


Cite this: *RSC Adv.*, 2021, 11, 39438

A systematic variational approach to band theory in a quantum computer

Kyle Sherbert,^a Frank Cerasoli^a and Marco Buongiorno Nardelli^{*ab}

Quantum computers promise to revolutionize our ability to simulate molecules, and cloud-based hardware is becoming increasingly accessible to a wide body of researchers. Algorithms such as Quantum Phase Estimation and the Variational Quantum Eigensolver are being actively developed and demonstrated in small systems. However, extremely limited qubit count and low fidelity seriously limit useful applications, especially in the crystalline phase, where compact orbital bases are difficult to develop. To address this difficulty, we present a hybrid quantum-classical algorithm to solve the band structure of any periodic system described by an adequate tight-binding model. We showcase our algorithm by computing the band structure of a simple-cubic crystal with one s and three p orbitals per site (a simple model for polonium) using simulators with increasingly realistic levels of noise and culminating with calculations on IBM quantum computers. Our results show that the algorithm is reliable in a low-noise device, functional with low precision on present-day noisy quantum computers, and displays a complexity that scales as $\Omega(M^3)$ with the number M of tight-binding orbitals per unit-cell, similarly to its classical counterparts. Our simulations offer a new insight into the "quantum" mindset and demonstrate how the algorithms under active development today can be optimized in special cases, such as band structure calculations.

Received 7th October 2021
Accepted 1st December 2021

DOI: 10.1039/d1ra07451b

rsc.li/rsc-advances

1 Introduction

The primary objective of computational chemistry is to calculate the eigenstates and eigenenergies of a chemical system. Increasing the accuracy of these calculations enables more accurate characterization of a wide range of chemical properties (ionization potential, equilibrium constants, *iR* spectra, *etc.*). However, the computational complexity of the most accurate algorithms scales exponentially with the number of basis orbitals, when using a classical computer. Therefore, large systems are typically treated approximately with a single-electron approximation, in which each electron independently interacts with an effective potential produced by all other electrons and atomic centers. Spin exchange forces can be treated through the Hartree-Fock self-consistent field method, and multi-body correlations can be treated through *ad hoc* correction terms as in density functional theory, but these methods fail when applied to highly-correlated systems. Therefore, computational chemists have begun to turn to *quantum* computers to reach higher levels of accuracy.

Quantum computers surpass Hartree-Fock by imposing fermionic statistics onto qubits and including electron correlation terms directly in the system Hamiltonian. Quantum Phase Estimation (QPE) is a quantum algorithm which extracts

eigenstate energies from a simulated system on a noise-resilient quantum computer at a low cost of classical computational resources.^{1,2} However, quantum circuits designed for the QPE algorithm tend to require longer coherence times than are available in the era of Noisy Intermediate-Scale Quantum (NISQ) devices, so recent efforts focus on hybrid algorithms which balance quantum and classical resource costs. In particular, a more popular algorithm for molecular ground-state energy calculations is the Variational Quantum Eigensolver (VQE).³⁻⁶ Many variants of VQE have arisen in recent literature, including methods such as Variational Quantum Deflation (VQD) capable of exploring excited states.⁷ However, resource costs for interesting systems still tend to exceed those currently available, and there is still some doubt whether error in hybrid algorithms can be made to converge sub-exponentially.⁸

Periodic systems, such as the crystalline phase of a molecule or protein, are an especially interesting arena in the development of quantum algorithms. On the one hand, translational symmetry over the entire crystal seems to offer extremely powerful ways to reduce the computational complexity.⁹ On the other, periodic systems are typically considered infinite in extent and thus appear to require an exceedingly *large* number of resources to adequately approximate. For example, one may simulate a periodic system of N unit cells, each consisting of M orbitals (see for example Cade *et al.*³⁶), where M is typically comparable to the number of orbitals considered in a single molecular simulation, but N is large enough to approximate

^aDepartment of Physics, University of North Texas, Denton, TX 76203, USA. E-mail: mbn@unt.edu

^bSanta Fe Institute, Santa Fe, NM 87501, USA


infinity. Alternatively, one may adopt a plane-wave basis, for which a quantum circuit is available to efficiently diagonalize the kinetic and potential operators directly.¹⁰ In either approach, the size of the basis, and therefore the number of qubits, must be very large to accurately represent the periodic system. The quantum resources required to simulate a crystal will thus tend to be many times larger than those required for a solitary molecule, and generally larger than the size of quantum computers available today.

In this work, we offer an easier transition to adapt quantum algorithms for periodic systems, by implementing correlation-free band structure calculations on NISQ-era quantum computers. Band structures are the fundamental toolbox of materials scientists in the characterization and discovery of the electronic properties of crystalline solids. Band theory adopts the single-electron approximation (as in Hartree-Fock), for which a periodic Hamiltonian becomes separable in reciprocal space, reducing the system at any particular momentum \mathbf{k} to the complexity of a single unit cell. In this way, the eigenstates of an electron with momentum \mathbf{k} can be efficiently calculated in a classical computer; the energies of each eigenstate along a path varying \mathbf{k} through reciprocal space form the band structure of the crystal. Integrating the band structure provides early insight into structural, electronic, optical, and thermal properties of the crystal.¹¹

Quantum algorithms to obtain classically verifiable results provide an invaluable tool in establishing foundational building blocks such as efficient measurement protocols and error mitigation,¹² as well as being generally easier to understand and replicated by researchers just breaking into quantum computation. With this motivation, we show how the single-electron approximation accommodates a systematic approach to efficiently apply VQE to solve the band structure of any periodic system. We will illustrate our procedure using an empirical tight-binding model for a simple cubic lattice, but our procedure is easily applied to any tight-binding Hamiltonian. For example, recent work in our group has demonstrated that “exact” tight-binding Hamiltonians can be derived from accurate electronic structure calculations using a procedure based on the projection of electronic wavefunctions on localized orbital bases (PAO-FLOW).¹³ PAO-FLOW is an advanced software tool that constructs tight-binding Hamiltonians from self-consistent electronic wavefunctions projected onto a set of atomic orbitals and provides numerous materials properties that otherwise would have to be calculated *via* approximate model approaches. Thus, an efficient quantum algorithm for the solution of the tight-binding Hamiltonian would provide an avenue for the calculation of advanced crystalline properties on a quantum computer. We do not expect our approach in this paper to offer immediate quantum advantage; rather, our purpose is to help chemists think the *quantum* way, motivating new, resource-efficient approaches to studying highly correlated systems. By considering the simplest available model, we can provide lower bounds on resource complexity and give insight into the practical difficulties chemists may expect when implementing quantum algorithms.

We have previously considered this topic, employing a VQE-based algorithm to iteratively calculate the band energies of

a tight-binding silicon model.¹⁴ We now apply recent developments in the literature^{7,15,16} to extend and improve upon our previous work. In Section 2, we briefly outline the essential ideas we have taken from the literature. In Section 3, we present our robust procedure for accurately calculating the band structure of any periodic system with a quantum computer. In Section 4, we demonstrate our procedure applied to a simple-cubic lattice, presenting data from a quantum simulator and preliminary results from IBM’s open-access `ibmq_athens` and `ibmq_santiago` cloud devices. In Section 5, we discuss the algorithmic complexity of our procedure and highlight the steps which may or may not be improved in later work.

2 Background

In this section we briefly outline some essential techniques actively studied in the quantum computing literature. In particular, while QPE provides a robust strategy for measuring eigenenergies with minimal classical resources, its performance suffers greatly from the imperfect fidelity of NISQ devices. As such, we will focus mostly on VQE and its close cousin VQD when measuring eigenspectra, applying QPE when available as an optional refinement.

2.1 Quantum phase estimation

In the QPE algorithm,^{1,2} a set of qubits (the “state register”) are first prepared into an eigenstate $|\psi\rangle$ of a unitary operator \hat{U} , such that $\hat{U}|\psi\rangle = e^{2\pi i\phi}|\psi\rangle$. The unitary operator U is then repeatedly applied as a controlled quantum circuit so that the phase shift ϕ is encoded into another set of qubits (the “readout register”). Measurements on the readout register give the binary expansion of ϕ . In molecular simulations one selects $\hat{U} \equiv \exp(i\hat{H}\tau)$, the operator which evolves a system with Hamiltonian \hat{H} by a unit time τ . If one first transforms \hat{H} to guarantee that all possible energies E fall within the interval $[0, 2\pi/\tau)$, the measured phases ϕ map directly onto the eigenenergies of \hat{H} . The algorithm is deterministic when the state register is prepared exactly into an eigenstate $|\psi\rangle$ and its eigenvalue ϕ has an exact binary expansion, but it retains some probability of success when both conditions are relaxed. Thus, QPE can be adapted to discover eigenstates and eigenvalues *a priori*, at the cost of additional rounds of measurement.

