
Quantum chemistry

Computational chemistry seeks to use computational methods to predict the physical properties and behaviors of atoms, molecules, and materials, and includes methods such as first-principles simulation, classical molecular dynamics, and cheminformatics. We will restrict our focus to first-principles simulations that treat chemical systems quantum mechanically (noting that these methods may be used within a larger workflow incorporating other techniques). Despite the apparent exponential cost of exact classical methods for this task, scientists have made incredible progress over the last century via increasingly sophisticated approximate methods. As a result, quantum chemistry is now a core part of several applications, including the analyses of chemistry experiments, the pharmaceutical drug discovery pipeline, and the optimization of materials for catalysts and batteries. Given the inherently quantum mechanical nature of these problems, it follows that several quantum algorithms have been proposed for computational chemistry [66]. In this chapter, we focus on simulating the electrons and vibrations of nuclei in molecules and materials. For further reviews of quantum computing for chemistry, we refer readers to [752, 229, 102, 783].

The authors are grateful to Ryan Babbush, Joshua Goings, and Ashley Montanaro for reviewing this chapter.

2.1 Simulating electrons in molecules and materials

Overview

We seek the energy eigenstates, thermal states (i.e., statistical ensembles of eigenstates at a given temperature), or dynamics corresponding to the Hamil-

tonian used to describe the electrons in molecules or material systems. The electrons interact with each other, in addition to fields produced by the nuclei and any external applied fields. In many systems it is appropriate to use the Born–Oppenheimer approximation, which treats the nuclei classically and fixes their spatial positions, separating the nuclear and electronic degrees of freedom.

Material systems can be described by a periodically repeating (i.e., translationally invariant) finite-size computational cell of interacting atoms. By simulating increasingly large computational cells and mitigating finite-size effects, we can extrapolate simulation results to the thermodynamic limit. This enables the measurement of bulk properties, such as magnetization, tensile strength, and thermal or electrical conductivity.

Actual end-to-end problem(s) solved

The Hamiltonian of a system consisting of K nuclei and η electrons interacting via the Coulomb interaction is (in atomic units)

$$H = - \sum_{i=1}^{\eta} \frac{(\nabla_i)^2}{2} - \sum_{I=1}^K \frac{(\nabla_I)^2}{2M_I} - \sum_{i,I} \frac{Z_I}{|r_i - R_I|} + \sum_{i \neq j} \frac{1}{2|r_i - r_j|} + \sum_{I \neq J} \frac{Z_I Z_J}{2|R_I - R_J|},$$

where ∇ is the gradient operator, r_i gives the position of the i -th electron, and R_I and Z_I give the position and charge of the I -th nucleus. This Hamiltonian can be discretized using a basis set $\{\phi_i(r)\}_{i=1}^N$ of electron spin orbital and $\{\chi_i(r)\}_{i=1}^M$ of nuclear orbital functions, or grid points, and can either be used with the time-dependent Schrödinger equation to simulate dynamics, or with the time-independent Schrödinger equation to obtain energy eigenstates. When simulating dynamics, it is necessary to use a basis set that is sufficiently flexible (or adaptive) to accurately describe the states at all times, as many chemical basis sets are highly optimized for ground state calculations and so are less suitable for dynamics calculations. It is often appropriate to make the Born–Oppenheimer approximation, fixing the positions of the nuclei, which are treated as classical point charges. The resulting electronic Hamiltonian at a fixed nuclear configuration is given by

$$H(\{R_I\}) = - \sum_i \frac{(\nabla_i)^2}{2} - \sum_{i,I} \frac{Z_I}{|r_i - R_I|} + \frac{1}{2} \sum_{i \neq j} \frac{1}{|r_i - r_j|} + V(\{R_I\}), \quad (2.1)$$

where $V(\{R_I\})$ is the constant offset from the nuclear repulsion energy. In this case, the Hamiltonian is discretized using a basis set $\{\phi_i(r)\}_{i=1}^N$ of electron spin orbital functions or grid points. For many molecules at room temperature, the ground state of the electronic structure Hamiltonian is a good approximation

for the thermal state $\rho \propto e^{-\beta H}$ (with $\beta = 1/k_B T$, where k_B is the Boltzmann constant and T the temperature), as the electronic energy levels are well separated with respect to $k_B T$. This can be contrasted with the vibrational structure of molecules, where vibrational energies are on the order of $k_B T$, and so excited states are also populated at room temperature.

The electronic eigenstates (or thermal states) often provide a good starting description of a wide range of system properties, which then can be corrected by, for example, vibrational, rotational, relativistic (e.g., spin-orbit coupling) or entropic contributions. Preparing the desired electronic state for a given nuclear configuration is typically the first step in learning properties of the system. We then measure the expectation values of observables with respect to these states. Moreover, the electronic response to weak or slowly varying perturbations can often be described by a sequence of static calculations, for example, linear response theory for radiation absorption, or Born–Oppenheimer molecular dynamics [745] where one iteratively solves the electronic time-independent Schrödinger equation to obtain the forces on the nuclei, whose positions can then be updated using Newton’s laws. Static calculations can be used to probe:

- Energy values (potentially across a range of nuclear configurations)—for electronic excitation energies at a fixed nuclear geometry, for determining molecular geometries by computing the electronic ground state energy at different geometries, and for finding reaction pathways and rates by computing energy differences between a sequence of geometries involved in a reaction. To obtain accurate predictions, the electronic energy values must be corrected by other contributions to the free energy.
- Determining transition probabilities between states—for reactions and optical properties.
- Differential changes in electronic energy in response to an applied field, for example, electronic or magnetic dipole moments, polarizability.
- Calculating forces on the nuclei, for use in molecular dynamics calculations—used in a range of applications, including protein folding and calculating drug molecule binding affinities.
- Orbital occupancies and correlation functions, as well as real- and imaginary-time Green’s functions.

Properties of interest for materials include the following:

- Energy densities for given system parameters, to determine phase diagrams.
- Bulk properties, such as magnetization, thermal or electrical conductivity, and tensile strength.

