

Efficient implementation of single particle Hamiltonians in exponentially reduced qubit space

Martin Plesch^{1,2}, Martin Friák^{3,4}, and Ijaz Ahamed Mohammad¹

¹Institute for Physics, Slovak Academy of Sciences, Bratislava, Slovakia

²Matej Bel University, Banská Bystrica, Slovakia

³Institute of Physics of Materials, v.v.i., Czech Academy of Sciences, Žitkova 22, CZ-616 00, Brno, Czech Republic

⁴Department of Condensed Matter Physics, Faculty of Science, Masaryk University, Kotlářská 2, CZ-611 37, Brno, Czech Republic

Current and near-term quantum hardware is constrained by limited qubit counts, circuit depth, and the high cost of repeated measurements. We address these challenges for solid-state Hamiltonians by introducing a logarithmic-qubit encoding that maps a system with N physical sites onto only $\lceil \log_2 N \rceil$ qubits while maintaining a clear correspondence with the underlying physical model. Within this reduced register, we construct a compatible variational circuit and a Gray-code-inspired measurement strategy whose number of global settings grows only logarithmically with system size. To quantify the overall hardware load, we introduce a volumetric efficiency metric that combines the number of qubits, circuit depth, and the number of measurement settings into a single measure, expressing the overall computation costs. Using this metric, we show that the total space-time sampling volume required in a variational loop can be reduced dramatically from N^2 to $(\log N)^3$ for a hardware-efficient ansatz, allowing an exponential reduction in time and size of the quantum hardware. These results demonstrate that large, structured solid-state Hamiltonians can be simulated on substantially smaller quantum registers with controlled sampling overhead and manageable circuit complexity, extending the reach of variational quantum algorithms on near-term devices.

1 Introduction

Determining the energy spectrum of molecules and condensed-matter systems lies at the heart of quantum chemistry and many-body physics. The problem can be formally expressed as finding the eigenvalues and eigenstates of a high-dimensional Hamiltonian, whose complexity grows exponentially with the number of interacting particles and orbitals. As a result, obtaining accurate electronic energies and wavefunctions remains one of the major computational bottlenecks in modern quantum science [1, 2].

Classical computational approaches, such as Full Configuration Interaction (FCI) [3], Coupled-Cluster (CC) theory [4], and Density Functional Theory (DFT) [5], have achieved remarkable success in describing a wide range of chemical and physical systems. However,

Ijaz Ahamed Mohammad: fyzijaz@savba.sk

their computational cost grows rapidly with system size. Even the widely used CCSD (Coupled-Cluster Singles and Doubles) method, which captures most of the electron correlation by truncating excitations to the single- and double-excitation manifold, scales as $\mathcal{O}(N^6)$ in the number of orbitals [4]. For larger or strongly correlated systems, such methods become intractable, necessitating approximate or model-specific simplifications.

Classically, eigenvalues and eigenstates are typically computed with sparse tools from linear algebra (e.g. the Lanczos algorithm, also known as the recursion method [6]). This is particularly relevant for semiconductors and nanostructures such as quantum dots which are often well described in a single-particle model [7–9]. In addition to scaling challenges, classical iterative eigensolvers such as the Lanczos algorithm suffer from numerical instability. In finite precision, the orthogonality of the Krylov subspace vectors gradually degrades, resulting in spurious or unstable eigenvalues [10]. Reorthogonalization techniques—such as Gram–Schmidt or block-based methods—can restore numerical stability but at the expense of significant memory overhead and increased runtime, especially at high spectral resolution [11, 12]. This trade-off between precision and scalability fundamentally limits classical simulations of large Hamiltonians.

These limitations motivate the exploration of quantum algorithms as a novel and promising approach for computing molecular and material energy spectra. Quantum hardware provides a natural platform to represent and manipulate exponentially large Hilbert spaces with polynomial resources. In particular, Variational Quantum Algorithms (VQAs) [13] have emerged as a promising class of methods that combine parameterized quantum circuits with classical optimization. Among them, the Variational Quantum Eigensolver (VQE) [14–19] stands out as a practical candidate for near-term devices: it minimizes the expectation value of a Hamiltonian to approximate ground state using shallow ansätze. For higher-energy states, extensions such as Variational Quantum Deflation (VQD) [20] further refine the energy spectrum by introducing orthogonality penalties against previously found eigenstates.

Similar to classical truncation schemes such as CCSD, one can restrict the quantum Hilbert space to a physically motivated subset of configurations. In this work, we focus on the single-particle Hamiltonian, which describes a single particle delocalized across N orbitals or lattice sites. This model naturally maps onto a tight-binding (TB) Hamiltonian of dimension N , which models electrons in a crystal using a localized orbital basis, with diagonal terms corresponding to on-site energies and off-diagonal terms representing hopping amplitudes between sites. TB is used to model band structures with a small set of physically informed parameters. Parameters can be obtained empirically by fitting band edges and effective masses (empirical TB) [21–23] or derived from first-principles frameworks such as density functional theory [24]. Despite its apparent simplicity, this approximation retains essential physical structure and provides a meaningful testbed for studying variational quantum algorithms under realistic hardware constraints.

While single-particle and tight-binding Hamiltonians do not by themselves constitute classically hard problems, they offer a useful benchmark for developing efficient quantum protocols. We view the present work as meaningful progress toward that broader objective, developing tools and procedures to be used on more complex systems, when quantum devices will become available in such a scale.

This single-particle approach for VQD was formulated in [25], and subsequent work by the present authors [26] significantly reduced the measurement overhead from an explicit $\mathcal{O}(N)$ scaling to a constant number of *three* global measurement settings, independent of the system size. In the present work, we take this approach a step further by showing how to realize the same class of Hamiltonians on a logarithmic number of qubits, thereby reducing

the required quantum resources even further. Although Hamiltonians of dimension N can, in principle, be represented using only $\log_2 N$ qubits, practical and resource-efficient implementations have received relatively little attention in the literature [25, 27, 28]. Here, we close this gap in knowledge by providing a logarithmic-register variational approach that extends the reach of single-particle simulations to substantially larger system sizes on limited hardware.

To compare the efficiency of hybrid quantum–classical algorithms, we introduce a volumetric measure that captures the actual resource demand of such procedures. In classical computation, two basic parameters are considered: memory demand (number of bits) and time (number of operations). These parameters are mostly independent, and each independently limits the feasibility of executing the algorithm: the algorithm must both fit into the computer’s memory and complete within a bearable time.

In the quantum world, the situation is different. While the size of the device (number of qubits) is likewise a limiting parameter, the notion of time is more subtle. First, the output of a quantum computer is quantum, and converting it into a classical, human-readable form requires measurements. The number of incompatible measurement bases affects the efficiency of the algorithm, since more bases require more independent runs. Moreover, as most quantum devices can, unlike classical ones, perform many gates in parallel, the depth of the circuit (number of layers) is more relevant than the total number of operations. Finally, because the outcomes of quantum measurements are stochastic, many repetitions are needed; on large devices, these repetitions can be parallelized, speeding up the process. Based on these differences, we define a measure given by the product of the device size, circuit depth, and number of incompatible measurement settings, and use it to highlight the performance of our approach.

