUC)
- Tom Deveraux (Stanford)
- Jim Freericks (Georgetown)
- Emanuel Gull (UMich)
- Barbara Jones (IBM Almaden)
- Austin Minnich (Caltech)
- Kirill Shtengel (UC Riverside)
- Steve White (UC Irvine)

Quantum Algorithms and Hardware

- Andrew Childs (UMD)
- Sergei Bravyi (IBM)
- Sophia Economou (Virginia Tech)
- Sabre Kais (Purdue)
- Guang Hao Low (Microsoft)
- Antonio Mezzacapo (IBM Yorktown)

Outreach, broader disciplines, education

- Daniel Crawford (Virginia Tech, MolSSI)
- Edgar Solomonik (UIUC)
- Takeshi Yamazaki (1QBit)

Other Accepted Participants: Christopher Chang, Alexandra Curtis, Sarom Leang, Mekena Metcalf, Anurag Mishra, Mario Motta, Petr Plechac, Ushnish Ray, Julia Rice, Yuan Su, Chong Sun, Miroslav Urbanek, Prakash Verma, Erika Ye.

Event Website: <https://sites.google.com/view/nsfworkshopquantumleap/>

1	Introduction	5
1.1	Summary of open problems and recommendations	6
1.1.1	Simulation challenges	7
1.1.2	Quantum algorithms	7
1.1.3	Software stack	9
1.1.4	Broader issues	9
1.2	What is a quantum computer?	10
1.2.1	Current quantum architectures	13
1.2.2	Building a circuit-based digital quantum computer . .	14
1.3	Broader issues	15
2	Simulation challenges in molecular and materials science	17
2.1	Quantum chemistry	17
2.2	Quantum molecular spectroscopy	20
2.3	Chemical quantum dynamics	22
2.4	Correlated electronic structure in materials	26
2.5	Dynamical quantum effects in materials	29
3	Challenges for quantum algorithms in quantum simulation	31
3.1	Overview of algorithms	31
3.2	Qubit representation of many-body systems	32
3.2.1	Ab initio electronic structure qubit representations . .	33
3.2.2	Electronic basis functions	36
3.2.3	Fermion-to-qubit mappings	38
3.2.4	Model and non-electronic problems	40
3.3	Quantum algorithms for ground and excited states	41
3.3.1	Preparing ground states along a prescribed path . . .	42

3.3.2	Variational state preparation and quantum eigensolver	44
3.3.3	Excited states	50
3.4	Phase estimation	51
3.5	Quantum algorithms for time evolution	54
3.5.1	Hamiltonian simulation problem	54
3.6	Finite-temperature algorithms	63
3.7	Hybrid quantum-classical methods	66
3.7.1	Quantum embedding methods	66
3.7.2	Other hybrid quantum-classical algorithms	68
3.7.3	Open questions	68
3.8	Benchmark systems	69
3.9	Reading out results	70
3.9.1	Equal-time measurements	70
3.9.2	Dynamical properties and Green's functions	71
4	Software stack	73
4.1	Quantum compilers and libraries	73
4.2	Classical simulation of (noisy) quantum circuits	74
4.3	Libraries for quantum simulation of chemistry and materials .	75
4.4	Interfacing and incorporating classical simulation software .	76
4.5	Summary	78
5	Broader issues	79
5.1	Diversity	79
5.2	Academia and industry	80
5.3	Education and training	81

5.4	Summary	82
-----	-------------------	----

1 Introduction

As we begin to reach the limits of classical computers, quantum computing has emerged as a technology that has captured the imagination of the scientific world and the public, as a potentially revolutionary advance.

The particular way in which quantum computing extends classical computing means that one cannot expect arbitrary tasks to be sped up by a quantum computer. However, early theoretical work showed that in certain structured problems – for example, the famous factoring problem underlying modern encryption – quantum computers offer exponential computational speedups over the best known classical algorithms [1]. Should these results translate broadly, quantum computing stands to transform society much like the invention of digital computing.

While for many years, the ability to execute quantum algorithms was only a theoretical possibility, recent advances in hardware mean that quantum computing devices now exist that can carry out quantum computation on a limited scale. Thus it is now a real possibility, and of central importance at this time, to assess the potential impact of quantum computers on real problems of interest.

One of the earliest and most compelling applications for quantum computers is Richard Feynman’s idea of simulating quantum systems with many degrees of freedom [2]. Such systems are found across chemistry, physics, and materials science. To precisely understand the potential impact of quantum computing in the scientific context of simulating systems of interest in chemistry, condensed matter physics, and materials science, we organized a workshop sponsored by the National Science Foundation, titled *Enabling the Quantum Leap: Quantum Algorithms for Quantum Chemistry and Materials*.

The workshop was attended by roughly 40 participants drawn from a wide range of institutions, from the disciplines of chemistry, physics, materials science, and computer science, and with representation from both academia and

industry, and consisted both of traditional presentations as well as multiple sessions of targeted small group discussion. At the beginning of the workshop, our stated goals were to:

1. Identify central problems in chemistry and materials science, in the areas of electronic structure, quantum statistical mechanics, and quantum dynamics, as target problems for quantum computers,
2. Identify areas for new quantum algorithms and critical aspects of current quantum algorithms to improve for applications to the above problems,
3. Bring together active and interested workers in the fields of quantum chemistry, quantum materials science, and quantum algorithms,
4. Explore how to educate a new generation of scientists who are trained both in the application areas of quantum simulation as well as in quantum algorithms,
5. Propose concrete mechanisms (e.g. funding) which will help advance the field and achieve convergence in efforts from different disciplines,
6. Interface academic and industry efforts in theoretical algorithms, software, and hardware, for chemistry and materials simulations,
7. Prepare a workshop report for the community which will provide recommendations on how to advance the field.

This is the report described in point 7. We have aimed to capture the flavor of the discussions and to provide context for the scientific questions; any omissions and errors are the sole responsibility of the authors. We hope that this report will prove useful in advancing the field of quantum simulations and algorithms for chemistry and materials science in the coming decade.

1.1 Summary of open problems and recommendations

Here, we summarize open problems and recommendations from all the forthcoming sections. Each class of open problems is accompanied by a set of keywords. We recommend referring to the appropriate sections for further details and context.

1.1.1 Simulation challenges

Quantum chemistry – The description of multi-reference quantum chemical problems at the same level of accuracy as single-reference quantum chemical problems (see [Fig. 2.1](#)). Quantitative modeling in large basis sets. Consideration of aspects beyond pure electronic structure, such as environment and dynamical effects. See [Sec. 2.1](#).

Quantum molecular spectroscopy and dynamics – Reaching the high accuracy of molecular spectroscopy. Non-rigid molecules. Basis representations for molecular spectroscopy. See [Sec. 2.2](#) and [Sec. 2.3](#).

Correlated electronic structure in materials – Strongly correlated electrons and spins, e.g. high-temperature superconductivity and its nearby phases, fractional quantum Hall effect and Moiré materials, interaction-driven phases in two-dimensional materials. Reaching the thermodynamic limit in quantum simulations and associated small energy scales that give rise to competing phases. See [Sec. 2.4](#).

Dynamical quantum effects in materials – Spectral functions, real-time non-equilibrium dynamics in strongly driven correlated materials. Quantum dynamics at long times and emergence of thermodynamics. See [Sec. 2.5](#).

1.1.2 Quantum algorithms

Qubit representations – Reducing redundancy in first quantization. Choice of electronic basis functions, handling the Coulomb cusp. Exploiting symmetries. Specialized representations for models. Basis representation for non-electronic-structure problems. Boson-to-qubit representations. See [Sec. 3.2](#).

Quantum algorithms for ground and excited states – Heuristic analysis of adiabatic state preparation and quantum imaginary time evolution in real problems. Better numerical and practical understanding of different ansatz in real problems. Relationship between different types of coupled cluster ansatz. Preparation methods for tensor networks. Representational power of deep tensor networks. Optimization in variational algorithms. Understanding and

improving robustness of variational methods against errors. Practical performance and relationship between different excited state methods. Proof-of-principle demonstrations. See [Sec. 3.3](#).

Phase estimation – Tradeoff between depth and spectral resolution. Error mitigation. Incorporation of better classical post-processing to extract more eigenvalues. Replacing Hamiltonian evolution by different types of evolution. See [Sec. 3.4](#).

Time evolution – Improving system-specific time evolution methods. Better bounds on the runtime using structure of the initial state and state-specific simulation methods. Quantum walks in place of time evolution. Practical reduction of runtime, improved compilation of time-evolution circuits. Space-time tradeoffs and depth reduction. Time evolution error mitigation without active error correction. Exploring connections to classical long-time evolution ideas such as long-time integrators. See [Sec. 3.5](#).

Finite-temperature algorithms – Benchmarking and deciding between many current proposed algorithms. More heuristic evaluation for problems of chemistry and physics, actual implementations on quantum devices. Avoiding long thermalization times and ancillae especially in the near-term era. Additional formulations for finite-temperature simulations. See [Sec. 3.6](#).

Hybrid quantum-classical methods – How to best adapt quantum algorithms to classical embedding frameworks. Hamiltonian representations of actions. More efficient evaluation of Green's functions and density matrices. New quantum-based embedding ideas, such as based on properties of quantum circuits. Better feedback between quantum and classical parts of the algorithm, improving classical optimization. See [Sec. 3.7](#).

Benchmark systems – Benchmarks to compare quantum algorithms against each other. Benchmarks to compare against the best classical methods. Including data for multiple observables. Systematic molecular and material benchmarks where complexity can be tuned. See [Sec. 3.8](#).

Reading out results – Reducing the number of measurements e.g. for variational methods. Reorganizing terms. Exploiting structure in Green's functions. Avoiding the need to re-prepare the initial state. See **Sec. 3.9**.

1.1.3 Software stack

Quantum compilers, libraries, circuit simulations – Developing libraries of common subroutines. Optimal compilation into gates. Fast general-purpose quantum circuit simulators. Specialized quantum circuit simulators. Approximate contraction methods, potentially using domain knowledge. Cross-platform operation of different quantum software libraries. Improved tools for near-term experiments. See **Sec. 4**.

Interfacing to classical simulation software – Standardizing interfaces; flexible middleware; better packaging of specialized codes; modularization of large codes. See **Sec. 4**.

1.1.4 Broader issues

Diversity – Accumulating best practices. Starting from leadership. Improving undergraduate internships. Engaging social sciences to quantify progress. See **Sec. 5**.

Academia and industry – How to engage with quantum hardware. Understanding intellectual property rights and negotiations. Balancing organizational sizes and power differential. Collective organizations for academia. See **Sec. 5**.

Education and training – Encouraging basic course requirements in quantum mechanics. Incorporating basic quantum information and quantum computing into introductory QM. Introducing quantum into introductory CS. Graduate level programs with cross-disciplinary course requirements. Support for multi-institutional efforts in education and curricula development. Working with industry in quantum education. See **Sec. 5**.

1.2 What is a quantum computer?

How is it relevant to quantum simulation?

A quantum computer is a device which computes by manipulating quantum information [3, 4, 5, 6, 7, 8]. The basic unit of quantum information, called a *qubit*, is synonymous with a two-level quantum system. Denoting the two basis states of a qubit as $|0\rangle, |1\rangle$, the general single-qubit state may be a superposition $c_0|0\rangle + c_1|1\rangle$ where c_0, c_1 are complex numbers. For n qubits, there are 2^n basis states, which can be enumerated as the bitstrings $|00 \dots 0\rangle, |00 \dots 1\rangle$, etc. To specify a general quantum state, one must specify a complex coefficient (amplitude) for each basis state, e.g. $c_{00\dots 0}, c_{00\dots 1}$ etc. The exponential number of amplitudes needed to specify the state of n qubits stands in contrast to the linear amount of information needed to encode a single bitstring, the state of n classical bits.

Measuring the qubits (in this basis, often referred to as the computational basis), collapses the state onto the measurement outcome, which is a single (random) bitstring x that appears with the probability $|c_x|^2$. Despite this probabilistic interpretation, the fact that quantum amplitudes can be complex prevents one from treating quantum states as classical probability distributions, which is at the heart of the quantum phenomenon of *entanglement*, the correlations in a system that cannot be generated by a classical distribution of states. The possibility of creating entangled states in a space of exponentially large dimension and manipulating these states by exploiting their constructive or destructive interference is the source of the computational power of a quantum computer.

A convenient way to describe the manipulation of quantum states in the context of quantum computation is via few-qubit operations called gates. The action of each gate can be viewed as a unitary time evolution of the n -qubit system under a suitable Hamiltonian that acts non-trivially only on a few qubits (usually one or two). Any quantum computation can be expressed by a sequence of elementary gates, called a *quantum circuit*, applied to an initial basis state, e.g. $|00 \dots 0\rangle$, and followed by the measurement of some set of the qubits.

The exponential separation between quantum and classical information does not simply mean that quantum computers can compute answers to problems “in parallel” with an exponential speedup, for example, by storing multiple different solutions in the many amplitudes. This is because reading out from a

quantum computer destroys the state, and thus to harness the power of quantum information, a quantum algorithm must coordinate interference between the amplitudes such that useful information can be read out with high confidence without many measurements.

The interest in quantum computing for quantum simulations of molecules and materials stems from the fact that in many cases, the chemistry and physics of molecules and materials is best described using quantum mechanics. Thus, the state of a many-particle molecule also encodes quantum information, and as the number of atoms increases, similarly can require an exponentially large number of classical bits to describe. This means that in the worst case, quantum simulation is exponentially hard on classical computers. This is the motivation for Feynman’s famous observation that “Nature isn’t classical, dammit, and if you want to make a simulation of nature, you’d better make it quantum mechanical” [2].

A moment’s reflection, however, suggests that the potential quantum advantage for a quantum computer in quantum simulation is nonetheless subtle. For example, if it were indeed impossible to say anything about how atoms, molecules, or materials behave, without using a quantum computer, there would be no disciplines of chemistry, condensed matter physics, or materials science! Decades of electronic structure and quantum chemistry simulations suggest that reasonable, and in some cases very high accuracy, solutions of quantum mechanics can be obtained by classical algorithms in practice. Quantum advantage in quantum simulation is thus problem-specific, and must be tied both to the types of questions that are studied, as well as the accuracy required.

We can look to theoretical quantum computer science to better understand the power of quantum computers in quantum simulation. The natural problem to solve on a quantum computer is the time evolution of a quantum system given some initial state,

$$\textit{Quantum dynamics: } i\partial_t|\Psi(t)\rangle = \hat{H}|\Psi(t)\rangle. \quad (1.1)$$

This problem is representative for the complexity class BQP, i.e. it is of polynomial cost on a quantum computer and believed to offer a clear separation from the classical case (an exponential quantum speedup has been rigorously proven only in the query complexity setting [10]). However, it is necessary to prepare the initial state, which may be difficult by itself. In particular, preparing a low-energy state may be challenging, which naturally leads to considering other

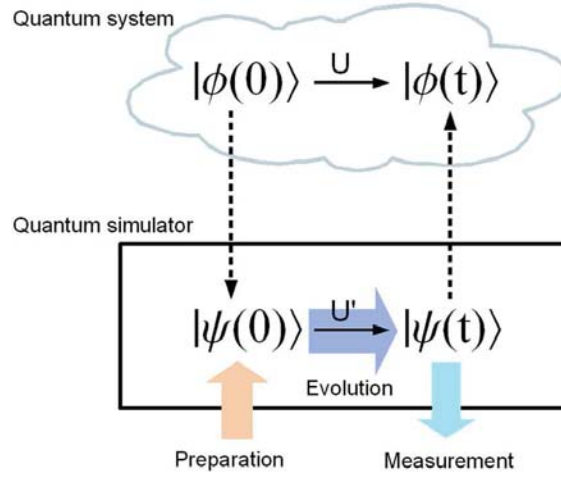


Figure 1.1: Schematic of a quantum simulation of quantum dynamics. A quantum simulator (bottom) is prepared in an initial state $|\psi(0)\rangle$, that is a representation of the initial state $|\phi(0)\rangle$ of the actual system of interest (top). The simulator is manipulated by a unitary transformation \hat{U}' that is an encoding of the real-time evolution of the system of interest, and the final simulator state $|\psi(t)\rangle$ is measured, yielding information about the dynamics of the original system. From [9].

important problems,

$$\textbf{Ground state: } \hat{H}|\Psi_0\rangle = E_0|\Psi_0\rangle; \quad E_0 = \min_{|\Psi\rangle} \langle\Psi|\hat{H}|\Psi\rangle \quad (1.2)$$

$$\textbf{Thermal averages: } \langle\hat{A}\rangle = \frac{\text{Tr}[\hat{A}e^{-\beta\hat{H}}]}{\text{Tr}[e^{-\beta\hat{H}}]} \quad (1.3)$$

Ground state determination lies in complexity class QMA [6], a class of problems not known to be efficiently solvable (in general) on a quantum computer. This also means that thermal averages cannot in general be computed efficiently on a quantum computer, since in the limit of zero temperature, this problem reduces to ground-state determination. Although it is the case that there are many physical ground state and thermal problems that are not so hard to solve in practice (as demonstrated by efficient classical simulation of many problems) and similarly many initial states of interest in quantum dynamics that are easy to prepare, the above tells us that in a rigorous sense, we do not have a straightforward theoretical guarantee that a quantum com-

putational advantage can be naturally achieved for the central problems in quantum simulation.

Given the limits to the guidance that can be provided by rigorous computational complexity results, it is clear that to understand quantum advantage in chemistry, condensed matter physics, and quantum materials science, we must be guided by actual empirical data in the form of numerical and theoretical experiments with quantum algorithms and quantum devices on simulation problems of interest. This requires making progress on both theoretical and practical questions of quantum simulations, ranging from the basic algorithms and choices of encoding and representation to issues of circuit compilation, readout, and mapping to specialized hardware. A central purpose of this report is to provide a community perspective on what the relevant chemical and materials problems are today; to give a snapshot of the limitations of classical methods for these problems; and in these contexts to understand the strengths, weaknesses, and bottlenecks of existing ideas for quantum algorithms, and where they need to be improved, both in terms of fundamental theoretical aspects as well as practical methods of implementation.

1.2.1 Current quantum architectures

The idea of using a quantum mechanical device to perform a computation was first considered in earnest by Richard P. Feynman in a famous lecture in 1982 [2]. His suggestion was to build a lattice of spins with tunable interactions. He conjectured that by appropriately tuning these interactions, such a system could be made to imitate the behavior of any other (bosonic) quantum system with the same dimensionality, and thus could serve as a way to compute the properties of some other system that one would like to study. This idea, which is often referred to as *analog quantum computation*, is still very much alive today and embodied in the field of cold atomic gases and related techniques, which have made great progress in simulating complicated physics of strongly correlated systems in a controlled environment. A general schematic of the idea is shown in Fig. 1.1.

However, such systems come with some limitations: for example, it appears unlikely that a system of spins can naturally emulate a system of fermions, or that a two-dimensional array of spins could emulate a three-dimensional material; in other words, a given analog quantum computer can only simulate a certain set of physical systems. Furthermore, it is itself an experiment with limited accuracy; while it may be easier to perform than experiments on the

original system of interest, the accuracy nonetheless remains limited. A more general approach, which is the focus of this report, is *digital quantum computation*. Here, very much in analogy to the classical computers we are used to, one considers a system of quantum registers – qubits – that are controlled through some set of instructions - the quantum gate operations. One can show that with a finite (and indeed relatively small) set of such gate operations, one can in principle generate arbitrary quantum evolution to arbitrary precision! In other words, every problem in the complexity class BQP can be mapped into a quantum program, i.e. the quantum circuit composed of such gates.

A third model of quantum computation is *adiabatic quantum computation*. Here, the computation is encoded into a time-dependent Hamiltonian, and the system is evolved slowly to track the instantaneous ground state of this Hamiltonian. It can be shown that this model is equivalent to circuit-based digital quantum computation [11], but it is usually considered to be less practical. However, a restricted version of it, adiabatic quantum optimization [12], has gained some popularity. Here, a classical optimization problem is encoded into a quantum Hamiltonian, to which one adds some non-commuting terms to endow the system with non-trivial quantum dynamics. If one then slowly turns off the quantum terms, the optimal solution to the classical problem should be obtained. In practice, however, one may have to go impractically slowly for this to be true; the detailed analysis of this approach is a complex problem that is not covered in this report.

1.2.2 Building a circuit-based digital quantum computer

Returning to circuit-based quantum simulation, what is the status of the field today? The natural enemy of quantum computation is decoherence of the qubit, i.e. the tendency of the stored quantum state to decay into a classical state. After decades of research, a number of qubit technologies [13, 14], for example superconducting qubits [15] and ion traps [16], have reached the point where small devices of a few dozen qubits can be sufficiently isolated from decoherence to execute non-trivial quantum algorithms of a few tens to hundreds of gates. This has been termed the *noisy intermediate-scale quantum* (NISQ) era [17]. With these devices, the first demonstrations of *quantum supremacy* appear to be on the horizon. For this, an artificial but well-defined problem that is intractable classically is solved on a quantum computer. However, these problems do not generally have practical relevance, and

a more important and challenging question is when a scientifically or commercially relevant problem can be solved on a quantum computer.

To address large-scale problems, it is necessary to correct errors that occur on physical qubits during the computation. While the no-cloning theorem prevents error correction by simple redundancy, it can be shown that **quantum error correction** is possible nonetheless by encoding a single qubit into an entangled state of many qubits [19], to use suitable measurement to detect errors occurring in the system and to apply suitable unitary transformations correcting such errors without disturbing the information encoded in the system (see Fig. 1.2). This leads to the very important distinction between **physical** and **logical** qubits. The latter are error-corrected and encoded in the state of many physical qubits. Quantum algorithms are performed on the logical qubits, and the error-correction scheme translates the operations on logical qubits into physical operations. This incurs significant overhead: depending on the error rate of the underlying physical qubits and the target error rate of the logical qubits, one may need hundreds or even thousands of physical qubits to realize a single logical qubit. Therefore, when evaluating the capabilities of some qubit platform with respect to an algorithm, one must be careful to include the cost of encoding logical qubits into physical qubits to achieve the required error rates.

1.3 Broader issues

Progress in quantum algorithms for quantum simulation requires training a scientific workforce. Quantum simulation research is, by its very nature, an interdisciplinary field. To carry out research in this discipline requires understanding quantum algorithms as well as chemistry and materials science. What should be the mechanism to educate such a workforce? How do we ensure best practices in hiring? How do we ensure collaboration between academic software and theory efforts and industry hardware and software groups? An important component of the workshop was also to discuss these issues and provide community feedback from both academia and industry on the desired best practices for new collaborations, workforce training, and engaging the scientific and broader population as a whole. We return to these questions in Sec. 5.

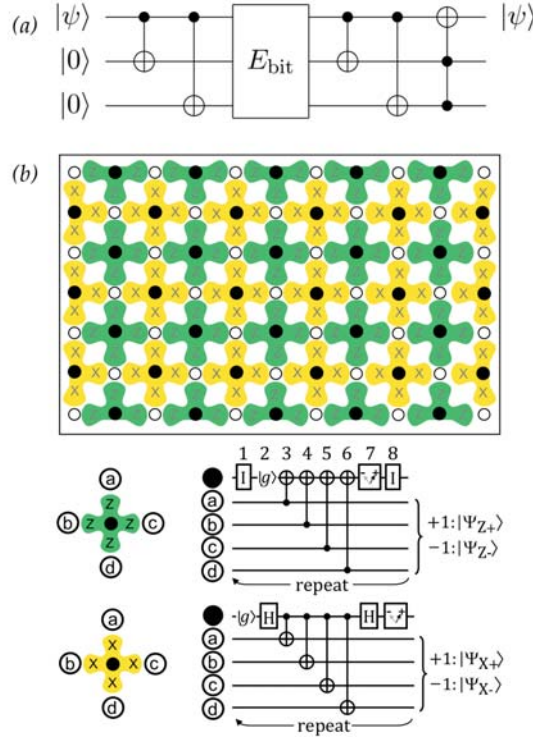


Figure 1.2: Examples of error correction schemes. (a) Quantum circuit implementing the bit-flip code. The state of a single logical qubit, $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is encoded in the state of three physical qubits as $\alpha|000\rangle + \beta|111\rangle$. The system is then sent through a channel $E_{\text{bit}} = \mathcal{E}^{\otimes 3}$, where \mathcal{E} flips a qubit with probability p (X error). Measuring the Z operator on the last two qubits permits to determine whether a qubit has been flipped by the channel, and to correct such an error without corrupting the transmitted state. (b) Schematic of the surface code for quantum error correction and quantum fault-tolerance. Top: a two-dimensional array implementation of the surface code, data qubits are open circles and measurement qubits are filled circles, with measure- Z qubits colored green (dark) and measure- X qubits colored orange (light). Middle: sequence of operations (left), and quantum circuit (right) for one surface code cycle for a measure- Z qubit, to detect sign flip errors. Bottom: similar, but for measure- X qubits. From [18].

2 Simulation challenges in molecular and materials science

In this section we describe a representative, but certainly non-exhaustive, set of scientific problems of relevance to quantum simulation.

2.1 Quantum chemistry

Quantum chemistry is concerned with determining the low-lying eigenstates of the electronic Hamiltonian of a molecule. The eigenstates are determined for fixed sets of nuclear positions, i.e. within the Born-Oppenheimer approximation. Determining the main features of the resulting potential energy surface, i.e. the electronic energy as a function of nuclear positions, its minima and saddle points, is key to understanding chemical reactivity, product distributions, and reaction rates.