- Particle densities and correlation functions between sites, as well as real- and imaginary-time Green's functions.

In order to understand how these observables vary as the system parameters (nuclear positions, atomic doping, temperature, applied field, etc.) are changed, the desired state may need to be prepared and measured a number of times.

Simulations of dynamics may be used to explicitly probe many of the equilibrium phenomena implicitly being probed above—for example, chemical reactions occurring at thermodynamic (quasi-)equilibrium—as well as additional nonperturbative or nonequilibrium phenomena that are difficult to implicitly describe as a sequence of static calculations, such as scattering from collisions, absorption of UV and X-ray radiation, charge-transfer dynamics, and optimal control. As a result, many of the same observables described above are still of interest and can be monitored as a function of time, including the following:

- Changes in kinetic or potential energy.
- Changes in particle densities or orbital occupancies.
- Changes in charge or spin densities.

Dominant resource cost/complexity

Mapping the problem to qubits: We discretize the electron positions using a basis of N spin orbital functions. For many basis sets, the discretization error decays as $1/N$ [488, 932] and is limited by the resolution of singularities in the Coulomb interaction from charge coalescence. A variety of functional forms have been considered for the electron orbitals (see Table 2.1 for a list of orbitals commonly considered in quantum computing). The optimal choice will be system dependent and must consider the following non-exhaustive list of factors:

- The resolution of the orbital, improved by matching the character of local vs. delocalized physics in the system to that of the orbital.
- The cost of computing the Hamiltonian, either in classical precomputation or (if required) coherently on a quantum device (see §Accessing the Hamiltonian, below).
- The properties of the resulting Hamiltonian (number of terms, norm, locality of terms, etc.) which determine the cost of accessing the Hamiltonian in algorithms.

The commonly used “Galerkin discretization scheme” discretizes the Hamiltonian via integrals over the basis functions, with one- and two-electron

Representation	First-quantized	Second-quantized
Gaussians	[74] ^a [416]	[1034]
Plane waves	[962]	[76]
Bloch/Wannier functions	Not yet studied	[561, 886]
Grids	[628]	[76, App. A]
Pseudo-spectral / Discrete variable representations	[601, 253, 287]	[76, 721]

Table 2.1 Representative references (chosen based on their discussion of their choice of representation) showing the use of different basis functions in quantum algorithms for the electronic structure problem.

^a This reference is not technically a first-quantized representation, as antisymmetry is stored in the operators rather than the wavefunction, but it stores states in an analogously compressed way to first-quantized representations.

integrals

$$\begin{aligned}
 h_{ij} &= \int dr \phi_i^*(r) \left(-\frac{(\nabla)^2}{2} - \sum_l \frac{Z_l}{|r - R_l|} \right) \phi_j(r) \\
 h_{ijkl} &= \int dr_1 dr_2 \frac{\phi_i^*(r_1) \phi_j^*(r_2) \phi_k(r_2) \phi_l(r_1)}{|r_1 - r_2|}.
 \end{aligned} \tag{2.2}$$

Hamiltonians defined using grids, pseudo-spectral representations, or discrete-variable representations are not obtained from integrals over basis functions, as specified in Eq. (2.2), and the values of h_{ij} , h_{ijkl} are instead defined using finite difference formulas, and/or by their values at discrete grid points. An attractive feature of the Galerkin discretization scheme is that the discretization error is strictly positive. We refer readers to [287, 76, 244] for a more complete discussion.

We can represent electronic states on a quantum computer using either first or second quantized representations.

- For η electrons in N spin orbitals, first quantization uses η registers, which each contain $\log_2(N)$ qubits; each register enumerates which orbital its corresponding electron is in, and the wavefunction must then be antisymmetrized to respect fermionic constraints [139]. The Hamiltonian of Eq. (2.1) in first quantization can be written as

$$H = \sum_{\alpha=1}^{\eta} \sum_{i,j=1}^N h_{ij} |i\rangle \langle j|_{\alpha} + \frac{1}{2} \sum_{\alpha \neq \beta} \sum_{i,j,k,l=1}^N h_{ijkl} |i\rangle \langle l|_{\alpha} \otimes |j\rangle \langle k|_{\beta},$$

where α, β index which of the electron registers the operators act on.

- In second quantization, antisymmetry is stored in the operators, which obey fermionic anticommutation relations: $\{a_i, a_j^{\dagger}\} = \delta_{ij} I$ and $\{a_i, a_j\} = \{a_i^{\dagger}, a_j^{\dagger}\} =$

0 (where δ_{ij} denotes the Kronecker delta symbol, I denotes the identity operator, and $\{u, v\} = uv + vu$). The Hamiltonian of Eq. (2.1) in second quantization can be written as

$$H = \sum_{i,j=1}^N h_{ij} a_i^\dagger a_j + \frac{1}{2} \sum_{i,j,k,l=1}^N h_{ijkl} a_i^\dagger a_j^\dagger a_k a_l.$$

A number of fermion-to-qubit mappings have been studied; see [752] for discussion. Under the commonly used Jordan–Wigner mapping we require N qubits, where each qubit stores the occupancy of the corresponding spin orbital. These mappings induce a mapping of the Hamiltonian (and other observables) to qubit operators.