Generally, \hat{H} is given as a weighted sum of non-commuting Pauli words (Section 3.1). An exact circuit for $\hat{U} \equiv \exp(i\hat{H}\tau)$ is not readily available, but it can be closely approximated by Suzuki-Trotter expansion, which factors \hat{U} into many small-time slices.¹⁷ The number of time slices scales polynomially with the accuracy required, and the depth of each time slice depends on the number of commuting groups in \hat{H} . For this reason, QPE is extremely susceptible to errors arising from the low gate fidelity and short coherence times which plague NISQ devices. Alternative approaches scale more favorably with error at the cost of additional ancilla qubits.^{18,19}



2.2 Variational quantum eigensolver

In the VQE algorithm,^{3,4} one begins with a Hamiltonian \hat{H} , represented as a weighted sum of non-commuting Pauli words, and a parameterized quantum circuit $\hat{V}(\theta)$, the “ansatz”, which prepares a set of qubits into an arbitrary state (Section 3.2). The ansatz is applied on an ensemble of states such that qubit measurements give the expectation value of each Pauli word in \hat{H} . The weighted sum of expectation values gives the energy $E(\theta)$ of the arbitrary state prepared – this procedure is called “operator estimation”.^{4,20} The parameters θ are varied by a classical optimization routine until E is minimized. According to the variational principle, this minimum is exactly the ground-state energy of the system when the ansatz $\hat{V}(\theta)$ is robust (*i.e.* it spans the full Hilbert space of the system), and if the classical optimization succeeds in producing the global minimum.

The algorithmic complexity of VQE depends on several factors. Measuring the expectation value of a Pauli word is a stochastic process, requiring a large number of measurements on the order of $O(\varepsilon^{-2})$ for an acceptable sampling noise ε . These ensembles are usually measured for each Pauli word in \hat{H} , although recent advances reduce the size of the ensemble by simultaneously measuring each commuting group of Pauli words³⁵ or by “classically shadowing”²¹ the quantum circuit to require only a logarithmic number of measurements. Like in QPE, the efficiency of the algorithm is determined by the complexity of \hat{H} , and every element of the ensemble requires a unique application of the ansatz, meaning that circuit depth and gate count should be kept minimal. The dimension of the ansatz also determines the efficiency and efficacy of the classical optimization. For all these reasons, VQE tends to be impractical for perfectly robust ansatz, and much of the literature focuses on methods for constructing effective ansatz accounting for system symmetries and hardware limitations.^{15,22–25} Because circuit depth and gate count are kept low, VQE is well-suited to NISQ devices.

2.3 Variational quantum deflation

Variational Quantum Deflation (VQD)⁷ is one approach for extending the VQE algorithm to explore excited states in addition to the ground-state. VQD begins as a typical VQE run to locate the ground state, and the ground-state parametrization θ_0 is recorded. The variational process is then repeated with an additional term in the optimization routine's cost function, which gives the overlap between the current ansatz and the ground-state, weighted by a factor β . States similar to the ground-state will be shifted into a higher effective energy, so that the optimization routine considers them unfavorable. Meanwhile, higher-energy eigenstates must be orthogonal to the ground-state, so their overlap contribution will be zero. Therefore, the next lowest energy that can be found is the first excited state. This process is repeated for each energy level, adding a new overlap term for each eigenenergy already found. Each overlap can be evaluated as the expectation value of a single commuting group of Pauli words in the Hamiltonian, so that the total number of additional measurements after finding M eigenvalues is $\Theta(M^2)$. If \hat{H} consists of $\Omega(M)$

commuting groups, measured for each of M energy levels, then the additional cost of the overlap circuits is negligible.

3 Method

Our objective is to calculate the band structure of a periodic system, as described by a tight-binding Hamiltonian \hat{H} of the form:

$$\hat{H} = \sum_{\alpha,\beta} t_{\alpha\beta} c_{\alpha}^{\dagger} c_{\beta} \quad (1)$$

Each c_j^{\dagger} and c_j represent a creation and annihilation (ladder) operator on an atomic orbital ϕ_j , centered on a coordinate \mathbf{r}_j in the crystal. The hopping parameters $t_{\alpha\beta}$ denote the energy cost of an electron transitioning from orbital ϕ_{β} to orbital ϕ_{α} . They are calculated from the overlap integrals between each pair of orbitals ϕ_{α} and ϕ_{β} , or they are selected to fit empirical observations. A general tight-binding Hamiltonian may also include multi-electron correlation such as $t_{\alpha\beta\gamma\delta} c_{\alpha}^{\dagger} c_{\beta}^{\dagger} c_{\gamma} c_{\delta}$, but we neglect these terms in this work.

Our strategy is to transform eqn (1) into reciprocal space and to apply VQD to solve for each eigenenergy at each momentum \mathbf{k} along the desired path through reciprocal space. When sufficient quantum resources are available, we refine each band energy with QPE. Our procedure for mapping a single-electron periodic system onto a set of qubits is derived in Section 3.1. The variational ansatz we have selected, suitable for any band structure calculation, is described in Section 3.2. Details of implementing the quantum algorithm are presented in Section 3.3. Finally, we provide a step-by-step schematic of our algorithm and its relation to VQE in Fig. 2.

3.1 Qubit mapping

The Hamiltonian in eqn (1) consists of ladder operators acting on atomic orbitals. The Hamiltonians appearing in the quantum algorithms of Section 2 consist of Pauli words acting on qubits. We define a “Pauli word” \hat{P}_i as an operator acting independently on each qubit with either the identity \hat{I} or one of the Pauli spin matrices \hat{X} , \hat{Y} , \hat{Z} . Pauli words are a natural choice for representing physical operators in a quantum computer because their expectation values can be readily measured¹⁴ and their unitary time evolution $\exp(i\hat{P}_i t)$ can be readily implemented as a quantum circuit.²⁶ Our goal in this section is to map our atomic orbitals onto a qubit basis, and our Hamiltonian to a weighted sum of Pauli words:

$$\hat{H} = \sum_{\alpha} \sum_{\beta} t_{\alpha\beta} c_{\alpha}^{\dagger} c_{\beta} \rightarrow \sum_i a_i \hat{P}_i \quad (2)$$

3.1.1 Qubit basis. The simplest conceivable mapping between orbitals and qubits is to identify each orbital with its own qubit. The qubit state $|1\rangle$ represents an occupied orbital, while $|0\rangle$ is empty. There are however an infinite number of orbitals in an infinite crystal, and quantum computers with an infinite number of qubits are beyond our engineering capabilities. We therefore reinterpret \hat{H} as the Hamiltonian of an



arbitrarily large supercell with periodic boundary conditions, consisting of N unit cells, each with M orbitals.

$$\hat{H} = \sum_{\nu}^{(N)} \sum_{\nu'}^{(N)} \sum_{\alpha=0}^{M-1} \sum_{\beta=0}^{M-1} t_{\alpha\beta}^{(\delta)} c_{\nu'\alpha}^{\dagger} c_{\nu\beta} \quad (3)$$

Hopping parameters are now dependent on the orbitals α , β and the displacement vector $\delta \equiv \mathbf{r}_{\nu'\alpha} - \mathbf{r}_{\nu\beta}$ between their atoms. As δ increases, $t_{\alpha\beta}^{(\delta)}$ tends to vanish, permitting a nearest-neighbor approximation in which one considers only a few of the smallest δ .

Eqn (3), when supplemented with two-electron correlation terms, is the form typically considered when applying quantum algorithms to periodic systems, requiring a total of MN qubits. In the single-electron approximation, however, we can reduce the size of the system to only M qubits by transforming into reciprocal space. Reciprocal space orbitals are characterized by their own ladder operators $\tilde{c}_{\mathbf{k}j}^{\dagger}$ and $\tilde{c}_{\mathbf{k}j}$, related to $c_{\nu j}^{\dagger}$ and $c_{\nu j}$ by Fourier transform:

$$c_{\nu'\alpha}^{\dagger} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}'} e^{i\mathbf{k}'\mathbf{r}_{\nu'\alpha}} \tilde{c}_{\mathbf{k}'\alpha}^{\dagger} \quad (4a)$$

$$c_{\nu\beta} = \frac{1}{\sqrt{N}} \sum_{\mathbf{k}} e^{-i\mathbf{k}\mathbf{r}_{\nu\beta}} \tilde{c}_{\mathbf{k}\beta} \quad (4b)$$

Substituting eqn (4) into eqn (3), we obtain

$$\hat{H} = \sum_{\mathbf{k}}^{(N)} \sum_{\mathbf{k}'}^{(N)} \sum_{\alpha=0}^{M-1} \sum_{\beta=0}^{M-1} \left(\frac{1}{N} \sum_{\nu} \sum_{\nu'} t_{\alpha\beta}^{(\delta)} e^{i(\mathbf{k}'\mathbf{r}_{\nu'\alpha} - \mathbf{k}\mathbf{r}_{\nu\beta})} \tilde{c}_{\mathbf{k}'\alpha}^{\dagger} \tilde{c}_{\mathbf{k}\beta} \right) \quad (5)$$

We simplify this sum by recalling $\mathbf{r}_{\nu'\alpha} = \mathbf{r}_{\nu\beta} + \delta$. Then $\mathbf{r}_{\nu'\alpha}$ becomes a common factor of each \mathbf{k} in the exponential, and we may exploit the orthogonality relation $\frac{1}{N} \sum_{\nu'} e^{i(\mathbf{k}' - \mathbf{k})\mathbf{r}_{\nu'\alpha}} = \delta_{\mathbf{k}'\mathbf{k}}$.