There exists a wide range of quantum chemical methods with different accuracy and speed tradeoffs, ranging from density functional methods that routinely treat thousands of atoms or more [20], to high-level many-electron wavefunction methods, such as coupled cluster theory, that can attain chemical accuracy of 1 kcal/mol and better, on systems of tens of atoms [21, 22]. However, most methods in quantum chemistry are most accurate for problems where there is a dominant electronic configuration, a subset of the quantum chemistry problem known as the *single-reference problem*. Single-reference quantum chemistry is found in the ground-states of many simple molecules (e.g. hydrocarbons), but in many molecular excited states, at stretched bond geometries, and in transition metal chemistry, multiple electronic configurations can come into play, which is referred to as *multi-reference quantum chemistry*. Despite much work (and progress) in extending quantum chemistry to multi-reference situations, the attainable accuracy in molecules with

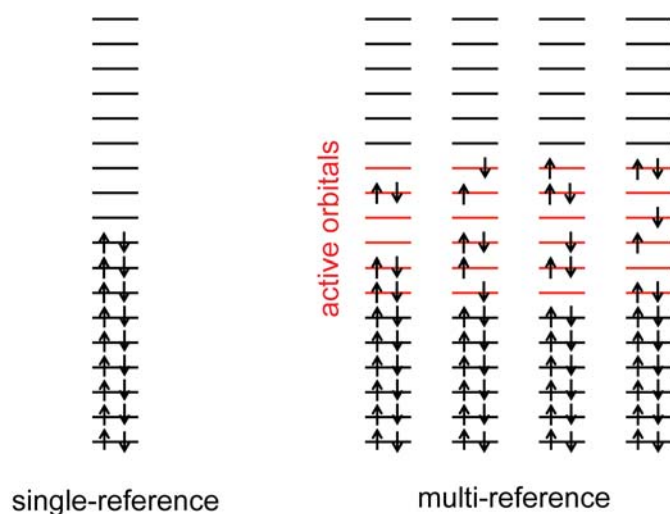


Figure 2.1: Single-reference (left) and multi-reference (right) wavefunctions. The former is qualitatively described by a single Slater determinant, the latter by a linear combination of a potentially large number of Slater determinants. Often such determinants correspond to different configurations of electrons in an “active space” of orbitals.

more than a few atoms is significantly lower than in the single-reference case. Some examples of multi-reference quantum chemical problems include:

- *The chemistry of enzyme active sites.* Such active sites can involve multiple coupled transition metals, famous examples being the 4 manganese ions in the oxygen evolving complex [25], or the 8 transition metals in the iron-sulfur clusters of nitrogenase [26], shown in Fig. 2.2. They pose some of the most complicated multi-reference quantum chemistry problems in the biological world. Combined theoretical and experimental studies, primarily at the level of density functional theory, have proven successful in unravelling many structural and electronic features of such enzyme active sites [27, 28, 29]. However, a detailed understanding of the interplay between spin-coupling and delocalization between metals, which requires true multi-reference quantum chemistry and is needed to interpret aspects of experimental spectroscopy, remains in its infancy [30, 24, 23, 23, 31, 32].
- *Transition metal nanocatalysts and surface catalysts.* Similarly to




Figure 2.2: Left: Iron-sulfur clusters associated with different oxidation states (P^N , Syn (a synthetic model of P^N), P^{1+} and P^{ox}) of the P cluster of the nitrogenase enzyme, from [23]. Right: Mn_4Ca core of the oxygen evolving complex of photosystem II, from [24].

enzyme active sites, simulating the mechanism of action of synthetic heterogeneous catalysts remains a major challenge. While density functional theory has been widely employed, predictions of even basic quantities such as the adsorption energy of small molecules are unreliable [33, 34]. While not all such problems are expected to be multi-reference in character, even the single-reference modeling of such chemistry, at a level significantly beyond density functional theory, is currently challenging or impossible. In addition, multi-reference effects are expected to play a role in certain catalysts, such as transition metal oxides, or at intermediate geometries in reaction pathways [35, 33, 36, 37].

- **Light harvesting and the vision process.** The photochemistry of conjugated organic molecules is the means by which Nature interacts with light. Some prominent examples of such natural conjugated systems include the carotenoid and chlorophyll pigments in the light-harvesting complex of plants [38, 39], as well as the rhodopsin system associated with vision [40, 41]. While describing the interaction with light is not purely a question of electronic structure, as it involves the quantum dynamical evolution of quantum states, the quantum chemical questions revolve around the potential energy surfaces of the ground- and excited-states, and the influence of the environment on the spectrum [42]. These questions are currently challenging due to the size of the systems in-

volved as well as the varying degree of single- and multi-reference character in many of the conjugated excited states [43, 44, 45].

The basic quantum simulation problem is the ground-state (or low-energy eigenstate) problem for the electronic Hamiltonians, and the basic metric is whether ground-state or low-energy eigenstate quantum algorithms yield more accurate energies (for comparable computational resources) than the best classical algorithms, for the problem sizes of interest. Proof-of-principle demonstrations could be carried out in simplified models of the above problems (e.g. in small active spaces of orbitals). However, to make real progress, one should also consider more quantitative models, which requires treating a large number of electrons in a large number of orbitals; at minimum, tens of electrons in hundreds of orbitals. This poses new challenges for ground-state algorithms, and raises questions of how best to represent and encode the resulting Hamiltonians and states. In addition, there are many aspects of the chemical problems beyond the modeling of the electronic wavefunctions, for example, to treat environmental, solvent, and dynamical effects [46]. These will require interfacing the quantum simulation with other classical simulation methods. Finally, although the above examples have focused on multireference and strongly correlated quantum chemistry, weakly correlated chemistry itself becomes hard to model classically when the number of degrees of freedom is very large [47]. These may also be interesting to target with quantum algorithms when sufficiently large quantum machines are available.

2.2 Quantum molecular spectroscopy

High-resolution gas-phase rovibrational spectroscopy provides an extremely precise experimental probe of molecular structure [48]. Such spectroscopy is important not only for the fundamental understanding of small molecules and the quantum control of atomic and molecular states, but also to provide insight into the basic chemical processes and species involved, for example, in atmospheric chemistry [49] and in astrochemistry [50]. In larger molecules, with more than a few atoms, even the low-energy rovibrational spectrum contains many peaks which cannot be interpreted without theoretical simulation [51]. The theoretical goal is to compute the eigenstates of the *nuclear Schrödinger equation* [52]. However, unlike the electronic structure problem, there are several challenges even in setting up the best form of the nuclear Schrödinger equation to solve. The first is that the nuclear Hamiltonian (in particular, the




Figure 2.3: Dominant families of configurations of fluxional CH_5^+ . From [60].

nuclear-nuclear interactions) are not known a priori because the interactions are mediated by the electrons. This nuclear potential energy term must instead be determined from quantum chemical calculations at fixed nuclear geometries and then fitted to an appropriate functional form; this requires a large number of high accuracy quantum chemistry calculations. The second is that nuclear ro-vibrational motion is often far from harmonic and not well approximated by simple mean-field theories, unlike many electronic structure problems. Thus there is a need for a proper choice of curvilinear nuclear coordinates that decreases coupling in the nuclear potential energy (e.g. in a harmonic system, normal modes are such a choice of coordinates) while retaining a simple form for the kinetic energy operator, and which also exposes the symmetry of the molecular system.

Once the nuclear Schrödinger equation has been properly formulated, one then faces the problem of representing the eigenstates. While methods such as diffusion Monte Carlo have made progress on vibrational ground-states [53], spectroscopy involves transitions to excited states. In this setting, tensor factorization [54, 55] and other approaches [56, 57] have been explored to approximate the rovibrational wavefunctions [58]. However, the high dimensionality and spectral congestion, requiring resolution between peaks on the scale of 1 wavenumber, proves extremely challenging [59, 55]. Some famous examples include:

- **Spectra of floppy molecules.** Floppy molecules are by their nature very anharmonic and thus far from a simple vibrational description. CH_5^+ is a prototypical floppy molecule (see Fig. 2.3): the five hydrogen atoms move around the central carbon and the molecule has almost no structure in the traditional sense [61, 62, 63, 64, 55].
- **Hydrogen bonded clusters.** Another vibrational problem with large anharmonicity is found in hydrogen bonded clusters, such as in the spec-

troscopy of water molecules and protonated water clusters [65, 66]. The hydrogen bond network is fluid and even small clusters can transition between many different minima on the potential energy surface [67]. Resolving the peaks and tunnelling splittings is important for interpreting water spectra in the atmosphere, as well as in understanding reaction mechanisms in water. Further, the spectroscopy of molecules with intermolecular hydrogen bonds, such as the malonaldehyde molecule, has also posed long-standing challenges for the field [68, 69, 70, 71].

From a quantum algorithms perspective, although there are similarities with the quantum chemistry problem (in particular one is interested in low-energy eigenstates), there are significant differences. One important difference is that the Hamiltonian is no longer of simple two-particle form due to the effective nuclear-nuclear interaction and typically includes important three- and four-mode terms. Also, one is often interested in an order of magnitude more states (e.g. hundreds of excited states) than in the electronic structure problem. All these features are sufficiently distinct from the usual quantum chemical scenarios that quantum algorithms are likely to require additional innovation to be useful in the nuclear problem. Some steps in this direction have recently appeared [72, 73, 48]. One simplification is that many nuclei are distinguishable avoiding the need to consider indistinguishable particles. The lack of a good mean-field starting point together with the various technical complications means that one can find relatively small systems (in terms of the Hilbert space size) where classical methods already have trouble; for example, a minimal quantum model of the CH_5^+ molecule can be formulated as a 12 dimensional problem with 10 basis functions per mode [64].

2.3 Chemical quantum dynamics

Chemical quantum dynamics is another important target for quantum simulation [74]. This field is concerned with modeling time-dependent electronic and nuclear quantum effects in molecules (as distinct from computing the time-independent electronic or nuclear eigenstates in quantum chemistry and quantum molecular spectroscopy). *Quantum molecular dynamics* is primarily concerned with the nuclear motion and describes the rates of chemical processes as well as the dynamical interaction of molecules with light, as involved in spectroscopy and quantum control. However, with the development of short

X-ray pulses direct experimental access to electron dynamics in molecules is now also available.

Currently these dynamical simulations are challenging. For example, nuclear motion is poorly described by mean-field theory and the classical limit is often a better starting point, but offers no zeroth order description of quantum effects. Thus classical simulations of quantum dynamics either invoke methods based on the classical limit that scale to large systems but which are difficult to systematically improve (such as approximate path integral methods [75], [76]), or methods which model the wave-function dynamics or the path integral accurately for a small number of degrees of freedom, but which are not scalable due to dimensionality or the dynamical sign problem [77, 78].

A subfield of quantum molecular dynamics, but one of important chemical interest, is the description of *non-adiabatic quantum effects* [74, 79]. At nuclear configurations where different electronic surfaces approach each other, the Born-Oppenheimer approximation can break down and the quantum behavior of the nuclei, coupled indirectly via the electrons, is enhanced. The associated quantum non-adiabatic effects govern non-radiative energy relaxation via the crossing of electronic surfaces (so-called conical intersections) and are thus central in describing energy transfer. The faithful description of non-adiabatic quantum effects requires the simultaneous treatment of quantum electrons and quantum nuclei. The complexity of this problem together with the often large system sizes where non-adiabatic effects are of interest means that current classical methods rely on simple heuristic approximations, such as surface hopping [80], for which a rigorous quantum formulation is lacking. Examples of relevant chemical problems in the area of chemical quantum dynamics include:

- *Proton coupled electron transfer* (PCET) [81]. PCET is known to be an important mechanism in catalysis and energy storage: electrons are transferred at lower overpotentials when thermodynamically coupled to proton transfer. Examples range from homogeneous catalysts [82, 83, 84, 85] and heterogeneous electrocatalysts [86, 87] to enzymes that perform PCET, including soybean lipoxygenase (that catalyzes the oxidation of unsaturated fatty acids) [88], photosystem II (in the tyrosine oxidation step) [89, 90] and the redox-leveling mechanism of catalytic water oxidation [87]. While semiclassical predictions are often good enough for describing electron transfer, the quantum nature of molecular dynamics is paramount with PCET because of the quantum nature of the proton [88, 91] as evidenced in kinetic isotope effects (referring to

the ratio of the proton to deuteron reaction rates) which can be larger than 80 [88]. As such, classical mechanics is not sufficient, and quantum simulations of PCET would be extremely helpful for making predictions in large (especially biological) systems.

- *Vibrational dynamics in complex environments.* For many systems of interest, vibrational spectroscopy is the key tool available for characterization. There is overlap with the problems in section 2.2 but the questions here focus on larger scale systems and condensed phase problems, where the line-shapes as well as frequencies are important, and the system size limits the use of fully quantum formalisms. From the librations of water to the high frequency motion of C-H bonds, it would appear that all of physical chemistry uses nuclear vibrational frequencies to characterize complex systems; for instance, it is routine nowadays to use Stark shifts of reporter molecules with large vibrational frequencies to characterize electric fields either within biological enzymes [92] or at electrochemical interfaces [93, 83]. Nevertheless, because of computational limitations, the standard approach today for modeling vibrational dynamics in large complex environments is to invoke a very old flavor of theory: Kubo theory [94]. One focuses on energy gaps and uses a semiclassical expansion of the line-shape. Obviously, this approach relies on diagonalization of the quantum subsystem and as such is limited to systems with only a few quantum states of interest; more generally, for systems with many interacting quantum states – for instance, H/J molecular aggregates [95] – one can calculate vibrational states only by brute force, which is severely limiting.
- *Plasmonic chemistry.* One of the most exciting areas today in physical chemistry is the possibility of using metal particles (with large cross sections) as a tool to absorb light and, with the resulting plasmonic excitations, initiate “plasmonic chemistry”. Already, there are a few examples in the literature of hot plasmonic chemistry, including the dissociation of hydrogen molecules [96]. To better understand this arena, however, one of the key questions is: how do we characterize plasmonic excitations? While classical descriptions of plasmons are easy to obtain, quantum descriptions are necessary if we are going to model quantum processes, e.g. electron transfer. And yet, by definition, plasmonic excitations are correlated excitations, where many electrons push against each other and force a collective motion of the electron cloud of the particle. Thus, perhaps not surprisingly, characterizing plasmons fully quantum mechanically has proven to be extremely difficult.




Figure 2.4: Qualitative phase diagram of cuprate high-temperature superconductors, which has challenged theory and simulation in condensed matter for decades. Figure adapted from Ref. [98].

While quantum dynamics is in principle an ideal simulation problem for a quantum computer, the quantum simulation of quantum molecular dynamics entails several practical challenges. Much of this may be viewed as an issue of representation. As in the problem of quantum molecular spectroscopy, the nuclear quantum Hamiltonian contains complicated interactions which must be tabulated or calculated on the fly [97]. In addition to this, the dynamical quantum state involves near continuum degrees of freedom, posing a challenge for the standard discretizations of Hilbert space considered in quantum algorithms. Finally, typical spectroscopic observables may be accessible to relatively simple treatments, not requiring the full fidelity of the quantum wavefunction evolution. These technical issues mean that implementing quantum molecular dynamics with a quantum advantage is likely to remain challenging in practice, despite the favorable theoretical complexity on a quantum device.

2.4 Correlated electronic structure in materials

The goal of electronic structure calculations in materials is to determine their low-energy properties. There is some overlap in methods and ideas between the materials electronic structure problem and the problem of quantum chemistry. When electron-electron interactions are weak, the low energy material properties can normally be described by computing the band structure using *density-functional theory* and one of the many popular density functionals, such as the local density approximation (LDA) or generalized gradient approximations (GGA). However, in some materials, commonly referred to as *strongly correlated*, the electron-electron interactions fundamentally alter the behavior and such an effective non-interacting description is no longer appropriate. A paradigmatic example are Mott insulators, which appear as conductors in band structure theory but in fact become insulating due to electron-electron interactions. While the mechanism behind the insulating behavior of Mott insulators is well-understood, for many other phenomena in strongly correlated systems the underlying microscopic mechanism is not fully known, let alone a quantitative and predictive theory of the associated physics with material specificity. Some famous examples of such problems include:

- Originally discovered in 1986, *high-temperature superconductivity* has eluded a complete theoretical explanation to date. The experimental phase diagram, which is sketched in Fig. 2.4, is accurately characterized experimentally across several materials. While the general properties of the superconducting phase itself are relatively well characterized, the mechanism driving superconductivity is not yet fully elucidated. Also, two nearby regimes, the *pseudogap* and *strange metal* phase, continue to puzzle theorists [98]. In both cases, their nature as well as their precise relation to the superconducting phase are not understood. The strange metal phase (also known as non-Fermi liquid) exhibits behavior inconsistent with a simple weakly interacting metal even at high energies, and has motivated a whole area of research on exotic metallic systems [99]. Meanwhile, the pseudogap phase exhibits several competing ordering tendencies [100], which are extremely challenging to resolve in numerical methods because most methods naturally favor a particular ordering pattern, thus making it challenging to disentangle physical effects from method biases.
- The non-Fermi liquid behavior exemplified in the strange-metal phase of cuprates appears also in other classes of materials, such as *heavy*

fermion compounds and *fermionic systems near criticality* [101, 102]. In general, these systems are amenable to classical simulation only at special points where quantum Monte Carlo methods do not suffer from the infamous sign problem [103].

- Since quantum fluctuations are enhanced in lower dimensions, *two dimensional systems* have long been of central interest in strongly correlated physics. Many material systems realize effectively two-dimensional physics, including two-dimensional electron gases in semiconductor heterostructures, where the *integer and fractional quantum Hall effect* were first discovered [104], layered materials (including cuprate high temperature superconductors), graphene [105] and transition-metal dichalcogenides (TMDs). More recently, it has been found that so-called *Moiré materials* exhibit rich phase diagrams due to strong interactions, including exotic superconductivity [106] and exotic topologically non-trivial phases. A paradigmatic example is twisted bilayer graphene [107], which consists of two graphene layers that are slightly twisted with respect to each other. This leads to a Moiré pattern with a very large unit cell, which effectively quenches the kinetic energy (i.e., leads to almost flat bands) and drastically enhances the effect of Coulomb interaction.
- *Frustrated spin systems* have long been an important topic especially for numerical simulations in condensed-matter physics. These systems potentially realize a host of high non-trivial phases, in particular topological and gapless spin liquids [108]. They have historically been the testbed for computational methods such as tensor networks and variational methods. As such, they appear as good test cases also for quantum simulations. Furthermore, recent developments in particular in materials with strong spin-orbit coupling have opened the door on a variety of new materials that may exhibit exotic topological phases, and in particular realize a non-Abelian spin liquid [109, 110].

While many methods have been developed to accurately include electron-electron interactions, their scope generally remains limited. For example, tensor network methods have revolutionized the study of one- and to a limited extent two-dimensional effective models for magnets and itinerant fermions. However, these methods so far have not been successfully applied to more realistic models, and in particular in three dimensions. On the other hand, quantum embedding methods such as the dynamical mean-field theory (DMFT) and its many cousins can capture interaction effects in three-dimensional systems, including multi-band systems. However, they require some approximations to




Figure 2.5: Illustration of pump-probe spectroscopy using X-ray probes. Figure adapted from Ref. [111].

the correlations of the state; for example, in its simplest form, DMFT disregards momentum dependence of the electron self-energy. While many improved variants of these embedding methods exist, their accuracy is often difficult to control, and so far they have not been applied to realistic models without further approximations. Finally, quantum Monte Carlo methods have been extremely successful for bosonic systems and unfrustrated spin systems, but the sign problem hinders their application to frustrated or fermionic systems (away from special points) without other uncontrolled approximations.

From the perspective of quantum algorithms, the materials electronic structure problem is both simpler and more difficult than the quantum chemistry problem. Some ways in which it is simpler include the fact that often very simple Hamiltonians describe the main physics, as well as the potential presence of translational invariance. A major way in which it is more complicated is the fact that one needs to treat systems approaching the thermodynamic limit, which involves a very large number of degrees of freedom. This not only increases the number of qubits required but also heavily impacts the circuit depth of algorithms, such as state preparation. The thermodynamic limit can also lead to small energy scales for excitations and energy differences between competing phases. For these reasons, it remains to be understood whether the quantum algorithms of relevance to quantum chemistry are those of relevance to materials electronic structure.

2.5 Dynamical quantum effects in materials

Many experiments on condensed-matter systems do not probe the equilibrium properties of the system, but rather dynamical properties. For example, the main workhorse of mesoscopic quantum physics is *electron transport*, i.e. the response of the system when it is coupled to electron reservoirs and a voltage is applied [112]. Likewise, material properties are often probed through scattering experiments, such as neutron scattering or angle-resolved photoemission spectroscopy (ARPES) [113], which probe dynamical properties such as the structure factor or spectral function. Going beyond spectral properties, the *non-equilibrium real-time dynamics* of quantum systems have increasingly come into focus, both because of experiments that can probe quantum dynamics at atomic scales and because of fundamental interest in equilibration of quantum systems. Experimental setups that can probe ultra-fast dynamics in materials include, for example, free-electron lasers [114, 115] as well as other pulsed laser systems. These allow the application of experimental techniques, such as pump-probe spectroscopy [116, 111], to provide novel insights into the behavior of correlated quantum systems.

On the other hand, *cold atomic gases* [117] provide a highly controllable environment that allows systematic exploration of quantum dynamics even in the strongly interacting regime [118], see Fig. 2.6. A key advantage is that one can engineer the evolution of the system to closely follow a target model; this approach is also referred to as analog quantum simulation (see also Sec. 1.2.1). However, classical simulation still plays a crucial role in establishing the accuracy of cold atom setups.

From a conceptual point of view, a central question has become *the connection between statistical mechanics and the dynamics of closed quantum systems*. The general goal is to put quantum statistical mechanics on a solid conceptual foundation as well as understanding the cases where it does not apply, such as many-body localized systems [120, 121, 122].

Numerically simulating all these systems has been a severe challenge. While many approaches to quantum dynamics exist, none are universally applicable. At short times, tensor-network methods can accurately describe the dynamics. However, in general the computational cost grows exponentially with the desired simulation time, thus severely limiting the timescales that can be resolved [123]. Conversely, at long times, the system is often effectively described by classical dynamics controlled by conservation laws. However, the interesting strongly correlated behavior is generally exhibited at intermediate




Figure 2.6: Cold atomic gases are controllable quantum systems that can be arranged in a lattice (a). They can be tuned to provide analog simulators of real crystal potentials (b). From Ref. [119].

times, inaccessible to the established classical methods. Quantum Monte Carlo methods can scale exponentially in time even for unfrustrated systems. Non-equilibrium dynamical mean-field theory [124] has emerged as a powerful method especially for systems in high dimensions, but requires uncontrolled approximations (both in the setup of the method and its numerical solution). Finally, in the regime of weak interactions, time-dependent density functional theory can be used, but likewise implies uncontrolled approximations [125].

3 Challenges for quantum algorithms in quantum simulation

3.1 Overview of algorithms

In the introduction we described three quantum problems that lie at the heart of chemistry and materials physics: the problem of *quantum dynamics*, representative of computational tasks that can be efficiently tackled by a quantum computer, aside from the initial state preparation; *quantum ground- and low-energy state determination*, central to quantum chemistry, condensed phase electronic structure, and quantum molecular spectroscopy; and thermal averages or *quantum statistical mechanics*, to describe finite-temperature chemistry and physics. In next sections, we survey the current status and theoretical and practical challenges to implement quantum algorithms for these problems, and some future directions.

Regardless of the problem we are studying, the first step in a quantum simulation is to choose a *representation* for the Hamiltonian and the states. Thus, in **Sec. 3.2** we first examine the possibilities for different qubit representations, and open questions in that area.

Quantum ground-state algorithms fall into several classes. *Quantum phase estimation* is a direct route to (near exact) eigenstate determination, but has been challenging to implement in the near-term era. A complementary technique is to prepare the exact ground-state via a prescribed “exact” evolution path, either in real time (*adiabatic state preparation*), or in imaginary time (*quantum imaginary time evolution*). On the other hand, *variational quantum algorithms* provide the possibility to introduce approximations with adjustable circuit depth, which are determined via a non-linear optimization, usually implemented in a hybrid-quantum classical loop. The variety

of algorithms currently explored for quantum ground- (and low-lying excited state) determination are discussed in [Sec. 3.3](#).

While polynomial time algorithms for *quantum time evolution* have been known for some time, algorithms with optimal asymptotic complexity (linear in time) as well as favorable scaling with error and low prefactors, remain an area of active research. Also, much work remains to be done to optimize Hamiltonian simulation algorithms for the Hamiltonians of interest in chemistry and materials science. Quantum time evolution is a fundamental building block in many quantum algorithms, such as phase estimation. The current status of quantum time evolution algorithms is summarized in [Sec. 3.5](#).

How best to simulate thermal states in chemical and materials science applications remains an open question. A wide variety of techniques have been discussed, ranging from eigenstate thermalization, to state preparation methods, to hybrid quantum-classical algorithms, though few have been implemented. The current status of thermal state methods, the prospects for implementing them, and other open questions are discussed in [Sec. 3.6](#).

Many quantum algorithms involve interfacing with classical data and classical algorithms. This can be to take advantage of classical optimization strategies, as in variational quantum algorithms. Another reason is to enable a multi-level/multi-scale representation of the problem. Quantum embedding provides a framework for such multi-scale quantum/classical hybrids, with the quantum representation of a subsystem coupled either to a classical environment, or another quantum representation via the exchange of classical data. We discuss the current status of *hybrid quantum-classical algorithms and quantum embedding* in particular in [Sec. 3.7](#).

Finally, an important consideration when developing improved quantum algorithms for real chemical and materials science problems is to establish *benchmark systems and results*, from the best available classical simulation data. The possibilities and prospects for such benchmarks are discussed in [Sec. 3.8](#).

3.2 Qubit representation of many-body systems

Many-body systems in chemistry and materials physics are composed of interacting electrons and atomic nuclei. An exact quantum mechanical treatment involves continuous variables such as the particles' positions and momenta.

To simulate such systems on a digital computer (either quantum or classical), the infinite-dimensional Hilbert space of a many-body system has to be truncated.

The most direct route is to define a finite set of *basis functions* and then to project the exact many-body Hamiltonian onto the chosen basis. The resulting discretized system is then expressed in terms of qubits. Depending on the problem, the Hamiltonian of interest may be different, e.g. in electronic structure it is the electronic Hamiltonian, while in molecular vibrational problems, it is the nuclear Hamiltonian. Alternatively, one can write down a simple form of the Hamiltonian a priori that contains the main interactions (a model Hamiltonian) with adjustable parameters. This latter approach is particularly popular in condensed matter applications. Finally, depending on the particles involved it may also be necessary to account for their fermionic or bosonic nature, in which case a suitable encoding of the statistics is required.

A choice of a good representation is important as it may affect the simulation cost dramatically. In this section we briefly summarize known methods for the qubit representation of many-body systems, discuss their relative merits, and outline important directions for future research.