Accessing the Hamiltonian: Quantum algorithms for both static and dynamic simulations require access to the Hamiltonian. This is typically provided by block-encoding or Hamiltonian simulation.¹ A common block-encoding strategy for chemistry Hamiltonians is the linear combinations of unitaries (LCU) block-encoding, whereby the Hamiltonian is expressed as a linear combination of unitary operators $\sum_{i=0}^{L-1} c_i U_i$ (e.g., U_i could be products of Pauli matrices), and the block-encoding is then realized using the oracles PREPARE and SELECT² that act on the main register and a $\lceil \log_2(L) \rceil$ ancilla system as

$$\begin{aligned} \text{PREPARE} |0^{\lceil \log_2(L) \rceil}\rangle &= \frac{1}{\sqrt{\lambda}} \sum_{j=0}^{L-1} \sqrt{|c_j|} |j\rangle \\ \text{SELECT} &= \sum_{j=0}^{2^{\lceil \log_2(L) \rceil} - 1} |j\rangle \langle j| \otimes \text{sign}(c_j) U_j, \end{aligned}$$

where $\lambda = \sum_{i=0}^{L-1} |c_i|$. Then the sequence $\text{PREPARE}^\dagger \cdot \text{SELECT} \cdot \text{PREPARE}$ is a $(\lambda, \lceil \log_2(L) \rceil, 0)$ -block-encoding of the Hamiltonian written as an LCU. The oracle SELECT can be implemented using the unary iteration method [75] or the approach of [1011]. The oracle PREPARE can be implemented by coherently loading coefficients stored in memory [75, 722, 140] or by computing coefficients on-the-fly using quantum arithmetic [601, 74, 73, 962]. In many cases, loading the coefficients for PREPARE is the bottleneck, with cost—in terms of the number of non-Clifford gates—scaling linearly with $\tilde{L}$, the number of unique coefficients in the Hamiltonian (this scaling can be reduced to

¹ Hamiltonian simulation is used to explicitly simulate dynamics, but can also be used implicitly to provide access to the Hamiltonian for use in static calculations, for example, in quantum phase estimation.

² To be precise for $j \notin \{0, 1, \dots, L-1\}$ we define $\text{sign}(c_j) U_j := I$.

$O(\tilde{L}^{1/2})$ using $O(\tilde{L}^{1/2})$ ancilla qubits [722]). As a result, a number of algorithms have reduced the quantity of data to load by using compressed representations of the Coulomb Hamiltonian, achieved through tensor factorizations [140, 212, 669, 886] (see [712] for a recent unifying perspective on these approaches).

State preparation: Simulating the behavior of electrons in molecules and materials reduces to the task of preparing a desired state and measuring observables. The state to be prepared is typically an energy eigenstate, a thermal state, or a time-evolved state.

- **Energy eigenstates:** In the following discussion, we refer to the overlap $\gamma = |\langle \psi | E_j \rangle|$ between a desired eigenstate $|E_j\rangle$ and a given initial state $|\psi\rangle$, and the minimum gap Δ between the desired energy eigenvalue and other energy eigenvalues. Below, we list several methods for preparing energy eigenstates, or approximations to them.
 - **Classical trial states:** Approximate eigenstates obtained from a classical calculation can be prepared as quantum trial states, including Slater determinant states [1031], linear combinations of D Slater determinants (with complexity $\tilde{O}(D)$ [394]– $O(ND)$ [998]), and matrix product states (MPSs) with bond dimension χ (with complexity $O(N\chi^2)$ [394, 144]). In [144] it was observed that MPS with modest bond dimension could have large overlaps with chemical systems of interest. Several of these methods have been adapted to simulations performed in first quantization [81, 551].
 - **Quantum trial states:** Parameterized quantum circuits, in conjunction with variational quantum algorithms, have been proposed for preparing approximate energy eigenstates (see §NISQ implementations, below). Like classical trial states, the states prepared by these circuits can be used as inputs to other quantum algorithms that further refine the initial state, such as eigenstate filtering, or quantum phase estimation. Initial resource estimates for models of materials systems can be found in [301].
 - **Eigenstate filtering:** Methods such as those in [689, 688] filter out undesired eigenstates using spectral window functions applied via quantum singular value transformation (QSVT) to a block-encoding of the Hamiltonian. The complexity to prepare the ground state (to infidelity ϵ , with failure probability less than θ) using this approach scales as

$$\tilde{O}\left(\frac{\alpha}{\gamma\Delta} \log(\theta^{-1}\epsilon^{-1})\right)$$

calls to an $(\alpha, m, 0)$ -block-encoding of the Hamiltonian, where $\alpha \geq \|H\|$ is a normalization factor of the block-encoding. For comparison to related

methods, we refer the reader to [688, 412]. These algorithms can also be adapted for the case where access to the Hamiltonian is provided by Hamiltonian simulation [358].

- Adiabatic state preparation (ASP): ASP can be used to prepare a target eigenstate (typically the ground state) by evolving from the corresponding easy-to-prepare eigenstate of an initial Hamiltonian $H(0)$ to the full Hamiltonian $H(1)$. Time evolution can be implemented using algorithms for Hamiltonian simulation. The total evolution time is typically chosen according to the heuristic

$$T \gg \max_{0 \leq s \leq 1} \frac{\|\frac{dH}{ds}\|}{\Delta(s)^2},$$

where s describes the adiabatic path $H(s)$ and $\Delta(s)$ is the spectral gap of $H(s)$. It is difficult to analytically bound this (highly system-dependent) complexity for molecular systems (see, e.g., [870]) motivating numerical studies on small molecules [1006, 644, 670, 966].

- Quantum phase estimation (QPE): The above techniques all provide methods of preparing approximate eigenstates, in some cases using promises on the gap Δ , or by exploiting pre-existing knowledge of the energy eigenvalue. Given an approximate eigenstate, and a unitary $U = f(H)$ that encodes the eigenspectrum of the Hamiltonian (with a known, classically invertible relationship f), we can use QPE to project into the desired eigenstate and provide an estimate of the eigenphase ϕ_i of U , which can then be converted into an estimate of the eigenenergy of H using $\phi_i = f(E_i)$. QPE makes

$$O\left(\gamma^{-2} \|f'(H)\|^{-1} \epsilon^{-1} \log(\theta^{-1})\right)$$

calls to the unitary $U(H)$ encoding the spectrum of the Hamiltonian, where $\gamma = |\langle \psi | E_j \rangle|$ is the overlap between the state $|\psi\rangle$ input to QPE and the desired energy eigenstate $|E_j\rangle$, θ is the failure probability, and ϵ is the desired precision in the eigenenergy of H . It is possible to improve the complexity to

$$O\left(\gamma^{-1} \|f'(H)\|^{-1} \epsilon^{-1} \log(\theta^{-1})\right)$$

using amplitude amplification if a sufficiently precise estimate of the eigenvalue is known, or to