Summing over $\delta_{\mathbf{k}'\mathbf{k}}$, we obtain $\hat{H} = \sum_{\mathbf{k}} \hat{H}_{\mathbf{k}}$, where

$$\hat{H}_{\mathbf{k}} \equiv \sum_{\alpha=0}^{M-1} \sum_{\beta=0}^{M-1} H_{\alpha\beta}(\mathbf{k}) \tilde{c}_{\mathbf{k}\alpha}^{\dagger} \tilde{c}_{\mathbf{k}\beta} \quad (6)$$

$$H_{\alpha\beta}(\mathbf{k}) \equiv \sum_{\delta} t_{\alpha\beta}^{(\delta)} e^{i\mathbf{k}\delta} \quad (7)$$

Each momentum \mathbf{k} contributes an independent subsystem with only M orbitals, whose eigenenergies may be solved independently. Classically, the values $H_{\alpha\beta}(\mathbf{k})$ in eqn (7) form an $M \times M$ Hermitian matrix whose eigenvalues can be efficiently calculated with standard linear algebraic techniques in $\Theta(M^3)$ time. This work instead considers how to calculate these eigenvalues the “quantum” way.

We focus on a specific $\hat{H}_{\mathbf{k}}$ for the remainder of this section, with the understanding that our procedure must be repeated for each momentum \mathbf{k} along the path of interest in reciprocal space. Eqn (6) has a form very similar to eqn (1), except that it

acts on reciprocal-space orbitals rather than atomic orbitals. We therefore adopt a “reciprocal-orbital” basis, in which each reciprocal-space orbital is identified with its own qubit.

3.1.2 Hamiltonian mapping. Having transformed our Hamiltonian into reciprocal space (eqn (6)), we must now consider mapping each ladder operator to a set of Pauli words. The ladder operators must satisfy the following:

$$\tilde{c}|0\rangle = 0 \quad \tilde{c}|1\rangle = |0\rangle \quad (8a)$$

$$\tilde{c}^{\dagger}|0\rangle = |1\rangle \quad \tilde{c}^{\dagger}|1\rangle = 0 \quad (8b)$$

Meanwhile, the Pauli spin operators \hat{X} , \hat{Y} , \hat{Z} act on a qubit's basis states in the following way:

$$\hat{X}|0\rangle = |1\rangle \quad \hat{X}|1\rangle = |0\rangle \quad (9a)$$

$$-i\hat{Y}|0\rangle = |1\rangle \quad i\hat{Y}|1\rangle = |0\rangle \quad (9b)$$

$$\hat{Z}|0\rangle = |0\rangle \quad -\hat{Z}|1\rangle = |1\rangle \quad (9c)$$

It is easy to verify that the following mapping suffices for a single qubit:

$$\tilde{c} \rightarrow \frac{1}{2}(\hat{X} + i\hat{Y}) \quad (10a)$$

$$\tilde{c}^{\dagger} \rightarrow \frac{1}{2}(\hat{X} - i\hat{Y}) \quad (10b)$$

In multi-electron systems, one typically adopts the Jordan-Wigner transformation, which retains the form of eqn (10) but appends a \hat{Z} operation on $\Theta(M)$ other qubits to enforce fermionic antisymmetry. Alternatively, one may adopt the Bravyi-Kitaev transformation, which requires operations on only $\Theta(\log M)$ qubits, but uses a non-intuitive basis and involves non-adjacent interactions more difficult to simulate on certain qubit architectures. We refer the reader to Seeley *et al.*²⁶ for an excellent introduction to both transforms. However, because we are considering single-electron systems, there are no other fermions to exchange with, and we may use eqn (10) directly, so that each ladder operator acts on only $\Theta(1)$ qubits. This significantly reduces the complexity of our Hamiltonian, as we shall presently see.

We may rewrite eqn (6) to exploit the Hermiticity of $\hat{H}_{\mathbf{k}}$.

$$\hat{H}_{\mathbf{k}} = \sum_{\alpha} H_{\alpha\alpha} c_{\alpha}^{\dagger} c_{\alpha} + 2 \sum_{\alpha} \sum_{\beta > \alpha} \text{Re}(H_{\alpha\beta} c_{\alpha}^{\dagger} c_{\beta}) \quad (11)$$

Since the transpose term $H_{\beta\alpha} c_{\beta}^{\dagger} c_{\alpha} = (H_{\alpha\beta} c_{\alpha}^{\dagger} c_{\beta})^{\dagger}$. Applying eqn (10) and noting $\hat{X}^2 = \hat{Y}^2 = \hat{I}$, $-i\hat{X}\hat{Y} = i\hat{Y}\hat{X} = \hat{Z}$:

$$\begin{aligned} \hat{H}_{\mathbf{k}} \rightarrow & \frac{1}{2} \sum_{\alpha} H_{\alpha\alpha} (\hat{I} - \hat{Z}_{\alpha}) \\ & + \frac{1}{2} \sum_{\alpha} \sum_{\beta > \alpha} \text{Re}(H_{\alpha\beta}) (\hat{X}_{\alpha} \hat{X}_{\beta} + \hat{Y}_{\alpha} \hat{Y}_{\beta}) \\ & + \frac{1}{2} \sum_{\alpha} \sum_{\beta > \alpha} \sum_{\gamma > \alpha\beta} \text{Im}(H_{\alpha\beta}) (\hat{Y}_{\alpha} \hat{X}_{\beta} - \hat{X}_{\alpha} \hat{Y}_{\beta}) \end{aligned} \quad (12)$$



Eqn (12) provides the weighted sum of Pauli words required in the quantum algorithms of Section 2.

Eqn (12) consists of $\Theta(M^2)$ Pauli words. The complexity of each algorithm in Section 2 is determined in part by the number of commuting groups in \hat{H} . In eqn (12), all terms of the form \hat{Z}_α , $\hat{X}_\alpha\hat{X}_\beta$, and $\hat{Y}_\alpha\hat{Y}_\beta$ each form commutative groups. Therefore, when \hat{H}_k has no imaginary part, the energy can be determined with just 3 rounds of measurement. When \hat{H}_k does have an imaginary part, we note that for fixed α , $\hat{Y}_\alpha\hat{X}_{\beta\gamma\alpha}$ and $\hat{X}_\alpha\hat{Y}_{\beta\gamma\alpha}$ each form commutative groups, so in general we have $\Theta(M)$ commuting groups. Finally, we note that each of these commuting groups are *qubit-wise* commutative, meaning that each index of all Pauli words in the set has either the same spin operator or the identity. This accommodates a particularly simple procedure for measuring expectation values of each set simultaneously, requiring no additional overhead in the measurement circuit.

3.2 Ansatz

The VQE and VQD algorithms require an ansatz – a parameterized quantum circuit $\hat{V}(\theta)$ preparing a trial state $|\Psi(\theta)\rangle \equiv \hat{V}(\theta)|0\rangle$ for energy measurements. A quantum circuit to span the full Hilbert space of M qubits requires $2(2^M - 1)$ parameters, and it will not generally have an efficient decomposition into one- and two-qubit gates. However, most applications to molecular simulation consider a system with fixed number of electrons. In the orbital basis, or in our reciprocal orbital basis, one need only consider that subset of Hilbert space spanned by the basis states whose Hamming weights match the number of electrons in our system. For example, in band structure calculations we consider just one electron, so we need only consider the space spanned by $|10\dots\rangle$, $|010\dots\rangle$, etc.

Gard *et al.*¹⁵ provide a procedure for generating variational ansatz which conserve particle number, which is particularly

simple when particle number is 1. We begin with M qubits labeled 0 through $M - 1$ in the state $|0\rangle$. First, we apply an \hat{X} gate to qubit 0, to set our ansatz with a single filled orbital. Then we apply the entangling parameterized A gate¹⁵ such that each qubit is entangled directly or indirectly with qubit 0 (see Fig. 1). This ansatz requires $M - 1$ A gates, each contributing two independent parameters, for a total of $\Theta(M)$ gates and parameters. The circuit is compatible with any quantum architecture exhibiting linear qubit connectivity and has a depth of $\Theta(M)$. Alternatively, in a fully connected device, the A gates could be applied with a “divide-and-conquer” strategy, reducing the circuit depth to $\Omega(\log M)$.