3.2.1 Ab initio electronic structure qubit representations

The main objective of electronic structure in chemistry and physics to understand the low-energy properties of the electronic structure Hamiltonian that describes a system of interacting electrons moving in the potential created by atomic nuclei [126, 127],

$$\hat{H} = \hat{H}_1 + \hat{H}_2, \quad \hat{H}_1 = \sum_{i=1}^K -\frac{1}{2}\Delta_i^2 + V(r_i), \quad \hat{H}_2 = \sum_{1 \leq i < j \leq K} \frac{1}{|r_i - r_j|}. \quad (3.1)$$

Here K is the number of electrons, r_i is the position operator of the i -th electron, Δ_i is the corresponding Laplacian, and $V(r)$ is the electric potential created by atomic nuclei at a point r . The term \hat{H}_1 includes the kinetic and the potential energy of non-interacting electrons while \hat{H}_2 represents the Coulomb repulsion. Here we ignore relativistic effects and employ the standard Born-Oppenheimer approximation to solve the electronic Hamiltonian for fixed nuclei positions.

Each electron is described quantum mechanically by its position $r_i \in \mathbb{R}^3$ and spin $\omega_i \in \{\uparrow, \downarrow\}$. Accordingly, a quantum state of K electrons can be specified by a wave function $\Psi(x_1, \dots, x_K)$, where $x_i = (r_i, \omega_i)$. The wave function must obey Fermi statistics, that is, $\Psi(x_1, \dots, x_K)$ must be anti-symmetric under exchanging any pair of coordinates x_i and x_j .

The first step of any quantum electronic structure simulation algorithm is to approximate the electronic Hamiltonian \hat{H} with a simpler simulator Hamiltonian that describes a system of interacting qubits. This is usually achieved by truncating the Hilbert space of a single electron to a finite set of basis functions ψ_1, \dots, ψ_N known as (spin) orbitals. For example, each orbital could be a linear combination of atom-centered Gaussian functions with a fixed spin orientation.

Electronic structure simulation algorithms based on the *first quantization* method [128, 129, 130] describe a system of K electrons using the Configuration Interaction (CI) space (in classical simulations, this would be called the Full Configuration Interaction space). This is a linear space spanned by all possible Slater determinant states that can be formed by distributing K electrons over N orbitals. The CI space has dimension $\binom{N}{K}$ and can be identified with the anti-symmetric subspace of $(\mathbb{C}^N)^{\otimes K}$.

The projection of the full electronic Hamiltonian \hat{H} onto the CI space has the form

$$\hat{H} = \sum_{i=1}^K \sum_{p,q=1}^N t_{pq} |p\rangle \langle q|_i + \sum_{1 \leq i \neq j \leq K} \sum_{p,q,r,s=1}^N u_{pqrs} |p\rangle \langle r|_i \otimes |q\rangle \langle s|_j. \quad (3.2)$$

Here $|p\rangle \equiv |\psi_p\rangle$ are the chosen spin-orbitals. The coefficients $t_{p,q}$ and u_{pqrs} are known as one- and two-electron integrals. For example,

$$t_{p,q} = \langle \psi_p | (-\frac{1}{2}\Delta + V) | \psi_q \rangle. \quad (3.3)$$

Likewise, u_{pqrs} is the matrix element of the Coulomb interaction operator $1/|r_1 - r_2|$ between anti-symmetrized versions of the states $\psi_p \otimes \psi_q$ and $\psi_r \otimes \psi_s$. Each copy of the single-electron Hilbert space \mathbb{C}^N is then encoded by a register of $\log_2 N$ qubits. This requires $n = K \log_2 N$ qubits in total. The CI Hamiltonian H includes multi-qubit interactions among subsets of $2 \log_2 N$ qubits. The full Hilbert space of n qubits contains many unphysical states that do not originate from the CI space. Such states have to be removed from simulation by enforcing the anti-symmetry condition. This can be achieved by adding suitable energy penalty terms to the CI Hamiltonian [131].

An important parameter that affects the runtime of quantum simulation algorithms is the sparsity of the simulator Hamiltonian. A Hamiltonian H is said to be d -sparse if the matrix of H in the standard n -qubit basis has at most d non-zero elements in each row (each column). For example, the runtime of simulation algorithms based on quantum signal processing [132] scales linearly with the sparsity d . The CI Hamiltonian H has sparsity $d \sim (KN)^2$. Thus the first-quantization method is well-suited for high-precision simulation of small molecules when the number of electrons $K = O(1)$ is fixed and the number of orbitals N is a large parameter. As one approaches the continuum limit $N \rightarrow \infty$, the number of qubits grows only logarithmically with N while the sparsity of H scales as $d \sim N^2$.

The *second quantization* approach often results in a simpler simulator Hamiltonian and requires fewer qubits, especially in the case when the filling fraction K/N is not small. This method is particularly well suited for quantum simulation algorithms [133] and has been experimentally demonstrated for small molecules [134]. Given a set of N orbitals ψ_1, \dots, ψ_N , the second-quantized simulator Hamiltonian is

$$H = \sum_{p,q=1}^N t_{pq} \hat{c}_p^\dagger \hat{c}_q + \frac{1}{2} \sum_{p,q,r,s=1}^N u_{pqrs} \hat{c}_p^\dagger \hat{c}_q^\dagger \hat{c}_r \hat{c}_s, \quad (3.4)$$

where \hat{c}_p^\dagger and \hat{c}_p are the creation and annihilation operators for the orbital ψ_p . The Hamiltonian H acts on the Fock space spanned by 2^N basis vectors $|n_1, n_2, \dots, n_N\rangle$, where $n_p \in \{0, 1\}$ is the occupation number of the orbital ψ_p . The advantage of the second quantization method is that the Fermi statistics is automatically enforced at the operator level. However, the number of electrons can now take arbitrary values between 0 and N . The simulation has to be restricted to the subspace with exactly K occupied orbitals. The second-quantized Hamiltonian H can be written in terms of qubits using one of the fermion-to-qubit mappings discussed in Section Sec. 3.2.3.

Within the above outline, there are several active areas of research. For example, one may wonder if the redundancy of the space in the first quantized representation can be reduced or completely avoided. Other questions include the choice of basis functions and fermion-to-qubit mappings. These are discussed in the next sections.

3.2.2 Electronic basis functions

The discretization of the electronic Hilbert space for a quantum simulation requires balancing two concerns. We need to represent the state with as few qubits, but also, retain maximal Hamiltonian sparsity. These requirements only partially align with those of classical many-particle quantum simulations. In the classical setting, a compact state is crucial due to the exponential Hilbert space, while Hamiltonian sparsity is less so; the choice of basis functions has historically been made so that matrix elements of the Hamiltonian (the one- and two-electron “integrals”) can be analytically evaluated [135, 126].

There are two families of basis functions in wide use in quantum chemistry and quantum materials science: atomic orbital *Gaussian bases* and *plane waves*. Gaussian bases are most commonly employed in molecular simulations due to their compactness, while plane waves are most often used in crystalline materials simulation, due to their intrinsic periodicity and ease of regularizing the long-range contributions of the Coulomb operator (which are conditionally convergent in an infinite system).

In Gaussian bases, linear combinations of Gaussian functions (referred to as simply Gaussian basis functions) are placed at the nuclear positions. As they are placed where the ground-state electron density is highest, they give a compact representation of the wavefunction for bound states, but the Hamiltonian is not sparse, with $O(N^4)$ second-quantized matrix elements. In a quantum algorithm, this leads to high gate counts even for simple quantum primitives such as a Trotter step.

Plane-waves offer greater simplicity as the accuracy of the basis is controlled by a single parameter, the kinetic energy cutoff. While the number of plane waves needed to reach a desired accuracy is larger than the number of Gaussian states, the Hamiltonian contains fewer terms ($O(N^3)$) due to momentum conservation. To reduce the number of required plane waves, it is essential to employ pseudopotentials to remove the sharp nuclear cusp [137, 138]. Furthermore, the asymptotic basis convergence of Gaussian and pseudopotential plane wave calculations is the same: the feature governing the rate of convergence is the wavefunction discontinuity or *electron-electron cusp* due to the singularity of the Coulomb interaction (see Fig. 3.1). In classical simulations, so-called explicit correlation methods can be used to remove the slow convergence due to the singularity [139, 140]. How to use such techniques with quantum computers has yet to be explored.

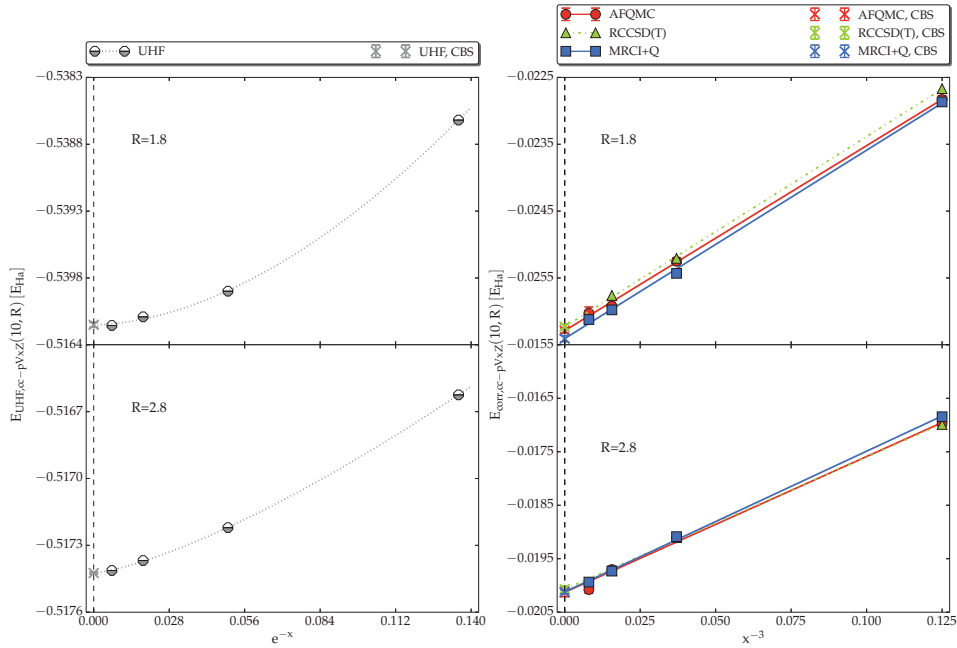


Figure 3.1: Extrapolation of the Hartree-Fock (left) and correlation energy (right) for chains of 10 hydrogen atoms, using cc-pVxZ Gaussian bases, and various methods from [136]. While the mean-field energy converges exponentially fast in the basis size (roughly given by x^3), the correlation energy converges as x^{-3} (inversely proportional to the number of basis functions) due to the electron-electron cusp.

The need to expose more sparsity in the Hamiltonian while retaining a reasonably compact wavefunctions is an active area of research in both classical and quantum algorithms. Recent ideas have included new types of basis function that return to a more grid-like real-space basis [141, 142, 143, 144, 145] where the Coulomb operator and thus Hamiltonian has only a quadratic number of terms, as well as factorizations of the Coulomb operator itself [146, 147]. The best choice of basis for a quantum simulation remains very much an open question.

3.2.3 Fermion-to-qubit mappings

Since the basic units of a quantum computer are qubits rather than fermions, any quantum simulation algorithm of fermions (e.g. for electronic structure) employs a suitable encoding of fermionic degrees of freedom into qubits. For example, the standard *Jordan-Wigner* mapping (sketched in Fig. 3.2) identifies each Fermi mode (orbital) with a qubit such that the empty and the occupied states are mapped to the qubit basis states $|0\rangle$ and $|1\rangle$ respectively. More generally, the Fock basis vector $|n_1, \dots, n_N\rangle$ is mapped to a qubit basis vector $|x_1\rangle \otimes \dots \otimes |x_N\rangle$, where each bit x_j stores a suitable partial sum (modulo two) of the occupation numbers n_1, \dots, n_N . The Jordan-Wigner mapping corresponds to $x_j = n_j$ for all j . This is not quite satisfactory since single-mode creation/annihilation operators contain Jordan-Wigner strings $(-1)^{n_1+\dots+n_j}$ which may be very non-local in terms of qubits. On the other hand, updating the qubit state x upon application of a single creation/annihilation operator requires a single bit flip (see Fig. 3.3). More efficient fermion-to-qubit mappings balance the cost of computing Jordan-Wigner strings and the bit-flip cost of updating the qubit state. For example, the *binary tree encoding* proposed by Kitaev and one of the authors [148] maps any fermionic single-mode operator (e.g. \hat{c}_p or \hat{c}_p^\dagger) to a qubit operator acting non-trivially on roughly $\log N$ or less qubits. Generalizations of this encoding were studied in [149, 150, 151]. As a consequence, the second-quantized Hamiltonian Eq. (3.4) expressed in terms of qubits becomes a linear combination of Pauli terms with weight at most $O(\log N)$. This is important in the context of VQE-type quantum simulations (see Sec. 3.3.2) since Pauli operators with an extensive weight (of order N) cannot be measured reliably in the absence of error correction.

A natural question is whether the number of qubits required to express a Fermi system can be reduced by exploiting symmetries such as the particle




Figure 3.2: Schematic of the Jordan-Wigner encoding for 3 spin-orbitals. Fermions are represented by blue spheres. Each Fock basis vectors (from the vacuum state, left, to the completely filled state, right) is mapped onto a 3-qubit state, with empty (filled) spin-orbitals corresponding to qubits in 0 (1).

number conservation or the point group symmetries of molecules. For example, zero temperature simulations often target only one symmetry sector containing the ground state. This motivates the study of *symmetry-adapted fermion-to-qubit mappings*. The goal here is to reduce the number of qubits required for the simulation without compromising the simple structure of the resulting qubit Hamiltonian (such as sparsity). The simplest case of \mathbb{Z}_2 symmetries is now well understood and the corresponding symmetry adapted mappings are routinely used in experiments [131, 152, 134, 153]. The $U(1)$ symmetry underlying particle number conservation was considered in [154, 131, 155].

Dimension counting shows that a system of N Fermi modes with exactly K particles can be mapped to roughly $\log_2 \binom{N}{K}$ qubits. However, it remains an open question whether this mapping can be chosen such that the resulting qubit Hamiltonian admits a sparse representation in some of the commonly used operator bases (such as the basis of Pauli operators) to enable applications in VQE. Mappings adapted to point group symmetries have been recently considered in [156]. It is also of great interest to explore fermion-to-qubit mappings adapted to approximate and/or emergent symmetries.

Alternatively, one may artificially introduce symmetries either to the original Fermi system or its encoded qubit version with the goal of simplifying the resulting qubit Hamiltonian. This usually requires redundant degrees of freedom such as auxiliary qubits or Fermi modes [148, 157, 158, 159]. In the case of lattice fermionic Hamiltonians such as the 2D Fermi Hubbard model or more general models defined on bounded degree graphs, such symmetry-adapted mappings produce a local qubit Hamiltonian composed of Pauli operators of constant weight independent of N .




Figure 3.3: Difference between Jordan-Wigner and parity encoding. The former encodes occupation numbers x_i on qubit states, the latter parities $p_i = \sum_{j < i} x_j \mod 2$. Other strategies, such as the binary-tree encoding, balance non-locality of occupation numbers and parities to achieve more efficient encodings.

3.2.4 Model and non-electronic problems

While much of the above has focused on ab initio quantum chemistry and electronic structure in quantum simulations, the diverse questions of chemistry and materials physics discussed in section 2 raise additional issues of representation. For example, model Hamiltonians avoid the problem associated with a choice of basis by restricting the Hamiltonian to a predetermined simple form on a lattice. The lattice structure permits specialized techniques, such as specific fermionic encodings. Developing specialized representations for model problems is of particular importance in simulating condensed matter systems.

Other kinds of non-electronic simulations may involve different requirements on the basis than electronic problems. For example, in quantum reactive scattering processes, there is little a priori information on the positions of the particles; instead various grid representations, often in non-Cartesian coordinate systems, are used [160, 161]. Alternatively, the particles of interest may be bosons which engender new encoding considerations. Relatively little attention has been paid to these questions so far in quantum algorithms.

3.3 Quantum algorithms for ground and excited states

There are many approaches to obtaining ground states or excited states on a quantum computer. State preparation procedures attempt to construct a circuit to prepare a state with as large as possible overlap with the desired eigenstate. One set of such procedures, which we review in Sec. 3.3.1 and which includes *adiabatic state preparation* and *quantum imaginary time evolution*, uses a prescribed evolution path. An alternative strategy is based on *variational methods* (often called variational quantum eigensolvers) and is reviewed in Sec. 3.3.2. Here, the preparation circuit itself is defined via the optimization of the energy with respect to parameters of the circuit.

Given some state with reasonably large overlap with the desired state, one can perform *quantum phase estimation*, which simultaneously projects the state onto an eigenstate of the Hamiltonian and obtains an estimate for the energy of this eigenstate. The error depends inversely on the simulation time. The probability of successfully projecting onto the desired state is given by the square overlap of the input state and the desired state, and it is thus necessary to use some other method (such as the state preparation procedures above) to prepare an input state with sufficient overlap with the desired state.

The various approaches come with different strengths and weaknesses. While phase estimation allows the deviation of the final state from an exact eigenstate (although not necessarily the desired eigenstate) to be systematically reduced, it can require deep circuits with many controlled gates that are challenging for devices with limited coherence and without error correction. Variational methods or quantum imaginary time evolution replace such circuits by a large number of potentially shorter simulations, which is expected to be easier to implement on near-term machines. However, if one does not measure the energy by phase estimation, but instead by expressing the Hamiltonian as a sum of multi-qubit Pauli operators and measuring the terms individually, the state preparation and measurements must be repeated many times, with the error converging only as the square root of the number of repetitions. Variational methods are also limited by the variational form and ability to solve the associated optimization problem, which may by itself represent a difficult classical optimization. Finally, adiabatic state preparation and quantum imaginary time evolution become inefficient for certain Hamiltonians.

State preparation methods are first discussed in **Sec. 3.3.1** and **Sec. 3.3.2**

(variational state preparation is discussed separately due to the large number of different types of ansatz). Considerations for excited states are discussed in **Sec. 3.3.3**. Phase estimation is reviewed in **Sec. 3.4**.

3.3.1 Preparing ground states along a prescribed path

Adiabatic state preparation

One general route to prepare the ground-state of a physical system on a quantum device is through adiabatic state preparation. This relies on the well-known *adiabatic theorem* [162, 163], which states that a system evolved under a time-dependent Hamiltonian will remain in the instantaneous ground state as long as the evolution is sufficiently slow and the spectrum of the Hamiltonian remains gapped. To make use of this, one chooses a Hamiltonian path $\hat{H}(\lambda)$, $0 \leq \lambda \leq 1$, such that the ground state of $\hat{H}(0)$ is easily prepared, while $\hat{H}(1)$ is the Hamiltonian whose ground state one wants to obtain. The system is then evolved under the time-dependent Schrödinger equation,

$$i \frac{d}{dt} |\Phi_t\rangle = \hat{H}(t/T) |\Phi_t\rangle, \quad 0 \leq t \leq T \quad (3.5)$$

In the limit $T \rightarrow \infty$ and if the spectrum of $\hat{H}_{t/T}$ is gapped for all t , the final state $|\Phi_T\rangle$ is the exact ground state of \hat{H} . Away from the adiabatic limit $T \rightarrow \infty$, corrections arise that depend on the instantaneous gap of $\hat{H}(\lambda)$ and the total time T . Furthermore, if degeneracies occur along the path, a different time-dependent Hamiltonian path (although with the same endpoints) must be chosen.

Thus while this approach is very general, its practical applicability is limited by the requirement of having to take the limit of large T , and to choose a path without degeneracies (see **Fig. 3.4**). The limit of large T may require deep circuits that may not be practical in near-term quantum machines. Analyzing the errors (and optimizing the path, for example by choosing an improved $f(s)$) is also challenging due to the dependence on the unknown spectrum of $\hat{H}_{t/T}$. Some of these questions have been studied in the more general context of adiabatic quantum computation [164]. However, while this is one of the first state preparation methods discussed for chemical systems [133], more heuristic work in this area for problems of interest to chemistry and physics is needed [165].




Figure 3.4: Adiabatic state preparation. Starting from an eigenstate of a simple Hamiltonian \hat{H}_0 and slowly switching on the interaction $\hat{H} - \hat{H}_0$ leads to an eigenstate of \hat{H} (left). Along paths with degeneracies, adiabatic state preparation can lead to the wrong eigenstate.

Quantum imaginary-time evolution

In classical simulations, one popular approach to prepare (nearly exact) ground-states is imaginary-time evolution, which expresses the ground-state as the long-time limit of the imaginary-time Schrödinger equation,

$$|\Psi\rangle = \lim_{\beta \rightarrow \infty} \frac{e^{-\beta \hat{H}} |\Phi\rangle}{\|e^{-\beta \hat{H}} \Phi\|} . \quad (3.6)$$

Imaginary time-evolution underlies the family of projector quantum Monte Carlo methods in classical algorithms [127].

To perform imaginary time evolution on a quantum computer, it is necessary to implement the (repeated) action of the short-imaginary-time propagator $e^{-\Delta\tau \hat{H}}$ on a state. Given a Hamiltonian that can be decomposed into geometrically local terms, $\hat{H} = \sum_m \hat{h}_m$, and a state $|\Psi\rangle$ with finite correlation length C , the action of $e^{-\Delta\tau \hat{h}_i}$ can be generated by a unitary $\hat{U} = e^{i\hat{A}}$ acting on $O(C)$ qubits surrounding those acted on by \hat{h}_i , i.e.

$$\frac{e^{-\Delta\tau \hat{h}_i} |\Psi\rangle}{\|e^{-\Delta\tau \hat{h}_i} \Psi\|} = \hat{U} |\Psi\rangle = e^{i\hat{A}} |\Psi\rangle , \quad (3.7)$$

where the coefficients of the Pauli strings in \hat{A} can be determined from local measurements of the qubits around \hat{h}_i . This is the idea behind the quantum

imaginary time evolution (QITE) algorithm [166]. Like adiabatic state preparation, quantum imaginary time evolution can in principle prepare exact states without the need for variational optimization. Also, the total length of imaginary time propagation to achieve a given error is determined by the spectrum of \hat{H} and the initial overlap, rather than by the spectrum of $\hat{H}_{t/T}$ along the adiabatic path. However, the method becomes inefficient in terms of the number of measurements and complexity of the operator \hat{A} if the domain C grows to be large along the imaginary time evolution path. In these cases, QITE can be used as a heuristic for approximate ground-state preparation, analogous to using adiabatic state preparation for fixed evolution time. While initial estimates in a limited set of problems show QITE to be resource efficient compared to variational methods due to the lack of an optimization loop [166], a better numerical understanding of its performance and cost across different problems, as well as the accuracy of inexact QITE in different settings, remains to be developed.

3.3.2 Variational state preparation and variational quantum eigensolver

A class of state preparation methods that have been argued to be particularly amenable to near-term machines is variational state preparation [167, 168, 169]. Here, similar to classical variational approaches, one chooses a class of ansatz states for the ground state of the Hamiltonian of interest. Generally speaking, such an ansatz consists of some initial state and a unitary circuit parametrized by some set of classical variational parameters. Applying this circuit to the initial state yields a guess for the ground state, whose energy is then evaluated. This yields an upper bound to the true ground state energy. One then varies the variational parameters to lower the energy of the ansatz state.

In choosing the class of ansatz states, one pursues several goals: on the one hand, it is crucial that the class contains an accurate approximation to the true ground state of the system. On the other hand, one desires a class of circuits that are easily executed on the available quantum computer, i.e. for a given set of available gates, connectivity of the qubits, etc. Finally, it is important for the classical optimization over the variational parameters to be well-behaved, so as to be able to find low-energy minima. While we cannot list all possible ansatz states below, we provide a representative sample.

Unitary coupled cluster

An early example of a particular class of ansatz states that has been suggested for applications in quantum chemistry is the unitary coupled-cluster (uCC) ansatz [170, 168, 171]

$$|\Psi_{uCC}\rangle = e^{\hat{T}-\hat{T}^\dagger} |\Psi_{HF}\rangle, \quad \hat{T} = \sum_{k=1}^d \sum_{\substack{i_1 \dots i_k \\ a_1 \dots a_k}} t_{i_1 \dots i_k}^{a_1 \dots a_k} \hat{c}_{a_k}^\dagger \dots \hat{c}_{a_1}^\dagger \hat{c}_{i_k} \dots \hat{c}_{i_1}. \quad (3.8)$$

Here, d denotes the maximum order of excitations in the uCC wavefunction (for example $d = 1, 2, 3$ for singles, doubles and triples respectively), $\hat{c}_{a_k}^\dagger \dots \hat{c}_{a_1}^\dagger (\hat{c}_{i_k} \dots \hat{c}_{i_1})$ are creation (destruction) operators relative to orbitals unoccupied (occupied) in the Hartree-Fock state, and t is a rank- $2k$ tensor, antisymmetric in the $a_k \dots a_1$ and $i_k \dots i_1$ indices. This choice of ansatz is motivated by the success of mean-field theory, which suggests that the density of excitations in the true wavefunction should be small relative to the mean-field state. Standard coupled cluster theory – written as $e^{\hat{T}} |\Psi_{HF}\rangle$ – is widely used in classical quantum chemistry but is challenging to implement on a quantum device, whereas the reverse is true for the unitary variant. Understanding the theoretical and numerical differences between standard and unitary coupled cluster is an active area of research [172, 173]. The variational quantum eigensolver algorithm applied to the unitary coupled-cluster Ansatz is depicted in Fig. 3.5

Hardware-efficient ansatz

The unitary coupled cluster ansatz involves non-local gate operations and is expensive to implement on near-term devices with limited qubit connectivity. An alternative variational approach, pursued e.g. in Ref. [134] and termed “hardware-efficient” there, is to tailor the ansatz specifically to the underlying hardware characteristics. The circuits considered in Ref. [134], sketched in Fig. 3.6, consist of alternating layers of arbitrary single-qubit gates and an entangling gate that relies on the intrinsic drift Hamiltonian of the system. While this drift Hamiltonian and thus the entangling gate is not known precisely, for the variational approach it is sufficient to know that the gate is reproducible. The variational parameters are only the rotations in the layer of single-qubit gates. While it is not guaranteed that such an ansatz contains a good approximation to the state of interest, it is an example of an adaption of a method to NISQ devices [134]. An application to BeH_2 is seen in Fig. 3.6.