$$O\left((\gamma^{-2} \Delta^{-1} + \epsilon^{-1}) \|f'(H)\|^{-1} \log(\theta^{-1})\right)$$

by exploiting knowledge of the gap Δ between the energy eigenstates to perform rejection sampling [139]. The unitary encoding the Hamil-

tonian is often chosen to be $U(H) \approx e^{iHt}$ approximated via quantum algorithms for Hamiltonian simulation. In this case, the approximation error in U must be balanced against the error from QPE. Alternatively, it is common to encode the Hamiltonian using a quantum walk operator $W(H)$ which acts like $e^{i\arccos(H/\alpha)}$ and can be implemented exactly via qubitization [841, 139]. The costs to implement $U(H)$ are inherited from the method used, based on the properties (e.g., commutativity, locality, number of terms, norm, cost of coherently calculating coefficients) of the Hamiltonian in the chosen spin orbital basis and representation. As indicated by the complexities presented above, QPE incurs an overhead from imperfect overlap with the target eigenstate, as well as from needing to suppress the failure probability of the method. We refer to Chapter 13 on QPE for a more detailed discussion of the latter issue, which can either be mitigated by repeating the calculation and using methods for probability amplification [792, 222], or by using window functions [143, 829, 456] to guarantee a desired confidence interval. The latter strategy appears to require fewer resources in practice, especially when considering the imperfect overlap with the target state [144].

- **Thermal states:** Several quantum algorithms have been proposed for preparing thermal states, also known as Gibbs states. The most efficient algorithms proceed by simulating the dissipative open system dynamics, and typically make repeated calls to a block-encoding of the Hamiltonian. The complexity of these methods for concrete electronic structure problems of interest has not yet been determined and depends on the spectral gap of the Lindbladian considered, at the desired temperature. Thermal states could be used as an approximation to the ground state, by choosing the temperature to be sufficiently low compared to the gap between the ground and first excited state [260]. Quantum algorithms for simulating open systems dynamics can also be adapted to directly prepare approximations to the ground state [351].
- **Time-evolved states:** A time-evolved state can be prepared using Hamiltonian simulation algorithms, which approximate the propagator to error ϵ (which provides an upper bound on the error in the resulting state). The cost of Hamiltonian simulation depends both on the algorithm used and the details of the Hamiltonian being simulated. Plane wave, grid, and pseudo-spectral (DVR) basis sets are well suited to simulations of dynamics, as they treat all points in space on an equal footing. For both the plane wave basis and pseudo-spectral DVR in first quantization, the scaling is [77, 962]

$$\tilde{O}\left(\eta^{8/3}N^{1/3}t + \eta \log(1/\epsilon)\right)$$

using Hamiltonian simulation in the interaction picture [718], or

$$\tilde{O}\left((\eta^{8/3}N^{1/3} + \eta^{4/3}N^{2/3})t + \eta \log(1/\epsilon)\right)$$

using qubitization with quantum signal processing [77, 962]. In the pseudo-spectral DVR, the cost scales as

$$\tilde{O}\left((\eta^{7/3}N^{1/3} + \eta^{4/3}N^{2/3})\frac{t^{1+o(1)}N^{o(1)}}{\epsilon^{o(1)}}\right)$$

using high-order product formulas [81].

Measuring observables: Many proposed algorithms consider the ground or excited state energy of the chemical system as the observable of interest. This can be calculated directly using QPE, as discussed above. QPE (and related methods) can also be adapted for other application, for example, calculating absorption spectra of molecules [395].

In other applications, it may be necessary to measure observables other than the energy. In a fault-tolerant computation, it can be preferable to measure these observables through phase-estimation-like approaches, rather than direct measurement averaging, as the former is asymptotically more efficient and can be made robust to logical errors through repetition and probability amplification. Measurement schemes have been developed which achieve this using overlap estimation [637] (see Section 14.2 on amplitude estimation) or the approach of [549, 49] based on the quantum gradient estimation algorithm of [430]. Both approaches require access to a state preparation unitary U_ψ , and its inverse.³ The algorithm based on overlap estimation can be formulated as performing amplitude estimation on U_O , a unitary block-encoding of the observable O with subnormalization factor α_O . The complexity to compute the expectation value to precision ϵ is $O(\alpha_O/\epsilon)$ calls to U_O and U_ψ (or the reflection $R_\psi = I - 2|\psi\rangle\langle\psi|$) and their inverses. This approach has been considered in the context of measuring correlation functions, density of states, and linear response properties (all in [853]), as well as energy gradients with respect to various parameters, which can be used to compute forces or dipole moments, and for which a range of estimation strategies are possible [805, 956].

The gradient-based algorithm simultaneously computes the value of M potentially noncommuting observables O_j by making $\tilde{O}(M^{1/2}/\epsilon)$ calls to U_ψ , $U_\psi^\dagger$ (or R_ψ) and either $\tilde{O}(M^{3/2}/\epsilon)$ calls to gates of the form e^{ixO_j} [549] or $\tilde{O}(M/\epsilon)$ calls to a block-encoding of the observables [49]. The algorithm also requires

³ Note that it can be substantially cheaper to directly execute the reflection $R_\psi = I - 2|\psi\rangle\langle\psi|$ used in both methods, rather than through the use of U_ψ , as the complexity of R_ψ does not depend on the overlap γ that appears in state preparation—see [688] for additional discussion.

$O(M \log(1/\epsilon))$ additional qubits. This approach has been considered in the context of measuring nuclear forces [805], fermionic reduced density matrices [549], and dynamic correlation functions [549].

Existing resource estimates

There are a large number of resource estimates for performing phase estimation to learn the ground state energies of molecular or material systems, which we list in Table 2.2 and Table 2.3. These resource estimates use compilation methods described in Part III on fault-tolerant quantum computing. We also note the existence of a software package that provides features for calculating the non-Clifford costs of QPE for the electronic structure problem [238].