Rather than assigning each orbital to its own qubit, we *could* assign each orbital to an individual basis state, requiring only $\Theta(\log M)$ qubits total. This “compact basis” is the approach of our previous work.¹⁴ While this is more efficient in the number of qubits, it must explore states with an arbitrary Hamming weight. The number of parameters required to span the space of interest is unchanged, and a suitable ansatz must be developed *ad hoc*. Additionally, the Hamiltonian for the compact basis consists of global operators acting on all qubits at once, and it would generally form the maximum number $3^{\log_2 M} = M^{\log_2 3}$ of commuting sets, requiring a less efficient measurement protocol. Reliance on a random ansatz and a global cost function made our previous procedure vulnerable to exponentially difficult optimizations induced by barren plateaus.^{8,27} Finally, the Hamiltonian for the compact basis is less-structured and more difficult to reduce based on symmetries in the Hamiltonian (for example our observation in Section 3.1 that a real \hat{H}_k results in $\Theta(1)$ rounds of measurement).

3.3 Band structure calculation

With our ansatz $\hat{V}(\theta)$ (Fig. 1) and qubit Hamiltonian \hat{H}_k (eqn (7) and (12)) prepared, we are ready to implement VQD for each

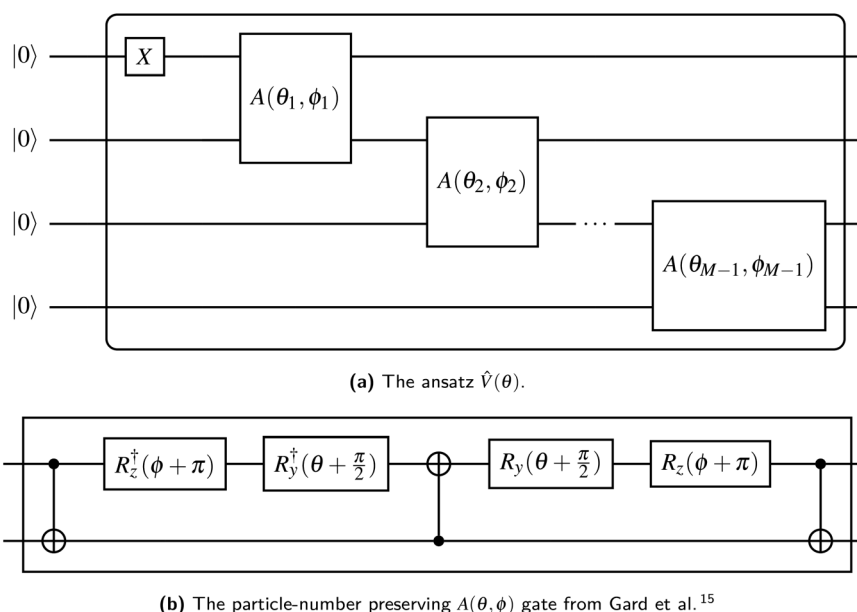


Fig. 1 The ansatz $V(\theta)$ suitable for any band structure calculation. Each qubit is initialized in the $|0\rangle$ state; the output is an arbitrary superposition of states with a single qubit in the $|1\rangle$ state.



momentum \mathbf{k} along a path through reciprocal space. This path is usually constructed from high-symmetry segments in the crystal's First Brillouin Zone, because this proves sufficient to calculate many properties of interest. As briefly described in Section 2, the idea is to vary the trial state prepared by our ansatz until the energy $E \equiv \langle \hat{H} \rangle$ is minimized. We repeat the optimization for each band energy, adding additional terms to the cost function proportional to the overlap between the trial state and each previously found eigenstate, weighted by the constant factor β .

The expectation values $\langle \hat{H} \rangle$ of a generic observable cannot be directly measured in the quantum computer. Rather, the expectation value of each Pauli word \hat{P}_i are measured independently, and the energy is evaluated from the weighted sum $\langle \hat{H} \rangle = \sum_i a_i \langle \hat{P}_i \rangle$, with weights a_i taken from eqn (12). Obtaining the Pauli expectation values $\langle \hat{P}_i \rangle$ is also somewhat indirect. First, the Pauli word \hat{P}_i should be transformed so that it contains only letters \hat{I} or \hat{Z} – let us refer to the modified Pauli word as \hat{Q}_i . In practice, the transformation is easily accomplished by applying a “basis rotation” gate to each qubit before measurement. Next, each qubit is measured to be in one of the two computational basis states $|0\rangle$ or $|1\rangle$. The bitstring obtained from concatenating the state of each qubit is itself an eigenstate of \hat{Q}_i , with eigenvalue $+1$ or -1 . This procedure is applied to a large *ensemble* of qubits, each prepared independently with the ansatz and basis rotation gates. The expectation value $\langle \hat{P}_i \rangle$ is the average of all the eigenvalues of \hat{Q}_i measured across the ensemble.

The ensemble necessarily has a finite size S , introducing an energy variance on the order of $\varepsilon^2 \in O(1/S)$. In practice, the ensemble is usually prepared in sequence, resetting a single register of qubits after each round of measurement, relegating the sampling error ε a parameter in the time complexity of any VQE-based algorithm. Fortunately, the same ensemble may be used to calculate the expectation values of any Pauli word which is *qubit-wise* commutative with \hat{P}_i . For simplicity, we assign $S = 8096$ for each commuting group in this work, although advanced methods exist which optimally distribute measurements to minimize the sampling error ε .²⁰

Many popular optimization routines (e.g. SLSQP, BFGS) are gradient-based, and they have difficulty converging to the correct value in the presence of sampling noise. Therefore, we use COBYLA, a simplex-based algorithm implemented in the *scipy* python package, which we have empirically noted to give good results. We randomly generate our initial guess for the parameters θ , and we use the default tolerance parameters implemented by *scipy*. These choices are by far the simplest, but they are by no means optimal, and our results may be improved greatly by a more careful choice of optimization routine.²⁸

Before we can implement the deflation procedure, we must select the constant β suitable for “deflating” each band energy. We do this with a systematic procedure, first *maximizing* the energy of our system to find the highest possible energy E_{\max} . We then minimize the energy to find E_0 and $\Delta \equiv E_{\max} - E_0$. In theory, $\beta = \Delta$ is a sufficiently high number to guarantee each eigenstate is projected sufficiently out of the optimization in

later steps. In practice, we take $\beta = 2\Delta$ to insure against errors in the sampling and optimization process.

Higgott *et al.*⁷ offer several strategies for computing the overlap, offering robustness against error at the cost of ancilla qubits or additional optimization steps. In this work, we choose the simplest, evaluating the overlap with an eigenstate $\Psi(\theta_i)$ by preparing the trial state $\Psi(\theta)$ and applying the *adjoint* circuit $\hat{V}_i^\dagger \equiv \hat{V}_i^\dagger(\theta_i)$. The probability of measuring the bitstring $0 \dots 0$ gives the overlap $|\langle \Psi(\theta) | \Psi(\theta_i) \rangle|^2$. In practice, the probability of measuring bitstring $0 \dots 0$ is equivalent to the expectation value of an operator $\Omega_0 \equiv \sum_i \hat{Q}_i$, the sum of all unique Pauli words spelled with the letters \hat{I} and \hat{Z} (e.g. $\hat{I}\hat{I}$, $\hat{I}\hat{Z}$, $\dots \hat{Z}\hat{Z}\hat{Z}$). All such operators are qubit-wise commutative and can be estimated with a single round of measurements. Therefore, we can implement the deflation procedure conveniently in the *qiskit* Python package provided by IBM, by solving for each band energy and then adding to our Hamiltonian the deflation operator $\beta \hat{V}_i \hat{\Omega}_0 \hat{V}_i^\dagger$.

Initializing the Hamiltonian $\hat{H}_0 \equiv \hat{H}_{\mathbf{k}}$, our procedure can be formally summarized as follows:

$$\theta_l \equiv \arg \min_{\theta} \langle 0 | \hat{V}_l^\dagger(\theta) \hat{H}_l \hat{V}_l(\theta) | 0 \rangle \quad (13)$$

$$\hat{V}_l \equiv \hat{V}_l(\theta_l) | 0 \rangle \quad (14)$$

$$E_l \equiv \langle 0 | \hat{V}_l^\dagger \hat{H}_l \hat{V}_l | 0 \rangle \quad (15)$$

$$H_{l+1} \equiv H_l + \beta \hat{V}_l \hat{\Omega}_0 \hat{V}_l^\dagger \quad (16)$$

Each E_l we find is recorded as the energy of the l th band at momentum \mathbf{k} , and we repeat the procedure for each \mathbf{k} in our selected path.