The paper is organized as follows. In Section II, we outline the theoretical framework of the Variational Quantum Eigensolver (VQE) and introduce the specific Single Excitation Subspace (SES) Hamiltonians. Section III discusses optimal state preparation strategies on a logarithmic number of qubits for such Hamiltonians, while Section IV deals with measurement settings, highlighting how its structure enables a significant reduction in its number. In Section V we show the exponential speedup of the presented algorithm in comparison with the previously used approaches and the Section VI concludes our findings.

2 Background

2.1 Variational Quantum Eigensolver

The Variational Quantum Eigensolver (VQE) [14] is a hybrid quantum-classical algorithm designed to estimate the ground-state energy of a given Hamiltonian. Its relatively low circuit depth and reliance on classical optimization make it particularly suitable for near-term quantum devices.

VQE functions by preparing a parameterized quantum state $|\psi(\boldsymbol{\theta})\rangle$ using a variational ansatz, typically constructed from layers of parameterized single-qubit rotations and entangling gates. The expectation value of the Hamiltonian H with respect to this state, $\langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle$, is estimated through quantum measurements and serves as the cost function for optimization.

A classical optimizer is employed to minimize this cost by iteratively updating the parameters $\boldsymbol{\theta}$, with the aim of converging to the ground state of the system. This hybrid loop continues until a convergence criterion is met, ideally resulting in a quantum state that closely approximates the true ground state of the target system.

In recent years, VQE has become a standard benchmark for variational quantum algorithms across various domains, including quantum chemistry, condensed matter physics, and combinatorial optimization [13]. Nonetheless, its performance is strongly influenced by the choice of ansatz, the effectiveness of the classical optimization strategy [29], and the presence of quantum noise. A very important aspect is also the implementation of the physical Hamiltonian, addressed in the paper.

2.2 Solid states SES Hamiltonians

The Single Excitation Subspace (SES), equivalent to single-particle approximation in tight-binding method, refers to a constrained region of the full Hilbert space in which only one excitation is present across all qubits or quantum modes at any given time. Mathematically, the SES consists of all quantum states that are linear combinations of basis states with a single excitation, such as $|100 \dots 0\rangle$, $|010 \dots 0\rangle$, ..., $|000 \dots 1\rangle$. This subspace is particularly relevant for simulating physical systems where the dynamics are dominated by a single particle or quasi-particle, such as an exciton, electron, or spin flip.

The SES approximation is powerful because it drastically reduces the dimensionality of the problem—from an exponentially large space of 2^n states for n qubits to only n states—making both classical and quantum simulations more tractable. This makes SES particularly useful for near-term quantum hardware, which is limited in qubit number and coherence time.

Although the SES describes a linear system, it remains relevant for electronic structure problems. In particular, the single-excitation subspace captures the physics of one-electron models, such as the calculation of molecular orbitals in the linear combination of atomic orbitals (LCAO) framework. The SES forms the basis for computing one-electron energy levels and can serve as a starting point for more advanced approaches. Even though the SES does not include electron–electron interactions, it provides a computationally efficient approximation to simulating the electronic properties of materials.

2.2.1 SES Ansatz

In an N -qubit system restricted to the Single Excitation Subspace (SES), the variational state must encompass all basis vectors $\{|e_j\rangle\}_{j=1}^N$ of Hamming weight one. A compact strategy to generate such states begins with a localized excitation, e.g. $|e_1\rangle = X_1|0\rangle^{\otimes N}$, and subsequently propagates it through the register by means of two-qubit entangling gates.

We use family of gates, $\hat{A}_{j,j+1}(\beta_j, \gamma_j)$, that act on neighboring qubits and mix the amplitudes of $|e_j\rangle$ and $|e_{j+1}\rangle$ [25]. By applying these gates sequentially across the chain, one constructs the trial state

$$|\psi(\boldsymbol{\theta})\rangle = \left(\prod_{j=1}^{N-1} \hat{A}_{j,j+1}(\beta_j, \gamma_j) \right) X_0|0\rangle^{\otimes N}. \quad (1)$$

After all $N - 1$ gates have been applied, the register occupies a general superposition of single-excitation states,

$$|\psi\rangle = \sum_{j=1}^N \alpha_j |e_j\rangle, \quad (2)$$

with coefficients $\alpha_j = |\alpha_j|e^{i\theta_j}$ capturing both amplitude ($|\alpha_j|$) and phase (θ_j) of the excitation at site j .




Figure 1: Circuit representation of the SES ansatz. The localized excitation is propagated across the register through a sequence of two-qubit entangling gates $\hat{A}_{j,j+1}$ acting on neighboring qubits.




Figure 2: Implementation of the $\hat{A}_{j,j+1}$ gate used in the SES ansatz. Each block is expressed in terms of three CNOTs and a sequence of parameterized single-qubit rotations R_y and R_z .

The Fig. 1 illustrates the sequential layout of the entangling gates, which effectively “sweep” the excitation across the system. Each gate $\hat{A}_{j,j+1}$ admits a standard decomposition into elementary operations: three two-qubit controlled gates (e.g. CNOTs), two R_z rotations, and two R_y rotations, as shown in Fig. 2.

2.2.2 Hamiltonian

In the single-excitation subspace (SES), the system evolves within the sector of the Hilbert space that contains exactly one excitation. Each computational basis state in this subspace corresponds to a single qubit being in the excited state $|1\rangle$, with all other qubits remaining in the ground state $|0\rangle$. The general SES Hamiltonian, which is mathematically equivalent to the tight-binding Hamiltonian commonly used in solid-state physics, includes both on-site energy terms h_{jj} and coherent hopping terms h_{jk} between different sites. It can be written as:

$$H = \sum_{k=1}^N h_{kk} |k\rangle \langle k| + \sum_{j \neq k} h_{jk} |j\rangle \langle k|, \quad (3)$$

where $|j\rangle$ denotes the state with the excitation localized at site j .

To express this Hamiltonian in terms of Pauli operators acting on the full qubit space, we map the diagonal terms using the projector $|1\rangle_j \langle 1| = \frac{1-Z_j}{2}$ and the off-diagonal hopping terms using the ladder operators: $|j\rangle \langle k| = \sigma_j^+ \sigma_k^-$, where $\sigma_j^\pm = \frac{1}{2}(X_j \mp iY_j)$. Decomposing

these into Pauli matrices leads to the full qubit Hamiltonian [25]:

$$\begin{aligned}
H = & \sum_{k=1}^N \frac{h_{kk}}{2} (1 - Z_k) \\
& + \sum_{j < k} \left(\frac{\text{Re}(h_{jk})}{2} (X_j X_k + Y_j Y_k) \right. \\
& \quad \left. + \frac{\text{Im}(h_{jk})}{2} (Y_j X_k - X_j Y_k) \right). \tag{4}
\end{aligned}$$

This form is Hermitian and preserves the excitation number, ensuring that time evolution remains within the SES. It also provides a convenient basis for implementing SES dynamics on quantum hardware using variational approaches.