There have been comparatively few studies of the logical resources required for the simulation of chemical dynamics. Recent work has computed the resources required to calculate the energy loss of charged particles moving through a medium (“stopping power”), as pertaining to nuclear fusion experiments [887]. End-to-end resource estimates were determined, including the costs of initial state preparation, measurement of observables, and repetitions across a range of parameters. The resource estimates for the end-to-end task ranged from roughly 2000 logical qubits and order- 10^{13} Toffoli gates to roughly 30,000 logical qubits and order- 10^{17} Toffoli gates.

Caveats

Existing resource estimates typically consider only a single run of phase estimation and assume that we have access to the desired energy eigenstate. As outlined above, both phase estimation and eigenstate filtering scale as $\Omega(\gamma^{-1}\Delta^{-1})$ when we have a lower bound on the gap. The “orthogonality catastrophe” suggests that the overlap of simple trial states with the desired eigenstate will decay exponentially as a function of system size. Although simulations will always be performed on finite-size systems, it is an open question as to when asymptotic behavior becomes problematic and whether initial states with sufficient overlaps can be prepared for systems of interest [998, 670, 144]. This issue may become more pressing for materials systems as we scale to the thermodynamic limit. In general, we know that the problem of finding the ground state of electronic structure Hamiltonians is QMA-hard [1035], but it is not yet known if these complexity-theoretic statements provide intuition for physically realistic Hamiltonians.

As noted above, to accurately resolve the system, a large basis set must be used, as the discretization error decays as $1/N$ where N is the number of spin orbitals considered. In practice, one typically repeats the calculation using increasingly accurate basis sets and then extrapolates to the continuum limit.

Molecule(s) & references	Algorithm	Logical qubits	T/Toffoli gates per sample	Number of samples
FeMo-co (nitrogen fixation) [870, 140, 212, 669, 1013, 238]	2nd Q, THC qubitization, Gaussians 2nd Q, randomized compilation, Gaussian	2196 [669] ~193 [1013]	3.2×10^{10} [669] $\sim 3 \times 10^{12}$ [1013]	7 ~600
Cytochrome P450 (biological drug metabolizing enzyme) [437]	2nd Q, THC qubitization, Gaussians	1434	7.8×10^9	7
Lithium-ion battery molecules [617, 962]	2nd Q, DF qubitization, Gaussians 1st Q, qubitization, plane waves	10^4 – 10^5 [617] 2000–3000 [962]	10^{12} – 10^{14} [617] 10^{11} – 10^{12} [962]	7 7
Chromium dimer [375]	2nd Q, sparse qubitization, Gaussians	~1300	$\sim 10^{10}$	7
Ruthenium catalyst (CO ₂ fixation) [212]	2nd Q, DF qubitization, Gaussians	~4000	$\sim 3 \times 10^{10}$	7
Ibuprofen (drug molecule) [158]	2nd Q, Sparse qubitization, Gaussians	2207	1.1×10^{10}	7
Molybdenum catalysts (nitrogen fixation) [113] Amyloid beta binding site fragment (Alzheimer's disease) [818]	2nd Q, DF qubitization, Gaussians	1000–8000	10^{11} – 10^{12}	3–9 ^a
Fullerene-encapsulated cyclic ozone (rocket fuel) [1029]	2nd Q, DF qubitization, Gaussians 2nd Q, DF qubitization, Gaussians 2nd Q, sparse qubitization, plane wave dual	5000 10^3 10^4	10^{14} 10^{12} – 10^{13} 10^{15}	7 7

Table 2.2 Logical resource estimates for quantum phase estimation (QPE) applied to a range of molecular systems to compute a single energy eigenvalue. We list state-of-the-art resource estimates and refer to the references therein for prior estimates. The presented gate counts are for a single run of the phase estimation circuit. The number of samples is determined by the desired maximum failure probability, taken here as 0.1. In most cases, the number of samples is calculated using [792, Lemma 1] which assumes that the median value is taken from a number of independent repetitions of QPE. The success probability of a single run of QPE is $8/\pi^2$ [186]. QPE must be run a number of times if the overlap $\gamma \leq 1$, in general contributing a multiplicative cost of $O(1/\gamma^2)$ (though this may be reduced, as described in the main text). “Algorithm” denotes the phase estimation unitary considered (e.g., Trotterization, Qubitization) as well as details about the quantization scheme (1st or 2nd), basis used, or factorization method used to compile the unitary (sparse [140], SF = single factorized [140], DF = double factorized [212], THC = tensor hypercontraction [669]). The resource estimates presented can be for different numbers of electrons and orbitals, accuracies, and can have differing classical simulation complexities. As such, the results may not be directly comparable, even within a single row of the table.

^a This resource estimate assumed an overlap of $\gamma < 1$, and a lower failure probability of 0.01.

Material(s) & references	Algorithm	Logical qubits	T /Toffoli gates per sample
Electron gas (prototypical model) [75, 630, 753, 962]	1st Q, qubitization, plane waves 2nd Q, qubitization [75]/Trotter [753], plane wave (dual)	1500–5000 [962] 100–1000 [75, 753]	10^9 – 10^{14} [962] 10^8 – 10^{11} [75, 753]
Lithium-ion battery materials [341, 934, 886, 142]	2nd Q, sparse/SF/DF/THC qubitization, Bloch orbitals 1st Q, qubitization, plane waves (w. pseudopotential)	10^5 – 10^6 [886] ~1000 [142]	10^{12} – 10^{14} [886] ~ 10^{14} [142]
Transition metal catalysts nickel/palladium oxide [561, 416]	2nd Q, sparse qubitization, Bloch/Wannier orbitals 1st Q, sparse qubitization, plane wave (dual)	10^4 – 10^5 [561] 17,505 [416]	10^{10} – 10^{13} [561] 10^{15} [416]
Magnesium/nickel alloys (corrosion resistant) [800]	2nd Q, qubitization, plane wave (dual)	9000–500,000	10^{14} – 10^{19}
Nitrogen vacancy center in diamond (quantum sensing) [562]	2nd Q, DF qubitization, plane wave (projector augmented-wave method)	5000–150,000	10^{12} – 10^{14}