Optimization routines do not always converge to the true minimum, and errors incurred early in the deflation procedure can propagate unfavorably to higher bands. Therefore, we include an optional QPE refinement to our algorithm, which applies QPE to each optimized state $\Psi_l \equiv \hat{V}_l | 0 \rangle$. QPE has the effect of selecting the dominant eigenstate of Ψ_l and giving the corresponding eigenenergy with high precision. Thus, as long as the optimization procedure is “good enough”, we may update our energy calculations with the result of the QPE experiment. We have used the iterative version of QPE implemented in *qiskit*. Details of the algorithm can be found in Dobšiček *et al.*²

4 Results

To demonstrate our procedure, we consider a basic model for a crystal in a simple cubic lattice structure (see Fig. 3). Each atom has s , p_x , p_y , and p_z orbitals ($M = 4$), with a large energy gap between the s and p orbitals, and comparatively small hopping parameters between neighboring s and p orbitals. This is the simplest three-dimensional periodic system accommodating multiple orbitals per atom, and the required number of qubits (one qubit per orbital, plus one ancilla qubit to implement QPE) fits nicely onto IBM's open-access five-qubit machines. It may also be considered a rough model for



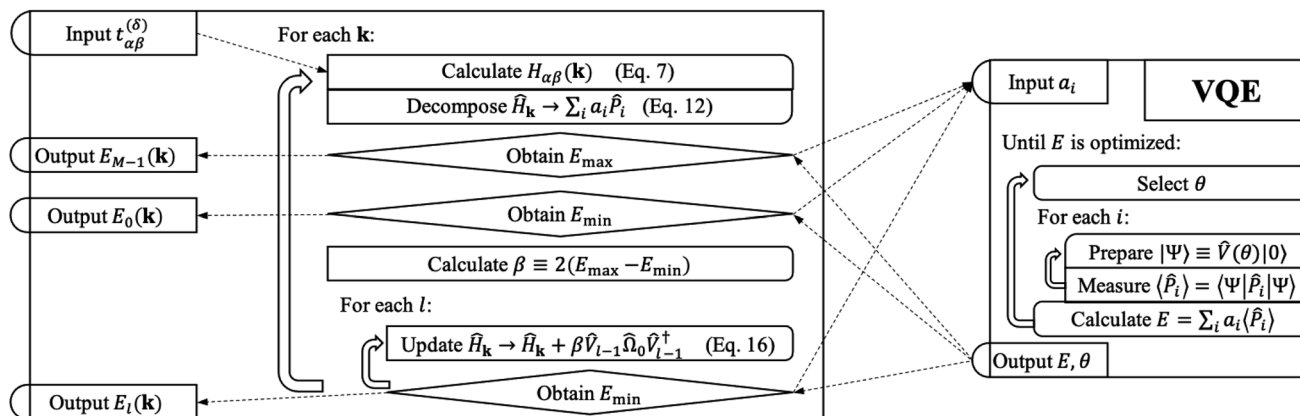


Fig. 2 A schematic of our algorithm and its relation to VQE. Our algorithm takes tight-binding parameters $t_{\alpha\beta}^{(\delta)}$ as input and outputs each band energy $E_l(\theta)$. Optionally, each band energy may be refined with QPE. The operator \mathcal{Q}_0 is the sum of all Pauli words spelled with letters l and Z . The operator V is the quantum circuit presented in Fig. 1.

elemental Polonium, although more accurate models should take into account relativistic effects and Coulomb interaction between orbitals located on the same atom.^{29,30} The exact eigenenergies of our model at specific k -points along a high-symmetry path are calculated using standard linear algebraic techniques to diagonalize the matrix elements in eqn (7). We

compare this band structure to the results from the quantum algorithm presented in Section 3 in four different levels of simulation:

- (1) Statevector – quantum operations are simulated with unitary matrices, and expectation values are calculated exactly.




Fig. 3 Statevector simulator – the band structure of a simple cubic lattice with s and p orbitals (right inset) along the high-symmetry path XMI through the lattice's. First Brillouin zone (left inset). Solid curves denote classical (exact) diagonalization. Diamonds denote the median optimization result from applying our method on a noiseless statevector simulator 32 times with a different random seed. Bars (only visible between the third and fourth bands at nearly-degenerate momenta) denote interquartile ranges. Hopping parameters are 2 eV between adjacent s and p orbitals and 2 eV between colinear p orbitals. Each s orbital has a self-energy of -14 eV to generate a large band gap.

(2) Sampling – expectation values are now calculated by sampling from a probability distribution.

(3) Noisy – quantum operations and measurements are now applied with an error rate drawn from real quantum devices.

(4) Calibrated – the same noisy simulator is used, but classical post-processing steps are applied to mitigate the error.

We also present preliminary results from IBM quantum devices.

4.1 Statevector simulator

To validate our algorithm's capability of producing the correct band structure, we model the state of an n qubit system as a complex statevector and quantum operations as unitary matrices acting on the Hilbert space spanned by the 2^n dimensional basis vectors. Expectation values are evaluated analytically. Such a simulation gives the ideal behavior of a quantum computer, with perfect qubit fidelity and no sampling variance. Fig. 3 summarizes the results of over 32 randomly-seeded optimization runs, marking the median optimization with a diamond and the interquartile range with a bar. For a few momenta where the third and fourth bands are very close together, optimization tends to locate the wrong eigenstate, giving a small variance in results. For every other

point, the diamond coincides perfectly with the classical solution and the bar is absent, demonstrating that our ansatz is robust, the deflation procedure is mathematically sound, and that our choice of optimization routine (COBYLA) is generally consistent in converging to the correct values on a smooth surface.

4.2 Sampling simulator

We now consider long-term viability of our procedure by retaining perfect qubit fidelity but simulating realistic measurement. The same unitary matrices as in the statevector simulator are applied to an ensemble of states, which are “measured” by sampling from the resulting probability distribution a finite number of times. While mathematically equivalent, the sampling noise resulting from the stochastic measurement process can make the energy surface bumpy, which can have a detrimental effect on the optimization step. We have selected the COBYLA optimization algorithm because it is resistant to these bumps; nevertheless, the anomalous variance observed at nearly degenerate points in the statevector simulator is now commonplace. Fig. 4 shows our results on the noiseless qubit simulator over 32 randomly-seeded runs, clearly marking the median (asterisk) and mean (X) for both