3 State preparation

While a constant number of measurement settings as previously shown in our [26] constitutes a very compelling reduction on the number of runs of the quantum device needed, the algorithm still relies on a large number of qubits N and a long series of $O(N)$ subsequent gates. Here we present a different approach: we exponentially reduce the number of qubits in the system and then carefully select the appropriate measurement bases to achieve additional savings in both qubit count and measurement resources. Such logarithmic reduction effectively compresses the original Hilbert space into a smaller, yet information-preserving subspace, ensuring that the essential dynamical and spectral properties of the Hamiltonian remain intact. When combined with optimized measurement strategies, this approach allows for an even more resource-efficient estimation of expectation values. Consequently, it becomes possible to simulate exponentially larger systems using only a modest number of physical qubits, substantially extending the reach of variational quantum algorithms on near-term devices.

Using a standard binary encoding, each of the N SES basis states is mapped to an n -qubit register with $n = \lceil \log_2 N \rceil$. The site label $k \in \{0, \dots, N-1\}$ is written in binary as a bit string $b(k) = (b_1, \dots, b_n)$, and the corresponding register state is $|b_1 \dots b_n\rangle = |\text{binary}(k)\rangle$.

This encoding strategy is most natural when the number of qubits matches an exact power of two, i.e., when $N = 2^n$. In this setting, each computational basis state of the n -qubit register corresponds directly to one of the N sites of the system, and no additional restructuring is required. When the system size lies between two consecutive powers of two, $2^{n-1} < N < 2^n$, the situation becomes a bit more complicated: the Hilbert space of the register then contains more basis states than are actually needed to represent the system. If left untreated, a variational algorithm would naturally explore this larger space, which may lead to parameter updates that leave the physical domain.

There are two possible strategies for addressing this mismatch: (i) modify the cost function by constructing a new, artificial Hamiltonian energetically punishing the newly introduced states, or (ii) restrict the variational circuit (ansatz) so that the evolution is confined to the subset of states corresponding to the true system's degrees of freedom.

3.1 Extended Hamiltonian with generic ansatz

The first approach involves enlarging the Hilbert space and extending the Hamiltonian accordingly. This is achieved by adding penalty terms that energetically suppress any population outside the physical subspace:

$$H_{\text{ex}} = H_o + \sum_{i=N+1}^{2^n} C_p \frac{1 - Z_i}{2}. \quad (5)$$

Here, H_o denotes the original Hamiltonian, and C_p is a large positive constant chosen such that any non-physical configuration incurs a prohibitively high energy cost. As a result, during variational optimization, amplitudes in the unphysical subspace are strongly suppressed, while the low-energy spectrum of the physical sector remains effectively unaltered. In practice, C_p must be selected significantly larger than the characteristic energy scales of H_o , ensuring that the optimizer has no incentive to explore forbidden regions of the Hilbert space. Once this extended Hamiltonian is constructed, one can proceed with standard reduction techniques as if the system naturally possessed 2^n valid states.

A key advantage of this method is that it allows the use of any n -qubit ansatz, such as a hardware-efficient ansatz [16], without imposing additional restrictions on the parameter space, since the penalty terms automatically exclude unphysical states from the optimization process. The price to pay is that a significantly (up almost double) Hilbert space needs to be researched, leading to an increased risk of redundant parameter symmetries, large barren plateaus, and poor gradient scaling, all of which hinder convergence of the optimization process. Even more importantly, all relevant hardware-efficient ansätze are biased in the sense that they naturally favor less complex states (as they start from a product state and use only a very limited number of entangling operations), while a large part of the Hilbert space remains inaccessible [30].

As there is no one-to-one physical correspondence between states in the reduced system and states in the original system (where all states are, in a sense, comparable), the algorithm may favor “simple” solutions over more complex ones, potentially leading to biased or incorrect results.

3.2 Binary encoded SES ansatz

The second approach avoids modifying the Hamiltonian and instead enforces the restriction at the level of the variational ansatz. The idea is to construct a parameterized quantum circuit whose support lies strictly within the physical subspace, such that amplitudes for non-physical states are identically zero throughout the optimization. Although this requires more careful circuit design, it has several attractive features. In particular, the ansatz can be engineered to preserve the same parameterization as the original ansatz defined for the unreduced Hamiltonian, which means that the numerical optimization landscape—and thus the physical results—remain consistent with those obtained from the full system.

We design a quantum circuit that operates within a logarithmically reduced Hilbert space using a combination of data and two ancillary qubits. The key idea is to transfer the amplitude encoded in a pair of ancillary qubits using the A gate (Eq. 2), to the appropriate data-register basis state, thus emulating the original single-excitation ansatz Fig. 1 without ever populating non-physical subspace.

To effectively reproduce the same state as generated by the original ansatz, we employ two ancillary qubits for the application of the A gate. Conditioned on one of these ancillary qubits, which encodes the amplitude information, we prepare the single-excitation subspace state in the data register. After this step, the ancillary qubit is reset, and the procedure

is repeated: the A gate is applied again to regenerate the amplitude information, which is then transferred to the data register through a sequence of controlled operations. This can be summarized into the following steps:

1. *Preparation of $|1\rangle$* : Only once, at the beginning of the procedure, the first qubit is initialized.
2. *Application of A gate*: The two ancillary qubits are updated via an A gate, which performs a controlled two-qubit rotation that mixes the ancillary states.
3. *State preparation on data register*: Controlled on the first ancilla qubit, the appropriate computational basis state $|i\rangle$ is prepared on the (logarithmically reduced) data register using binary encoding.
4. *Unflagging ancilla qubit*: A multi-controlled Toffoli gate is used to reset the ancilla back to $|0\rangle$, conditioned on the data register being in state $|i\rangle$.

The points 2 to 4 are repeated for each $i < N$, sequentially building up a coherent superposition on the data register. This procedure is illustrated in the Fig. 3. This construction guarantees that only valid states from the single-excitation subspace are ever populated.

To efficiently implement this ansatz we use shifted binary encoding. Concretely, each site label $k \in \{0, \dots, N-1\}$ is stored as the binary string of $k+1 \pmod{2^n}$, i.e. $|k\rangle \mapsto |b_1 b_2 \dots b_n\rangle = |\text{binary}(k+1)\rangle$. This choice makes the initialized data register $|0 \dots 0\rangle$ distinct from the first prepared state ($k=0$ encodes to $|0 \dots 01\rangle$). For instance, an eight-state SES uses three qubits with the mapping $k=0 \rightarrow |001\rangle$, $k=1 \rightarrow |010\rangle, \dots, k=7 \rightarrow |000\rangle$, as summarized in Table 1.