Table 2.3 Logical resource estimates for quantum phase estimation (QPE) applied to a range of material systems. We list state-of-the-art resource estimates and refer to the references therein for prior estimates. The presented gate counts are for a single run of the phase estimation circuit. The number of samples for all listed systems is 7, determined by the desired maximum failure probability, taken here as 0.1 (calculated using [792, Lemma 1] which assumes that the median value is taken from a number of incoherent repetitions of QPE. The success probability of a single run of QPE is $8/\pi^2$ [186]). QPE must be run a number of times if the overlap $\gamma \leq 1$, in general contributing a multiplicative cost of $O(1/\gamma^2)$ (though this may be reduced, as described in the main text). “Algorithm” denotes the phase estimation unitary considered (e.g., Trotterization, qubitization) as well as details about the quantization scheme (1st or 2nd), basis used, or factorization method used to compile the unitary (sparse [140], SF = single factorized [140], DF = double factorized [212], THC = tensor hypercontraction [669]). The resource estimates presented can be for different numbers of electrons and orbitals, and can have differing classical simulation complexities. As such, the results may not be directly comparable, even within a single row of the table.

Many quantum resource estimates consider active spaces with a small number of active orbitals, and so underestimate the resources required to achieve sufficiently accurate results to be informative. It is an active area of research to develop methods for increasing the resolution without increasing the basis set size, such as perturbative approaches, downfolding techniques, and embedding theories.

The end-to-end applications typically solved in the electronic structure problem can require between tens (e.g., structure determination, low temperature properties) and millions (e.g., biologically relevant molecular dynamics) of energy evaluations—each with different Hamiltonian parameters that may require preparing a new state to be measured. For example, a recent analysis of quantum algorithms applied to pharmaceutical chemistry [896] highlighted that to calculate the binding affinity between a drug molecule and its target (free energy differences) requires sampling a range of thermodynamic configurations, resulting in millions to billions of single-point energy evaluations. This introduces a large overhead when preparing a different state for each configuration and measuring its energy [805], although alternative approaches may provide more favorable scaling [941].

Comparable classical complexity and challenging instance sizes

The cost of exact diagonalization of the electronic structure Hamiltonian scales exponentially with the number of electrons and basis set size. As such, classical approaches to the electronic structure problem typically utilize a range of approximations that reduce their complexity to polynomial in the system size but introduce a (potentially uncontrolled) deviation from the exact ground state, leading to a bias in energy estimates and/or the expectation values of other observables. Approaches include Hartree–Fock (HF), density functional theory (DFT), perturbation theory, configuration interaction (CI) methods, coupled cluster (CC) methods, quantum Monte Carlo (QMC) techniques, and tensor network approaches. The cheapest approaches can be applied to thousands of orbitals but can be qualitatively inaccurate for strongly correlated systems. The most expensive approaches are more effective for strongly correlated systems, but their higher computational cost limits their applicability to roughly 100 spin orbitals. For example, [437] found that a density matrix renormalization group (DMRG) calculation performed on an 86 spin orbital active space of the cytochrome P450 enzyme molecule referenced in Table 2.2 required around 50 hours, using 32 threads, 48 gigabytes of RAM, and 235 gigabytes of disk memory. We also refer to [1046] for a comparison of 20 first-principles many-body electronic structure methods applied to a test set of seven transition metal atoms and their ions and monoxides.

Due to their extended nature, material systems are most commonly targeted with DFT. DFT can be applied to systems with thousands of electrons and orbitals but can lead to uncontrolled energy bias in strongly correlated systems. QMC and tensor network methods have been successfully applied to prototypical models of material systems and are becoming increasingly practical for more realistic models. We refer to [666, 784, 785, 900] for benchmarks of classical electronic structure methods on hydrogen chains and Hubbard models scaling to the thermodynamic limit, which act as simplified models for real materials.

Many of the techniques discussed above for computing ground and excited states of chemical systems have been extended to explicitly simulate the time dynamics of the electronic Hamiltonian. These include time-dependent HF, real-time time-dependent DFT, and time-dependent CI & CC methods. In general, the errors from the approximations made in these approaches are larger than for their static counterparts. We refer readers to [682] for a more detailed discussion of classical methods for real-time time-dependent electronic structure theory.

Speedup

It is nontrivial to determine the speedup of quantum algorithms for the electronic structure problem over their classical counterparts. If we consider the subtask of determining energy eigenstates, then for speedup greater than polynomial to be achieved, we require:

- The ability to prepare a trial state with nonexponentially vanishing overlap with the ground state as the system size increases.
- Polynomially scaling (with system size) classical algorithms having an exponential growth in their approximation parameter (e.g., bond dimension, number of excitations) as the system size increases.

Whether these two requirements can coexist in systems of interest is an active area of research [670, 144]. Even if exponential speedups are not available, it may be the case that quantum algorithms provide polynomial speedups over exact classical algorithms—and potentially over approximate classical algorithms [252].

From a complexity-theoretic viewpoint, we know that simulating the dynamics of a quantum system is a BQP-complete problem [705]. Combined with the observed difficulty of classically simulating the time evolution of electronic structure Hamiltonians, this may be taken as evidence for the possibility of an exponential speedup when simulating dynamics. In [81], quantum algorithms for simulating the fully correlated dynamics of electrons in a pseudo-

spectral DVR or plane-wave basis [601, 77, 962] were compared against classical methods for mean-field dynamics. Large polynomial speedups were observed, ranging from superquadratic to seventh power in the salient parameters, depending on the relation between N and η .

NISQ implementations

Solving the electronic structure problem is one of the most widely studied and promoted NISQ applications. The primary NISQ approach is the variational quantum eigensolver (VQE). There have been a number of experimental demonstrations on small molecules, for example, [595, 442], as well as proposals to simulate material systems [1068, 738, 301]. Related methods, such as quantum computing–assisted QMC methods [550] have also been developed. Nevertheless, current device noise rates are too high to enable the running of circuits sufficiently deep that they can outperform classical electronic structure methods, and the number of circuit repetitions required to measure energy expectation values can be impractically large [440]. As such, there are several challenges that must be overcome if heuristic NISQ approaches are to scale to classically intractable system sizes and provide advantage over classical methods. There have also been proposals to simulate the electronic structure problem using analog quantum simulators [55], though to the best of our knowledge, these have not yet been experimentally demonstrated and are limited by the high-precision requirements of the electronic structure problem.

