

Ground-State Energy and Related Properties Estimation in Quantum Chemistry with Linear Dependence on the Number of Atoms

Taehee Ko

*School of Computational sciences**Korea Institute for Advanced Study*

South Korea

kthmomo@kias.re.kr

Xiantao Li

*Department of Mathematics**Pennsylvania State University*

USA

xiantao.li@psu.edu

Chunhao Wang

*Department of Computer Science**and Engineering**Pennsylvania State University*

USA

cwang@psu.edu

Abstract—At the heart of quantum chemistry and materials science lies the critical task of estimating ground-state properties. We present a quantum algorithm for this task by quantizing the density-functional theory (DFT). A key aspect of implementing DFT faithfully is the requirement for self-consistent calculations, which involve repeated diagonalizations of the Hamiltonian. This procedure, however, creates a significant bottleneck, as a classical algorithm generally demands a computational complexity that grows cubically with the number of electrons, restricting the scalability of DFT for tackling large-scale problems that involve complex chemical environments and microstructures. This article presents the first quantum algorithm that has provided substantial speedup for the ground state computation, by improving the complexity to one with a *linear* scaling with the number of *atoms*. The algorithm leverages the exponential speedup by the quantum singular value transformation to generate a quantum circuit to encode the density-matrix, followed by an efficient estimation method for the output electron density, which constitutes a simple hybrid approach for achieving self-consistency.

XL's research is supported by the National Science Foundation Grants DMS-2111221. CW acknowledges support from the National Science Foundation grant CCF-2238766 (CAREER). Both XL and CW were supported by a seed grant from the Institute of Computational and Data Science (ICDS) and the National Science Foundation Grants CCF-2312456. TK is supported by a KIAS Individual Grant CG096001 at Korea Institute for Advanced Study.

Moreover, the algorithm produces the ground state Hamiltonian, from which the ground state energy and band structures can be efficiently computed. The proposed framework is accompanied by a rigorous error analysis that establishes the convergence and quantifies various sources of error and the overall computational complexity. The combination of efficiency and precision opens new avenues for exploring large-scale physical systems.

Index Terms—quantum algorithm, quantum chemistry, density functional theory, quantum simulation

I. INTRODUCTION

Estimating ground-state energy and related properties of a quantum system are of ultimate interest in quantum chemistry and quantum computation. Classically, one of the breakthroughs in computational chemistry and material science is the development of the density-functional theory (DFT) [22]. The theory is founded on the observation that the electronic structures are fully determined by the underlying electron density $n(\mathbf{r}); \mathbf{r} \in \mathbb{R}^3 \rightarrow [0, +\infty)$ which, thanks to the ground-breaking work of Kohn and Sham [26], can be represented through an auxiliary system of non-interacting electrons with an effective Kohn-Sham Hamiltonian $\mathcal{H}$. Consequently, the task of simulating a problem in $\mathbb{R}^{3N}$ is reduced to one in $\mathbb{R}^3$, thus drastically reducing

the problem size. Meanwhile, the electron-electron interactions are captured by an exchange-correlation energy functional of $n(\mathbf{r})$, which is part of the Hamiltonian operator. The electron density must be computed self-consistently to meet the self-consistent field (SCF) requirement [37]. However, direct SCF computation amounts to calculating many eigenvalues of a large-dimensional matrix, for which the computational cost typically scales cubically with the dimension [54]. Such scaling has been the major limiting factor for large-scale DFT calculations e.g., perovskite materials [20], high-entropy alloys [8], two-dimensional materials with a twist angle [53], and biomaterials [12].

This paper aims to demonstrate a quantum advantage by proposing quantum algorithms for estimating ground-state energy based on DFT. Instead of explicitly computing the eigenvalues and eigenvectors, we apply the quantum eigenvalue transformation and construct a quantum circuit to generate the density-matrix in DFT, which exhibits an exponential speedup in terms of the number of electrons. The electron density is then extracted from the diagonals of the density-matrix. What makes the algorithm appealing is that the overall gate and query complexity only scales linearly with the dimension of the problem, which we will compress down to be proportional to the number of atoms. To enable further speedup, we devise an efficient SCF iteration algorithm, where only some components of the electron density need to be updated at each step, thereby significantly reducing the measurement cost. Our theoretical analysis shows the convergence of the iteration methods, and our numerical results indicate that the new SCF algorithm is much more efficient than the classical approach.