SES state k	Basis $ \cdot\rangle$	Binary $ b_1 b_2 b_3\rangle$
0	$ 00000001\rangle$	$ 001\rangle$
1	$ 00000010\rangle$	$ 010\rangle$
2	$ 00000100\rangle$	$ 011\rangle$
3	$ 00001000\rangle$	$ 100\rangle$
4	$ 00010000\rangle$	$ 101\rangle$
5	$ 00100000\rangle$	$ 110\rangle$
6	$ 01000000\rangle$	$ 111\rangle$
7	$ 10000000\rangle$	$ 000\rangle$

Table 1: Mapping of SES basis states to their basis and 3-qubit shifted binary encodings.

The resource need of the proposed ansatz is governed by the number of modules needed to span the single-excitation subspace (SES) and by the cost of each module. Starting from N SES basis states, we encode the excitation index into $n = \lceil \log_2 N \rceil$ qubits; one module is invoked per valid label, for a total of N modules.

Step 2 applies the A gate on the two ancillary qubits, which can be decomposed into three CNOT gates as shown in Fig. 2. In Step 3, the desired computational basis state $|i\rangle$ is prepared on the data register, controlled on a single ancilla qubit. This corresponds to a single-control, multi-target X operation, decomposable on $O(n)$ (and on average $n/4$) CNOT gates. The last step, unflagging the ancilla, is the most resource-intensive. It requires implementing an n -controlled Toffoli gate conditioned on the data register. With




Figure 3: Two concatenated modules of the subspace-constrained ansatz. In each module, a two-qubit gate A acts on the flag-selector pair (a_0, a_1) , followed by a SWAP. Conditioned on the flag, the data register is prepared in $|i\rangle$ (Module 1) and $|i+1\rangle$ (Module 2), and a multi-controlled toffoli operation targets a_1 .

one additional clean ancilla, such an n -controlled X can be realized using $(2n - 3)$ Toffoli-type blocks [31]. Each of these blocks can be implemented as a standard CCX gate, costing six CNOTs [32]. Hence, the CNOT cost in this step scales linearly with n , i.e., $O(n)$. All in all, the total CNOT cost per module scales as $O(n)$. As the full ansatz consists of N modules, the overall cost scales as $O(Nn) = O(N(\log N))$.

4 Measurement settings

In our previous work [26], for a general SES ansatz we showed that the measurement protocol can be implemented in an optimal manner: all information needed to evaluate the SES Hamiltonian can be obtained using only three global measurement settings, rather than the $O(N)$. Here, we present a subtly different, yet equivalent, formulation that will be convenient for the upcoming subsection 4.2 on binary-encoded measurement protocol. For clarity, we briefly revisit the correlators involved and the organization of measurements at the level of global settings. This reformulation preserves the three-setting measurement cost and makes explicit the structure we later exploit for implementing efficient measurement protocol in the binary encoded system.

For SES Hamiltonians expressed in Eq. (4), the correlators $\langle X_j Y_k \rangle$ and $\langle Y_j X_k \rangle$ differ only by sign, and thus can also be combined into a single measurement setting. More precisely,

$$\langle X_j X_k \rangle = \langle Y_j Y_k \rangle = 2|\alpha_j||\alpha_k| \cos(\theta_k - \theta_j), \quad (6)$$

where $|\alpha_j|$, $|\alpha_k|$ and θ_j , θ_k denote the magnitudes and phases associated with the basis states $|e_j\rangle$ and $|e_k\rangle$.

Similarly,

$$\langle X_j Y_k \rangle = 2|\alpha_j||\alpha_k| \sin(\theta_k - \theta_j), \quad (7)$$

$$\langle Y_j X_k \rangle = -2|\alpha_j||\alpha_k| \sin(\theta_k - \theta_j). \quad (8)$$

Eqs. (6)–(8) show that the Hamiltonian’s expectation value can be fully determined once, for every pair of basis states $|e_j\rangle$ and $|e_k\rangle$, one knows the absolute magnitude of their overlap together with the cosine and sine of their relative phase.

Using the relations in Eqs. (6)–(8), and as described in Appendix A, one can evaluate the cost function or Hamiltonian expectation value as

$$\begin{aligned}
E = & \sum_{k=l}^N h_{kk} |\alpha_k|^2 \\
& + \sum_{j=1}^N \sum_{k>j}^N \left[(h_{jk} + h_{kj}) |\alpha_j| |\alpha_k| \cos(\theta_k - \theta_j) \right. \\
& \quad \left. + \frac{h_{jk} - h_{kj}}{2i} |\alpha_j| |\alpha_k| \sin(\theta_k - \theta_j) \right].
\end{aligned} \tag{9}$$

4.1 Measurement protocol

The extraction of amplitudes and phases from quantum measurements follows naturally from the structure of the variational state in the single-particle picture as described below.

4.1.1 Obtaining amplitudes from Z -measurements

In particular, determining the site occupations—or equivalently, the magnitudes of the expansion coefficients—is a straightforward first step, as it relies solely on measurements in the computational basis. These measurements directly reveal how the probability density of the particle is distributed across the available sites.

In the single-excitation subspace expressed in Eq. (2), the occupation probability of site j is $p_j = |\alpha_j|^2$. With computational-basis readout, identifying $|e_j\rangle$ with $Z_j = -1$ and all other qubits with $Z = +1$, the single-qubit expectation satisfies

$$|\alpha_j| = \sqrt{\frac{1 - \langle Z_j \rangle}{2}}. \tag{10}$$

Knowing all $|\alpha_j|$ fixes the magnitudes.

4.1.2 Nearest-neighbour phase differences

Once the amplitudes are known, the next step is to extract phase differences between neighbouring qubits. To this end, the qubits are arranged in a cyclic structure as shown in the Fig. 4 and nearest-neighbour two-qubit correlators are measured:

$$\langle X_j X_{j+1} \rangle = 2|\alpha_j| |\alpha_{j+1}| \cos(\theta_{j+1} - \theta_j), \tag{11}$$

$$\langle X_j Y_{j+1} \rangle = 2|\alpha_j| |\alpha_{j+1}| \sin(\theta_{j+1} - \theta_j), \tag{12}$$

where $(j, j+1)$ are neighbouring sites. The relative phase difference between basis states j and $j+1$ is then directly obtained as

$$\theta_{j+1} - \theta_j = \arctan\left(\frac{\langle X_j Y_{j+1} \rangle}{\langle X_j X_{j+1} \rangle}\right). \tag{13}$$




Figure 4: Circular arrangement of the basis states $|j\rangle, |j+1\rangle, \dots, |j+7\rangle$ used to extract phase differences. The blue edge shows the local relative phase $\delta\theta_j = \theta_{j+1} - \theta_j$, while the red arc illustrates the cumulative phase difference across multiple steps, $\theta_{j+4} - \theta_{j+2} = \sum_{r=2}^3 (\theta_{j+r+1} - \theta_{j+r})$.