Outlook

Simulating the behavior of electrons in molecules and materials has repeatedly been identified as one of the most promising applications of quantum computing. Nevertheless, the discussion above highlights several challenges for current quantum approaches to become practical. Most notably, after incorporating the costs of initial state preparation and measuring observables, using larger active spaces to capture dynamic correlation, and including algorithmic repetitions to account for nonzero failure probabilities and sampling a range of parameters, a large number of logical qubits and total T /Toffoli gates may be required. The success of approximate classical methods for a wide range of chemical problems sets a high bar for quantum simulations to achieve advantage and encourages continued focus on resource estimations for end-to-end applications.

2.2 Simulating vibrations in molecules and materials

Overview

We seek the energy eigenstates, thermal states (i.e., statistical ensembles of eigenstates at a given temperature), or dynamics corresponding to the Hamiltonian that describes the vibrations of the nuclei in a molecule or material around their equilibrium positions. This Hamiltonian contains the kinetic energy of the nuclei and the effective potential that they move on, which is determined by the electronic potential energy surface (i.e., the electronic energy expressed as a function of the nuclear coordinates). It is also possible to consider non-adiabatic couplings between the vibrational degrees of freedom and electronic (“vibronic”) or rotational (“ro-vibrational”) degrees of freedom.

Actual end-to-end problem(s) solved

Classically solving the Schrödinger equation while treating electrons and nuclei on an equal footing has prohibitively high computational cost for all but the smallest systems. For systems where it is valid to separate the electronic and nuclear motions (the Born–Oppenheimer approximation), we can imagine the nuclei moving on the electronic potential energy surface (PES). For molecules composed of light atoms (where relativistic effects can be neglected), the vibrations of the nuclei around their equilibrium positions provide a first-order correction to the electronic energies and influence photo-emission/absorption properties. For a system with K classical nuclei at equilibrium positions $\{R_I\}$, the vibrational Hamiltonian can be written as

$$H = - \sum_I \frac{\nabla_I^2}{2M_I} + V_e(\{R_I\}),$$

where $V_e(\{R_I\})$ denotes the nuclear potential determined by the electronic potential energy surface, obtained by first solving the electronic Hamiltonian for a range of nuclear positions. Computing vibrational eigenstates can be made classically tractable by modeling V_e as a harmonic potential, which reduces the problem to solving a number of coupled quantum harmonic oscillators. The harmonic approximation can also be used when simulating vibronic transitions between vibrational energy levels on different PESs. However, due to the coordinate change between the PESs—which acts as a squeezing and displacement operation on the vibrational modes—exact simulation is #P-hard, and therefore inefficient for both classical and quantum algorithms. Nevertheless, vibronic spectra can be efficiently approximated using classical [552, 811] and quantum [553] algorithms in many regimes of interest.

To accurately describe nonrigid molecules or highly excited vibrational states, additional anharmonic terms are required in the potential. These can be obtained by performing higher-order fits of the potential, for example, by expanding the potential V_e to degree d . Computing accurate solutions of this Hamiltonian is prohibitively costly for many systems of interest. Probing certain phenomena, such as vibronic spectra, internal conversion, intersystem crossings, and conical intersections, additionally requires the consideration of vibrations on multiple PESs and may require a description of the nonadiabatic couplings between the different PESs (which must often be explicitly determined [828]). We seek to prepare eigenstates or thermal states, or simulate the dynamics of the anharmonic vibrational Hamiltonian, and then measure the expectation values of observables with respect to these states. Properties of interest include the following:

- The vibrational energy at the minimum of the PES, which provides a first-order correction to the electronic energies—for calculating excitation energies, determining stable molecular structures, or finding reaction pathways and rates.
- Determining transition probabilities between states and transition dipole moments—for calculating infrared/Raman spectra between vibrational levels of the same electronic state or vibronic spectra between vibrational levels of different electronic states.
- Measuring the occupancy of vibrational modes as a function of time following excitation, to understand vibrational energy transfer and relaxation in chemical systems (i.e., internal conversion and intersystem crossings).

Thermal states $\rho \propto e^{-\beta H}$ (with $\beta = 1/k_B T$, where k_B is the Boltzmann constant and T the temperature) are often of greater interest in the vibrational case than in the electronic case; vibrational energies are on the order of $k_B T$ and so excited vibrational states are populated even at room temperature. This can be contrasted with the electronic structure problem, where the larger electronic energy gaps of many molecules mean that ground states are typically of primary interest at room temperature.

Dominant resource cost/complexity

A molecule with K atoms has $M = 3K - 6$ ($M = 3K - 5$ for linear molecules) vibrational modes. Each vibrational mode excitation is treated as a distinguishable particle and so the wavefunction does not need to be explicitly symmetrized. The Hamiltonian is discretized using a basis set of vibrational modal functions $\{\chi_i\}_{i=1}^N$, for example, the truncated eigenfunctions of the quantum

harmonic oscillator Hamiltonian, pseudo-spectral (Fourier) discrete variable representations, or grids.

The vibrational wavefunction can be stored in a first-quantized representation using $M \log(N)$ qubits, where the basis function of each vibrational mode is specified in binary (or an equivalent representation, such as the Gray code [897]). Alternatively, one can use second-quantized representation using MN qubits [813].