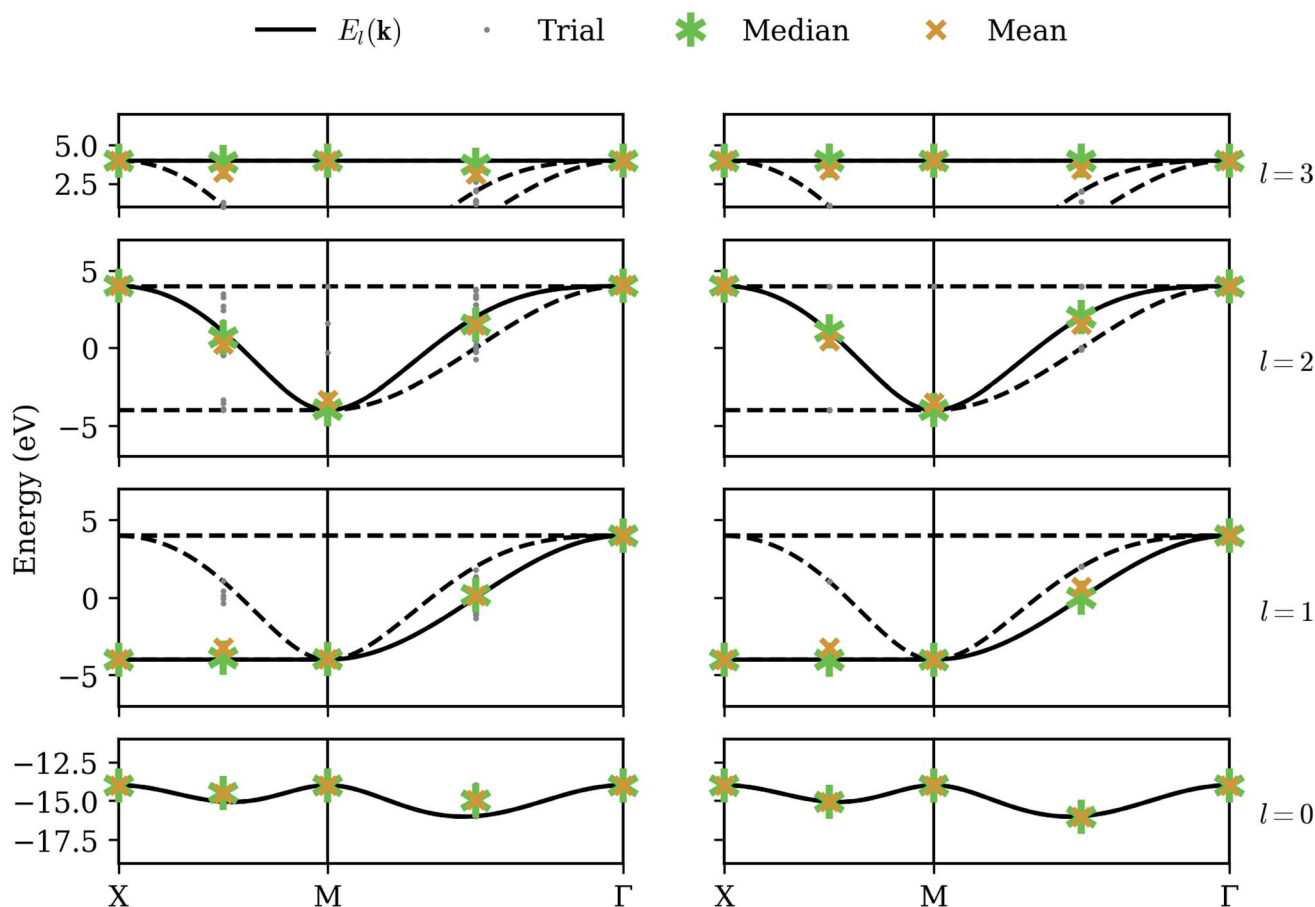


Fig. 4 Sampling simulator – our method applied in the presence of sampling noise (high-fidelity qubits). The left column shows raw optimization results; the right column shows the energy obtained by QPE refinement. Gray dots denote the results from each of 32 trials, with each band given its own row. The asterisk and X denote the median and mean, respectively.

optimization (left) and QPE refinement (right). The smaller dots denote the results of individual trials.

The optimization results are extremely accurate and precise on the high-symmetry momenta but deviate slightly on the intermediate points. In fact, the high-symmetry points in this particular model each happen to have matrices $H_{\alpha\beta}(\mathbf{k})$ (eqn (7)) which are already diagonalized, and the resulting cost function yields a well-behaved surface which is reliably optimized, even in the presence of noise. Averaged results on the intermediate points still tend to be quite good, but individual trials can exhibit a large variance. However, the optimization does succeed in finding a point *close enough* to a correct eigenstate that the QPE refinement consistently extracts the dominant eigenvalue with high precision. The median QPE results prove to be as accurate as is permitted by the finite binary expansion calculated by the algorithm.

4.3 Noisy simulator

We now consider the realistic application of our procedure on present-day quantum computers, which suffer from relatively short coherence times and are vulnerable to a number of error sources. This makes practical computations extremely difficult, even in systems requiring relatively few qubits. We model

environmental effects by simulating a dephasing channel, randomly introducing a bit-flip and/or phase-flip on each qubit after each unitary operation.³¹ We also model errors in the measurement process by randomly introducing a bit-flip with low probability prior to sampling each probability distribution. Fig. 5 shows our results on a simulator emulating the error rates characteristic of IBM's *ibmq_athens* quantum computer. Qubit noise has a clearly negative impact on the quality of results. Lowest-band optimization results tend to suffer a large systematic shift, characteristic of coherent noise in a quantum computer. Additionally, while average QPE refinement often improves energy estimates, its results are now clustered with some variance around each nearby band, and on occasion (*e.g.* the third band between *M* and *Γ*) the mean optimization result is more accurate. This is symptomatic of the long circuit requirements for QPE, and supports the widely-believed notion that variational algorithms are better suited to NISQ devices.

4.4 Calibrated simulator

Although automated error correction procedures, based on redundant qubit registers and applied during calculation, are the most promising path toward practical quantum computation, several classical post-processing methods have already

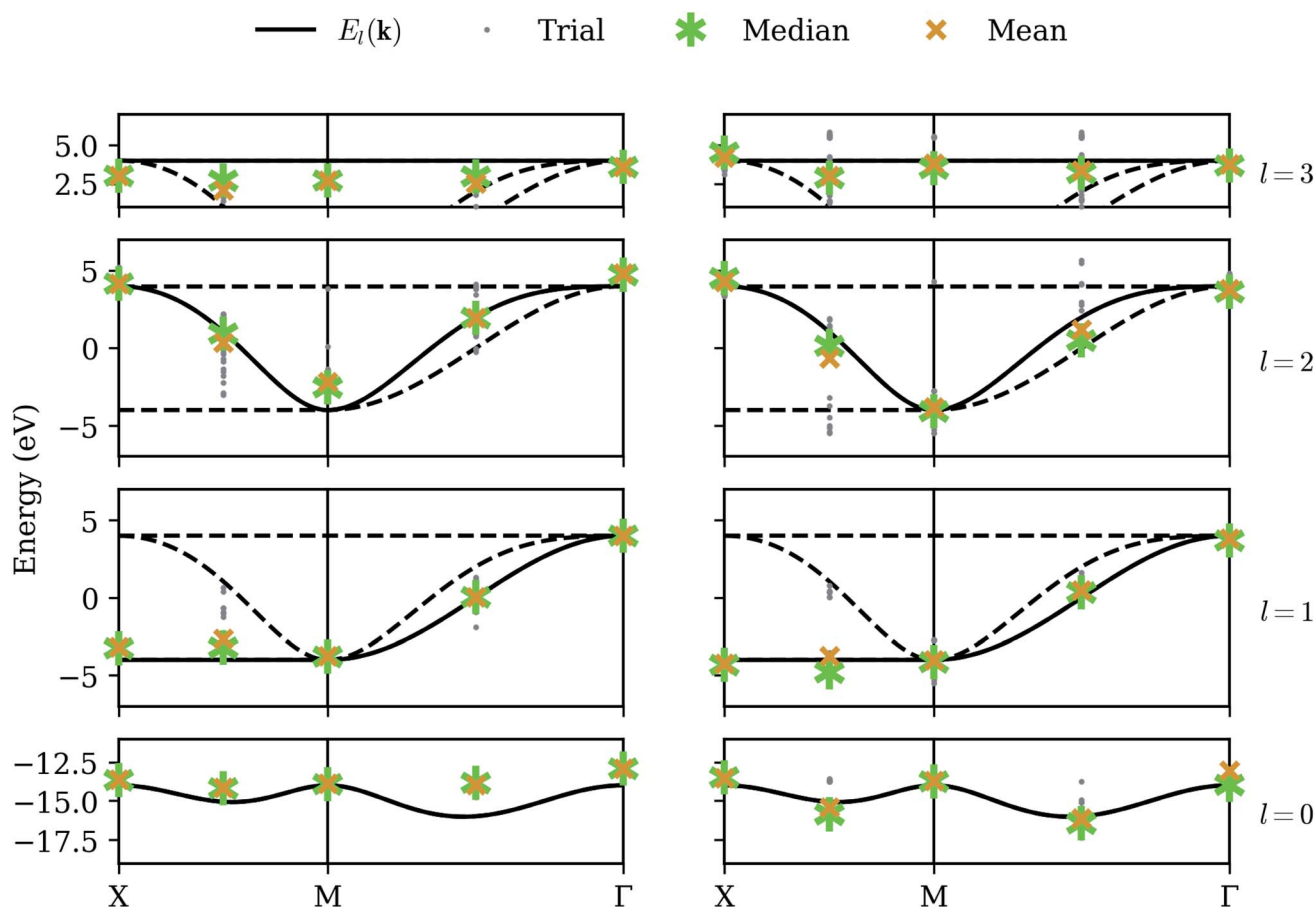


Fig. 5 Noisy simulator – our method applied while simulating low-fidelity qubits, without calibration. The left column shows raw optimization results; the right column shows the energy obtained by QPE refinement. Gray dots denote the results from each of 32 trials, with each band given its own row. The asterisk and X denote the median and mean, respectively.



proven successful in mitigating error. Errors modeled by bit-flips in the measurement process ("readout error") can be mitigated in part or entirely, at the cost of additional calibration circuits.³² Errors modeled by a dephasing channel as each unitary operation is applied ("gate error") tend to result in systematic distortions of the energy surface, as we have seen in the optimization results of Fig. 5. These distortions can be mitigated by applying Zero-Noise Extrapolation (ZNE)¹⁶ in which the same measurements are repeated several times, each time modifying the circuit to incur *more* noise. These measurements can then be extrapolated to a hypothetical circuit with zero noise, using Richardson extrapolation or a similar method.