4.1.3 Phase differences between non-neighbouring qubits

For qubits that are not directly connected, relative phase differences can be reconstructed by summing over paths through the cyclic graph as shown in the Fig. 4. For example, if j and k are separated by intermediate qubits $j+1, j+2, \dots, k-1$, then

$$\theta_k - \theta_j = \sum_{m=j}^{k-1} (\theta_{m+1} - \theta_m), \quad (14)$$

where each nearest-neighbour difference $\theta_{m+1} - \theta_m$ is obtained from Eq. (13).

Global measurement settings :

In summary, the full reconstruction requires only three global measurement settings to obtain the amplitudes $|\alpha_i|$ and all relative phases $\theta_j - \theta_k$ for arbitrary j, k :

- $\mathcal{M}_Z$: One round of measurements in the computational basis, yielding

$$|\alpha_j|, |\alpha_j|^2, \quad \text{for all } j = 1, \dots, N.$$

- $\mathcal{M}_{XX}$: One round of measurements in the X basis, yielding nearest-neighbour correlations

$$\langle X_j X_{j+1} \rangle, \quad \text{for neighbouring pairs } (j, j+1).$$

- $\mathcal{M}_{XY}$: One round of measurements in an alternating pattern of X and Y , e.g. $(X, Y, X, Y, \dots)$, yielding nearest-neighbour mixed correlators

$$\langle X_j Y_{j+1} \rangle, \quad \text{for neighbouring pairs } (j, j+1).$$

From these three settings, all nearest-neighbour phase differences are obtained via Eq. (13), and phase differences between arbitrary qubit pairs follow by summing along paths in the cycle. Thus, amplitudes and relative phases of the entire state in the single-excitation subspace can be fully reconstructed using only three measurement settings.

4.1.4 Handling vanishing amplitudes

The difficulty with this protocol arises when Eq. (13) fails, i.e., when one of the amplitudes $|\alpha_j|$ or $|\alpha_k|$ vanishes for a given pair (j, k) . In such cases, both the numerator and denominator of the ratio become zero, leaving the angle undefined and thereby obstructing the calculation of phase differences along paths involving that pair. To address this, one can simply omit states with vanishing amplitudes ($|\alpha_j| \approx 0$) and reconstruct the cyclic graph using only non-vanishing states. This is reasonable, as states with negligible amplitudes contribute insignificantly to the expectation value of the Hamiltonian.

Any qubit with close to zero amplitude simply forms a disconnected arm of the circle and can be ignored, since such qubits contribute nothing to the total energy and do not interfere with the determination of relative phases among the remaining qubits.

The measurement strategy introduced here differs from the one in our earlier work [26], where the phase differences were extracted in an alternative, but equivalent, way. There, the relative phase between basis states was encoded in the complex two-qubit correlator

$$C_{jk} = \langle X_j X_k \rangle + i \langle X_j Y_k \rangle = 2 |a_j| |a_k| e^{i\varphi_{jk}}, \quad (15)$$

whose argument directly yields the phase difference φ_{jk} between sites j and k . Using this quantity, phase differences for opposite-parity pairs (j, k) can be determined from just the $\mathcal{M}_{XX}$ and $\mathcal{M}_{XY}$ global settings. Once those odd-parity phases are fixed, the remaining even-parity relations follow recursively from the odd-parity relations, so that all relative phases are recovered without introducing any additional measurement settings. For more detailed discussion on recursive construction of phase differences, please refer to the article [26].

4.2 Efficient Measurement Settings in Binary encoded system

In the binary encoding, the first step in characterizing the variational state is to recover the amplitudes, $|\alpha_i|$. This task can be accomplished very simply: a single global measurement of all qubits in the computational (Z) basis produces a distribution over bitstrings. Since each codeword $|b(j)\rangle$ is uniquely associated with a basis state of the single-excitation subspace, the corresponding outcome probability is exactly $|\alpha_j|^2$. In other words, the magnitudes of the amplitudes can be directly obtained from the frequency statistics of a single all- Z measurement setting.

The more challenging step is to determine the relative phases between the amplitudes. To achieve this, we introduce a family of operators that selectively couple two chosen codewords. For any two indices j, k , define

$$O_X^{(j,k)} = \bigotimes_{\ell \in D_{jk}} X_\ell \otimes \bigotimes_{\ell \in S_{jk}} P_\ell^{(b_\ell(j))}, \quad (15)$$

where the index sets are given by

$$D_{jk} = \{\ell : b_\ell(j) \neq b_\ell(k)\}, \quad (16)$$

$$S_{jk} = \{\ell : b_\ell(j) = b_\ell(k)\}. \quad (17)$$

Here D_{jk} denotes the qubits on which the two binary strings $b(j)$ and $b(k)$ differ, while S_{jk} denotes the qubits where they coincide. On the differing positions, X operators are applied, whereas on the coinciding positions we project onto the appropriate computational basis state by using

$$P_\ell^{(0)} = \frac{1 + Z_\ell}{2}, \quad P_\ell^{(1)} = \frac{1 - Z_\ell}{2}. \quad (18)$$




Figure 5: The protocol consists of quantum measurements to extract amplitudes and relative phases, followed by classical post-processing to reconstruct the total energy.

By construction, the operator $O_X^{(j,k)}$ swaps the two codewords $|b(j)\rangle$ and $|b(k)\rangle$ while annihilating all other codewords.

However, this operator is not unique, since different pairs of codewords can correspond to the same pair of difference and sum patterns (D, S) . In such cases, the operator O_X couples multiple pairs of basis states simultaneously, and its measured expectation value becomes a sum over several independent contributions rather than isolating a single transition. For instance, in a two-qubit encoding, the operator $O_X = X \otimes X$ acts nontrivially on both the pairs $(00, 11)$ and $(01, 10)$. Consequently, its expectation value contains overlapping information from both transitions, preventing the unique reconstruction of individual relative phases from the measurement outcomes alone.

To eliminate this ambiguity, the operator can be restricted to act only between codeword pairs (i, j) that satisfy $|D_{ij}| = 1$ —that is, pairs differing by a single bit flip. Under this condition, the operator O_X becomes unique, as each such pair corresponds to a distinct

single-qubit transition. The measured expectation value then takes the form

$$\langle O_X^{(j,k)} \rangle = 2|\alpha_j||\alpha_k| \cos(\theta_k - \theta_j), \quad (19)$$

which directly yields the cosine of the relative phase difference between amplitudes α_j and α_k .

This uniqueness can be systematically ensured by arranging the codewords along a cyclic *Gray code* path [33–35], where successive n -bit strings differ by exactly one qubit flip. Following such an ordering establishes a one-to-one correspondence between each operator O_X and an edge in the Gray code graph, guaranteeing that every measurement isolates a single transition in the encoded space. An example of this Gray-coded construction is depicted in Fig. 6.




Figure 6: Cyclic representation of the 3-qubit Gray code sequence, where adjacent states differ by a single bit flip.