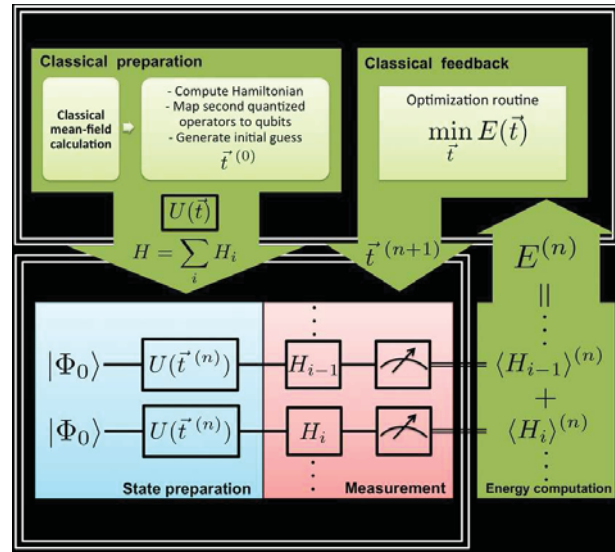


Figure 3.5: Workflow of the variational quantum eigensolver algorithm. The classical optimization routine adds expectation values of the Hamiltonian Pauli terms to calculate the energy and estimates new values for the unitary parameters. The process is repeated until convergence. From [174].

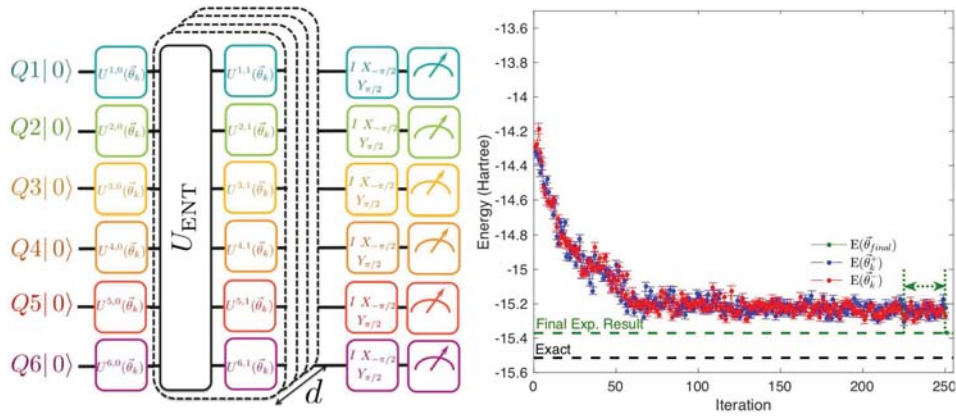


Figure 3.6: Left: Hardware-efficient quantum circuit for trial state preparation and energy estimation, shown here for 6 qubits. The circuit is composed of a sequence of interleaved single-qubit rotations, and entangling unitary operations U_{ENT} that entangle all the qubits in the circuit. A final set of post-rotations prior to qubit readout are used to measure the expectation values of the terms in the qubit Hamiltonian, and estimate the energy of the trial state. Right: energy minimization for the six-qubit Hamiltonian describing BeH_2 . Adapted from [134].

Adapt-VQE ansatz

In the adapt-VQE scheme, a collection of operators \hat{A}_i (operator pool) is chosen in advance, and the ground state is approximated by

$$|\Psi_{\text{adapt-VQE}}\rangle = e^{\theta_n \hat{A}_n} \dots e^{\theta_1 \hat{A}_1} |\Psi_{\text{HF}}\rangle, \quad (3.9)$$

Given a current parameter configuration $\theta_1 \dots \theta_n$, the commutator of the Hamiltonian with each operator in the pool is measured to obtain the gradient of the energy $\langle \Psi_{\text{adapt-VQE}} | \hat{H} | \Psi_{\text{adapt-VQE}} \rangle$ with respect to the parameters θ . Repeating this multiple times and averaging over the obtained samples gives the gradient of the expectation value of the Hamiltonian with respect to the coefficient of each operator. The ansatz is improved by adding the operator \hat{A}_i with the largest gradient to the left end of the ansatz with a new variational parameter, thereby increasing n . The operation is repeated until convergence of the energy [175]. Numerical simulations, for example for short hydrogen chains, show that adapt-VQE can improve over the unitary coupled cluster approach in terms of the accuracy reached for a given circuit depth.

Tensor networks

Tensor networks are a class of variational states which construct the global wavefunction amplitude from tensors associated with local degrees of freedom. They specify a class of quantum states that can be represented by an amount of classical information proportional to the system size. There are two main families of tensor networks: those based on *matrix product states* and *tree tensor network states* (MPS, TTNS) [176, 177] (also known in the numerical multi-linear algebra community as the tensor-train decomposition [178, 179] and the tree-structured hierarchical Tucker representation [180, 181]) and their higher dimensional analogs, *projected entangled pair states* (PEPS) [182]; and those based on the *multi-scale entanglement renormalization ansatz* (MERA) [183]. Because of their success in representing low-energy states in classical simulations, they are a natural class of variational wavefunctions to try to prepare in a quantum algorithm for low-energy states. These tensor networks are schematically depicted in Fig. 3.7.

There are many analogies between tensor network algorithms and quantum circuits. This analogy can be exploited to develop an efficient preparation mechanism for these states on a quantum computer. By recognizing that




Figure 3.7: Graphical representation of (a) matrix product states, (b) projected-entangled pair states. The boxes and circles represent tensors of numbers. (c) Construction of deep multi-scale entanglement renormalization ansatz (deep MERA). The figure shows an isometry (one component of the MERA ansatz) being constructed with exponential dimension via quantum circuits. Adapted from [189] and [187].

the tensors in an MPS or a MERA can be associated with a block of unitaries (with the bonds between tensors playing the role of circuit lines in a quantum circuit) it is possible to prepare an MPS or MERA state on a quantum computer [184, 185, 186]. Because the dimensions of the associated tensor grow exponentially with the depth of the quantum circuit associated with it, it is possible to prepare certain tensor networks with large bond dimension on a quantum computer that presumably cannot be efficiently simulated classically; an example of this is the so-called *deep MERA* [187].

There are many open questions in the area of tensor networks and quantum computing. For example, preparing PEPS on a quantum computer appears to be much less straightforward than preparing a matrix product state or a MERA [188]. Similarly, although “deep” tensor network states can only be efficiently simulated on a quantum computer, their additional representational power over classically efficient tensor networks for problems of physical or chemical interest is poorly understood.

Other considerations

Besides the choice of ansatz state, the computational challenges of variational methods and VQE are twofold:

- Potentially, a very large number of measurements must be performed to accurately estimate the energy. Indeed, the scaling is quadratically

worse than when using quantum phase estimation. We will discuss this point further in [Sec. 3.9](#).

- The optimization of the variational parameters may be very difficult, in particular if the energy exhibits a very non-trivial dependence on classical parameters with many local minima, and if gradient information is not easily available. For some discussion of optimization algorithms in this context, see Refs. [\[190, 191\]](#).

Some key advantages of the VQE approach are that it can often be carried out with a large number of independent, short quantum simulations. This is more suitable to NISQ machines than the long coherent circuits required for approaches based on quantum phase estimation, which has been demonstrated in several experiments [\[134, 152\]](#). Furthermore, the approach is more resilient against certain types of errors. For example, as mentioned already above, it is generally not necessary to know exactly what circuit is executed for some variational parameters as long as it is reproducible; therefore, systematic coherent tuning errors of the qubits (for example systematic deviations between the desired and the actually applied single-qubit rotations) do not adversely affect the results. In addition to studying the robustness of VQE against errors, it has become a very active field to develop techniques that mitigate such physical errors. Such approaches promise to reduce the impact of errors on near-term machines before error correction becomes available. For work in this direction, see Refs. [\[190, 192, 193, 194, 195, 196, 197\]](#).

3.3.3 Excited states

While much of the above discussion of state preparation and variational algorithms has focused on ground-states, most of the same methods can also be used with minor extensions for excited states. For example, adiabatic state preparation can be used to prepare an excited state, so long as it is connected to the initial state without a vanishing gap.

In the area of variational methods, it is often useful to choose the excited state ansatz to be related to that of the ground-state, since at low-energies much of the physics is the same. This is widely used in classical simulations and essentially the same ideas have been ported to the quantum algorithm setting. For example, in the *quantum subspace expansion* (QSE) [\[198\]](#), the excited state is made via the ansatz $|\Psi'\rangle = \sum_{\alpha} c_{\alpha} \hat{E}_{\alpha} |\Psi\rangle$, where $\{E_{\alpha}\}$ is a set of “excitation” operators and $|\Psi\rangle$ is the ground-state constructed in VQE.

In QSE, one needs to measure all the subspace matrix elements $\langle \Psi | \hat{E}_\alpha^\dagger \hat{E}_\beta | \Psi \rangle$, $\langle \Psi | \hat{E}_\alpha^\dagger \hat{H} \hat{E}_\beta | \Psi \rangle$, thus the number of measurements grows quadratically with the subspace. In the **quantum Lanczos** method, the QITE algorithm is used to construct the subspace $\{e^{-\lambda \hat{H}} | \Psi \rangle, e^{-2\lambda \hat{H}} | \Psi \rangle, \dots\}$ and the special structure of this space means that all subspace matrix elements can be constructed with a number of measurements that grows only linearly with the size of the subspace [166]. Alternatively, one can fix the coefficients c_α and reoptimize the quantum circuit in the variational method; this is the basis of the multi-state VQE method; other similar ideas have also been proposed. Connections between quantum subspaces and error correction have been explored in [199].

The above methods compute total energies of excited states, which have to be subtracted from the ground state energy to give the excitation energies of the system. A method to directly access excitation energies is desirable. One route to achieve this goal is provided by the **equation-of-motion** (EOM) approach, also widely used in classical simulations [200] and recently extended to quantum computing [201]. In the EOM approach, excitation energies are obtained as

$$\Delta E_n = \frac{\langle \Psi | [\hat{O}_n, \hat{H}, \hat{O}_n^\dagger] | \Psi \rangle}{\langle \Psi | [\hat{O}_n, \hat{O}_n^\dagger] | \Psi \rangle} \quad (3.10)$$

where Ψ is an approximation to the ground state (such as the VQE ansatz) and \hat{O}_n is an excitation operator expanded on a suitable basis. The variational problem of finding the stationary points of ΔE_n leads to a generalized eigenvalue equation, the solutions of which are the excited-state energies.

3.4 Phase estimation

Quantum Phase Estimation (QPE) is a crucial step in many quantum algorithms. In the context of quantum simulation, QPE enables high-precision measurements of the ground and excited energy levels. This is achieved by preparing a trial initial state $|\psi(0)\rangle$ that has a non-negligible overlap with the relevant eigenvector of the target Hamiltonian \hat{H} and applying a quantum circuit that creates a superposition of time evolved states $|\psi(t)\rangle = e^{-i\hat{H}t}|\psi(0)\rangle$ over a suitable range of the evolution times t . In the simplest case, known as the iterative QPE [6, 133] and sketched in Fig. 3.8, the final state is a superposition of the initial state itself and a single time-evolved state,

$$\frac{1}{\sqrt{2}}(|0\rangle \otimes |\psi(0)\rangle + e^{i\theta}|1\rangle \otimes |\psi(t)\rangle). \quad (3.11)$$




Figure 3.8: Quantum circuit for iterative QPE. The first two single-qubit gates bring the ancilla to the state $\frac{1}{\sqrt{2}} (|0\rangle + e^{i\theta}|1\rangle)$. The controlled- $U(t)$ gate, with $U(t) = e^{-it\hat{H}}$, brings the register into the state (3.11). The last Hadamard gate is needed to measure in the X basis.

Here one ancillary control qubit has been added that determines whether each gate in the quantum circuit realizing $e^{-i\hat{H}t}$ is turned on (control is 1) or off (control is 0). The extra phase shift θ coordinates interference between the two computational branches such that useful information can be read out with high confidence. Finally, the control qubit is measured in the so-called X -basis, $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ and the measurement outcome $b \in \{+1, -1\}$ is recorded.

The iterative QPE works by performing many runs of the above subroutine with a suitable choice of parameters t , θ in each run and performing a classical post-processing of the observed measurement outcomes. The ancillary control qubit stays alive only over the duration of a single run since its state is destroyed by the measurement. However, the remaining qubits that comprise the simulated system stay alive over the entire duration of the QPE algorithm. More specifically, let E_α and $|\psi_\alpha\rangle$ be the eigenvalues and eigenvectors of \hat{H} such that $\hat{H} = \sum_\alpha E_\alpha |\psi_\alpha\rangle\langle\psi_\alpha|$. The trial state can be expanded in the eigenbasis of \hat{H} as $|\psi(0)\rangle = \sum_\alpha c_\alpha |\psi_\alpha\rangle$. Then the joint probability distribution describing measurement outcomes $b_1, \dots, b_N = \pm 1$ observed in N runs of QPE has the form

$$\text{Pr}(b_1, \dots, b_N) = \sum_\alpha |c_\alpha|^2 \prod_{i=1}^N \frac{1}{2} (1 + b_i \cos(\theta_i - E_\alpha t_i)). \quad (3.12)$$

Here θ_i and t_i are the phase shift and the evolution time used in the i -th run and $b_i = \pm 1$ is the observed measurement outcome. This has the same effect as picking an eigenvector of \hat{H} at random with the probability $|c_\alpha|^2$ and then running QPE on the initial state $|\psi(0)\rangle = |\psi_\alpha\rangle$. Accordingly, QPE aims at estimating a *random* eigenvalue E_α sampled from the probability distribution $|c_\alpha|^2$. We note that a suitably formalized version of this problem with the trial state $|00\dots 0\rangle$ and a local Hamiltonian \hat{H} composed of few-qubit interactions

is known to be BQP-complete [202]. In that sense, QPE captures the full computational power of quantum computers and any quantum algorithm can be expressed as a special case of QPE. A common application of QPE is the task of estimating the smallest eigenvalue $E_0 = \min_{\alpha} E_{\alpha}$. This requires a trial state $|\psi(0)\rangle$ that has a non-negligible overlap with the true ground state of \hat{H} to ensure that the minimum of a few randomly sampled eigenvalues E_{α} coincides with E_0 . For example, in the context of molecular simulations, $|\psi(0)\rangle$ is often chosen as the Hartree-Fock approximation to the ground state. Such a state is easy to prepare as it can be chosen to correspond to a standard basis vector, see Sec. 3.2.3.

The problem of obtaining a good estimate of the eigenvalue E_{α} based on the measured outcomes b_1, \dots, b_N is an active research area, see [203, 204, 205, 206]. Assuming that the trial state has a constant overlap with the ground state of \hat{H} , the smallest eigenvalue E_0 can be estimated using QPE within a given error ϵ using $N \sim \log(1/\epsilon)$ runs such that each run evolves the system over time at most $O(1/\epsilon)$. As discussed in Sec. 3.5, the evolution operator $e^{-i\hat{H}t}$ (as well as its controlled version) can be approximated by a quantum circuit of size scaling almost linearly in t (neglecting logarithmic corrections). Thus QPE can achieve an approximation error ϵ at the computational cost roughly $1/\epsilon$ even if the trial state has only a modest overlap with the ground state. This should be contrasted with VQE algorithms that have cost at least $1/\epsilon^2$ due to sampling errors and where the trial state must be a very good approximation to the true ground state, see Sec. 3.3.2. On the other hand, QPE is much more demanding in terms of the required circuit depth and the gate fidelity. It is expected that quantum error correction will be required to implement QPE in a useful way (e.g. to outperform VQE in ground-state determination).

QPE also has a single-run (non-iterative) version where the time evolution of the simulated system is controlled by a multi-qubit register and the X -basis measurement of the control qubit is replaced by the Fourier basis measurement [207]. The iterative version of QPE has the clear advantage of requiring fewer qubits. It also trades quantum operations required to realize the Fourier basis measurement for classical postprocessing, thereby reducing the overall quantum resource cost of the simulation.

Since QPE is used ubiquitously in a variety of quantum applications, it is crucial to optimize its performance. Below we list some open problems that are being actively investigated; see Ref. [206] for a recent review.

- Given limitations of near-term quantum devices, of particular interest

are tradeoffs between the depth of the QPE circuit and its spectral-resolution power as well as its sensitivity to noise. It was shown [206] that the computational cost of QPE interpolates between $1/\epsilon$ and $1/\epsilon^2$ as the depth (measured by the number of runs N per iteration of the algorithm) is reduced from $O(1/\epsilon)$ to $O(1)$. A particular version of QPE with a tunable depth that interpolates between the standard iterative QPE and VQE was proposed in [208].

- Several methods have been proposed for mitigating experimental errors for VQE-type simulations [194, 196, 197]. Such methods enable reliable estimation of expected values of observables on a given trial state without introducing any overhead in terms of extra qubits or quantum gates. Generalizing such error mitigation methods to QPE is a challenging open problem since QPE performs a non-trivial postprocessing of the measurement outcomes that goes beyond computing mean values.
- Classical post-processing methods that enable simultaneous estimation of multiple eigenvalues are highly desirable [206].
- Finally, a natural question is whether the time evolution operator $e^{-i\hat{H}t}$ in QPE can be replaced by some other functions of \hat{H} that are easier to implement [209, 210].

3.5 Quantum algorithms for time evolution

3.5.1 Hamiltonian simulation problem

It was recognized early on [211, 212] that a quantum computer can be programmed to efficiently simulate the unitary time evolution of almost any physically realistic quantum system. The time evolution of a quantum system initialized in a given state $|\psi(0)\rangle$ is governed by the Schrödinger equation

$$i\frac{d|\psi(t)\rangle}{dt} = \hat{H}|\psi(t)\rangle \quad , \quad t \geq 0 \quad , \quad (3.13)$$

where H is the system's Hamiltonian. Since any fermionic or spin system can be mapped to qubits, see **Sec. 3.2**, below we assume that H describes a system of n qubits. By integrating Eq. (3.13) for a time-independent Hamiltonian one obtains the time-evolved state

$$|\psi(t)\rangle = e^{-it\hat{H}}|\psi(0)\rangle \quad . \quad (3.14)$$

A quantum algorithm for Hamiltonian simulation takes as input a description of \hat{H} , the evolution time t , and outputs a quantum circuit \hat{U} that approximates the time evolution operator $e^{-it\hat{H}}$ within a specified precision ϵ , that is,

$$\|\hat{U} - e^{-it\hat{H}}\| \leq \epsilon. \quad (3.15)$$

More generally, the circuit \hat{U} may use some ancillary qubits initialized in the $|0\rangle$ state. The simulation cost is usually quantified by the runtime of the algorithm (the gate count of \hat{U}) and the total number of qubits. Applying the circuit \hat{U} to the initial state $|\psi(0)\rangle$ provides an ϵ -approximation to the time-evolved state $|\psi(t)\rangle$. The final state $\hat{U}|\psi(0)\rangle$ can now be measured to access dynamical properties of the system such as time-dependent correlation functions. The time evolution circuit \hat{U} is usually invoked as a subroutine in a larger enveloping algorithm. For example, the quantum phase estimation method employs a controlled version of \hat{U} to measure the phase accumulated during the time evolution, see [Sec. 3.4](#). The enveloping algorithm is also responsible for preparing the initial state $|\psi(0)\rangle$.

While practical applications are concerned with specific Hamiltonian instances, quantum simulation algorithms apply to general classes of Hamiltonians satisfying mild technical conditions that enable a quantum algorithm to access the Hamiltonian efficiently. For example, a Hamiltonian can be specified as a linear combination of elementary interaction terms denoted $\hat{V}_1, \dots, \hat{V}_L$ such that

$$\hat{H} = \sum_{i=1}^L \alpha_i \hat{V}_i, \quad \|\hat{V}_i\| \leq 1. \quad (3.16)$$

Here α_i are real coefficients and $\|\hat{V}_i\|$ is the operator norm (the maximum magnitude eigenvalue). In the case of [local Hamiltonians](#) [\[212, 6\]](#), each term \hat{V}_i acts non-trivially only a few qubits. This includes an important special case of lattice Hamiltonians where qubits are located at sites of a regular lattice and the interactions \hat{V}_i couple small subsets of nearest-neighbor qubits. Molecular electronic Hamiltonians mapped to qubits assume the form Eq. (3.16), where \hat{V}_i are tensor products of Pauli operators which may have a super-constant weight. This situation is captured by the [Linear Combination of Unitaries](#) (LCU) model [\[213\]](#). It assumes that each term \hat{V}_i is a black-box unitary operator that can be implemented at a unit cost by querying an oracle (more precisely, one needs a “select- \hat{V} ” oracle implementing a controlled version of $\hat{V}_1, \dots, \hat{V}_m$). The LCU model also assumes an oracle access to the coefficients α_i , see [\[213\]](#) for details. Alternatively, a quantum algorithm can access the Hamiltonian through a subroutine that computes its matrix elements. A

Hamiltonian \hat{H} is said to be d -sparse [214] if the matrix of \hat{H} in the standard n -qubit basis has at most d non-zero entries in a single row or column. The sparse Hamiltonian model assumes that positions and values of the nonzero entries can be accessed by querying suitable oracles [215]. Most physically realistic quantum systems can be mapped to either local, LCU, or sparse qubit Hamiltonians such that the corresponding oracles are realized by a short quantum circuit. As described below, the runtime of quantum simulation algorithms is controlled by a dimensionless parameter T proportional to the product of the evolution time t and a suitable norm of the Hamiltonian. One can view T as an effective evolution time. A formal definition of T for various Hamiltonian models is as follows.

$$T = \begin{cases} t \max_i |\alpha_i| & \text{(Local)} \\ t \sum_i |\alpha_i| & \text{(LCU)} \\ td \|H\|_{\max} & \text{(Sparse)} \end{cases} \quad (3.17)$$

Here $\|H\|_{\max}$ denotes the maximum magnitude of a matrix element.

It is strongly believed that the Hamiltonian simulation problem is hard for classical computers. For example, Ref. [216] showed that *any* problem solvable on a quantum computer can be reduced to solving an instance of a suitably formalized Hamiltonian simulation problem with a local Hamiltonian. Technically speaking, the problem is BQP-complete [6]. All known classical methods capable of simulating general quantum systems of the above form require resources (time and memory) exponential in n . On the other hand, Feynman's original insight [211] was that a quantum computer should be capable of simulating many-body quantum dynamics efficiently, such that the simulation runtime grows only polynomially with the system size n and the evolution time T . This intuition was confirmed by Lloyd who gave the first quantum algorithm for simulating local Hamiltonians [212]. The algorithm exploits the *Trotter-Suzuki product formula*

$$e^{-it(\hat{H}_1 + \hat{H}_2 + \dots + \hat{H}_L)} \approx \left(e^{-it\hat{H}_1} e^{-it\hat{H}_2} \dots e^{-it\hat{H}_L} \right)^{\frac{t}{\delta}}. \quad (3.18)$$

By choosing a sufficiently small Trotter step δ one can approximate the evolution operator e^{-iHt} by a product of few-qubit operators describing evolution under individual interaction terms. Each few-qubit operator can be easily implemented by a short quantum circuit. The runtime of Lloyd's algorithm scales as [212, 217]

$$O(L^3 T^2 \epsilon^{-1}), \quad (3.19)$$

where T is the effective evolution time for local Hamiltonians, see Eq. (3.17). Importantly, the runtime scales polynomially with all relevant parameters.

Lloyd's algorithm was a breakthrough result demonstrating that quantum computers can indeed provide an exponential speedup over the best known classical algorithms for the task of simulating time evolution of quantum systems. However, it was quickly realized that the runtime of Lloyd's algorithm is unlikely to be optimal. Indeed, since any physical system simulates its own dynamics in a real time, one should expect that a universal quantum simulator can attain a runtime scaling only linearly with t . Moreover, for any realistic Hamiltonian composed of short-range interactions on a regular lattice, one should expect that the simulation runtime is linear in the space-time volume nt . Clearly, the scaling Eq. (3.19) falls far behind these expectations.

Algorithmic tools

The last decade has witnessed several improvements in the runtime scaling based on development of new algorithmic tools for Hamiltonian simulation. Most notably, a breakthrough work by Berry et al. [215, 213, 218] achieved an exponential speedup over Lloyd's algorithm with respect to the precision ϵ . A powerful algorithmic tool introduced in [215] is the so-called *LCU lemma* [215]. It shows how to construct a quantum circuit that implements an operator $\hat{U}' = \sum_{i=1}^M \beta_i \hat{U}_i$, where β_i are complex coefficients and \hat{U}_i are black-box unitary operators. Assuming that U' is close to a unitary operator, the lemma shows that U' can be well approximated by a quantum circuit of size roughly $M \sum_i |\beta_i|$ using roughly $\sum_i |\beta_i|$ queries to the oracle implementing $\hat{U}_1, \dots, \hat{U}_M$ (and their inverses). The simulation algorithm of Ref. [213] works by splitting the evolution into small intervals of length τ and using the truncated Taylor series approximation $e^{-i\tau\hat{H}} \approx \sum_{m=0}^K (-i\tau\hat{H})^m/m! \equiv \hat{U}_\tau$. Accordingly, $e^{-it\hat{H}} = (e^{-i\tau\hat{H}})^{t/\tau} \approx (\hat{U}_\tau)^{t/\tau}$. Substituting the LCU decomposition of H into the Taylor series one obtains an LCU decomposition of \hat{U}_τ . For a suitable choice of the truncation order K , the truncated series \hat{U}_τ is close to a unitary operator. Thus \hat{U}_τ can be well approximated by a quantum circuit using the LCU lemma. The runtime of this simulation algorithm, measured by the number of queries to the Hamiltonian oracles, scales as [213]

$$\frac{T \log(T/\epsilon)}{\log \log(T/\epsilon)}. \quad (3.20)$$

Here T is the effective evolution time for the LCU Hamiltonian model, see Eq. (3.17). This constitutes a square-root speedup with respect to T and an exponential speedup with respect to the precision compared with Lloyd's algorithm.