II. PROBLEM STATEMENTS AND SUMMARY OF RESULTS

DFT finds the ground state property by solving the following auxiliary eigenvalue problem [26],

$$\begin{aligned}\mathcal{H}[n]|\psi_j\rangle &= E_j|\psi_j\rangle, \mathcal{H}[n] := -\frac{\nabla^2}{2} + V[n](\mathbf{r}), \\ V[n](\mathbf{r}) &:= V_H[n](\mathbf{r}) + V_{xc}[n](\mathbf{r}) + V_{\text{ext}}, \\ \int \psi_i(\mathbf{r})^* \psi_j(\mathbf{r}) d\mathbf{r} &= \delta_{ij},\end{aligned}\tag{1}$$

where E_j 's are the Kohn-Sham eigenvalues and ψ_j 's are the associated wavefunctions. The notation $[\cdot]$ indicates a dependence on the function $n(\mathbf{r})$. In (1), the first term in the Hamiltonian $\mathcal{H}$ is the one-electron kinetic energy. $V_H[n](\mathbf{r})$ is the Hartree potential, which is a functional of n . More precisely, this potential can be obtained by solving the Poisson equation.

The self-consistent field (SCF) in DFT asserts that the electron density $n(\mathbf{r})$ that enters the Hamiltonian $\mathcal{H}$ has to be the same as the electron density determined from the eigenvalues and eigenvectors of $\mathcal{H}$. The output density is denoted by $F(\mathbf{n})$, i.e., $F[n(\mathbf{r})] = \sum_j f(E_j)|\psi_j(\mathbf{r})|^2$, e.g., Fermi-Dirac distribution $f(E) = 1/(1 + e^{\beta(E-\mu)})$. Directly running a fixed-point iteration to update n with $F[n]$ often fails to convergence, a computational difficulty that later motivated many mixing schemes [5], [21], [47], [31]. Nevertheless, the most important issue is that the cost of computing $F[n]$ is often dominated by the computation of the eigenvalue problem in Eq. (1). To formulate a concrete problem, we consider a spatial discretization of Eq. (1) at some grid points, where the input and output density are denoted by $\mathbf{n}$ and $F(\mathbf{n})$, respectively, with dimension equal to the number of grid points.

Problem 1 (Estimating $F(\mathbf{n})$). *Assume that the Hamiltonian $\mathcal{H}$ in Eq. (1) is approximated by an s -sparse Hermitian matrix H on a set of grid points. Suppose we are given the sparse-access oracle for H as in Eqs. (4) and (5). Given any $\epsilon > 0$, the goal is to obtain an unbiased estimate for the updated electron density $\tilde{F}(\mathbf{n}) \in \mathbb{R}^{N_I}$ at N_I grid points, such that $\|\mathbb{E}[\tilde{F}(\mathbf{n})] - F(\mathbf{n})\| < \epsilon$.*

Due to the inherent nature of quantum computation, the estimated electron density $\tilde{F}$ is subject to random measurement outcomes. The expectation here $\mathbb{E}[\cdot]$ denotes the corresponding average.

Theorem 1 (Informal version of Theorem 3). *There is a quantum algorithm that solves Problem 1 by outputting a nearly unbiased approximate electron density $\tilde{F}(\mathbf{n})$ such that $\|\tilde{F}(\mathbf{n}) - F(\mathbf{n})\| < \epsilon$ with probability at least $1 - \delta$ and $\|\mathbb{E}[\tilde{F}(\mathbf{n})] - F(\mathbf{n})\| < \zeta\epsilon + \delta$ with some adjustable parameter ζ . Under the assumptions above, the algorithm involves a query complexity $\tilde{O}(\frac{N_I}{\zeta\epsilon})$.¹*

This quantum algorithm produces the density $\tilde{F}(\mathbf{n})$ whose expectation is close to the true one $F(\mathbf{n})$, which is then incorporated into the SCF iteration. We say that the algorithm converges, if the electron density converges to some $\mathbf{n}_*$ (thus $\mathbf{n}_* = F(\mathbf{n}_*)$), which leads to our second problem,