An analogous construction yields the sine terms. If in the definition of the swap operator we replace X by Y on the differing qubit, we obtain

$$O_Y^{(j,k)} = \bigotimes_{\ell \in D_{jk}} Y_\ell \otimes \bigotimes_{\ell \in S_{jk}} P_\ell^{(b_\ell(j))}, \quad (20)$$

whose expectation value gives

$$\langle O_Y^{(j,k)} \rangle = 2|\alpha_j||\alpha_k| \sin(\theta_k - \theta_j). \quad (21)$$

Thus, measuring O_Y provides the missing sine information.

Similar to Eq. (13), from the Eqs. (19)–(21), one can obtain the phase difference between the neighbouring qubits using the following equation.

$$\theta_{j+1} - \theta_j = \arctan \left(\frac{\langle X_j Y_{j+1} \rangle}{\langle X_j X_{j+1} \rangle} \right). \quad (22)$$

The phase difference of the remaining pairs can be calculated by summing along the path of gray code curve, by using the Eq. (14). After which one can calculate the energy of the quantum state for Hamiltonian H using Eq. (9).

Since there are n possible positions for the single flip in the Gray code, we need n distinct X -based settings for the cosine terms and n corresponding Y -based settings for the sine terms. In summary, the complete characterization of the encoded variational

state requires only $2n + 1$ distinct measurement settings: one global all- Z measurement to determine the magnitudes, n settings of O_X operators to extract all cosines, and n settings of O_Y operators to extract all sines. This Gray-code-based construction ensures uniqueness of each operator, prevents collisions of pairs with same D and S , and achieves an efficient scaling of the measurement overhead with the number of qubits n in the binary encoded system. The overall workflow is identical to Fig. 5; the only difference is the cost: $2n + 1$ settings instead of three for the SES ansatz.

5 Resource efficiency

In classical computation, based on the current hardware architecture consisting of a combination of (large) memory and localized processing unit, the costs of the computation are defined independently by two parameters. One is the size of the data register, defining the minimal size of the computer needed. This limit is unavoidable and, for example, also limiting the size of a quantum state that can be simulated classically. The other parameter is the number of operations that need to be performed, defining the time costs of the algorithm. For deterministic programs one run suffices, thus even a computer with large data register cannot help to reduce the time costs of an algorithm with high number of operations needed.

For quantum computers the situation is significantly different, as summarized below:

- While the size of a quantum computer is a similar limitation than in classical space, one can in principle encode exponentially more information into the same size of the device
- As in most existing scalable quantum infrastructure the number of possible parallel operations scales with the number of qubits, the depth (number of layers) of the quantum circuit is more relevant than the actual number of operations to be performed.
- As quantum computers are stochastic by nature, runs must be repeated and statistics gathered from measurements. Consequently, increasing the size of the device can reduce runtime via parallelization.
- Unlike the classical case, a single measurement will not reveal all the information about the system. Thus the number of measurement settings needed inherently enters into the time costs of the algorithm, as it increases the running time.

Based on these differences, we introduce a volumetric efficiency metric for quantum devices

$$\mathcal{E} = (\text{qubit width}) \times (\text{circuit depth}) \times (\text{number of measurement settings}), \quad (23)$$

which captures the total space-time-sampling volume that a hardware platform must support during its run. This measure is especially well suited for VQE applications, as all parameters of the circuit (size, depth and measurement settings) can be adjusted.

Let us now summarize the results for the above presented algorithms. The original single-excitation encoding uses N qubits, depth $D_{\text{SES}} = O(N)$ from the linear chain of two-qubit gates by application of A gates, and three global measurement settings for extracting amplitudes and relative phases. The corresponding volumetric cost thus scales as

$$\mathcal{E}_{\text{Original}} = N \times O(N) \times 3 = O(N^2). \quad (24)$$

Resource	Original SES	Binary-encoded SES		
		<i>Hardware-efficient</i>	<i>Full</i>	<i>Gray-code</i>
Qubits	N		$n = \lceil \log_2 N \rceil$	
Measurement settings	3		$2n+1$	
Circuit depth (CNOT)	$\mathcal{O}(N)$	$\mathcal{O}(n)$	$\mathcal{O}(N)$	$\mathcal{O}(Nn)$
Volumetric cost $\mathcal{E}$	$\mathcal{O}(N^2)$	$\mathcal{O}((\log N)^3)$	$\mathcal{O}(N(\log N)^2)$	$\mathcal{O}(N(\log N)^3)$

Table 2: Comparison of asymptotic resource scaling. The three Binary-encoded SES columns correspond to (i) Hardware-efficient ansatz, (ii) a general preparation scheme, and (iii) our Gray-code measurement scheme. Here N is the number of SES basis states and $n = \lceil \log_2 N \rceil$ is the number of qubits in the compressed register.

In the logarithmic-qubit encoding presented above, the width reduces to n qubits. The depth is different depending on the approach chosen. For hardware efficient ansatz, it can be as low as $\mathcal{O}(n)$, while for the most general preparation scheme [36] it is $\mathcal{O}(N)$ and when using the preparation mimicking original SES parametrization it increases to $\mathcal{O}(Nn)$. The number of global measurement settings is then independent on the preparation protocol and is $2n + 1 = \mathcal{O}(n)$. Thus the overall volumetric cost becomes

$$\begin{aligned}
\mathcal{E}_{\text{HE}} &= n \times \mathcal{O}(n) \times \mathcal{O}(n) = \mathcal{O}((\log N)^3) \\
\mathcal{E}_{\text{Full}} &= n \times \mathcal{O}(N) \times \mathcal{O}(n) = \mathcal{O}(N(\log N)^2) \\
\mathcal{E}_{\text{Gray}} &= n \times \mathcal{O}(Nn) \times \mathcal{O}(n) = \mathcal{O}(N(\log N)^3).
\end{aligned} \tag{25}$$

These results are depicted also in Tab. 2.

Consider two illustrative qubit sizes of interest. For $N = 1024$ sites, plausibly reachable on near-term hardware, the binary-encoded approach yields approximate speedups relative to the original SES implementation of 1000 with a hardware-efficient (HE) ansatz and 10 with a general ansatz over the full Hilbert space, while the ansatz that mimics the original SES preparation has essentially the same volumetric cost as the original SES.

On the other hand, for $N = 1048576$, already of practical interest, the corresponding approximate speedups relative to the original SES implementation are 10^8 for a hardware-efficient ansatz, 2.6×10^3 for a general ansatz over the full Hilbert space, and 130 for the SES-mimicking ansatz. For intuition, a computation that would take about a year with the original SES mapping could be reduced to a few days with the SES-mimicking ansatz, a few hours with a general ansatz, and to a fraction of a second with a hardware-efficient ansatz—although here the chance of a failure would probably be rather high. These estimates also ignore hardware size: the original mapping would require on the order of a million physical qubits, whereas our binary encoding needs only $n = \lceil \log_2 N \rceil = 20$ qubits in this example.