An important algorithmic tool proposed by Childs [219] is converting a Hamiltonian into a *quantum walk*. The latter is a unitary operator W that resembles the evolution operator $e^{-it\hat{H}}$ with a unit evolution time t . For a suitable normalization of \hat{H} , the quantum walk operator \hat{W} has eigenvalues $e^{\pm i \arcsin(E_\alpha)}$, where E_α are eigenvalues of \hat{H} . The corresponding eigenvectors of \hat{W} are simply related to those of \hat{H} . Unlike the true evolution operator, the quantum walk \hat{W} can be easily implemented using only a few queries to the oracles describing the Hamiltonian \hat{H} , e.g. using the LCU or the sparse models. To correct the discrepancy between \hat{W} and the true evolution operator, Low and Chuang [220] proposed the *Quantum Signal Processing* (QSP) method. One can view QSP as a compiling algorithm that takes as input a black-box unitary operator W , a function $f : \mathbb{C} \rightarrow \mathbb{C}$, and outputs a quantum circuit that realizes $f(\hat{W})$. Here it is understood that $f(\hat{W})$ has the same eigenvectors as \hat{W} while each eigenvalue z is mapped to $f(z)$. The circuit realizing $f(\hat{W})$ is expressed using controlled- \hat{W} gates and single-qubit gates on the control qubit. Remarkably, it can be shown that the Low and Chuang algorithm is optimal for the sparse Hamiltonian model [220]. Its runtime, measured by the number of queries to the Hamiltonian oracles, scales as

$$T + \frac{\log(1/\epsilon)}{\log \log(1/\epsilon)} \quad (3.21)$$

where T is the effective evolution time for the sparse model, see Eq. (3.17). This scaling is optimal in the sense that it matches previously known lower bounds [218, 215]. We note that simulation methods based on the sparse and LCU Hamiltonian models have been recently unified using a powerful framework known as *qubitization* [221]. It provides a general recipe for converting a Hamiltonian into a quantum walk using yet another oracular representation of a Hamiltonian known as a block encoding [221].

Algorithms based on the quantum walk (such as the QSP) or truncated Taylor series may not be the best choice for near-term applications since they require many ancillary qubits.

In contrast, the original Lloyd algorithm [212] and its generalizations based on *higher order product formulas* [222] require only as many qubits as needed to express the Hamiltonian. In addition, such algorithms are well-suited for simulating lattice Hamiltonians where qubits are located at sites of a regular D -dimensional grid and each elementary interaction V_i couples a few qubits located nearby. Lattice Hamiltonians contain $L = O(n)$ elementary interactions. Each few-qubit operator that appears in a product formula approximating $e^{-it\hat{H}}$ can be expressed using a few gates that couple nearest-neighbor




Figure 3.9: Trotter-like time evolution (top) and divide-and-conquer strategy for lattice systems (bottom). From [224] and [225] respectively.

qubits. The corresponding quantum circuit can be easily implemented on a device whose qubit connectivity graph is a D -dimensional grid.

Quite recently, Childs and Su [223] revisited simulation algorithms based on product formulas and demonstrated that their performance is better than what one could expect from naive error bounds. More precisely, an order- p product formula approximates the evolution operator e^{-iHt} under a Hamiltonian $\hat{H} = \hat{A} + \hat{B}$ with a simpler operator that involves time evolutions under Hamiltonians A and B such that the approximation error scales as t^{p+1} in the limit $t \rightarrow 0$. Childs and Su [223] showed that the gate complexity of simulating a lattice Hamiltonian using order- p product formulas scales as $(nT)^{1+1/p} \epsilon^{-1/p}$ which shaves off a factor of n from the best previously known bound. Here T is the effective evolution time for the local Hamiltonian model.

A major breakthrough in simulating lattice Hamiltonians has been recently achieved by Haah et al. [225]. This work introduced a new class of product formulas based on the *divide-and-conquer* strategy and showed that lattice

Hamiltonians can be simulated with gate complexity $\tilde{O}(nT)$, where \tilde{O} hides logarithmic corrections.

This result confirms the physical intuition that the cost of simulating lattice Hamiltonians scales linearly with the space-time volume. The algorithm of Ref. [225], sketched in Fig. 3.9, approximates the full evolution operator by dividing the lattice into small (overlapping) subsystems comprising $O(\log n)$ qubits each and simulating time evolution of the individual subsystems.

The errors introduced by truncating the Hamiltonian near the boundaries are canceled by alternating between forward and backward time evolutions. The error analysis is based on a skillful application of the *Lieb-Robinson bound* [226, 227] that controls how fast information can propagate across the system during the time evolution.

Finally, we note that while product formulas achieve a better scaling with system size than LCU or QSP methods by exploiting commutativity of Hamiltonian terms, they suffer from worse scaling with the simulation time and the error tolerance. A recent approach of *multiproduct formulas* [228] combines the best features of both. From Trotter methods, it inherits the simplicity, low-space requirements of its circuits, and a good scaling with system size. From LCU, it inherits the optimal scaling with time and error, up to logarithmic factors. Essentially, the work [228] shows how a certain type of high-order product formula can be implemented with a polynomial gate cost in the order (scaling as p^2). In contrast, the standard Trotter-Suzuki formulas scale exponentially with the order (scaling as 5^p).

Open problems

The existing simulation methods such as QSP are optimal in terms of the query complexity. However their runtime may or may not be optimal if one accounts for the cost of implementing the Hamiltonian oracles by a quantum circuit. Indeed, in the case of lattice Hamiltonians, *system-specific simulation strategies* that are not based on oracular models are known to achieve better runtime [225]. Improved system-specific simulation methods are also available for quantum chemistry systems. For example Ref. [143] achieved a quadratic reduction in the number of interaction terms present in molecular Hamiltonians by treating the kinetic energy and the potential/Coulomb energy operators using two different sets of basis functions – the plane wave basis and its dual. The two bases are related by the fermionic version of the

Fourier transform which admits a simple quantum circuit [143]. This simplification was shown to reduce the depth of simulation circuits (parallel runtime) based on the Trotter-Suzuki and LCU decompositions [143]. One may anticipate that further improvements can be made by exploiting system-specific information. It is an interesting open question whether molecular Hamiltonians can be simulated in depth scaling poly-logarithmically with the size of the electronic basis. More recent advanced sparse Hamiltonian simulation algorithms can exploit prior knowledge of other Hamiltonian norms to get better performance. For instance, the algorithm of Ref. [229] has an effective time $T = \sqrt{d\|H\|_{\max}\|H\|_1}$, which is a tighter bound than the one displayed in Eq. (3.17). This can be further improved to $T = t\sqrt{d}\|H\|_{1\rightarrow 2}$ which is a tighter bound than both, see [230] for details. However, these scaling improvements come with a larger constant factor that has not been thoroughly characterized.

In many situations the Hamiltonian to be simulated can be written as $H = H_0 + V$, where the norm of H_0 is much larger than the norm of V , while the time evolution generated by H_0 can be “fast-forwarded” by calculating the evolution operator analytically. For example, H_0 could represent the kinetic energy in the momentum basis or, alternatively, potential energy in the position basis. *Hamiltonian simulation in the interaction picture* [231] is an algorithmic tool proposed to take advantage of such situations. It allows one to pay a logarithmic cost (instead of the usual linear cost) with respect to the norm of H_0 . This tool has been used to simulate chemistry in the plane wave basis with $O(N^2)$ gates in second quantization, where N is the number of plane waves [231]. In the first-quantization, the gate cost scales as $O(N^{1/3}K^{8/3})$, where K is the number of electrons [232].

Likewise, better bounds on the runtime may be obtained by exploiting the structure of the initial state $\psi(0)$. In many applications, $\psi(0)$ is the Hartree-Fock approximation to the true ground state. As such, the energy of $\psi(0)$ tends to be small compared with the full energy scale of \hat{H} and one may expect that the time evolution is confined to the low-energy subspace of \hat{H} . How to develop state-specific simulation methods and improved bounds on the runtime is an intriguing open problem posed in [233].

An interesting alternative to quantum simulations proposed by Poulin et al. [209] is using the quantum walk operator \hat{W} described above directly in the quantum phase estimation method to estimate eigenvalues of \hat{H} . This circumvents errors introduced by the Trotter or LCU decomposition reducing the

total gate count. It remains to be seen whether other physically relevant quantities such as time dependent correlation functions can be extracted from the quantum walk operator sidestepping the unitary evolution which is relatively expensive compared with the quantum walk.

While the asymptotic runtime scaling is of great theoretical interest, practical applications are mostly concerned with specific Hamiltonian instances and constant approximation error (e.g. four digits of precision). To assess the practicality of quantum algorithms for specific problem instances and compare their runtime with that of state-of-the-art classical algorithms, one has to examine *compiling methods* that transform a high-level description of a quantum algorithm into a quantum circuit that can be run on particular hardware. This motivates development of compilers tailored to quantum simulation circuits. For example, Hastings et al. [234] examined Trotter-Suzuki type simulation of chemical Hamiltonians mapped to qubits using the Jordan-Wigner transformation. An improved compiling method was proposed reducing the runtime by a factor $O(N)$, where N is the number of orbitals. Childs et al. [235] and Campbell [217] recently proposed a randomized compiler tailored to Trotter-Suzuki simulations. It improves the asymptotic runtime scaling Eq. (3.19) to $L^{5/2}T^{3/2}\epsilon^{-1/2}$ and achieves almost 10^5 speedup for simulation of small molecule quantum chemistry Hamiltonians [235, 217]. Different types of compilers may be needed for NISQ devices and fault-tolerant quantum machines [236]. Indeed, in the context of NISQ devices, arbitrary single-qubit gates are cheap and the simulation cost is dominated by the number of two-qubit gates. In contrast, the cost of implementing an error-corrected quantum circuit is usually dominated by the number of non-Clifford gates such as the T-gate or the CCZ-gate. Since large-scale quantum simulations are expected to require error correction, developing compiling algorithms minimizing the T-gate count is vital. For example, Low et al. [237] recently achieved a square-root reduction in the T-count for a state preparation subroutine employed by the LCU simulation method. Compiling algorithms that minimize the number of generic single-qubit rotations (which are expensive to implement fault-tolerantly) were investigated by Poulin et al. [209]. This work considered implementation of the quantum walk operator W associated with a lattice Hamiltonian that contains only a few distinct parameters (e.g. translation invariant models).

Estimating resources required to solve practically important problems is linked with the study of *space-time tradeoffs* in quantum simulation [238, 239, 240]. For example, in certain situations the circuit depth (parallel runtime) can be reduced at the cost of introducing ancillary qubits [241] and/or

using intermediate measurements and feedback [242]. Of particular interest are methods for reducing the size of quantum simulation circuits using “dirty” ancilla — qubits whose initial state is unknown and which must be restored to their original form upon the completion of the algorithm [241, 243, 237], see also [244]. For example, if a Hamiltonian simulation circuit is invoked as a subroutine from a larger enveloping algorithm, the role of dirty ancillas may be played by data qubits borrowed from different parts of the algorithm.

The Hamiltonian simulation problem has a purely classical version. It deals with the ODE $\dot{\mathbf{p}} = -\partial H/\partial \mathbf{q}$, $\dot{\mathbf{q}} = \partial H/\partial \mathbf{p}$, where \mathbf{p}, \mathbf{q} are canonical coordinates of a classical Hamiltonian system. Numerical algorithms for integrating classical Hamiltonian dynamics known as *symplectic integrators* have a long history [245, 246] and are widely applied in simulations of molecular dynamics [247]. One may ask whether new advances in quantum or classical Hamiltonian simulation algorithms can be made by an exchange of ideas between these two fields.

Finally, the problem of *simulating Hamiltonians in the presence of noise*, without active error correction in the NISQ era, is an open problem. The primary question in this context is whether noise tolerant methods, analogous to variational algorithms in ground-state energy optimization, can also be developed for Hamiltonian simulation.

3.6 Finite-temperature algorithms

How a quantum computer can be used to simulate experiments on quantum systems in thermal equilibrium is an important problem in the field of quantum simulation. Early algorithms for the simulation of Gibbs states [248, 249, 250] were based on the idea of coupling the system of interest to a set of ancillary qubits and letting the system and bath together evolve under a joint Hamiltonian, thus mimicking the physical process of thermalization. The main disadvantages of thermalization-based methods are the presence of additional ancillary qubits defining the bath states, and the need to perform time evolution under \hat{H} for a thermalization time t that could be very long.

More recent proposals have focused on ways to generate finite-temperature observables without long system-bath thermalization times and/or large ancilla bath representations. For example, Ref. [251] showed how to realize the

imaginary time evolution operator $e^{-\beta\hat{H}/2}$ using Hamiltonian simulation techniques. Applying a suitable version of the Hubbard-Stratonovich transformation the authors of Ref. [251] obtained a representation $e^{-\beta\hat{H}/2} = \sum_{\alpha} c_{\alpha} U_{\alpha}$, where c_{α} are real coefficients and U_{α} are unitary operators describing the time evolution under a Hamiltonian $\hat{H}^{1/2}$. Although the square-root $\hat{H}^{1/2}$ is generally not easily available, Ref. [251] showed how to realize it for Hamiltonians composed of few-qubit positive semidefinite terms using an ancillary system. Applying state-of-the-art methods to simulate time evolution under $\hat{H}^{1/2}$ and the LCU Lemma (see Section 3.5.1) to realize the desired linear combination of unitaries, Ref. [251] obtained a quantum algorithm for preparing the thermal Gibbs state with gate complexity $\beta^{1/2} 2^{n/2} \mathcal{Z}^{-1/2}$, where \mathcal{Z} is the quantum partition function. Here we ignored a prefactor scaling poly-logarithmically with β and the inverse error tolerance. A closely related but slightly less efficient algorithm was discussed in Ref. [252].

Several more heuristic quantum algorithms for finite-temperature simulations have been proposed recently. However, most of these algorithms are challenging to analyze mathematically and generally do not have performance guarantees. For example, *variational ansatz states for the Gibbs state* that can be prepared with simple circuits have been proposed to bypass possibly long thermalization times. One example is the *product spectrum ansatz* [253] (PSA), where a shallow unitary circuit applied to a product thermal state is chosen to minimize the free energy of the system.

A different avenue is to sample from the Gibbs state rather than generate it explicitly on the quantum computer. For example, *quantum Metropolis sampling* [254] samples from the Gibbs state in an analog of classical Metropolis sampling, using phase estimation on a random unitary applied to the physical qubits to “propose” moves, and an iterative amplification procedure to implement the “rejection”. Much like the classical Metropolis algorithm, the fixed point of this procedure samples the Gibbs state. Alternatively, the *quantum minimally entangled typical thermal state (METTS) algorithm* [166] samples from the Gibbs state using imaginary time evolution applied to pure states, implemented via the quantum imaginary time evolution algorithm. One strength of quantum METTS is that it uses only the physical qubits of the system and potentially shallow circuits, thus making it feasible even in the NISQ era, where it has been demonstrated on quantum hardware for small spin systems. An application to the Heisenberg model is shown in Fig. 3.10.

Another alternative is to work within the microcanonical ensemble. This




Figure 3.10: Application of the quantum METTS algorithm to the 6-site 1D Heisenberg model. Energy is shown as a function of inverse temperature β , using unitaries acting on $D = 2 \dots 5$ spins. From [166].

is the basis of the *minimal effective Gibbs ansatz (MEGA)* [255], which attempts to generate pure states within the energy window corresponding to a microcanonical ensemble (for example, using phase estimation). The basic challenge is to ensure that the energy window is chosen according to the desired temperature. The MEGA method estimates the temperature from asymptotic properties of the ratio of the greater and lesser Green's functions, which in principle can be measured on the quantum device using the techniques in Sec. 3.9.

While there are many different proposed techniques for estimating observables of thermal states which all appear quite plausible on theoretical grounds, little is known about their heuristic performance, and almost none have been tested on real devices. In this sense, thermal state algorithms lag greatly behind those for ground-states for problems of interest in chemistry and materials science. To identify the best way forward, heuristic benchmarking for systems of relevance in physics and chemistry will be of major importance. The problem of benchmarking is further discussed in Sec. 3.8.

3.7 Hybrid quantum-classical methods

3.7.1 Quantum embedding methods

Embedding algorithms use a divide-and-conquer strategy to break a large quantum simulation into smaller pieces that are more amenable to simulation. The properties of the original model and the reduced models are related to each other in an (ideally) self-consistent fashion. These methods are popular both in condensed matter physics to study correlated electronic materials, where they reduce the problem of solving a bulk fermionic lattice model to that of studying a simpler Anderson-like impurity model, as well as in molecular applications, to reduce the computational scaling of methods in the simulation of complicated molecules.

In classical simulations, there are many flavors of quantum embedding. These can be grouped roughly by the choice of variable used to communicate between regions; *dynamical mean-field theory* (DMFT) works with the Green's function and self-energy [256, 257, 258]; *density matrix embedding theory* with the one-body density matrix [259, 258]; *density functional embedding* via the electron density [260, 258], and other methods, such as *QM/MM*, *ONIOM*, and *fragment MO* methods communicate via the electrostatic potential [261].

There has been growing interest in quantum embedding methods in the quantum information community, with the quantum computer playing the role of the quantum mechanical solver for the fragment/impurity problem (see Fig. 3.11). Ref. [224] suggested that a small quantum computer with a few hundred qubits could potentially speed up material simulations based on the DMFT method, and proposed a quantum algorithm for computing the Green's function of a quantum impurity model. Kreula et al. [263] subsequently proposed a proof-of-principle demonstration of this algorithm. Similarly, Rubin [264] and Yamazaki et al [262] have explored the potential of DMET in conjunction with a quantum computer for both condensed-phase lattice models as well as for large molecular calculations. Ground states of quantum impurity models and their structural properties in the context of quantum algorithms have been analyzed in Ref. [265].

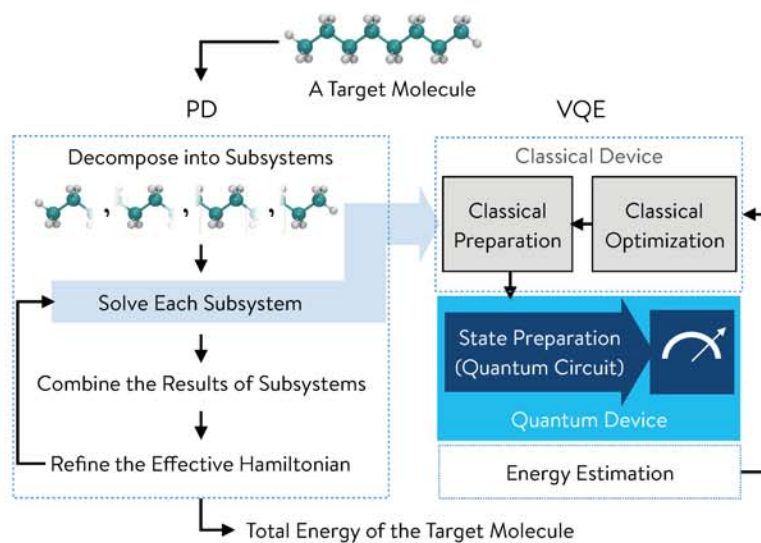


Figure 3.11: Illustration of hybrid quantum-classical embedding for a quantum chemistry system. Fragmentation of the original problem is performed on a classical computing device, and the more complex task of simulating each subproblem is handled by the quantum algorithm. From [262].

3.7.2 Other hybrid quantum-classical algorithms

Beyond standard quantum embedding, there are many other possibilities for hybrid quantum-classical algorithms. Variational quantum eigensolvers have previously been discussed for eigenstate problems, see [Sec. 3.3.2](#). One can also classically postprocess a quantum simulation of an eigenstate to improve it, in an analog of post-Hartree-Fock and post-complete-active-space methods, as has been explored in [\[266\]](#). In quantum molecular dynamics, it is natural to use a quantum computer to propagate the wavefunction or density matrix subject to motion of the nuclei, as is done today in classical Born-Oppenheimer dynamics. Similarly, the use of quantum optimizers and quantum annealers, for example to assist classical conformational search [\[267\]](#), can also be viewed as types of quantum-classical hybrids.

3.7.3 Open questions

It is clear that one will rely on hybrid quantum-classical algorithms for many years to come, and there remain many open questions. One is how to best adapt quantum algorithms within existing quantum-classical frameworks. For example, Green's function embedding methods are generally formulated in terms of actions rather than Hamiltonians; unfolding into a bath representation, consuming additional qubits, is currently required. More compact representations of the retarded interactions suited for quantum simulation should be explored (see for example, Ref. [\[268\]](#) for a related proposal to generate effective long-range interactions). Similarly one should explore the best way to evaluate Green's functions or density matrices, minimizing the coherence time and number of measurements. It is also possible that new kinds of embedding frameworks should be considered. For example, *quantum-quantum embedding* algorithms within the circuit model of quantum computation have been proposed to simulate large-scale "clustered" quantum circuits on a small quantum computer [\[269, 270\]](#). These are circuits that can be divided into small clusters such that there are only a few entangling gates connecting different clusters. Another promising class of embedding algorithms known as *holographic quantum simulators* was also recently proposed in [\[271, 272\]](#). Such algorithms enable the simulation of 2D lattice models on a 1D quantum computer by converting one spatial dimension into time. Whether it is useful to incorporate a classical component into such quantum-quantum frameworks clearly needs to be explored.

Another question is how best to implement the *feedback between the quantum and classical parts* of the algorithm. For example, as already discussed in the variational quantum eigensolver, such optimizations require the evaluation of approximate gradients on the quantum device, and to noisy optimization with limited gradient information on the classical device. Improving both aspects is clearly needed. Improving the performance of the quantum-classical interface also involves software engineering and building libraries between quantum and classical simulation software. These questions are discussed in section 4.

3.8 Benchmark systems

Researchers working on quantum simulation algorithms would greatly benefit from having access to well-defined benchmarks. Such benchmarks help the community by defining common conventions (e.g. choosing specific bases) and curating the best results.

There are two types of benchmarks to develop.

- the first is a benchmark that allows quantum algorithms (possibly on different hardware) to be compared against each other. For example, such problems could include a test suite of molecular Hamiltonian simulation problems for some specific choice of electron basis functions, fermion-to-qubit mapping, evolution time, and the desired approximation error.
- The second is curated data from the best classical methods for specific problems and well-defined Hamiltonians. Wherever possible, the data should not only include ground-state energies, but also excited-states and other observables, and if exact results are not available, an estimate of the precision should be given.

In the near-term era, suitable candidates for benchmarking may include molecular or material science Hamiltonians that can be expressed with about 50 or fewer qubits. Some promising candidates discussed in the literature include lattice spin Hamiltonians [273], and models of correlated electrons such as the 2D Fermi-Hubbard model [274], the uniform electron gas (jellium) [143], and the Haldane pseudo-potential Hamiltonian that models FQHE systems [275, 276, 277]. While some aspects of these models are easy to solve classically for systems of moderate size, others remain difficult, providing room for quantum advantage.

Naturally, there is a wide range of molecular or materials problems that could be chosen, a small number of which are highlighted in section 2. One relevant factor to check when constructing a Hamiltonian benchmark problem is to verify that it is indeed difficult to simulate classically [278]. Ideally, there should be a way to maximally tune the “complexity” for classical simulation, which then defines a natural setting for demonstrating quantum supremacy. Commonly, a way to tune the Hamiltonian (in model Hamiltonians) is to change the parameters of the Hamiltonian directly. In more realistic settings, one may change the size or geometry of the system or the chemical identity of the atoms.

3.9 Reading out results

3.9.1 Equal-time measurements

It is of course an essential part of any quantum computation to perform a measurement on the final state and thus read out the result of the computation. Conventionally, this is achieved through projective measurements of individual qubits in the computational basis. More complicated operators can be measured through standard techniques, for example by first applying a unitary rotation to the state that maps the operator onto Pauli- Z , or by using an ancillary qubit.

For very complicated operators, however, this can become quite resource-intensive. Consider for example measuring the expectation value of the Hamiltonian, which is required e.g. in the VQE (see **Sec. 3.3.2**). This can be done by writing the Hamiltonian as a sum of products of Pauli operators and measuring each one individually. Each measurement must be repeated a sufficient number of times to collect accurate statistics. Since for many applications, the number of operators grows quite quickly with the number of qubits (for example as N^4 in typical quantum chemistry applications) and the state may have to be prepared anew after a projective measurement, it is important to organize the terms in such a way that the number of operations to achieve a desired accuracy is minimized. Some work in this direction appears in Sec. V of Ref. [134], as well as more recently in Refs. [279, 280]. Nonetheless, the need to converge a large number of measurements to high precision presents a practical problem in many applications, particularly in many hybrid quantum-classical algorithms.




Figure 3.12: Quantum circuit to measure $\langle \Psi | \hat{A}_i^\dagger(t) \hat{B}_j | \Psi \rangle$. Here \hat{A}_i, \hat{B}_j are unitaries, $\hat{A}_i^\dagger(t) = \hat{U}(t) \hat{A}_i^\dagger \hat{U}^\dagger(t)$ is the time evolution of \hat{A}_i^\dagger under $\hat{U}(t) = e^{-it\hat{H}}$, and $\sigma_- = |0\rangle\langle 1| = \frac{X+iY}{2}$. Adapted from [281].

3.9.2 Dynamical properties and Green's functions

Much of the experimentally relevant information about a system, for example as obtained in scattering experiments such as optical or X-ray spectroscopy or neutron scattering, is encoded in dynamical properties. Access to these properties allows for a more direct comparison between theoretical predictions and experiments, thus allowing to infer microscopic information that is difficult to extract solely from the experiment.

A convenient way to capture this information is via the single- or few-particle *Green's functions*, which can be simply related to time-correlation functions of observables, such as the dynamic structure factors and dipole-dipole correlation functions. For example, the particle and hole components of the single-particle Green's function in real time ($t \geq 0$) are given by

$$G_{\alpha\beta}^p(t) = \langle \psi | \hat{c}_\alpha(t) \hat{c}_\beta^\dagger(0) | \psi \rangle \quad G_{\alpha\beta}^h(t) = \langle \psi | \hat{c}_\alpha(t)^\dagger \hat{c}_\beta(0) | \psi \rangle, \quad (3.22)$$

where α, β can be spin, orbital or site indices, $|\psi\rangle$ is the quantum state of interest (for example the ground state), and $\hat{c}_\alpha^{(\dagger)}(t) = e^{it\hat{H}} \hat{c}_\alpha^{(\dagger)} e^{-it\hat{H}}$. These can be measured by decomposing the fermion creation and annihilation operators into unitary combinations and applying standard techniques, see e.g. Ref. [281, 224] (a quantum circuit for measuring products of unitaries, one of which time-evolved, is shown in Fig. 3.12).