Preparing the desired eigenstate or thermal state, or simulating the dynamics can be achieved using a range of quantum algorithms, including quantum phase estimation, quantum singular value transformation, adiabatic state preparation, variational quantum algorithms, Gibbs sampling, and Hamiltonian simulation. These methods are discussed in more detail for the electronic Hamiltonian, as the explicit costs of many of these methods have not yet been determined for simulating vibrations. Nonetheless, many of the same high-level considerations apply. The complexities of subroutines to prepare eigenstates and extract observables are determined by the following observations:

- (i) All methods scale as $\Omega(1/\epsilon)$ to measure the desired observable to an error of $\pm\epsilon$. For the energy, we typically seek $\epsilon \sim (1\text{--}10) \text{ cm}^{-1} \approx (4.56 \times 10^{-6})\text{--}(4.56 \times 10^{-5})$ Hartree.⁴ For comparison, the largest matrix elements in the vibrational Hamiltonian (the harmonic couplings) are typically on the order of 1000 cm^{-1} , and there are $O(M)$ such terms [899]. As such, the ratio $\|H\|_1/\epsilon$ that features multiplicatively in the complexity of quantum phase estimation (at least, variants based on qubitization) can be on the order of 10^4 (or larger) for modest system sizes with $M \approx 100$.
- (ii) To date, only product formula-based methods have been quantitatively studied in the context of providing coherent access to the vibrational Hamiltonian. If expanding the Hamiltonian as a linear combination of Pauli operators, the number of terms grows as $O(M^d N^{2d})$ for a degree d of anharmonic terms considered in the Hamiltonian (often at least 4th order). An alternative approach is to consider the Hamiltonian discretized by a real-space grid or Fourier pseudo-spectral basis, where the position and momentum operators can be easily applied [732, 731].

Existing resource estimates

To date, there have been no end-to-end resource estimates for the vibrational structure problem. In terms of initial steps in this direction, [897] considered

⁴ Due to the close historical ties with spectroscopy, in vibrational chemistry it is common to see energies expressed as wavenumbers. Interconversion can be performed using the Planck relation.

the resources required to map vibrational operators to qubit operators, [899] compared the number and magnitude of terms in vibrational Hamiltonians to those in electronic structure Hamiltonians, and [995] estimated the number of terms and Trotter steps required to perform quantum phase estimation on polyynes molecules.

Caveats

Quantum algorithms and many (but not all [1092]) classical algorithms for simulating the vibrations of nuclei require the availability of an electronic PES, from classical calculations. For a grid-based interpolation of the multi-dimensional PES with h points per dimension, we require $O(h^M)$ PES evaluations. Nevertheless, several interpolation techniques and adaptive methods have been developed to obtain high-accuracy PESs, at lower costs. Moreover, a few molecules with classically challenging vibrational spectra have been identified with classically easy-to-compute electronic PESs [899].

There has been less work on the number of vibrational basis states required to achieve a given accuracy than in the electronic case. While rigorous results exist for more simple bosonic Hamiltonians [991], the truncation level N has not yet been established for anharmonic potentials.

Comparable classical complexity and challenging instance sizes

A hierarchy of approximate classical methods has been developed for computing vibrational eigenstates, which trade increased accuracy for increased cost. Vibrational states with a multireference nature—which are required to describe vibrational resonances that arise due to near degeneracies between different vibrational eigenstates, resulting from anharmonicities in the PES—require more accurate (and thus costly) methods. Moreover, nonrigid molecules require a higher-degree approximation of the PES, leading to an increased cost for classical methods—and potentially increasing the complexity of the resulting eigenstates. For such challenging systems, accurate classical results have been obtained for molecules with $G = 20$ – 30 atoms [237, 84, 989, 99].

Recently, a number of quantum-inspired classical algorithms (see [811, 810] and references therein) have been developed for classically approximating the results of Gaussian boson sampling experiments. These experiments have been proposed as analog simulators of harmonic vibrational phenomena, including vibronic spectra and vibrational dynamics (see §NISQ implementations, below). This reduces the regime where quantum advantage may be possible and reinforces the necessity of considering anharmonicities in the vibrational Hamiltonian [811].

For a review of classical algorithms for simulating coupled vibrational and electronic degrees of freedom, we refer to [1092, 597]. Commonly used algorithms include multiconfigurational time-dependent Hartree (MCTDH) and the *ab initio* multiple spawning method (AIMS).

Speedup

The speedup for quantum algorithms in computing vibrational eigenstates depends on the overlap and error convergence of classical trial states to the true eigenstate and energy. This has yet to be determined for systems of interest. Nevertheless, in spectroscopic calculations that start from a classically easy-to-prepare state, the overlap between the initial state and a desired excited state becomes a quantity of interest—corresponding to the absorption intensity—rather than a limiting factor on the calculation. For example, in [898] it was proposed to use quantum phase estimation to project from the initial state into other eigenstates with probability given by the squared overlap between the states. However, while a single (exponentially costly) classical diagonalization of the vibrational Hamiltonian would provide complete access to the entire vibrational spectrum, a large number of repetitions of the quantum algorithm would be required to reconstruct the spectrum.

As discussed above, the development of recent quantum-inspired algorithms [811, 810] has reduced the prospect of achieving quantum advantage for vibrations in harmonic potentials, motivating the need to include anharmonicities in the vibrational Hamiltonian, or to consider more complex models that nonadiabatically couple vibrational and electronic degrees of freedom. We refer to [597] for a discussion of the classical complexity of simulating such models.

NISQ implementations

There have been proposals to apply variational algorithms to solve the vibrational structure problem [750, 813, 897, 899], but additional developments are required in order to implement sufficiently deep circuits to surpass classical methods, and the number of circuit repetitions required to measure energy observables is likely a bottleneck [899].

There have also been several proposals and experimental demonstrations for using analog quantum simulators to simulate molecular vibrations. Analog simulations have been performed for zero and finite temperature vibronic spectra [553, 1015] as well as vibrational dynamics [946]. We note that these approaches use harmonic approximations for the vibrational potential, and can be approximated efficiently by classical algorithms in some regimes [811]. There have also been analog quantum simulations of systems with coupled electronic

and vibrational degrees of freedom (typically via a linear vibronic coupling model) including simulations of conical intersections [1016, 1036, 1001] and vibrational assisted energy transfer [443].

Outlook

Further work is required to identify target systems that are challenging to simulate classically, but that may be amenable to quantum algorithms. In addition, existing quantum algorithms need to be further optimized for the accuracy required in vibrational structure problems and the form of the vibrational Hamiltonian. This will enable resource estimates for end-to-end applications, such as estimating vibrational spectra or simulating vibrational dynamics.