6 Conclusion

We have introduced an efficient framework for implementing solid-state Hamiltonians on quantum hardware by encoding the physical subspace into an exponentially smaller number of qubits. By reformulating the single-particle problem within a binary-encoded register and designing an ansatz that preserves the parameterization of the original SES circuit, we achieved an accurate representation of the physical wavefunction using only about $\log_2 N$ qubits. The proposed encoding compresses the Hilbert space exponentially while remaining

fully compatible with variational algorithms such as the VQE, thus enabling simulations of larger systems on near-term quantum hardware.

From a resource standpoint, the logarithmic-qubit encoding adds only a polylogarithmic overhead in circuit depth and measurement settings, while reducing the qubit width exponentially. The resulting volumetric cost improves significantly, meaning that large Hamiltonians can be simulated on dramatically smaller registers with comparable circuit complexity.

The proposed scheme provides a practical, hardware-efficient route for simulating solid-state Hamiltonians in logarithmically reduced qubit spaces. The same principles—compact encodings, structured ansätze, and measurement-efficient protocols—can be extended to more complex systems such as two-excitation Hamiltonians, as long as they occupy only a small fraction of the whole Hilbert space, allowing scalable variational simulations on near-term quantum devices. Studying noise robustness, extensions to correlated multi-electron sectors, as well as the resulting measurement complexity in compact encodings are promising directions for future work.

Acknowledgments

The authors acknowledge Ivana Miháliková’s helpful discussions and feedback.

M.F. acknowledges financial supports from (i) the Czech Academy of Science (the *Praemium Academiae* and the Strategy AV21, in particular the program “AI: Artificial Intelligence for Science and Society”) and (ii) the Ministry of Education, Youth and Sports of the Czech republic (project No. LUC25028 within the INTER-EXCELLENCE II program, a subprogram INTER-COST - LUC25). I.A.M. acknowledges the support of the VEGA project No. 2/0055/23 and Research and Innovation Authority project 09I03-03-V04-00685. M.P. acknowledges the support of project Research and Innovation Authority project 09I03-03-V04-00425.

A Expectation of the Hamiltonian in the SES

Consider an SES state of the form

$$|\Psi\rangle = \sum_{k=1}^N \alpha_k |k\rangle, \quad \alpha_k = |\alpha_k| e^{i\theta_k}, \quad (26)$$

where $|k\rangle$ denotes a single-excitation basis state and α_k is expressed in polar form with modulus $r_k = |\alpha_k|$ and phase θ_k .

The Hamiltonian restricted to this subspace is

$$H = \sum_k h_{kk} |k\rangle\langle k| + \sum_{j \neq k} h_{jk} |j\rangle\langle k|, \quad h_{kj} = h_{jk}^*. \quad (27)$$

The expectation value of H is

$$\begin{aligned} E &= \langle \Psi | H | \Psi \rangle \\ &= \sum_k h_{kk} |\alpha_k|^2 + \sum_{j \neq k} h_{jk} \alpha_j^* \alpha_k. \end{aligned} \quad (A1) \quad (28)$$

Since each off-diagonal term appears twice, we group them as

$$E = \sum_k h_{kk} |\alpha_k|^2 + \sum_{j < k} \left(h_{jk} \alpha_j^* \alpha_k + h_{kj} \alpha_k^* \alpha_j \right). \quad (29)$$

Now substitute $\alpha_j^* \alpha_k = r_j r_k e^{i(\theta_k - \theta_j)}$ and define the relative phase difference $\Delta_{jk} = \theta_k - \theta_j$. This gives

$$\begin{aligned} h_{jk} e^{i\Delta_{jk}} + h_{kj} e^{-i\Delta_{jk}} &= (h_{jk} + h_{kj}) \cos \Delta_{jk} \\ &\quad - \frac{h_{jk} - h_{kj}}{i} \sin \Delta_{jk}. \end{aligned} \quad (30)$$

Using Hermiticity ($h_{kj} = h_{jk}^*$), this separates into real and imaginary parts. The final compact expression for the expectation value is

$$\begin{aligned} E &= \sum_{k=1}^N h_{kk} |\alpha_k|^2 \\ &\quad + 2 \sum_{j=1}^N \sum_{k>j}^N \left[\operatorname{Re}(h_{jk}) |\alpha_j| |\alpha_k| \cos(\theta_k - \theta_j) \right. \\ &\quad \left. - \operatorname{Im}(h_{jk}) |\alpha_j| |\alpha_k| \sin(\theta_k - \theta_j) \right]. \end{aligned} \quad (31)$$

This is the desired expression, with diagonal terms from h_{kk} and off-diagonal contributions separated into real and imaginary parts of h_{jk} .

References

- [1] Trygve Helgaker, Poul Jorgensen, and Jeppe Olsen. *Molecular electronic-structure theory*. John Wiley & Sons, 2013.
- [2] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.

- [3] Peter J Knowles and Nicholas C Handy. A new determinant-based full configuration interaction method. *Chemical Physics Letters*, 111(4-5):315–321, 1984. DOI: [10.1016/0009-2614\(84\)85513-X](https://doi.org/10.1016/0009-2614(84)85513-X).
- [4] Rodney J Bartlett and Monika Musiał. Coupled-cluster theory in quantum chemistry. *Reviews of Modern Physics*, 79(1):291–352, 2007. DOI: [10.1103/RevModPhys.79.291](https://doi.org/10.1103/RevModPhys.79.291).
- [5] Robert G Parr and Yang Weitao. Density-functional theory of atoms and molecules. 1995.
- [6] H Teng, T Fujiwara, T Hoshi, T Sogabe, S-L Zhang, and S Yamamoto. Efficient and accurate linear algebraic methods for large-scale electronic structure calculations with nonorthogonal atomic orbitals. *Physical Review B*, 83(16):165103, 2011. DOI: [10.1103/PhysRevB.83.165103](https://doi.org/10.1103/PhysRevB.83.165103).
- [7] Jakub Valdhans and Petr Klenovský. Combined tight-binding and configuration interaction study of unfolded electronic structure of the g color center in si. *Phys. Rev. B*, 112:165413, Oct 2025. DOI: [10.1103/8d5x-779h](https://doi.org/10.1103/8d5x-779h). URL <https://link.aps.org/doi/10.1103/8d5x-779h>.
- [8] Jun-Qiang Lu, HT Johnson, Vaishno Devi Dasika, and RS Goldman. Moments-based tight-binding calculations of local electronic structure in InAs/GaAs quantum dots for comparison to experimental measurements. *Applied Physics Letters*, 88(5), 2006. DOI: <https://doi.org/10.1063/1.2171473>.
- [9] Alexander Mittelstädt, Andrei Schliwa, and Petr Klenovský. Modeling electronic and optical properties of III–V quantum dots—selected recent developments. *Light: Science & Applications*, 11(1):17, 2022. DOI: [10.1038/s41377-021-00700-9](https://doi.org/10.1038/s41377-021-00700-9).
- [10] Christopher C Paige. Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix. *IMA Journal of Applied Mathematics*, 18(3):341–349, 1976. DOI: [10.1093/imamat/18.3.341](https://doi.org/10.1093/imamat/18.3.341).
- [11] Beresford N Parlett and David S Scott. The Lanczos algorithm with selective orthogonalization. *Mathematics of Computation*, 33(145):217–238, 1979. DOI: [10.1090/S0025-5718-1979-0514820-3](https://doi.org/10.1090/S0025-5718-1979-0514820-3).
- [12] Taisuke Ozaki. $O(N)$ Krylov-subspace method for large-scale *ab initio* electronic structure calculations. *Physical Review B*, 74(24):245101, 2006. DOI: [10.1103/PhysRevB.74.245101](https://doi.org/10.1103/PhysRevB.74.245101).
- [13] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9).
- [14] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):4213, 2014. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [15] Dmitry A Fedorov, Bo Peng, Niranjana Govind, and Yuri Alexeev. VQE method: A short survey and recent developments. *Materials Theory*, 6(1):2, 2022. DOI: [10.1186/s41313-021-00032-6](https://doi.org/10.1186/s41313-021-00032-6).
- [16] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017. DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [17] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H Booth, et al. The variational