Fig. 6 shows our results on a noisy simulator, applying readout calibration and ZNE for each energy evaluation during the optimization. ZNE offers noticeable improvement in the highest and lowest bands (calculated independently), but appears less impactful on the intermediate bands (calculated after deflation), perhaps even *increasing* variance in the third band. This may be explained by noting that ZNE is designed to

assuage *systematic* error, and this is what we tend to observe when we can rely on the variational principle, where energies cannot in principle be measured below the ground-state energy. This is not *always* true because our energy estimates are linear combinations of stochastically evaluated Pauli expectation values, and on occasion we do observe trials which appear above the highest band, but these points are relatively rare, and the average values on the highest and lowest bands are shifted inwards. However, the deflation circuits \hat{V}_0 are somewhat different for each trial, depending on exactly what eigenstate was selected for the lowest band, and this, coupled with the coherent qubit error, has the effect of inducing a *random* noise on the intermediate bands. This explains why the intermediate bands seem to suffer a larger variance but reasonable average values. We note that many other error mitigation techniques besides readout calibration and ZNE have been proposed in the literature, and our results can likely be improved greatly by implementing more of them. Nevertheless, the best solution to combat random error remains averaging over more and more trials.

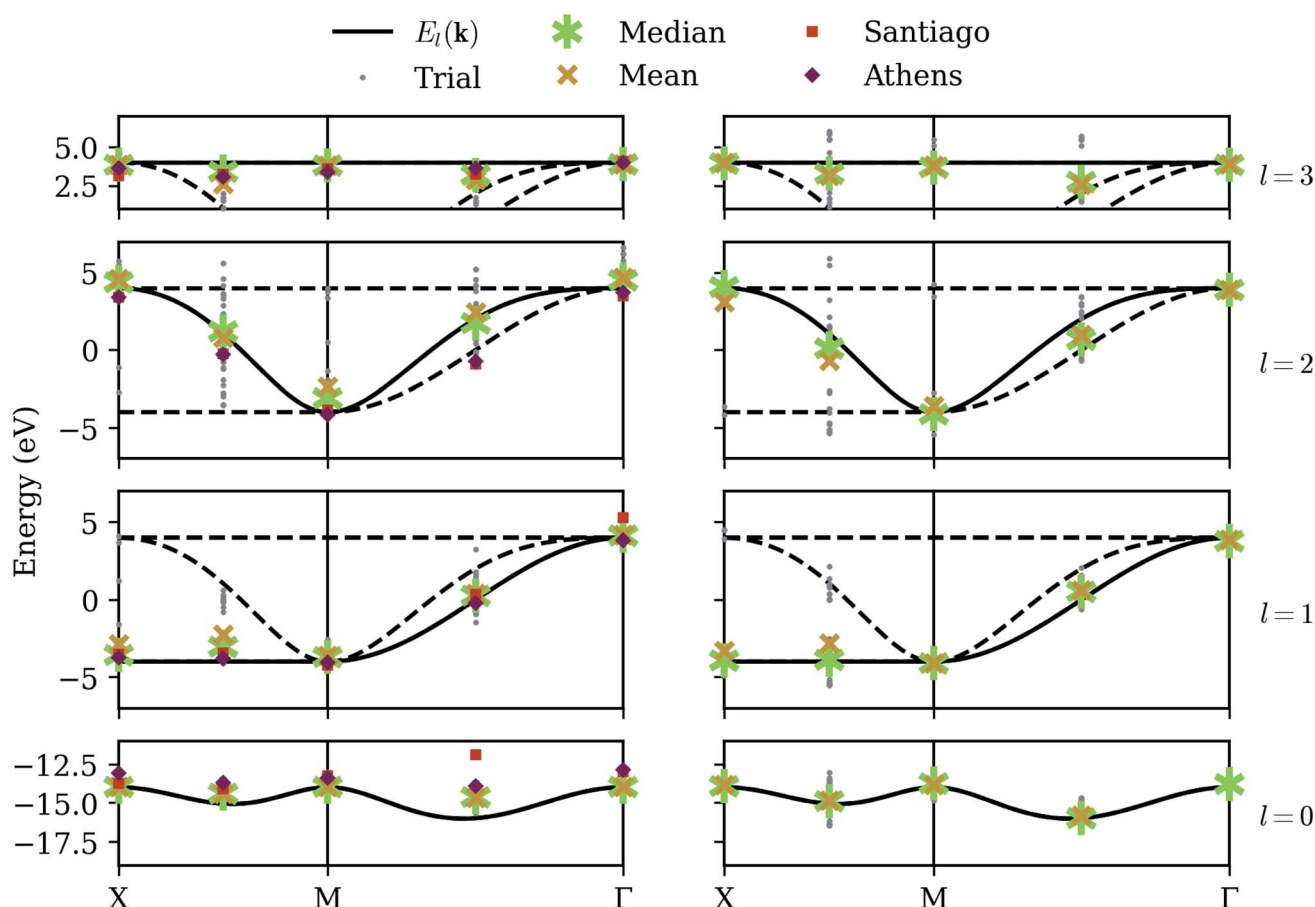


Fig. 6 Calibrated simulator – our method applied while simulating low-fidelity qubits, along with rudimentary calibration. The left column shows raw optimization results; the right column shows the energy obtained by QPE refinement. Gray dots denote the results from each of 32 trials, with each band given its own row. The asterisk and X denote the median and mean, respectively. The squares and diamonds on the left denote the energies measured on quantum devices ibmq_santiago and ibmq_athens respectively, using the least-error optimization results obtained with the calibrated simulation data. Device architecture constrains the length of quantum circuits, and QPE results for either device could not be obtained.

In addition to statistics from a calibrated simulator, Fig. 6 also shows data from the IBM devices `ibmq_athens` and `ibmq_santiago`. These are calibrated energy measurements of the eigenstates given by the least-error optimization runs on the (calibrated) noisy simulator. Results are generally consistent with the simulator, but our error mitigation is evidently even less effective on the real devices, and we note that `ibmq_santiago` is especially ill-behaved at certain points. This may be due to less effective thermal isolation from its environment at the hardware level. Finally, implementing the controlled-unitary operations necessary for the QPE procedure on a linear architecture introduces an overwhelming amount of overhead in the form of additional SWAP gates, making the QPE refinement part of our algorithm completely intractable on these devices.

5 Discussion

We have presented an application of VQD to calculate the band structure of a periodic system. This algorithm is hypothetically successful in producing accurate results on a device with low noise and is functional to a limited extent on current NISQ devices. In this section, we carefully analyze the complexity of the algorithm. The classical approach to band structure includes up to eqn (7), at which point the calculated values $H_{\alpha\beta}(\mathbf{k})$ are arranged into a Hermitian matrix. This matrix can be diagonalized using row-reduction or a similar technique in $\Theta(M^3)$ steps, where M is the number of atomic orbitals per unit cell. This is the standard against which we must compare our quantum algorithm.

Quantum resources are employed in the VQD phase of our algorithm during the operator estimation procedure, for every evaluation of the energy $E \equiv \langle \hat{H}_{\mathbf{k}} \rangle$. Each application of the ansatz from Fig. 1 requires M qubits, $\Theta(M)$ entangling gates, and has a depth between $\Theta(\log M)$ and $\Theta(M)$ layers, depending on qubit architecture. The Hamiltonian in eqn (12) has $\Theta(M)$ commuting groups, even including additional terms from the deflation procedure (eqn (16)). Our implementation requires an ensemble size of $O(\varepsilon^{-2})$ for each commuting group in \hat{H} to obtain an expectation value accurate within ε , but since ε does not scale with M , we omit it in the present analysis. The ensemble states may be prepared sequentially, for a worst-case (linear architecture) execution time on the order of $\Theta(M^2)$. Alternatively, the ensemble states may be prepared in parallel, decreasing execution time at the cost of additional qubits. In the best case, implementing “classical shadowing”²¹ reduces the number of required measurements to $\Theta(\log M)$, and a fully-connected architecture permits a circuit depth as low as $\Theta(\log M)$, bringing our algorithm into a sub-polynomial quantum resource requirement. However, the operator estimation procedure is still bounded by the number of Pauli words $\Theta(M^2)$ when measurement results are assembled into the energy $E(\theta) = \sum_i a_i \langle \Psi(\theta) | \hat{P}_i | \Psi(\theta) \rangle$.

Operator estimation is repeated for each function evaluation in the optimization procedure. The number of function evaluations required depends on the optimization routine selected and the shape of the energy surface, so it is difficult to estimate. Quantum variational algorithms are notoriously vulnerable to

so-called barren plateaus, regions in the parameter surface with a vanishing gradient expected to result in an exponential number of function evaluations for successful convergence.^{8,27} However, research into barren plateaus has focused mainly on densely-packed ansatz which alternate between parameterized rotation gates and entanglement among *each qubit*. The ansatz we have presented has a more constrained structure which rotates and entangles only two qubits at a time. Additionally, Cerezo *et al.*⁸ found that local cost functions such as the Hamiltonian we have employed are much more resistant to barren plateaus compared to global cost functions. Therefore, we conjecture that the number of function evaluations required in our algorithm may be expected to scale polynomially with the number of ansatz parameters, in our case $\Theta(M)$, provided sufficient error mitigation to suppress noise-induced barren plateaus, which are independent of the ansatz.³³ Thus, we include a factor of $\Theta(M^c)$, where $c \geq 1$ depends on the optimization. The optimization is repeated for each of M energy levels; therefore, the VQD phase of our algorithm has a total run-time on the order of $\Theta(M^{3+c})$.