Problem 2 (Determining the ground state electron density). *Given an input electron density $n_0(\mathbf{r})$, such that $\|n_*(\mathbf{r}) - n_0(\mathbf{r})\| < \gamma$, with sufficiently small γ . Determine an estimate $\hat{n}(\mathbf{r})$ such that $\|n_*(\mathbf{r}) - \hat{n}(\mathbf{r})\| < \epsilon$.*

The main contribution of this work is summarized as follows,

- We formulate the DFT-based computation as an approximation of a matrix function – the density-matrix to leverage the quantum singular value transformations to efficiently prepare the density-matrix $\Gamma = f(H)$ on the circuit.
- We ensure the self-consistency by constructing a hybrid algorithm, where the quantum algorithm produces the density-matrix $\Gamma = f(H)$, while the classical algorithm employs an SCF iteration to provide updated values for the electron density to reprogram the quantum algorithm by updating H at the next step.
- We prove the theoretical convergence for the hybrid algorithm, and our analysis takes into account the function approximation error, measurement error, and iteration error.

¹We use $\tilde{O}$ to neglect poly-logarithmic factors.

- We propose a flexible implementation of the hybrid algorithm, where only some components of the electron density are updated to reduce the sampling complexity. Motivated by the stochastic algorithms in machine learning, we incorporate the unbiased amplitude amplification [41] to mitigate algorithmic bias. This idea of improving hybrid quantum/classical algorithms can be of independent interest.

We first state the theoretical result regarding the convergence of the hybrid algorithm

Theorem 2 (Informal version of Theorem 5). *Under the same Assumption, there is a hybrid quantum-classical algorithm that solves Problem 2 by outputting an approximate ground state electron density $\hat{n}(\mathbf{r})$ such that $\|n_*(\mathbf{r}) - \hat{n}(\mathbf{r})\| < \epsilon$, with probability at least $1 - p_f$ for any given $p_f \in (\frac{\|n_0 - n_*\|^2}{\gamma^2}, 1)$. Neglecting logarithmic factors, and under the assumptions above, the algorithm involves $\tilde{O}(\frac{N_I}{\zeta\epsilon})$ queries to the Hamiltonian H .*

Comparison to classical algorithms for DFT-based computation In classical computing, the most expensive part of typical DFT implementations to compute the electron density is the step of solving the Kohn-Sham equations, which is equivalent to finding eigenpairs corresponding to the Hamiltonian matrix. Many efficient techniques have been proposed over the last two decades, such as polynomial filtering methods [3], [54], [32], direct energy minimization [49], [51] and spectrum slicing type methods [43], [27]. For numerical implementations of DFT, the readers are referred to [33], [28], [29], and a large collection of software packages [19], [52], [32], [44], [13], [45], [34]. The complexity of these algorithms typically scales cubically with the number of electrons: $\mathcal{O}(N_e^3)$. On the other hand, there are many classical algorithms for electron structure calculations that exhibit linear with respect to the number of atoms [46], [9], [18], [14], [15]. Such scaling is obtained with locality assumptions, also known as the “nearsightedness”, which is not assumed in this paper. Furthermore, it is difficult to quantify the error from linear-scaling algorithms. Finally, self-consistency is usually enforced in terms of charges, rather than the values of

the electron density in the original DFT formalism.

Notations. We use bold fonts for vectors, e.g., $\mathbf{r}$, and the entries will be labeled in parenthesis, e.g., $\mathbf{r}(j)$ being the j th entry of $\mathbf{r}$. $D \subset \mathbb{R}^3$ will be used to denote the physical domain. $n(\mathbf{r}) : D \rightarrow [0, +\infty)$ is a function representing the electron density. As a summary of the notations, N_e and N_a are respectively the number of electrons and the number of atoms. In the numerical discretization, D_δ is a set of grid points in D , with N_I being the number of grid points, which is often comparable to N_a . We choose $N_I = M = 2^m$ to map functions defined at the grid points to quantum states in a m -qubit system. $H \in \mathbb{R}^{N_I \times N_I}$ is the Hamiltonian represented at the grid points.

III. REAL-SPACE DISCRETIZATION

To solve Eq. (1), we assume that the Hamiltonian operator is properly discretized in a three-dimensional domain D by a finite-difference method [2] with grid size δ . Note that such a discretization generates an error typically of order δ^{2p} with a finite-difference stencil of width p . Thus one can choose $\delta = \epsilon^{1/2p}$ to control the discretization error to within ϵ .