However, this procedure is expensive as it requires separate calculations for each time t , the particle/hole, real/imaginary and potentially orbital/spatial components of the Green's function, and finally a potentially large number of repetitions for each to achieve some desired target accuracy. While some initial work has shown how to improve this slightly by avoiding having to re-prepare the initial state every time (which could be prohibitive), further improvements and developments of alternatives, would be very valuable.

It should be noted that the behavior of the Green's function is very constrained both at short times (due to sum rules) and at long times (where the decay is governed by the longest time-scale of the system). Therefore, of primary interest is the regime of intermediate times where classical methods are most difficult to apply. However, in these other regimes, the additional structure could potentially be employed to reduce the computational effort also on a quantum computer.

4 Software stack

4.1 Quantum compilers and libraries

Until recently, programming a quantum computer was generally done by writing out an algorithm as a sequence of quantum gates. If one were to compare this to classical computers, it is as if one programmed in logical gates, such as NAND. In practice, almost all programming on classical computers is done in high-level languages that hide such low-level details from the programmer, thus drastically increasing productivity as well as portability of code between different types of processors. The step of compilation turns high-level code into instructions for the specific processor, and achieving the same in the realm of quantum computation is one goal of quantum compilers.

To this end, a variety of approaches based on either instruction sets embedded into existing programming frameworks or independent programming languages aimed directly at quantum computers have been developed, along with tools to compile the programs into a target gate set or classically simulate the circuit (see [Sec. 4.2](#) and [Fig. 4.1](#)). In addition to high-level programming languages, an important aspect are libraries of commonly used subroutines, such as quantum phase estimation or arithmetic operations. Domain specific libraries are discussed in [Sec. 4.3](#).

An important challenge in quantum compilation is the optimal translation of a given quantum circuit into native gates. In the case of a NISQ quantum computer, where operations are performed directly on physical qubits, this is the set of physical operations; in the case of a fault-tolerant quantum computer, it is the set of protected gates offered by the error correction code. An important class of target gates for fault-tolerant schemes are Clifford+ T circuits, i.e. circuits consisting of Clifford gates and $\pi/4$ rotations, referred to as T gates. Compiling into these has received a great deal of attention, see e.g. Refs. [\[283, 284\]](#).

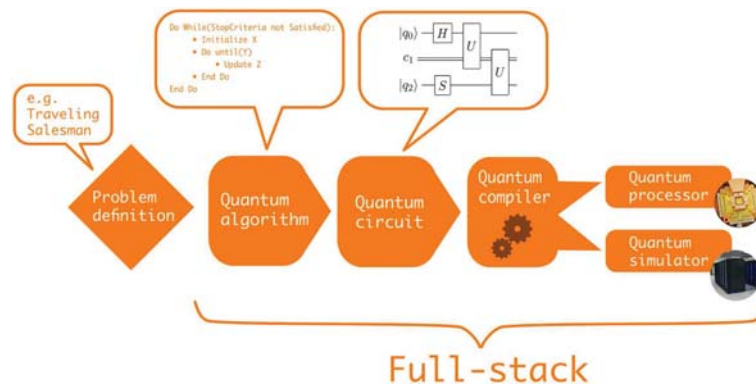


Figure 4.1: Typical workflow on a digital quantum simulator. A problem is first defined at a high-level and a suitable quantum algorithm is chosen. The quantum algorithm is then expressed as a quantum circuit, which in turn needs to be compiled to a specific quantum gate set. Finally, the quantum circuit is executed on a quantum processor or with a quantum circuit simulator. From [282].

4.2 Classical simulation of (noisy) quantum circuits

Testing out quantum algorithms by running their circuits on classical hardware has been of utmost importance to the field of quantum algorithms in recent years. Theoretical analysis by itself usually only bounds the asymptotic scaling of quantum algorithms, and as such is agnostic to prefactors that may well change which algorithm is preferable in practice. Furthermore, many algorithms, such as VQE, are heuristic and their performance cannot be guaranteed; in these cases, numerical experiments on classical hardware are the only way to assess the practical usefulness of the algorithms.

To facilitate such numerical tests, fast *quantum circuit simulators* can be used. In the general case, these scale exponentially with the number of qubits, although for very shallow circuits or special subclasses of gate sets (such as Clifford gates) more efficient algorithms may exist. Given this exponential scaling, up to 30 qubits can be simulated on desktop computers, and supercomputers can study systems of the order of 50 qubits. Several academic and industry groups are working on making high-performance simulators, both for desktop computers and massively parallel or cloud-based supercomputers, widely available to the public. At the same time, better methods are becoming avail-

able for simulating special classes of circuits, for example circuits that consist of a large number of Clifford gates and few non-Clifford gates (such as arbitrary rotations), which can be simulated in a time that is exponential only in the number of non-Clifford gates [285], or shallow circuits which can be evaluated more efficiently using a variety of techniques [286, 287, 288, 289, 290].

Another direction is the development of *approximate circuit simulation methods*, based e.g. on tensor networks [291, 292] or machine-learning approaches [293]. Expected aspects of NISQ circuits such as sparse connectivity of qubits, low depth, and noise, raise the potential viability of effective approximate simulation. Further, simulation of certain quantum algorithms for chemistry and materials may be effective for physical systems with limited correlation. Such simulations can help ascertain when quantum computation can provide more accurate solutions than classical approaches. The presence of noise in NISQ circuits can make them easier to simulate classically [294]. However, rigorous modeling of noise effects carries its own computational challenges. Standard methods for *simulation of noise* work by evolving a density matrix [295], which entails a quadratic increase in memory footprint. Approximate noise models such as the Pauli twirling approximation eliminate this overhead [296, 297]. Development of effective simulation mechanisms for noise and understanding its effect is essential to understanding the capabilities of quantum algorithms in the NISQ era.

4.3 Libraries for quantum simulation of quantum chemistry and quantum materials science

It is also essential to develop software to enable domain scientists in chemistry and physics to deploy or develop quantum computing methods without needing a detailed knowledge of all aspects of the quantum algorithms. There has been significant effort to develop packages and tools of this kind, some examples being OpenFermion [298], QISKit [299], and the Microsoft Quantum Development Kit (QDK) chemistry library [300]. Currently the packages have some overlapping capabilities but are different in scope and emphasis; for example, OpenFermion focuses on quantum chemistry applications; QISKit Aqua is a library of crossdomain quantum algorithms, which enables domain-specific users to run quantum algorithms across four areas:

quantum chemistry, artificial intelligence, optimization and finance; and the Microsoft QDK encompasses the Q# quantum programming language and libraries, while the Microsoft QDK chemistry library is a software stack that interfaces with the NWChem computational chemistry package [301] through a standardized “Broombridge” format, with the goal of automating the full quantum simulation of chemistry [302].

Given the proliferation of packages, domain specific languages, and programming interfaces, it is important to develop tools enabling *cross-platform development* and to develop *basic standards*. There are many industry efforts to develop user interfaces and GUI’s on top of the lower libraries. Improved tools for facilitating near-term experiments, including incorporating error mitigation or scheduling measurements for variational methods for specific hardware implementations, should also be developed.

4.4 Interfacing and incorporating classical simulation software

It is also important to properly interface to and incorporate classical simulation software. The types of interface will differ depending on the physical problem, but include

1. *classical preprocessing*, e.g. the generation of Hamiltonian matrix elements (integrals) or providing guess wave-functions for state preparation
2. *interfaces* for hybrid quantum-classical algorithms, e.g. for geometry optimization, orbital optimization, Born-Oppenheimer dynamics, and quantum embedding
3. *classical post-processing*, to interpret results or for further classical simulations, e.g. active space perturbation theory calculations, noise reduction, extrapolations, etc.

Distinct challenges arise in different applications. For example, the molecular quantum chemistry programs can typically provide molecular orbital Hamiltonian matrix elements to quantum algorithm software, but many condensed phase codes do not even compute such matrix elements as they are not needed in density functional calculations. Further, although such basic matrix elements are available in molecular codes, there are different conventions with

respect to basis function ordering, sign and normalization conventions, etc. We note also that many simulation programs are constructed in a relatively tightly coupled fashion making it hard to input information back into the program such as Green's functions or density matrices.

The software ecosystem also varies widely between different fields. While quantum chemistry and ab initio condensed phase electronic structure programs are often large multi-developer packages covering many simulation methods, and include both commercial and academic codes with large user communities, other simulation areas, such as lattice model condensed matter simulations, often involve more specialized programs focusing on a single technique. It is necessary to not only support the development of such specialized software but also the distribution, packaging, and education regarding its usage, in order for it to be used by the diverse set of researchers in the quantum algorithms area.

In general, it should be unnecessary for a quantum algorithms developer to learn about all the implementation quirks of classical simulation software. However, despite efforts by organizations such as the Molecular Software Sciences Institute ¹ to standardize where possible, it is also unrealistic to expect all codes to adopt and support the same conventions. A more fruitful strategy is to develop *dedicated middleware* between classical simulation software and quantum simulation software. Such middleware is already widely used to manage workflow, for example in the Materials Genome project (pymatgen) [303] or the Atomistic Simulation Environment (ASE) [304]. Middleware is also actively being developed in industry, for example in the Microsoft-PNNL collaboration, or in the QISKit and OpenFermion interfaces to open-source quantum chemistry packages such as Psi4 and PySCF. We note that these efforts have largely focused on quantum chemistry, and should be extended to other simulation problems.

Overall, the requirements of classical software to support the development of quantum algorithms overlap strongly with the general agenda of advancing simulation software infrastructure: standardizing interfaces; developing flexible middleware; better packaging of specialized codes; modularization of large codes. We expect therefore that resources devoted to supporting emerging quantum algorithms in classical software will benefit the wider classical simulation community as well.

¹<https://molssi.org>

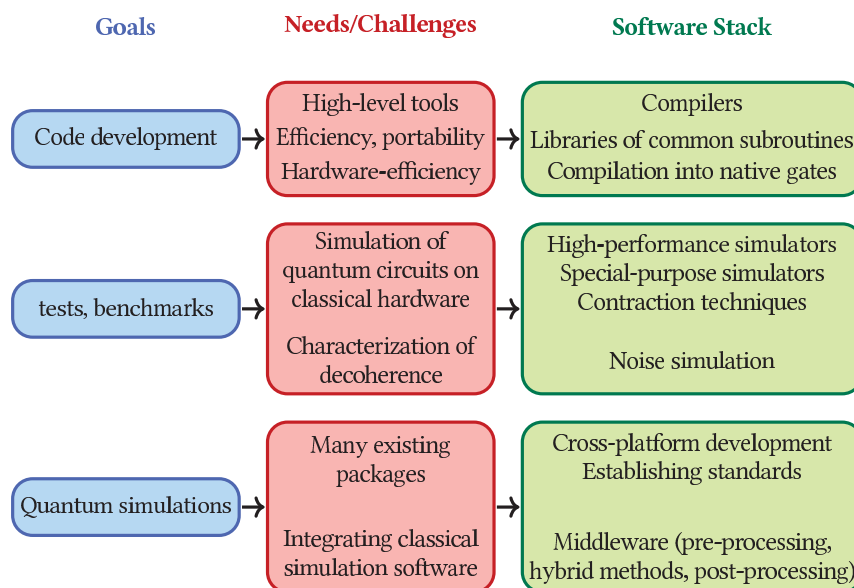


Figure 4.2: Development of quantum software for contemporary quantum simulation of quantum chemistry and materials science is motivated by research goals such as developing new quantum algorithms, benchmarking existing quantum algorithms on classical hardware, and performing simulations of molecules and materials. These goals entail specific technical needs and challenges (such as the development of compilers and libraries, of efficient classical simulators of quantum circuits in ideal and noisy conditions, and middleware to interface existing quantum packages between each other and with classical simulation software) that the community can address by the development of a quantum software stack.

4.5 Summary

In this section, we discussed some important goals of present-day quantum simulations, identified needs and challenges raised by these goals, and outlined a set of applications, utilities and routines – a software stack – to help the community advance efficiently. These items are summarized in the diagram in **Fig. 4.2**.

5 Broader issues

5.1 Diversity

Maintaining and enhancing diversity is critical to creating a large and qualified workforce. Given the rapid growth in the demand for qualified workers in the areas of quantum computing and quantum simulation, broadening participation is a time-critical problem. How to achieve diversity in all senses – gender, ethnicity, backgrounds, differently-able, geography, etc. – ultimately requires increasing the diversity of students who obtain chemistry, physics, math, and computer science degrees. Unfortunately, in some subdisciplines, recent years have seen the number of women and underrepresented minorities pursuing graduate degrees decrease rather than increase, highlighting the severity of the problem.

There are clearly no simple solutions. However, significant time and investment into this problem has been made by professional organizations (such as the American Chemical Society (ACS) and American Physical Society (APS) via their respective committees on diversity and women) as well as by industry (e.g. in the APS/IBM internship program ¹). It is important to adopt a data-driven approach and not to reinvent the wheel. For example, the set of *effective practices* documents produced by the APS ², or the *diversity resources* maintained by the ACS ³, are requisite reading. A common finding is that efforts should be organized around gathering and implementing and assessing best practices. Rather than list out all best practices, we mention particular ones highlighted during the workshop.

An important ingredient for a positive climate is the makeup of the leadership. In particular, it is important to have a chair/dean/high level manager

¹<https://www.aps.org/programs/women/scholarships/ibm/index.cfm>

²<https://www.aps.org/programs/women/reports/cswppractices/index.cfm>

³<https://www.acs.org/content/acs/en/membership-and-networks/acs/welcoming/diversity/diversity-resources.html>

who communicates well with faculty/students/employees, and with females and minorities in these groups in particular. The leadership needs to listen to concerns and to care enough to act both publicly and privately to effect needed changes, and to clarify behavior that will not be tolerated. The leadership also plays a central role in including diversity concerns in searches. Consequently, a potential avenue to ensure best practices are implemented is to focus on training of the key individuals in leadership positions.

Another question is when to start focusing on encouraging students and addressing the leaky pipeline. While public schools usually do not have time to introduce programs of any depth during the school year, during the undergraduate years there are more options. In particular, summer internships in STEM subjects for women and minorities have the potential for great impact. This is not solely because of the scientific training they provide, but because they introduce the student to mentors who play a central role in encouraging students to continue in the discipline. Ensuring that such mentors receive the appropriate training as described above may be one way to maximize the impact of internships. Local programs such as at universities can target promising students at the high school and junior high level and bring them in for shorter or longer periods during the summer for enrichment programs. Such university mentors of high school students can be encouraged and even nurtured as examples of best practices.

Given the diverse opinions and the various uncertainties in this area, it is also important to generate and maintain statistics to assess outcomes. With the limited resources that typical academic groups possess, it may be beneficial to engage efforts in the social sciences. Indeed, the early stage in the development of the quantum simulation workforce can be seen as a unique opportunity for innovation in diversity initiatives and for engagement between the social and physical sciences. Institutionalizing the metrics and assessment to simplify the process for future efforts will be valuable beyond the field of quantum simulation and quantum science.

5.2 Academia and industry

The development of quantum algorithms, hardware, and software for quantum simulation is currently a joint effort between academia and industry. This is a non-traditional situation and introduces new aspects to manage given the different motivations and interests of the various parties. For example, much

of the development of quantum hardware is focused in industry. What is the model for academia to engage with quantum hardware? What will be the funding model for this interaction? Will such machines be purchased by federal agencies for academic use?

Another challenge in the academia/industrial interaction is represented by the differing policies regarding *intellectual property* (IP). While freedom to publish is paramount to academia, industry is protective of intellectual property. The scope and ownership of possible IP needs to be clearly defined at the beginning of a collaboration, but in general academic collaborators do not know about IP rules or patent policies. It will be important for academics to educate themselves regarding these matters. The requisite material should be developed and provided online either by industry or in academic/industrial partnerships.

Finally, the academic/industry interface is complicated by the different organizational sizes. Many industrial efforts collate many scientists and would constitute multi-principal-investigator (multi-PI) efforts in an academic setting. Interacting as a single PI with such a large group presents a large disparity in size and power, which affects the academic interaction. This issue is also important regarding questions of standards and directions for the field. Ensuring that academics are represented by *collective organizations*, rather than only as individuals, will be important in ameliorating the imbalance in size and power in discussion and negotiations about scientific issues and directions in the field.

5.3 Education and training

Training students and researchers to work at the intersection of quantum chemistry, quantum materials science, and quantum algorithms requires new perspectives. The field needs the “right people”, but who these people are, and how to create a community of them, remains unclear.

In the area of undergraduate and graduate education, the relevant disciplines are physics, chemistry, computer science, and mathematics. We note that there is increasing demand in industry to hire people trained for quantum algorithms and quantum simulations at the bachelor’s and master’s level. There are deficiencies within all these undergraduate curricula with respect to training; typical chemistry and physics curricula leave little room for course

work in computer science and mathematics, and vice versa. While one way to enhance undergraduate education is via *internship programs* (see above) and *summer schools*, within the congested undergraduate curriculum it is most practical to condense the relevant information into existing courses rather than designing new curricula. Encouraging basic course requirements in introductory quantum mechanics, and ensuring that such a course contains concepts of quantum information as well as the usual chemical and physics applications is one possibility. Another is to include some aspects of quantum into introductory computer science (CS) and CS+X courses that are increasingly popular among chemistry and physics students. It is important not to waste effort and *developing shared curricula* in federally funded efforts should be encouraged.

At the graduate level, there is more room for *interdisciplinary education and programs*. One model to look towards is the Department of Energy Computational Science Graduate Fellowship program ⁴, which imposes strict program of study requirements for students of all academic backgrounds, requiring two graduate-level courses in three areas: a scientific field, mathematics, and computer science. We note that at the graduate level, it is important not only to cover traditional theoretical quantum computer science, but also more practical aspects including quantum software and hardware engineering. Given that expertise in these areas are spread across institutions multi-institutional efforts in the area of education and curricula development are a potential possibility for support. We note that industry has taken a proactive role in educational efforts into quantum computing in recent years, and new model of industry/academia collaboration in education should be developed.

5.4 Summary

In this Section, we explored some issues facing research activity in academia and at the interface with industry. We identified the need to build a diverse and inclusive workforce, of interacting in an informed manner with research in the private sector, and of increasing the presence of quantum mechanics and interdisciplinary requirements in various kinds of undergraduate and graduate courses. These issues are classified and summarized in **Fig. 5.1**.

⁴<https://www.krellinst.org/csgf/>

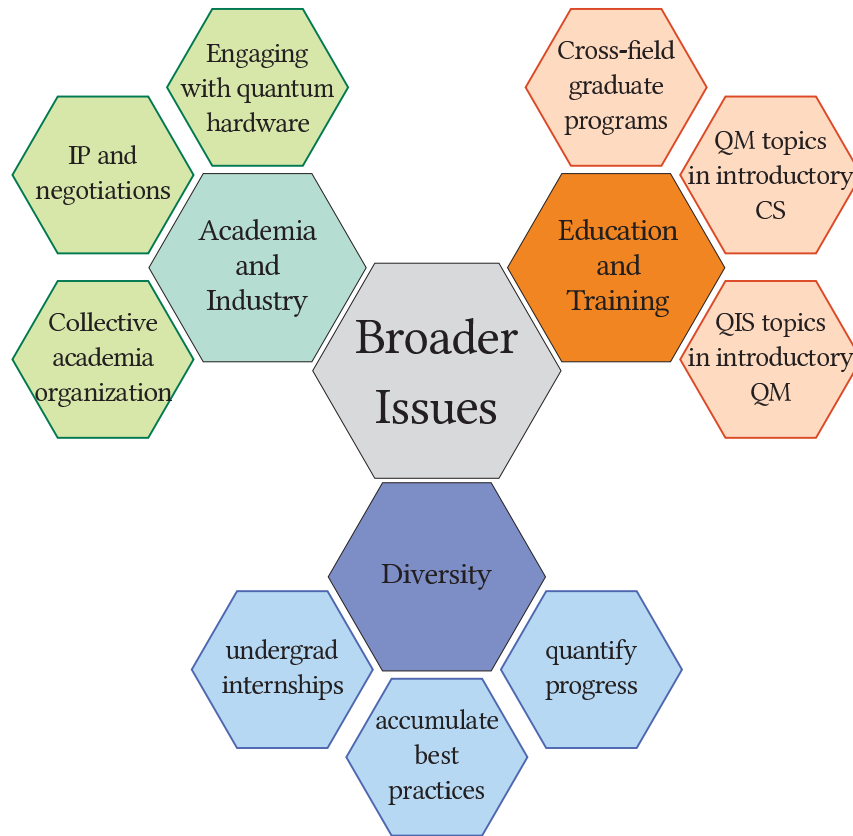


Figure 5.1: Issues discussed in **Sec. 5** related to diversity in research and the workforce, interactions between academia and industry, and education and training of undergraduate and graduate students. QM, CS, QIS and IP stand for quantum mechanics, computer science, quantum information science and intellectual property respectively.

Bibliography

- [1] Artur Ekert and Richard Jozsa. Quantum computation and Shor's factoring algorithm. *Reviews of Modern Physics*, 68(3):733, 1996.
- [2] Richard P Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6):467–488, 1982.
- [3] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000.
- [4] N.D. Mermin. *Quantum Computer Science: An Introduction*. Cambridge University Press, 2007.
- [5] G. Benenti, G. Casati, and G. Strini. *Principles of Quantum Computation and Information: Basic tools and special topics*. Principles of Quantum Computation and Information. World Scientific, 2004.
- [6] A.Y. Kitaev, A. Shen, M.N. Vyalyi, and M.N. Vyalyi. *Classical and Quantum Computation*. Graduate studies in mathematics. American Mathematical Society, 2002.
- [7] N.S. Yanofsky and M.A. Mannucci. *Quantum Computing for Computer Scientists*. Cambridge University Press, 2008.
- [8] D.C. Marinescu and G.M. Marinescu. *Approaching Quantum Computing*. Pearson/Prentice Hall, 2005.
- [9] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153–185, Mar 2014.
- [10] Michele Mosca and Artur Ekert. The hidden subgroup problem and eigenvalue estimation on a quantum computer. In *NASA International Conference on Quantum Computing and Quantum Communications*, pages 174–188. Springer, 1998.

- [11] Dorit Aharonov, Wim Van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008.
- [12] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science*, 292(5516):472–475, 2001.
- [13] Alexander I Johnson, Fhokrul Islam, Carlo M Canali, and Mark R Pederson. A multiferroic molecular magnetic qubit. *The Journal of Chemical Physics*, 151(17):174105, 2019.
- [14] Benjamin Kintzel, Michael Böhme, Junjie Liu, Anja Burkhardt, Jakub Mrozek, Axel Buchholz, Arzhang Ardavan, and Winfried Plass. Molecular electronic spin qubits from a spin-frustrated trinuclear copper complex. *Chemical communications*, 54(92):12934–12937, 2018.
- [15] G Wendin. Quantum information processing with superconducting circuits: a review. *Reports on Progress in Physics*, 80(10):106001, sep 2017.
- [16] Colin D Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 2019.
- [17] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [18] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.
- [19] Barbara M. Terhal. Quantum error correction for quantum memories. *Rev. Mod. Phys.*, 87:307–346, Apr 2015.
- [20] Robert O Jones. Density functional theory: Its origins, rise to prominence, and future. *Reviews of modern physics*, 87(3):897, 2015.
- [21] Rodney J Bartlett and Monika Musiał. Coupled-cluster theory in quantum chemistry. *Reviews of Modern Physics*, 79(1):291, 2007.
- [22] Manuel Sparta and Frank Neese. Chemical applications carried out by local pair natural orbital based coupled-cluster methods. *Chemical Society Reviews*, 43(14):5032–5041, 2014.

- [23] Zhendong Li, Sheng Guo, Qiming Sun, and Garnet Kin-Lic Chan. Electronic landscape of the p-cluster of nitrogenase as revealed through many-electron quantum wavefunction simulations. *Nature Chemistry*, pages 1–8, 2019.
- [24] Yuki Kurashige, Garnet Kin-Lic Chan, and Takeshi Yanai. Entangled quantum electronic wavefunctions of the Mn_4CaO_5 cluster in photosystem ii. *Nature chemistry*, 5(8):660, 2013.
- [25] Clyde W Cady, Robert H Crabtree, and Gary W Brudvig. Functional models for the oxygen-evolving complex of photosystem ii. *Coordination chemistry reviews*, 252(3-4):444–455, 2008.
- [26] Louis Noodleman, Timothy Lovell, Tiqing Liu, Fahmi Himo, and Rhonda A Torres. Insights into properties and energetics of iron–sulfur proteins from simple clusters to nitrogenase. *Current opinion in chemical biology*, 6(2):259–273, 2002.
- [27] Per EM Siegbahn. Structures and energetics for O_2 formation in photosystem II. *Accounts of Chemical Research*, 42(12):1871–1880, 2009.
- [28] Kyle M Lancaster, Michael Roemelt, Patrick Ettenhuber, Yilin Hu, Markus W Ribbe, Frank Neese, Uwe Bergmann, and Serena DeBeer. X-ray emission spectroscopy evidences a central carbon in the nitrogenase iron-molybdenum cofactor. *Science*, 334(6058):974–977, 2011.
- [29] Javaria Batool, Torsten Hahn, and Mark R Pederson. Magnetic signatures of hydroxyl-and water-terminated neutral and tetra-anionic Mn_{12} -Acetate. *Journal of computational chemistry*, 40(26):2301–2308, 2019.
- [30] Sandeep Sharma, Kantharuban Sivalingam, Frank Neese, and Garnet Kin-Lic Chan. Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics. *Nature chemistry*, 6(10):927, 2014.
- [31] Vijay Gopal Chilkuri, Serena DeBeer, and Frank Neese. Ligand field theory and angular overlap model based analysis of the electronic structure of homovalent iron–sulfur dimers. *Inorganic Chemistry*, 2019.
- [32] Lili Cao, Octav Caldararu, and Ulf Ryde. Protonation and reduction of the FeMo cluster in nitrogenase studied by quantum mechanics/molecular mechanics (QM/MM) calculations. *Journal of Chemical Theory and Computation*, 14(12):6653–6678, 2018.