An optional QPE phase may be implemented to estimate the eigenvalue to an arbitrary binary precision t .² The implementation of QPE we have used requires $M + 1$ qubits and $\Omega(t)$ rounds of measurement (see Dobšiček *et al.*² for a tighter bound). Each round applies a quantum circuit approximating a unitary operator $\hat{U}_j = \exp(i\hat{H}\tau_j)$. Each time slice in the Suzuki–Trotter expansion of \hat{U}_j on a linear architecture requires an entangling gate count of $\Theta(M^3)$ and a circuit depth of $\Theta(M^2)$.³⁴ QPE is repeated for each of M energy levels, setting the best case run-time of the QPE phase of our algorithm on the order of $\Theta(M^3)$. Note also that the simulation time τ_j scales exponentially with the accuracy of the phase estimation procedure, and the number of time-slices must scale accordingly to maintain an accurate \hat{U}_j . Thus, QPE tends to incur too much overhead for practical application on present-day NISQ devices.

Altogether, evaluating the band energies for each momentum \mathbf{k} requires $\Omega(M^3)$ time steps, comparable to the classical approach. Even with a “perfect” optimizer in which the optimal parameters θ_l are produced instantly ($c = 0$), the complexities of operator estimation and QPE alone exhibit the same scaling as classical diagonalization and incur significantly greater overhead from the finite accuracy ε . While in this form band structure calculations are not a strong candidate for quantum advantage, quantum computers *are* expected to provide a superior edge when including electron correlation terms such as $t_{\alpha\beta\gamma\delta} c_{\alpha}^{\dagger} c_{\beta}^{\dagger} c_{\gamma} c_{\delta}$ in the Hamiltonian, which introduce factors of exponential complexity in the classical approach. However, such terms also appear to force us to abandon several simplifications we have made. First, transforming into reciprocal space no longer enables H to be separated into subsystems of size M , meaning many more qubits are required to accurately simulate a periodic system. Second, considering multiple electrons forces us to adopt a qubit mapping which enforces fermionic antisymmetry, greatly increasing the number of commuting groups in the Hamiltonian. Finally, our ansatz dimension, entangling gates, and circuit depth can no longer remain linear in the number of qubits while simultaneously



remaining robust. Our hope is that this work will inspire similar simplifications to those we have made here, while remaining applicable to highly-correlated systems.

6 Conclusion

In this work, we have presented a systematic algorithm for evaluating band structures on a quantum computer. We have demonstrated the viability of implementing this algorithm in noiseless qubit systems, and we have demonstrated several of the difficulties faced when implementing it on present-day NISQ devices. Given the analogy to the classical band structure problem, our algorithm evidently generalizes to solving the eigenvalues of any Hermitian matrix. Finally we have demonstrated how state-of-the-art quantum algorithms can be applied with drastically lower resource requirements to correlation-free crystalline systems and motivated similar approaches for highly-correlated systems less accessible to classical computing.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

We thank Oliviero Andreussi, Itay Hen, Rosa Di Felice, Marco Fornari, Ilaria Siloi, Virginia Carnevali, and Anooja Jayaraj for useful discussions. We acknowledge support from the U.S. Department of Energy through the grant Q4Q: Quantum Computation for Quantum Prediction of Materials and Molecular Properties (DE-SC0019432). We are also grateful to IBM for providing quantum computing hardware and software.

References

- 1 A. Aspuru-Guzik, A. D. Dutoi, P. J. Love and M. Head-Gordon, *Science*, 2005, **309**(5741), 1704–1707.
- 2 M. Dobšiček, G. Johansson, V. Shumeiko and G. Wendin, *Phys. Rev. A*, 2007, **76**(3), 030306.
- 3 A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik and J. L. O'Brien, *Nat. Commun.*, 2014, **5**, 4213.
- 4 J. R. McClean, J. Romero, R. Babbush and A. Aspuru-Guzik, *New J. Phys.*, 2016, **18**, 023.
- 5 A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow and J. M. Gambetta, *Nature*, 2017, **549**, 242–246.
- 6 M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio and P. J. Coles, *Nat. Rev. Phys.*, 2021, **3**, 625–644.
- 7 O. Higgott, D. Wang and S. Brierley, *Quantum*, 2019, **3**, 156.
- 8 M. Cerezo, A. Sone, T. Volkoff, L. Cincio and P. J. Coles, *Nat. Commun.*, 2021, **12**, 1791.
- 9 D. Z. Manrique, I. T. Khan, K. Yamamoto, V. Wichitwechkarn and D. M. Ramo, arXiv: 2008.08694, 2020.
- 10 R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven and G. K.-L. Chan, *Phys. Rev. X*, 2018, **8**(1), 011044.
- 11 G. Grosso and G. Parravicini, *Solid State Physics*, Elsevier Science, 2000.
- 12 Google AI Quantum, *Science*, 2020, **369**, 1084–1089.
- 13 M. Buongiorno Nardelli, F. T. Cerasoli, M. Costa, S. Curtarolo, R. D. Gennaro, M. Fornari, L. Liyanage, A. R. Supka and H. Wang, *Comput. Mater. Sci.*, 2018, **143**, 462–472.
- 14 F. T. Cerasoli, K. Sherbert, J. Słostokawińska and M. Buongiorno Nardelli, *Phys. Chem. Chem. Phys.*, 2020, **22**(38), 21816–21822.
- 15 B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou and E. Barnes, *Npj Quantum Inf.*, 2020, **6**, 10.
- 16 T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari and W. J. Zeng, *2020 IEEE International Conference on Quantum Computing and Engineering*, QCE, 2020.
- 17 N. Hatano and M. Suzuki, in *Quantum Annealing and Other Optimization Methods*, ed. A. Das and B. K. Chakrabarti, Springer Berlin Heidelberg, 2005, pp. 37–68.
- 18 D. W. Berry, A. M. Childs, R. Cleve, R. Kothari and R. D. Somma, *Phys. Rev. Lett.*, 2015, **114**(9), 090502.
- 19 A. Kalev and I. Hen, *Quantum*, 2021, **5**, 156.
- 20 N. C. Rubin, R. Babbush and J. McClean, *New J. Phys.*, 2018, **20**, 053020.
- 21 H.-Y. Huang, R. Kueng and J. Preskill, *Nat. Phys.*, 2020, **16**, 1050–1057.
- 22 J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. J. Love and A. Aspuru-Guzik, *Quantum Sci. Technol.*, 2018, **4**, 014008.
- 23 I. G. Ryabinkin, T.-C. Yen, S. N. Genin and A. F. Izmaylov, *J. Chem. Theory Comput.*, 2018, **14**, 6317–6326.
- 24 H. R. Grimsley, S. E. Economou, E. Barnes and N. J. Mayhall, *Nat. Commun.*, 2019, **10**, 3007.
- 25 T. C. Yen, R. A. Lang and A. F. Izmaylov, *J. Chem. Phys.*, 2019, **151**, 164111.
- 26 J. T. Seeley, M. J. Richard and P. J. Love, *J. Chem. Phys.*, 2012, **137**, 224109.
- 27 J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush and H. Neven, *Nat. Commun.*, 2018, **9**, 4812.
- 28 W. Lavrijsen, A. Tudor, J. Muller, C. Iancu and W. De Jong, *Proceedings – IEEE International Conference on Quantum Computing and Engineering*, QCE, 2020, vol. 2020, pp. 267–277.
- 29 B. I. Min, J. H. Shim, M. S. Park, K. Kim, S. K. Kwon and S. J. Youn, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2006, **73**, 3–6.
- 30 A. Silva and J. van Wezel, *SciPost Phys.*, 2018, **4**(6), 28.
- 31 M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, Cambridge University Press, USA, 10th edn, 2011.
- 32 P. J. Karalekas, N. A. Tezak, E. C. Peterson, C. A. Ryan, M. P. da Silva and R. S. Smith, *Quantum Sci. Technol.*, 2020, **5**, 024003.
- 33 S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio and P. J. Coles, *Nat. Commun.*, 2021, **12**, 6961.
- 34 I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan and R. Babbush, *Phys. Rev. Lett.*, 2018, **120**(11), 110501.
- 35 T. C. Yen, V. Verteletskyi and A. F. Izmaylov, *J. Chem. Theory Comput.*, 2020, **16**(4), 2400–2409.
- 36 C. Cade, L. Mineh, A. Montanaro and S. Stanisic, *Phys. Rev. B*, 2020, **102**(23), 235122.