Within the discretization, the electron density at the grid points is expressed as a vector $\mathbf{n}$. Following the Hamiltonian operator in Eq. (1), we can express the matrix H as follows,

$$H(\mathbf{n}) = -\frac{1}{2}\nabla_\delta^2 + V_\delta(\mathbf{n}), \quad (2)$$

where ∇_δ^2 approximates the kinetic energy operator [2]. V_δ , which enters the Hamiltonian through the diagonals, is the potential evaluated at the grid points and it collects all the potential terms in the Hamiltonian operator. In terms of the matrix H from the finite-difference approximation, we can define the density-matrix $\Gamma \in \mathbb{C}^{2^m \times 2^m}$: $\Gamma = f(H)$. Thus, we generalize the continuous fixed-point problem to a discrete one,

$$\mathbf{n} = f(H(\mathbf{n})), \quad \mathbf{n} \in \mathbb{R}^{N_I}. \quad (3)$$

Let $N_I = 2^m$ be the number of grid points. Denote by D_Δ , the set of the grid points. Fig. 1 illustrates the ground-state electron density from a DFT calculation, satisfying (3).




Fig. 1: A quick illustration of the electron density of Silicon Nitride on its planar subsystem. The data is obtained by DFT calculations using an example from [23].

The matrix H from the spatial discretization is usually sparse. In quantum computing, the sparsity implies that the matrix is *efficiently row/column computable*. To access H , we assume we have access to a procedure O_S that can perform the following mapping:

$$O_S : |i\rangle |k\rangle \mapsto |i\rangle |r_{i_k}\rangle, \quad (4)$$

where r_{i_k} is the k -th nonzero entry of the i -th row of H . In addition, O_H can also perform the following mapping:

$$O_H : |i\rangle |j\rangle |0\rangle \mapsto |i\rangle |j\rangle |H(i, j)\rangle. \quad (5)$$

One key ingredient of our quantum algorithm is block encoding. If a $(m + a)$ -qubit unitary U_A satisfies $\|A - \alpha(\langle 0^{\otimes a} | \otimes I)U_A(|0^{\otimes a}\rangle \otimes I)\|_2 \leq \epsilon$, then it is called an (α, a, ϵ) -block-encoding of A . With the oracles O_S and O_H , the block encoding can be constructed as a unitary with the upper-left block being proportional to H , [17, Lemma 27]

$$U_H = \begin{pmatrix} H & \cdot \\ \cdot & \cdot \end{pmatrix}. \quad (6)$$

The implementation of O_S and O_H depends on the number of distinct values of non-zero entries of H in (3). For this, we need to store $\mathcal{O}(N_I)$ parameters in quantum random-access memory (QRAM) in

order to update the diagonals of H input oracle. The gate complexity for implementing the addressing scheme of such QRAM is $O(N_I)$. But the circuit depth is $\mathcal{O}(\log N_I)$ [36]. The input oracles O_S and O_H for H can be implemented as a procedure that reads data in the QRAM. Another alternative is to construct such oracles without QRAM, but rely on circuit architectures for structured matrices [6].

IV. PREPARING THE DENSITY-MATRIX

Since H is Hermitian, one can use the spectral map and approximate the density-matrix Γ by polynomial approximations of the Fermi-Dirac function. To apply this technique to the density-matrix Γ , we rescale the Hamiltonian matrix

$$f(H) = \left(1 + e^{\beta\left(\frac{\lambda_+ + \lambda_-}{2} - \mu\right)} e^{\hat{\beta}\tilde{H}}\right)^{-1}, \quad (7)$$

where

$$\hat{\beta} = \frac{\lambda_+ - \lambda_-}{2}\beta, \tilde{H} = \frac{2}{\lambda_+ - \lambda_-} \left[H - \frac{\lambda_+ + \lambda_-}{2} I \right]. \quad (8)$$

Here λ_- and λ_+ are some lower and upper bounds of the eigenvalues of H . The scaling is simply to map the eigenvalues of H to the interval $[-1, 1]$. Noticing that $\sigma(\tilde{H}) \subset [-1, 1]$, we apply the Chebyshev approximation of the density matrix. The following lemma, as in [24, Remark 4.8], quantifies the quality of the approximation

Lemma 1. *For a given inverse temperature β , the degree of the Chebyshev expansion to approximate $f(H)$, up to a precision ϵ , is at least,*

$$\ell = \Theta \left(\log_r \frac{1}{\epsilon} \right). \quad (9)$$

Here $r \in \left(1, \frac{c(\beta) + \sqrt{c(\beta)^2 + 4}}{2}\right)$ with $c(\beta) = \frac{4\pi}{\lambda_+ - \lambda_-} \frac{1}{\beta}$.