- quantum eigensolver: a review of methods and best practices. *Physics Reports*, 986: 1–128, 2022. DOI: [10.1016/j.physrep.2022.08.003](https://doi.org/10.1016/j.physrep.2022.08.003).
- [18] Ijaz Ahamed Mohammad, Matej Pivoluska, and Martin Plesch. Meta-optimization of resources on quantum computers. *Scientific Reports*, 14(1):10312, 2024. DOI: [10.1038/s41598-024-59618-y](https://doi.org/10.1038/s41598-024-59618-y).
- [19] Ivana Miháliková, Matej Pivoluska, Martin Plesch, Martin Friák, Daniel Nagaj, and Mojmir Šob. The cost of improving the precision of the variational quantum eigensolver for quantum chemistry. *Nanomaterials*, 12(2):243, 2022. DOI: [10.3390/nano12020243](https://doi.org/10.3390/nano12020243).
- [20] Oscar Higgott, Daochen Wang, and Stephen Brierley. Variational quantum computation of excited states. *Quantum*, 3:156, 2019. DOI: [10.22331/q-2019-07-01-156](https://doi.org/10.22331/q-2019-07-01-156).
- [21] Jean-Marc Jancu, Reinhard Scholz, Fabio Beltram, and Franco Bassani. Empirical *spds** tight-binding calculation for cubic semiconductors: General method and material parameters. *Physical Review B*, 57(11):6493, 1998. DOI: [10.1103/PhysRevB.57.6493](https://doi.org/10.1103/PhysRevB.57.6493).
- [22] Anh-Luan Phan, Alessandro Pecchia, Alessia Di Vito, and Matthias Auf der Maur. Empirical tight-binding method for large-supercell simulations of disordered semiconductor alloys. *Physica Scripta*, 99(7):075903, 2024. DOI: [10.1088/1402-4896/ad4f65](https://doi.org/10.1088/1402-4896/ad4f65).
- [23] P Vogl, Harold P Hjalmarson, and John D Dow. A semi-empirical tight-binding theory of the electronic structure of semiconductors. *Journal of Physics and Chemistry of Solids*, 44(5):365–378, 1983. DOI: [10.1016/0022-3697\(83\)90064-1](https://doi.org/10.1016/0022-3697(83)90064-1).
- [24] Yaohua Tan, Michael Povolotskyi, Tillmann Kubis, Yu He, Zhengping Jiang, Gerhard Klimeck, and Timothy B Boykin. Empirical tight binding parameters for GaAs and MgO with explicit basis through dft mapping. *Journal of Computational Electronics*, 12(1):56–60, 2013. DOI: [10.1007/s10825-013-0436-0](https://doi.org/10.1007/s10825-013-0436-0).
- [25] Kyle Sherbert, Frank Cerasoli, and Marco Buongiorno Nardelli. A systematic variational approach to band theory in a quantum computer. *RSC Advances*, 11(62): 39438–39449, 2021. DOI: [10.1039/D1RA07451B](https://doi.org/10.1039/D1RA07451B).
- [26] Michal Krejčí, Lucie Krejčí, Ijaz Ahamed Mohammad, Martin Plesch, and Martin Friák. Minimum measurements quantum protocol for band structure calculation. *arXiv preprint arXiv:2511.04389*, 2025.
- [27] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C Benjamin, and Xiao Yuan. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, 2020. DOI: [10.1103/RevModPhys.92.015003](https://doi.org/10.1103/RevModPhys.92.015003).
- [28] Kyle Sherbert, Anooja Jayaraj, and Marco Buongiorno Nardelli. Quantum algorithm for electronic band structures with local tight-binding orbitals. *Scientific Reports*, 12(1):9867, 2022. DOI: [10.1038/s41598-022-13627-x](https://doi.org/10.1038/s41598-022-13627-x).
- [29] Ijaz Ahamed Mohammad, Yury Chernyak, and Martin Plesch. HOPSO: A robust classical optimizer for VQE. *arXiv preprint arXiv:2508.13651*, 2025.
- [30] A generic pure state on n qubits lives in a an exponentially big Hilbert space and, up to normalization and global phase, requires $\mathcal{O}(2^n)$ real parameters to specify. By contrast, a hardware-efficient ansatz (HEA) with fixed layout and limited depth contains only $\mathcal{O}(pn)$ continuous parameters (for p layers) and uses a restricted set of entangling operations. Starting from a simple product state such as $|0\rangle^{\otimes n}$ and applying a shallow, architecture-constrained circuit therefore explores only a tiny subset of the full state space. This induces an intrinsic bias toward low-complexity states and leaves most of Hilbert space effectively inaccessible at practical depths.
- [31] Tanuj Khattar and Craig Gidney. Rise of conditionally clean ancillae for efficient

- quantum circuit constructions. *Quantum*, 9:1752, 2025. DOI: [10.22331/q-2025-05-21-1752](https://doi.org/10.22331/q-2025-05-21-1752).
- [32] Vivek V. Shende and Igor L. Markov. On the cnot-cost of toffoli gates. *Quantum Info. Comput.*, 9(5):461–486, May 2009. ISSN 1533-7146. URL <https://dl.acm.org/doi/10.5555/2011791.2011799>.
- [33] Edgard N Gilbert. Gray codes and paths on the n-cube. *The Bell System Technical Journal*, 37(3):815–826, 1958. DOI: [10.1002/j.1538-7305.1958.tb03887.x](https://doi.org/10.1002/j.1538-7305.1958.tb03887.x).
- [34] Frank Gray. Pulse code communication, 1953. Filed 1947.
- [35] Carla D. Savage. A survey of combinatorial Gray codes. *SIAM Review*, 39(4):605–629, 1997. DOI: [10.1137/S0036144595295272](https://doi.org/10.1137/S0036144595295272).
- [36] Martin Plesch and Časlav Brukner. Quantum-state preparation with universal gate decompositions. *Physical Review A*, 83(3):032302, 2011. DOI: [10.1103/PhysRevA.83.032302](https://doi.org/10.1103/PhysRevA.83.032302).