- [33] L Schimka, J Harl, A Stroppa, A Grüneis, M Marsman, F Mittendorfer, and G Kresse. Accurate surface and adsorption energies from many-body perturbation theory. *Nature materials*, 9(9):741, 2010.
- [34] Marçal Capdevila-Cortada, Zbigniew Łodziana, and Núria López. Performance of dft+ u approaches in the study of catalytic materials. *ACS Catalysis*, 6:8370, 2016.
- [35] Jens K Nørskov, Frank Abild-Pedersen, Felix Studt, and Thomas Bligaard. Density functional theory in surface chemistry and catalysis. *Proceedings of the National Academy of Sciences*, 108(3):937–943, 2011.
- [36] Alec M Wodtke. Electronically non-adiabatic influences in surface chemistry and dynamics. *Chemical Society Reviews*, 45(13):3641–3657, 2016.
- [37] Jens Kehlet Nørskov, Thomas Bligaard, Jan Rossmeisl, and Claus H Christensen. Towards the computational design of solid catalysts. *Nature chemistry*, 1(1):37, 2009.
- [38] Yuan-Chung Cheng and Graham R Fleming. Dynamics of light harvesting in photosynthesis. *Annual review of physical chemistry*, 60:241–262, 2009.
- [39] Tomáš Polívka and Villy Sundström. Ultrafast dynamics of carotenoid excited states- from solution to natural and artificial systems. *Chemical Reviews*, 104(4):2021–2072, 2004.
- [40] Susanne Hahn and Gerhard Stock. Quantum-mechanical modeling of the femtosecond isomerization in rhodopsin. *The Journal of Physical Chemistry B*, 104(6):1146–1149, 2000.
- [41] Tadeusz Andruniów, Nicolas Ferré, and Massimo Olivucci. Structure, initial excited-state relaxation, and energy storage of rhodopsin resolved at the multiconfigurational perturbation theory level. *Proceedings of the National Academy of Sciences*, 101(52):17908–17913, 2004.
- [42] Francesco Segatta, Lorenzo Cupellini, Marco Garavelli, and Benedetta Mennucci. Quantum chemical modeling of the photoinduced activity of multichromophoric biosystems. *Chemical reviews*, 2019.
- [43] Paul Tavan and Klaus Schulten. Electronic excitations in finite and infinite polyenes. *Physical Review B*, 36(8):4337, 1987.

- [44] Weifeng Hu and Garnet Kin-Lic Chan. Excited-state geometry optimization with the density matrix renormalization group, as applied to polyenes. *Journal of chemical theory and computation*, 11(7):3000–3009, 2015.
- [45] Andreas BJ Parusel and Stefan Grimme. A theoretical study of the excited states of chlorophyll a and pheophytin a. *The Journal of Physical Chemistry B*, 104(22):5395–5398, 2000.
- [46] Jacopo Tomasi, Benedetta Mennucci, and Roberto Cammi. Quantum mechanical continuum solvation models. *Chemical reviews*, 105(8):2999–3094, 2005.
- [47] Joost VandeVondele, Urban Borstnik, and Jurg Hutter. Linear scaling self-consistent field calculations with millions of atoms in the condensed phase. *Journal of chemical theory and computation*, 8(10):3565–3573, 2012.
- [48] Christopher S. Wang, Jacob C. Curtis, Brian J. Lester, Yaxing Zhang, Yvonne Y. Gao, Jessica Freeze, Victor S. Batista, Patrick H. Vaccaro, Isaac L. Chuang, Luigi Frunzio, Liang Jiang, S. M. Girvin, and Robert J. Schoelkopf. Quantum simulation of molecular vibronic spectra on a superconducting bosonic processor. *arXiv preprint arXiv:1908.03598*, 2019.
- [49] Veronica Vaida. Spectroscopy of photoreactive systems: Implications for atmospheric chemistry. *The J. Phys. Chem. A*, 113(1):5–18, 2008.
- [50] Ulrike Heiter, Karin Lind, Martin Asplund, Paul S Barklem, Maria Bergemann, Laura Magrini, Thomas Masseron, Š Mikolaitis, JC Pickering, and MP Ruffoni. Atomic and molecular data for optical stellar spectroscopy. *Physica Scripta*, 90(5):054010, 2015.
- [51] Vincenzo Barone, Malgorzata Biczysko, and Cristina Puzzarini. Quantum chemistry meets spectroscopy for astrochemistry: Increasing complexity toward prebiotic molecules. *Accounts of chemical research*, 48(5):1413–1422, 2015.
- [52] Joel M Bowman, Tucker Carrington, and Hans-Dieter Meyer. Variational quantum approaches for computing vibrational energies of polyatomic molecules. *Molecular Physics*, 106(16-18):2145–2182, 2008.

- [53] Alexandra Viel, Maurício D. Coutinho-Neto, and Uwe Manthe. The ground state tunneling splitting and the zero point energy of malonaldehyde: A quantum Monte Carlo determination. *The Journal of Chemical Physics*, 126(2):024308, 2007.
- [54] M.H. Beck, A. Jäckle, G.A. Worth, and H.-D. Meyer. The multiconfiguration time-dependent hartree (MCTDH) method: a highly efficient algorithm for propagating wavepackets. *Physics Reports*, 324(1):1 – 105, 2000.
- [55] Robert Wodraszka and Uwe Manthe. CH_5^+ : Symmetry and the entangled rovibrational quantum states of a fluxional molecule. *The Journal of Physical Chemistry Letters*, 6(21):4229–4232, Nov 2015.
- [56] Ove Christiansen. Vibrational structure theory: new vibrational wave function methods for calculation of anharmonic vibrational energies and vibrational contributions to molecular properties. *Phys. Chem. Chem. Phys.*, 9:2942–2953, 2007.
- [57] Ove Christiansen. Selected new developments in vibrational structure theory: potential construction and vibrational wave function calculations. *Phys. Chem. Chem. Phys.*, 14:6672–6687, 2012.
- [58] Xiao-Gang Wang and Tucker Carrington Jr. New ideas for using contracted basis functions with a lanczos eigensolver for computing vibrational spectra of molecules with four or more atoms. *The Journal of chemical physics*, 117(15):6923–6934, 2002.
- [59] Xiao-Gang Wang and Tucker Carrington Jr. Vibrational energy levels of CH_5^+ . *The Journal of chemical physics*, 129(23):234102, 2008.
- [60] Dominik Marx and Michele Parrinello. CH_5^+ : The cheshire cat smiles. *Science*, 284(5411):59–61, 1999.
- [61] Takeshi Oka. Taming CH_5^+ , the “enfant terrible” of chemical structures. *Science*, 347(6228):1313–1314, 2015.
- [62] Edmund T White, Jian Tang, and Takeshi Oka. CH_5^+ : the infrared spectrum observed. *Science*, 284(5411):135–137, 1999.
- [63] Xinchuan Huang, Anne B McCoy, Joel M Bowman, Lindsay M Johnson, Chandra Savage, Feng Dong, and David J Nesbitt. Quantum deconstruction of the infrared spectrum of CH_5^+ . *Science*, 311(5757):60–63, 2006.

- [64] Xiao-Gang Wang and Tucker Carrington Jr. Calculated rotation-bending energy levels of CH_5^+ and a comparison with experiment. *The Journal of chemical physics*, 144(20):204304, 2016.
- [65] Oriol Vendrell, Michael Brill, Fabien Gatti, David Lauvergnat, and Hans-Dieter Meyer. Full dimensional (15-dimensional) quantum-dynamical simulation of the protonated water-dimer III: Mixed Jacobi-valence parametrization and benchmark results for the zero point energy, vibrationally excited states, and infrared spectrum. *The Journal of chemical physics*, 130(23):234305, 2009.
- [66] Qi Yu and Joel M Bowman. Classical, thermostated ring polymer, and quantum VSCF/VCI calculations of IR spectra of H_7O_3^+ and H_9O_4^+ (eigen) and comparison with experiment. *The Journal of Physical Chemistry A*, 123(7):1399–1409, 2019.
- [67] Kun Liu, MG Brown, C Carter, RJ Saykally, JK Gregory, and DC Clary. Characterization of a cage form of the water hexamer. *Nature*, 381(6582):501, 1996.
- [68] Yimin Wang, Bastiaan J Braams, Joel M Bowman, Stuart Carter, and David P Tew. Full-dimensional quantum calculations of ground-state tunneling splitting of malonaldehyde using an accurate ab initio potential energy surface. *The Journal of chemical physics*, 128(22):224314, 2008.
- [69] Mauricio D Coutinho-Neto, Alexandra Viel, and Uwe Manthe. The ground state tunneling splitting of malonaldehyde: Accurate full dimensional quantum dynamics calculations. *The Journal of chemical physics*, 121(19):9207–9210, 2004.
- [70] Thorsten Hammer and Uwe Manthe. Intramolecular proton transfer in malonaldehyde: Accurate multilayer multi-configurational time-dependent Hartree calculations. *The Journal of Chemical Physics*, 134(22):224305, 2011.
- [71] Markus Schröder, Fabien Gatti, and Hans-Dieter Meyer. Theoretical studies of the tunneling splitting of malonaldehyde using the multiconfiguration time-dependent hartree approach. *The Journal of Chemical Physics*, 134(23):234307, 2011.
- [72] Nicolas P. D. Sawaya and Joonsuk Huh. Quantum algorithm for calculating molecular vibronic spectra. *The Journal of Physical Chemistry Letters*, 10(13):3586–3591, Jul 2019.

- [73] Sam McArdle, Alexander Mayorov, Xiao Shan, Simon Benjamin, and Xiao Yuan. Digital quantum simulation of molecular vibrations. *Chem. Sci.*, 10:5725–5735, 2019.
- [74] Samuel M. Greene and Victor S. Batista. Tensor-train split-operator fourier transform (TT-SOFT) method: Multidimensional nonadiabatic quantum dynamics. *Journal of Chemical Theory and Computation*, 13(9):4034–4042, Sep 2017.
- [75] Ian R Craig and David E Manolopoulos. Quantum statistics and classical mechanics: Real time correlation functions from ring polymer molecular dynamics. *The Journal of chemical physics*, 121(8):3368–3373, 2004.
- [76] Jianshu Cao and Gregory A Voth. The formulation of quantum statistical mechanics based on the feynman path centroid density. iv. algorithms for centroid molecular dynamics. *The Journal of chemical physics*, 101(7):6168–6183, 1994.
- [77] Dmitrii E Makarov and Nancy Makri. Path integrals for dissipative systems by tensor multiplication. condensed phase quantum dynamics for arbitrarily long time. *Chemical physics letters*, 221(5-6):482–491, 1994.
- [78] Haobin Wang and Michael Thoss. Multilayer formulation of the multiconfiguration time-dependent hartree theory. *The Journal of chemical physics*, 119(3):1289–1299, 2003.
- [79] Reinhard J. Maurer, Mikhail Askerka, Victor S. Batista, and John C. Tully. Ab initio tensorial electronic friction for molecules on metal surfaces: Nonadiabatic vibrational relaxation. *Phys. Rev. B*, 94:115432, Sep 2016.
- [80] Mario Barbatti. Nonadiabatic dynamics with trajectory surface hopping method. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 1(4):620–633, 2011.
- [81] My Hang V Huynh and Thomas J Meyer. Proton-coupled electron transfer. *Chemical Reviews*, 107(11):5004–5064, 2007.
- [82] Roc Matheu, Mehmed Z. Ertem, Jordi Benet-Buchholz, Eugenio Coronado, Victor S. Batista, Xavier Sala, and Antoni Llobet. Intramolecular proton transfer boosts water oxidation catalyzed by a Ru complex. *Journal of the American Chemical Society*, 137(33):10786–10795, Aug 2015.

- [83] Aimin Ge, Pablo E. Videla, Gwendolynne L. Lee, Benjamin Rudshiteyn, Jia Song, Clifford P. Kubiak, Victor S. Batista, and Tianquan Lian. Interfacial structure and electric field probed by in situ electrochemical vibrational stark effect spectroscopy and computational modeling. *The Journal of Physical Chemistry C*, 121(34):18674–18682, Aug 2017.
- [84] Ting Wang, Gary W. Brudvig, and Victor S. Batista. Study of proton coupled electron transfer in a biomimetic dimanganese water oxidation catalyst with terminal water ligands. *Journal of Chemical Theory and Computation*, 6(8):2395–2401, Aug 2010.
- [85] Muhamed Amin, Leslie Vogt, Serguei Vassiliev, Ivan Rivalta, Mohammad M. Sultan, Doug Bruce, Gary W. Brudvig, Victor S. Batista, and M. R. Gunner. Electrostatic effects on proton coupled electron transfer in oxomanganese complexes inspired by the oxygen-evolving complex of photosystem II. *The Journal of Physical Chemistry B*, 117(20):6217–6226, May 2013.
- [86] Jinquan Chen, Kaifeng Wu, Benjamin Rudshiteyn, Yanyan Jia, Wendu Ding, Zhao-Xiong Xie, Victor S. Batista, and Tianquan Lian. Ultrafast photoinduced interfacial proton coupled electron transfer from CdSe quantum dots to 4,4'-bipyridine. *Journal of the American Chemical Society*, 138(3):884–892, Jan 2016.
- [87] Kuo Liao, Mikhail Askerka, Elizabeth L. Zeitler, Andrew Bocarsly, and Victor Batista. Electrochemical reduction of aqueous imidazolium on Pt(111) by proton coupled electron transfer. *Topics in Catalysis*, 58:23–29, 02 2014.
- [88] Elizabeth Hatcher, Alexander V Soudackov, and Sharon Hammes-Schiffer. Proton-coupled electron transfer in soybean lipoxygenase. *Journal of the American Chemical Society*, 126(18):5763–5775, 2004.
- [89] Martin Sjödin, Stenbjörn Styring, Björn Åkermark, Licheng Sun, and Leif Hammarström. Proton-coupled electron transfer from tyrosine in a tyrosine- ruthenium- tris-bipyridine complex: Comparison with tyrosine oxidation in photosystem ii. *Journal of the American Chemical Society*, 122(16):3932–3936, 2000.
- [90] Eduardo M. Sproviero, José A. Gascón, James P. McEvoy, Gary W. Brudvig, and Victor S. Batista. A model of the oxygen-evolving center of

- photosystem II predicted by structural refinement based on EXAFS simulations. *Journal of the American Chemical Society*, 130(21):6728–6730, May 2008.
- [91] Claudio Carra, Nedialka Iordanova, and Sharon Hammes-Schiffer. Proton-coupled electron transfer in a model for tyrosine oxidation in photosystem II. *Journal of the American Chemical Society*, 125(34):10429–10436, 2003.
- [92] Gerold U Bublitz and Steven G Boxer. Stark spectroscopy: applications in chemistry, biology, and materials science. *Annual review of physical chemistry*, 48(1):213–242, 1997.
- [93] Joel G Patrow, Shayne A Sorenson, and Jahan M Dawlaty. Direct spectroscopic measurement of interfacial electric fields near an electrode under polarizing or current-carrying conditions. *The Journal of Physical Chemistry C*, 121(21):11585–11592, 2017.
- [94] Abraham Nitzan. *Chemical dynamics in condensed phases: relaxation, transfer and reactions in condensed molecular systems*. Oxford university press, 2006.
- [95] James P Riehl and Frederick S Richardson. Circularly polarized luminescence spectroscopy. *Chemical Reviews*, 86(1):1–16, 1986.
- [96] Shaunak Mukherjee, Florian Libisch, Nicolas Large, Oara Neumann, Lisa V Brown, Jin Cheng, J Britt Lassiter, Emily A Carter, Peter Nordlander, and Naomi J Halas. Hot electrons do the impossible: plasmon-induced dissociation of H₂ on Au. *Nano letters*, 13(1):240–247, 2012.
- [97] Ivan Kassal, Stephen P Jordan, Peter J Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, 2008.
- [98] Bernhard Keimer, Steven A Kivelson, Michael R Norman, Shinichi Uchida, and J Zaanen. From quantum matter to high-temperature superconductivity in copper oxides. *Nature*, 518(7538):179, 2015.
- [99] Sung-Sik Lee. Recent developments in non-fermi liquid theory. *Annual Review of Condensed Matter Physics*, 9:227–244, 2018.
- [100] Eduardo Fradkin, Steven A. Kivelson, and John M. Tranquada. Colloquium: Theory of intertwined orders in high temperature superconductors. *Rev. Mod. Phys.*, 87:457–482, May 2015.

- [101] Piers Coleman. Heavy fermions: Electrons at the edge of magnetism. *Handbook of Magnetism and Advanced Magnetic Materials*, 2007.
- [102] Qimiao Si and Frank Steglich. Heavy fermions and quantum phase transitions. *Science*, 329(5996):1161–1166, 2010.
- [103] Zi-Xiang Li and Hong Yao. Sign-problem-free fermionic quantum monte carlo: Developments and applications. *Annual Review of Condensed Matter Physics*, 10:337–356, 2019.
- [104] David Tong. Lectures on the quantum hall effect. *arXiv preprint arXiv:1606.06687*, 2016.
- [105] A. H. Castro Neto, F. Guinea, N. M. R. Peres, K. S. Novoselov, and A. K. Geim. The electronic properties of graphene. *Rev. Mod. Phys.*, 81:109–162, Jan 2009.
- [106] Yuan Cao, Valla Fatemi, Shiang Fang, Kenji Watanabe, Takashi Taniguchi, Efthimios Kaxiras, and Pablo Jarillo-Herrero. Unconventional superconductivity in magic-angle graphene superlattices. *Nature*, 556(7699):43, 2018.
- [107] Allan H MacDonald. Trend: Bilayer graphene’s wicked, twisted road. *Physics*, 12:12, 2019.
- [108] Leon Balents. Spin liquids in frustrated magnets. *Nature*, 464(7286):199–208, 2010.
- [109] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006.
- [110] Y. Kasahara, K. Sugii, T. Ohnishi, M. Shimozawa, M. Yamashita, N. Kurita, H. Tanaka, J. Nasu, Y. Motome, T. Shibauchi, and Y. Matsuda. Unusual thermal hall effect in a Kitaev spin liquid candidate α - RuCl_3 . *Phys. Rev. Lett.*, 120:217205, May 2018.
- [111] Michele Buzzi, Michael Först, Roman Mankowsky, and Andrea Cavalleri. Probing dynamics in quantum materials with femtosecond x-rays. *Nature Reviews Materials*, 3(9):299, 2018.
- [112] Yuli V Nazarov and Yaroslav M Blanter. *Quantum transport: introduction to nanoscience*. Cambridge university press, 2009.
- [113] Andrea Damascelli. Probing the electronic structure of complex systems by ARPES. *Physica Scripta*, T109:61, 2004.

- [114] BD Patterson, R Abela, HH Braun, U Flechsig, R Ganter, Y Kim, E Kirk, A Oppelt, M Pedrozzi, S Reiche, et al. Coherent science at the SwissFEL X-ray laser. *New Journal of Physics*, 12(3):035012, 2010.
- [115] SP Weathersby, G Brown, Martin Centurion, TF Chase, Ryan Coffee, Jeff Corbett, JP Eichner, JC Frisch, AR Fry, M Gühr, et al. Mega-electron-volt ultrafast electron diffraction at slac national accelerator laboratory. *Review of Scientific Instruments*, 86(7):073702, 2015.
- [116] Martin C Fischer, Jesse W Wilson, Francisco E Robles, and Warren S Warren. Invited review article: pump-probe microscopy. *Review of Scientific Instruments*, 87(3):031101, 2016.
- [117] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. *Reviews of modern physics*, 80(3):885, 2008.
- [118] Tim Langen, Remi Geiger, and Jörg Schmiedmayer. Ultracold atoms out of equilibrium. *Annu. Rev. Condens. Matter Phys.*, 6(1):201–217, 2015.
- [119] Markus Greiner and Simon Fölling. Condensed-matter physics: Optical lattices. *Nature*, 453(7196):736, 2008.
- [120] Rahul Nandkishore and David A Huse. Many-body localization and thermalization in quantum statistical mechanics. *Annu. Rev. Condens. Matter Phys.*, 6(1):15–38, 2015.
- [121] Luca D’Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics. *Advances in Physics*, 65(3):239–362, 2016.
- [122] Dmitry A Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn. Colloquium: Many-body localization, thermalization, and entanglement. *Reviews of Modern Physics*, 91(2):021001, 2019.
- [123] Sebastian Paeckel, Thomas Köhler, Andreas Swoboda, Salvatore R Manmana, Ulrich Schollwöck, and Claudius Hubig. Time-evolution methods for matrix-product states. *arXiv preprint arXiv:1901.05824*, 2019.
- [124] Hideo Aoki, Naoto Tsuji, Martin Eckstein, Marcus Kollar, Takashi Oka, and Philipp Werner. Nonequilibrium dynamical mean-field theory and its applications. *Rev. Mod. Phys.*, 86:779–837, Jun 2014.
- [125] Miguel Marques, Angel Rubio, Eberhard KU Gross, Kieron Burke, Fernando Nogueira, and Carsten A Ullrich. *Time-dependent density functional theory*, volume 706. Springer Science & Business Media, 2006.

- [126] Trygve Helgaker, Poul Jorgensen, and Jeppe Olsen. *Molecular electronic-structure theory*. John Wiley & Sons, 2014.
- [127] Richard M Martin, Lucia Reining, and David M Ceperley. *Interacting electrons*. Cambridge University Press, 2016.
- [128] Daniel S Abrams and Seth Lloyd. Simulation of many-body fermi systems on a universal quantum computer. *Physical Review Letters*, 79(13):2586, 1997.
- [129] Ryan Babbush, Dominic W Berry, Yuval R Sanders, Ian D Kivlichan, Artur Scherer, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially more precise quantum simulation of fermions in the configuration interaction representation. *Quantum Science and Technology*, 3(1):015006, 2017.
- [130] Ian D Kivlichan, Nathan Wiebe, Ryan Babbush, and Alán Aspuru-Guzik. Bounding the costs of quantum simulation of many-body physics in real space. *Journal of Physics A: Mathematical and Theoretical*, 50(30):305301, 2017.
- [131] Sergey Bravyi, Jay M Gambetta, Antonio Mezzacapo, and Kristan Temme. Tapering off qubits to simulate fermionic hamiltonians. *arXiv preprint arXiv:1701.08213*, 2017.
- [132] Guang Hao Low and Isaac L Chuang. Optimal hamiltonian simulation by quantum signal processing. *Physical review letters*, 118(1):010501, 2017.
- [133] Alán Aspuru-Guzik, Anthony D Dutoi, Peter J Love, and Martin Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005.
- [134] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242, 2017.
- [135] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.
- [136] Mario Motta, David M Ceperley, Garnet Kin-Lic Chan, John A Gomez, Emanuel Gull, Sheng Guo, Carlos A Jiménez-Hoyos, Tran Nguyen Lan, Jia Li, Fengjie Ma, et al. Towards the solution of the many-electron problem in real materials: equation of state of the hydrogen chain with state-of-the-art many-body methods. *Physical Review X*, 7(3):031059, 2017.

- [137] Richard M Martin and Richard Milton Martin. *Electronic structure: basic theory and practical methods*. Cambridge university press, 2004.
- [138] Dominik Marx and Jürg Hutter. *Ab initio molecular dynamics: basic theory and advanced methods*. Cambridge University Press, 2009.
- [139] Wim Klopper, Frederick R Manby, Seiichiro Ten-No, and Edward F Valeev. R12 methods in explicitly correlated molecular electronic structure theory. *International Reviews in Physical Chemistry*, 25(3):427–468, 2006.
- [140] Liguao Kong, Florian A Bischoff, and Edward F Valeev. Explicitly correlated R12/F12 methods for electronic structure. *Chemical reviews*, 112(1):75–107, 2011.
- [141] E. Miles Stoudenmire and Steven R. White. Sliced basis density matrix renormalization group for electronic structure. *Phys. Rev. Lett.*, 119:046401, Jul 2017.
- [142] Steven R White and E Miles Stoudenmire. Multisliced gausslet basis sets for electronic structure. *Physical Review B*, 99(8):081110, 2019.
- [143] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan. Low-depth quantum simulation of materials. *Physical Review X*, 8(1):011044, 2018.
- [144] Steven R White. Hybrid grid/basis set discretizations of the schrödinger equation. *The Journal of chemical physics*, 147(24):244102, 2017.
- [145] Lin Lin, Jianfeng Lu, Lexing Ying, and E Weinan. Adaptive local basis set for Kohn–Sham density functional theory in a discontinuous galerkin framework i: Total energy calculation. *Journal of Computational Physics*, 231(4):2140–2154, 2012.
- [146] Mario Motta, Erika Ye, Jarrod R McClean, Zhendong Li, Austin J Minnich, Ryan Babbush, and Garnet Kin Chan. Low rank representations for quantum simulation of electronic structure. *arXiv preprint arXiv:1808.02625*, 2018.
- [147] Edward G Hohenstein, Robert M Parrish, and Todd J Martínez. Tensor hypercontraction density fitting. I. quartic scaling second-and third-order Møller-Plesset perturbation theory. *The Journal of chemical physics*, 137(4):044103, 2012.