Polynomial approximations of the density-matrix Γ are not new. In fact, it has been used in [10] for DFT. But in this classical algorithm, the matrix multiplications will introduce significant computational overhead. In contrast, the quantum singular value transformation (QSVT) [17] can efficiently prepare the density-matrix with a complexity that does not depend on the matrix dimension explicitly.

According to the property of matrix norms, the condition that $|H_{i,j}| \leq 1$ for $i, j \in [m]$ is automatically satisfied due to the scaling in Eq. (8). Under this condition, QSVT builds a block-encoding of the following matrix function for $x \in [-1, 1]$,

$$U_{p_\ell(H)} = \begin{pmatrix} p_\ell(H) & \cdot \\ \cdot & \cdot \end{pmatrix}, \quad p_\ell(x) \approx \frac{1}{2(1 + \exp(\hat{\beta}x))}, \quad (10)$$

where $\hat{\beta}$ is defined in Eq. (8). This is summarized as follows,

Lemma 2 ([17, Theorem 31]). *Let U_H be a block encoding of H and $p_\ell(x) \in \mathbb{R}[x]$ be a polynomial of degree ℓ such that $\sup_{x \in [-1, 1]} |p_\ell(x)| \leq \frac{1}{2}$. Then there is a quantum circuit that implements a block encoding of $p_\ell(H)$, $U_{p_\ell(H)}$, with ℓ application of U_H and $U_H^\dagger$, one application of controlled- U_H gate, and $\mathcal{O}(\ell)$ other one- and two-qubit gates.*

V. ESTIMATING THE ELECTRON DENSITY

Recall that the electron density at different locations corresponds to the diagonals of $f(H)$:

$$F(\mathbf{n})(j) = \text{tr}(\rho_j f(H)), \quad \rho_j := |\mathbf{r}_j\rangle\langle\mathbf{r}_j|, \quad (11)$$

where $j \in [N_I]$, with $\mathbf{r}_j$ being a grid point in D_Δ .

The QSVT uses the polynomial approximation $f(H) \approx p_\ell(H)$, and it provides an approximate block encoding of $p_\ell(H)$. Thus, we use the following estimator for $F(\mathbf{n})$,

$$\hat{F}(\mathbf{n})(j) = 2 \text{tr}(\rho_j p_\ell(H)). \quad (12)$$

We have treated $p_\ell(H)$ as observables. To measure them, one may consider the amplitude amplification (AA) method [40], especially because we have prepared the density matrix as a block encoding. However, this approach can create a bias that complicates the hybrid algorithm. To circumvent this issue, we use the more recent algorithm [41] that is equipped with a controllably small bias. To apply the unbiased amplitude estimation [41], we let U_ρ be X_j where X_j is the Pauli X gate acting on the j -th qubit in the first register. In light of Lemma 1 together with [41, Theorems 6 and 23], we immediately have,

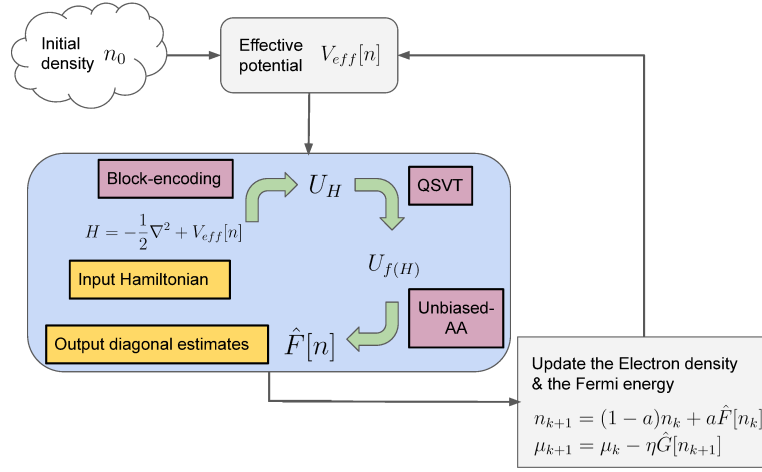


Fig. 2: An illustration of the hybrid algorithm: the blue box highlights the operations involved in the quantum algorithm, while the gray box shows the update formulas for the electron density and chemical potential on classical computers.