- [148] S. B. Bravyi and A. Y. Kitaev. Fermionic quantum computation. *Annals of Physics*, 298:210–226, May 2002.
- [149] Vojtěch Havlíček, Matthias Troyer, and James D Whitfield. Operator locality in the quantum simulation of fermionic models. *Physical Review A*, 95(3):032332, 2017.
- [150] Jacob T Seeley, Martin J Richard, and Peter J Love. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *The Journal of chemical physics*, 137(22):224109, 2012.
- [151] Andrew Tranter, Sarah Sofia, Jake Seeley, Michael Kaicher, Jarrod McClean, Ryan Babbush, Peter V Coveney, Florian Mintert, Frank Wilhelm, and Peter J Love. The Bravyi–Kitaev transformation: Properties and applications. *International Journal of Quantum Chemistry*, 115(19):1431–1441, 2015.
- [152] Peter JJ O’Malley, Ryan Babbush, Ian D Kivlichan, Jonathan Romero, Jarrod R McClean, Rami Barends, Julian Kelly, Pedram Roushan, Andrew Tranter, Nan Ding, et al. Scalable quantum simulation of molecular energies. *Physical Review X*, 6(3):031007, 2016.
- [153] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P Lanyon, Peter Love, Ryan Babbush, et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Physical Review X*, 8(3):031022, 2018.
- [154] Nikolaj Moll, Andreas Fuhrer, Peter Staar, and Ivano Tavernelli. Optimizing qubit resources for quantum chemistry simulations in second quantization on a quantum computer. *Journal of Physics A: Mathematical and Theoretical*, 49(29):295301, 2016.
- [155] Mark Steudtner and Stephanie Wehner. Fermion-to-qubit mappings with varying resource requirements for quantum simulation. *New Journal of Physics*, 20(6):063010, 2018.
- [156] Sean A Fischer and Daniel Gunlycke. Symmetry configuration mapping for compact representation of quantum chemistry on quantum computers. *arXiv preprint arXiv:1907.01493*, 2019.
- [157] RC Ball. Fermions without fermion fields. *Physical review letters*, 95(17):176407, 2005.

- [158] Frank Verstraete and J Ignacio Cirac. Mapping local Hamiltonians of fermions to local Hamiltonians of spins. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(09):P09012, 2005.
- [159] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D Whitfield. Superfast encodings for fermionic quantum simulation. *arXiv preprint arXiv:1810.05274*, 2018.
- [160] Daniel T Colbert and William H Miller. A novel discrete variable representation for quantum mechanical reactive scattering via the S-matrix Kohn method. *The Journal of chemical physics*, 96(3):1982–1991, 1992.
- [161] Russell T Pack and Gregory A Parker. Quantum reactive scattering in three dimensions using hyperspherical (APH) coordinates. theory. *The Journal of chemical physics*, 87(7):3888–3921, 1987.
- [162] M Born and V. Fock. Beweis des adiabatensatzes. *Zeitschrift für Physik*, 51:165–180, 1928.
- [163] A. Messiah. *Quantum mechanics, Vol. 1*. North-Holland, 1965.
- [164] Tameem Albash and Daniel A Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1):015002, 2018.
- [165] Libor Veis and Jiří Pittner. Adiabatic state preparation study of methylene. *The Journal of chemical physics*, 140(21):214111, 2014.
- [166] Mario Motta, Chong Sun, Adrian Teck Keng Tan, Matthew J. O’ Rourke, Erika Ye, Austin J. Minnich, Fernando G. S. L. Brandao, and Garnet Kin-Lic Chan. Quantum imaginary time evolution, quantum lanczos, and quantum thermal averaging. *arXiv preprint arXiv:1901.07653*, 2019.
- [167] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- [168] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213, Jul 2014. Article.
- [169] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.

- [170] Rodney J Bartlett, Stanislaw A Kucharski, and Jozef Noga. Alternative coupled-cluster Ansätze II. the unitary coupled-cluster method. *Chemical physics letters*, 155(1):133–140, 1989.
- [171] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, Abhinav Kandala, Antonio Mezzacapo, Peter Müller, Walter Riess, Gian Salis, John Smolin, Ivano Tavernelli, and Kristan Temme. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology*, 3(3):030503, jun 2018.
- [172] Bridgette Cooper and Peter J. Knowles. Benchmark studies of variational, unitary and extended coupled cluster methods. *The Journal of Chemical Physics*, 133(23):234102, 2010.
- [173] Gaurav Harsha, Toru Shiozaki, and Gustavo E Scuseria. On the difference between variational and unitary coupled cluster theories. *The Journal of chemical physics*, 148(4):044107, 2018.
- [174] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, 2018.
- [175] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications*, 10(1):3007, 2019.
- [176] Steven R White. Density matrix formulation for quantum renormalization groups. *Physical review letters*, 69(19):2863, 1992.
- [177] M. Fannes, B. Nachtergaele, and R. F. Werner. Finitely correlated states on quantum spin chains. *Comm. Math. Phys.*, 144:443–490, 1992.
- [178] Ivan V Oseledets and Eugene E Tyrtyshnikov. Breaking the curse of dimensionality, or how to use svd in many dimensions. *SIAM Journal on Scientific Computing*, 31(5):3744–3759, 2009.
- [179] Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- [180] Wolfgang Hackbusch and Stefan Kühn. A new scheme for the tensor representation. *Journal of Fourier analysis and applications*, 15(5):706–722, 2009.

- [181] Lars Grasedyck. Hierarchical singular value decomposition of tensors. *SIAM Journal on Matrix Analysis and Applications*, 31(4):2029–2054, 2010.
- [182] F. Verstraete and J. I. Cirac. Renormalization algorithms for quantum-many body systems in two and higher dimensions. *arXiv preprint arXiv:cond-mat/0407066*, 2004.
- [183] G. Vidal. Class of quantum many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101:110501, Sep 2008.
- [184] C Schön, K Hammerer, Michael M Wolf, J Ignacio Cirac, and E Solano. Sequential generation of matrix-product states in cavity QED. *Physical Review A*, 75(3):032311, 2007.
- [185] William James Huggins, Piyush Patil, Bradley Mitchell, K Birgitta Whaley, and Miles Stoudenmire. Towards quantum machine learning with tensor networks. *Quantum Science and technology*, 2018.
- [186] Jin-Guo Liu, Yi-Hong Zhang, Yuan Wan, and Lei Wang. Variational quantum eigensolver with fewer qubits. *arXiv preprint arXiv:1902.02663*, 2019.
- [187] Isaac H Kim and Brian Swingle. Robust entanglement renormalization on a noisy quantum computer. *arXiv preprint arXiv:1711.07500*, 2017.
- [188] Martin Schwarz, Kristan Temme, and Frank Verstraete. Preparing projected entangled pair states on a quantum computer. *Physical review letters*, 108(11):110502, 2012.
- [189] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, 2014.
- [190] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *arXiv preprint arXiv:1509.04279*, 2015.
- [191] Zhi-Cheng Yang, Armin Rahmani, Alireza Shabani, Hartmut Neven, and Claudio Chamon. Optimizing variational quantum algorithms using pontryagin’s minimum principle. *Phys. Rev. X*, 7:021027, May 2017.

- [192] Jarrod R. McClean, Mollie E. Kimchi-Schwartz, Jonathan Carter, and Wibe A. de Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A*, 95:042308, Apr 2017.
- [193] Y. Li and S. C. Benjamin. Efficient variational quantum simulator incorporating active error minimisation. *arXiv preprint arXiv:1611.09301*, 2016.
- [194] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. Error mitigation for short depth quantum circuits. *arXiv preprint arXiv:1612.02058*, 2016.
- [195] Suguru Endo, Simon C. Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Phys. Rev. X*, 8:031027, Jul 2018.
- [196] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O’Brien. Low-cost error mitigation by symmetry verification. *Phys. Rev. A*, 98:062339, Dec 2018.
- [197] Sam McArdle, Xiao Yuan, and Simon Benjamin. Error-mitigated digital quantum simulation. *Phys. Rev. Lett.*, 122:180501, May 2019.
- [198] James I Colless, Vinay V Ramasesh, Dar Dahlen, Machiel S Blok, ME Kimchi-Schwartz, JR McClean, J Carter, WA De Jong, and I Siddiqi. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Physical Review X*, 8(1):011021, 2018.
- [199] Jarrod R. McClean, Zhang Jiang, Nicholas C. Rubin, Ryan Babbush, and Hartmut Neven. Decoding quantum errors with subspace expansions. *arXiv preprint arXiv:1903.05786*, 2019.
- [200] D. J. Rowe. Equations-of-motion method and the extended shell model. *Rev. Mod. Phys.*, 40:153–166, Jan 1968.
- [201] Marc Ganzhorn, Daniel J. Egger, Panagiotis Kl. Barkoutsos, Pauline Ollitrault, Gian Salis, Nikolaj Moll, Andreas Fuhrer, Peter Mueller, Stefan Woerner, Ivano Tavernelli, and Stefan Filipp. Gate-efficient simulation of molecular eigenstates on a quantum computer. *Preprint*, 2018.
- [202] Pawel Wocjan and Shengyu Zhang. Several natural BQP-complete problems. *arXiv preprint quant-ph/0606179*, 2006.
- [203] Krysta M Svore, Matthew B Hastings, and Michael Freedman. Faster phase estimation. *arXiv preprint arXiv:1304.0741*, 2013.

- [204] Nathan Wiebe and Chris Granade. Efficient Bayesian phase estimation. *Physical review letters*, 117(1):010503, 2016.
- [205] Shelby Kimmel, Guang Hao Low, and Theodore J Yoder. Robust calibration of a universal single-qubit gate set via robust phase estimation. *Physical Review A*, 92(6):062315, 2015.
- [206] Thomas Eugene O’Brien, Brian Tarasinski, and Barbara Terhal. Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments. *New Journal of Physics*, 2019.
- [207] Daniel S Abrams and Seth Lloyd. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Physical Review Letters*, 83(24):5162, 1999.
- [208] Marcello Benedetti, Erika Lloyd, and Stefan Sack. Parameterized quantum circuits as machine learning models. *arXiv preprint arXiv:1906.07682*, 2019.
- [209] David Poulin, Alexei Kitaev, Damian S Steiger, Matthew B Hastings, and Matthias Troyer. Quantum algorithm for spectral measurement with a lower gate count. *Physical review letters*, 121(1):010501, 2018.
- [210] Dominic W Berry, Mária Kieferová, Artur Scherer, Yuval R Sanders, Guang Hao Low, Nathan Wiebe, Craig Gidney, and Ryan Babbush. Improved techniques for preparing eigenstates of fermionic Hamiltonians. *npj Quantum Information*, 4(1):22, 2018.
- [211] Richard Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, June 1982.
- [212] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [213] Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114(9):090502, 2015.
- [214] Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC ’03, pages 20–29, New York, NY, USA, 2003. ACM.

- [215] D. W. Berry, A. M. Childs, and R. Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. *arXiv preprint arXiv:1501.01715*, 2015.
- [216] Daniel Nagaj and Pawel Wocjan. Hamiltonian quantum cellular automata in one dimension. *Physical Review A*, 78(3):032311, 2008.
- [217] Earl Campbell. A random compiler for fast hamiltonian simulation. *arXiv preprint arXiv:1811.08017*, 2018.
- [218] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma. Exponential improvement in precision for simulating sparse Hamiltonians. *arXiv preprint arXiv:1312.1414*, 2013.
- [219] Andrew M. Childs. On the relationship between continuous- and discrete-time quantum walk. *Communications in Mathematical Physics*, 294(2):581–603, 2010.
- [220] G. Hao Low and I. L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *arXiv preprint arXiv:1606.02685*, 2016.
- [221] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019.
- [222] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders. Efficient quantum algorithms for simulating sparse Hamiltonians. *Communications in Mathematical Physics*, 270:359–371, March 2007.
- [223] Andrew M Childs and Yuan Su. Nearly optimal lattice simulation by product formulas. *arXiv preprint arXiv:1901.00564*, 2019.
- [224] Bela Bauer, Dave Wecker, Andrew J Millis, Matthew B Hastings, and Matthias Troyer. Hybrid quantum-classical approach to correlated materials. *Physical Review X*, 6(3):031045, 2016.
- [225] Jeongwan Haah, Matthew Hastings, Robin Kothari, and Guang Hao Low. Quantum algorithm for simulating real time evolution of lattice hamiltonians. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 350–360. IEEE, 2018.
- [226] Elliott H Lieb and Derek W Robinson. The finite group velocity of quantum spin systems. In *Statistical mechanics*, pages 425–431. Springer, 1972.

- [227] Matthew B Hastings and Tohru Koma. Spectral gap and exponential decay of correlations. *Communications in mathematical physics*, 265(3):781–804, 2006.
- [228] Guang Hao Low, Vadym Kliuchnikov, and Nathan Wiebe. Well-conditioned multiproduct Hamiltonian simulation. *arXiv preprint arXiv:1907.11679*, 2019.
- [229] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by uniform spectral amplification. *arXiv preprint arXiv:1707.05391*, 2017.
- [230] Guang Hao Low. Hamiltonian simulation with nearly optimal dependence on spectral norm. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 491–502. ACM, 2019.
- [231] Guang Hao Low and Nathan Wiebe. Hamiltonian simulation in the interaction picture. *arXiv preprint arXiv:1805.00675*, 2018.
- [232] Ryan Babbush, Dominic W Berry, Jarrod R McClean, and Hartmut Neven. Quantum simulation of chemistry with sublinear scaling to the continuum. *arXiv preprint arXiv:1807.09802*, 2018.
- [233] Jeongwan Haah. Plenary talk at the 22nd Annual Conference on Quantum Information Processing (QIP), Boulder Colorado, Jan. 2019.
- [234] Matthew B. Hastings, David Wecker, Bela Bauer, and Matthias Troyer. Improving quantum algorithms for quantum chemistry. *Quantum Information and Computation*, 15, 2015.
- [235] Andrew M Childs, Aaron Ostrander, and Yuan Su. Faster quantum simulation by randomization. *arXiv preprint arXiv:1805.08385*, 2018.
- [236] Yunseong Nam and Dmitri Maslov. Low-cost quantum circuits for classically intractable instances of the Hamiltonian dynamics simulation problem. *npj Quantum Information*, 5(1):44, 2019.
- [237] Guang Hao Low, Vadym Kliuchnikov, and Luke Schaeffer. Trading t-gates for dirty qubits in state preparation and unitary synthesis. *arXiv preprint arXiv:1812.00954*, 2018.
- [238] Charles H Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989.
- [239] Alex Parent, Martin Roetteler, and Krysta M Svore. Reversible circuit compilation with space constraints. *arXiv preprint arXiv:1510.00377*, 2015.

- [240] Giulia Meuli, Mathias Soeken, Martin Roetteler, Nikolaj Björner, and Giovanni De Micheli. Reversible pebbling game for quantum memory management. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 288–291. IEEE, 2019.
- [241] Adriano Barenco, Charles H Bennett, Richard Cleve, David P DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical review A*, 52(5):3457, 1995.
- [242] Daniel Gottesman and Isaac L Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(6760):390, 1999.
- [243] Thomas Häner, Martin Roetteler, and Krysta M Svore. Factoring using $2n+2$ qubits with toffoli based modular multiplication. *arXiv preprint arXiv:1611.07995*, 2016.
- [244] Harry Buhrman, Richard Cleve, Michal Koucký, Bruno Loff, and Florian Speelman. Computing with a full memory: catalytic space. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 857–866. ACM, 2014.
- [245] Paul J Channell and Clint Scovel. Symplectic integration of Hamiltonian systems. *Nonlinearity*, 3(2):231, 1990.
- [246] Ronald D Ruth. A canonical integration technique. *IEEE Trans. Nucl. Sci.*, 30(CERN-LEP-TH-83-14):2669–2671, 1983.
- [247] Stephen K Gray, Donald W Noid, and Bobby G Sumpter. Symplectic integrators for large scale molecular dynamics simulations: A comparison of several explicit methods. *The Journal of chemical physics*, 101(5):4062–4072, 1994.
- [248] Barbara M. Terhal and David P. DiVincenzo. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A*, 61:022301, Jan 2000.
- [249] David Poulin and Pawel Wocjan. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.*, 103:220502, Nov 2009.
- [250] Arnau Riera, Christian Gogolin, and Jens Eisert. Thermalization in nature and on a quantum computer. *Phys. Rev. Lett.*, 108:080402, Feb 2012.

- [251] Anirban Narayan Chowdhury and Rolando D Somma. Quantum algorithms for gibbs sampling and hitting-time estimation. *arXiv preprint arXiv:1603.02940*, 2016.
- [252] David Poulin and Pawel Wocjan. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters*, 103(22):220502, 2009.
- [253] John Martyn and Brian Swingle. Product spectrum ansatz and the simplicity of thermal states. *arXiv:1812.01015*, 2018.
- [254] K Temme, TJ Osborne, Karl G Vollbrecht, D Poulin, and Frank Verstraete. Quantum metropolis sampling. *Nature*, 471(7336):87–90, 2011.
- [255] Jeffrey Cohn, Khadijeh Sona Najafi, Forest Yang, Barbara Jones, and James K. Freericks. Minimal effective gibbs ansatz (MEGA): A simple protocol for extracting an accurate thermal representation for quantum simulation. *arXiv preprint arXiv:1812.03607*, 2018.
- [256] G. Kotliar, S. Savrasov, K. Haule, V. Oudovenko, O. Parcollet, and C. Marianetti. Electronic structure calculations with dynamical mean-field theory. *Reviews of Modern Physics*, 78(3):865, 2006.
- [257] Antoine Georges. Strongly correlated electron materials: Dynamical mean field theory and electronic structure. *AIP Conference Proceedings*, 715(1):3–74, 2004.
- [258] Qiming Sun and Garnet Kin-Lic Chan. Quantum embedding theories. *Accounts of chemical research*, 49(12):2705–2712, 2016.
- [259] Gerald Knizia and Garnet Kin-Lic Chan. Density matrix embedding: A simple alternative to dynamical mean-field theory. *Physical review letters*, 109(18):186404, 2012.
- [260] Tomasz Adam Wesolowski and Arieh Warshel. Frozen density functional approach for ab initio calculations of solvated molecules. *The Journal of Physical Chemistry*, 97(30):8050–8053, 1993.
- [261] Hans Martin Senn and Walter Thiel. QM/MM methods for biomolecular systems. *Angewandte Chemie International Edition*, 48(7):1198–1229, 2009.

- [262] Takeshi Yamazaki, Shunji Matsuura, Ali Narimani, Anushervon Saidmuradov, and Arman Zaribafiyani. Towards the practical application of near-term quantum computers in quantum chemistry simulations: A problem decomposition approach. *arXiv preprint arXiv:1806.01305*, 2018.
- [263] Juha M Kreula, Laura García-Álvarez, Lucas Lamata, Stephen R Clark, Enrique Solano, and Dieter Jaksch. Few-qubit quantum-classical simulation of strongly correlated lattice fermions. *EPJ Quantum Technology*, 3(1):11, 2016.
- [264] Nicholas C Rubin. A hybrid classical/quantum approach for large-scale studies of quantum systems with density matrix embedding theory. *arXiv preprint arXiv:1610.06910*, 2016.
- [265] Sergey Bravyi and David Gosset. Complexity of quantum impurity problems. *Communications in Mathematical Physics*, 356(2):451–500, 2017.
- [266] Tyler Takeshita, Nicholas C Rubin, Zhang Jiang, Eunseok Lee, Ryan Babush, and Jarrod R McClean. Increasing the representation accuracy of quantum simulations of chemistry without extra quantum resources. *arXiv preprint arXiv:1902.10679*, 2019.
- [267] DJJ Marchand, M Noori, A Roberts, G Rosenberg, B Woods, U Yildiz, M Coons, D Devore, and P Margl. A variable neighbourhood descent heuristic for conformational search using a quantum annealer. *arXiv preprint arXiv:1811.06999*, 2018.
- [268] Javier Argüello-Luengo, Alejandro González-Tudela, Tao Shi, Peter Zoller, and J Ignacio Cirac. Analog quantum chemistry simulation. *arXiv preprint arXiv:1807.09228*, 2018.
- [269] Sergey Bravyi, Graeme Smith, and John A Smolin. Trading classical and quantum computational resources. *Physical Review X*, 6(2):021043, 2016.
- [270] Tianyi Peng, Aram Harrow, Maris Ozols, and Xiaodi Wu. Simulating large quantum circuits on a small quantum computer. *arXiv preprint arXiv:1904.00102*, 2019.
- [271] Isaac H Kim. Holographic quantum simulation. *arXiv preprint arXiv:1702.02093*, 2017.
- [272] Isaac H Kim. Noise-resilient preparation of quantum many-body ground states. *arXiv preprint arXiv:1703.00032*, 2017.

- [273] Andrew M Childs, Dmitri Maslov, Yunseong Nam, Neil J Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018.
- [274] Dave Wecker, Matthew B Hastings, and Matthias Troyer. Progress towards practical quantum variational algorithms. *Physical Review A*, 92(4):042303, 2015.
- [275] Li Chen and Alexander Seidel. Algebraic approach to the study of zero modes of Haldane pseudopotentials. *Physical Review B*, 91(8):085103, 2015.
- [276] Gerald V Dunne. Slater decomposition of Laughlin states. *International Journal of Modern Physics B*, 7(28):4783–4813, 1993.
- [277] Andrea Di Gioacchino, Luca Guido Molinari, Vittorio Erba, and Pietro Rotondo. Unified Fock space representation of fractional quantum Hall states. *Physical Review B*, 95(24):245123, 2017.
- [278] Zhendong Li, Junhao Li, Nikesh S Dattani, CJ Umrigar, and Garnet Kin-Lic Chan. The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations. *The Journal of chemical physics*, 150(2):024302, 2019.
- [279] Andrew Jena, Scott Genin, and Michele Mosca. Pauli partitioning with respect to gate sets. *arXiv preprint arXiv:1907.07859*, 2019.
- [280] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F Izmaylov. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *arXiv preprint arXiv:1907.03358*, 2019.
- [281] Rolando Somma, Gerardo Ortiz, Emanuel Knill, and James Gubernatis. Quantum simulations of physics problems. *International Journal of Quantum Information*, 1(02):189–206, 2003.
- [282] Mark Fingerhuth, Tomáš Babej, and Peter Wittek. Open source software in quantum computing. *PloS one*, 13(12):e0208561, 2018.
- [283] Vadym Kliuchnikov, Alex Bocharov, Martin Roetteler, and Jon Yard. A framework for approximating qubit unitaries. *arXiv preprint arXiv:1510.03888*, 2015.
- [284] Michael Beverland, Earl Campbell, Mark Howard, and Vadym Kliuchnikov. Lower bounds on the non-clifford resources for quantum computations. *arXiv preprint arXiv:1904.01124*, 2019.

- [285] Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by clifford gates. *Physical review letters*, 116(25):250501, 2016.
- [286] Edwin Pednault, John A Gunnels, Giacomo Nannicini, Lior Horesh, Thomas Magerlein, Edgar Solomonik, and Robert Wisnieff. Breaking the 49-qubit barrier in the simulation of quantum circuits. *arXiv preprint arXiv:1710.05867*, 2017.
- [287] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J Bremner, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595, 2018.
- [288] Zhao-Yun Chen, Qi Zhou, Cheng Xue, Xia Yang, Guang-Can Guo, and Guo-Ping Guo. 64-qubit quantum circuit simulation. *Science Bulletin*, 63(15):964–971, 2018.
- [289] Igor L Markov, Aneeqa Fatima, Sergei V Isakov, and Sergio Boixo. Quantum supremacy is both closer and farther than it appears. *arXiv preprint arXiv:1807.10749*, 2018.
- [290] Riling Li, Bujiao Wu, Mingsheng Ying, Xiaoming Sun, and Guangwen Yang. Quantum supremacy circuit simulation on sunway taihulight. *arXiv preprint arXiv:1804.04797*, 2018.
- [291] Akira SaiToh. ZKCM: a C++ library for multiprecision matrix computation with applications in quantum information. *Computer Physics Communications*, 184(8):2005–2020, 2013.
- [292] Kieran Woolfe. Matrix product operator simulations of quantum algorithms. Technical report, University of New South Wales Melbourne United States, 2015.
- [293] Bjarni Jónsson, Bela Bauer, and Giuseppe Carleo. Neural-network states for the classical simulation of quantum computing. *arXiv preprint arXiv:1808.05232*, 2018.
- [294] Xun Gao and Luming Duan. Efficient classical simulation of noisy quantum computation. *arXiv preprint arXiv:1810.03176*, 2018.
- [295] Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. *arXiv preprint quant-ph/9806029*, 1998.

- [296] Michael R Geller and Zhongyuan Zhou. Efficient error models for fault-tolerant architectures and the Pauli twirling approximation. *Physical Review A*, 88(1):012314, 2013.
- [297] Nicolas PD Sawaya, Mikhail Smelyanskiy, Jarrod R McClean, and Alán Aspuru-Guzik. Error sensitivity to environmental noise in quantum circuits for chemical state preparation. *Journal of chemical theory and computation*, 12(7):3097–3108, 2016.
- [298] Jarrod R McClean, Ian D Kivlichan, Kevin J Sung, Damian S Steiger, Yudong Cao, Chengyu Dai, E Schuyler Fried, Craig Gidney, Brendan Gimby, Pranav Gokhale, et al. OpenFermion: the electronic structure package for quantum computers. *arXiv preprint arXiv:1710.07629*, 2017.
- [299] Qiskit. <https://qiskit.org>. Accessed: 2019-09-23.
- [300] Guang Hao Low, Nicholas P Bauman, Christopher E Granade, Bo Peng, Nathan Wiebe, Eric J Bylaska, Dave Wecker, Sriram Krishnamoorthy, Martin Roetteler, Karol Kowalski, et al. Q# and NWChem: Tools for scalable quantum chemistry on quantum computers. *arXiv preprint arXiv:1904.01131*, 2019.
- [301] Marat Valiev, Eric J Bylaska, Niranjan Govind, Karol Kowalski, Tjerk P Straatsma, Hubertus JJ Van Dam, Dunyou Wang, Jarek Nieplocha, Edoardo Apra, Theresa L Windus, et al. NWChem: A comprehensive and scalable open-source solution for large scale molecular simulations. *Computer Physics Communications*, 181(9):1477–1489, 2010.
- [302] Nicholas P Bauman, Eric J Bylaska, Sriram Krishnamoorthy, Guang Hao Low, Nathan Wiebe, Christopher E Granade, Martin Roetteler, Matthias Troyer, and Karol Kowalski. Downfolding of many-body hamiltonians using active-space models: Extension of the sub-system embedding sub-algebras approach to unitary coupled cluster formalisms. *The Journal of chemical physics*, 151(1):014107, 2019.
- [303] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffrey Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L Chevrier, Kristin A Persson, and Gerbrand Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.
- [304] Ask Hjorth Larsen, Jens Jørgen Mortensen, Jakob Blomqvist, Ivano E Castelli, Rune Christensen, Marcin Dulak, Jesper Friis, Michael N

Groves, Bjørk Hammer, Cory Hargus, et al. The atomic simulation environment — a Python library for working with atoms. *Journal of Physics: Condensed Matter*, 29(27):273002, 2017.