Theorem 3. For each $j \in [N_I]$ and any $\epsilon, \zeta, \delta > 0$, there is a quantum algorithm that outputs $\tilde{F}(\mathbf{n})(j)$ as an approximate unbiased estimate of $F(\mathbf{n})(j)$ satisfying that

$$\tilde{F}(\mathbf{n})(j) \in [0, 1], \quad (13)$$

$$\mathbb{P} \left[\left| \tilde{F}(\mathbf{n})(j) - F(\mathbf{n})(j) \right| < \epsilon \right] > 1 - \delta, \quad (14)$$

$$\left| \mathbb{E}[\tilde{F}(\mathbf{n})(j)] - F(\mathbf{n})(j) \right| \leq \zeta \epsilon + \delta. \quad (15)$$

It uses $\mathcal{O} \left(\frac{1}{\epsilon} \left(\log \left(\frac{1}{\delta} \right) + \frac{1}{\zeta} \right) \right)$ queries to $U_{p_\ell(H)}$ and $U_{p_\ell(H)}^\dagger$, and X_j gate.

Notice that the parameter ζ controls the unbiasedness of the estimator, which is determined by the degree of polynomial sampling within the unbiased AA in [41]. The proof is provided in Appendix X.

VI. HYBRID ALGORITHMS AND THE OVERALL COMPLEXITY

To perform SCF iterations, we use the estimate of the electron density to interface with fixed-point iterations on a classical computer. Overall, this constitutes a hybrid algorithm for implementing DFT: One iteration on a classical computer produces a new electron density at the grid points

in D_Δ , followed by the simple mixing scheme, as illustrated in Fig. 2. We establish theoretical results with the following assumptions,

Assumptions. Given the electron density at N_I grid points, the potential V in the Hamiltonian matrix (esp. V_H) can be evaluated with precision ϵ with cost $\mathcal{O}(N_I)$, excluding logarithmic factors (i.e. an efficient evaluation of potential via a local formulation of the electrostatics in [16]). In addition, we make the standard assumption that the mapping $F(\mathbf{n})$ is locally-contractive in that there exist some weighted vector norm $\|\cdot\|_w$ and some w -dependent $c \in (0, 1)$ such that

$$\|F(\mathbf{n}) - F(\mathbf{n}')\|_w \leq c \|\mathbf{n} - \mathbf{n}'\|_w, \quad (16)$$

for all $\mathbf{n}, \mathbf{n}' \in B_\gamma(\mathbf{n}_*)$ which denotes the ball centered at the fixed point $\mathbf{n}_*$ with radius γ . This condition can be obtained from mixing schemes [31], [24], [7]. Finally, for the estimator $\tilde{F}(\mathbf{n})(j)$ defined in (3), we assume that there exist $\sigma > 0$ and $\sigma_\infty > 0$ such that, for any $\mathbf{n} \in B_\gamma(\mathbf{n}_*)$ and each $j \in [N_I]$, the following bounds hold,

$$\text{var} [\tilde{F}(\mathbf{n})] \leq \sigma^2, \quad \text{var} [\tilde{F}(\mathbf{n})(j)] \leq \sigma_\infty^2. \quad (17)$$

This means that near the fixed point $\mathbf{n}_*$, the variance of noise involved in the density estimation is uniformly bounded, which is reasonable as long as the overall quantum noise is controlled. In the context of classical stochastic SCF calculations, a similar assumption is used [24].

Next, we establish a stochastic stability property to quantify the effect of measurement noise and analyze the convergence of hybrid SCF methods. More importantly, the stability analysis allows us to quantify the overall query complexity. Recall that we denote by $\tilde{F}(\mathbf{n})$ the vector in $\mathbb{R}^{N_I}$, whose components are defined in Theorem 3.

We first consider the SCF iteration with simple mixing $\mathbf{n}_{k+1} = (1 - a)\mathbf{n}_k + a\tilde{F}(\mathbf{n}_k)$ with $\tilde{F}(\mathbf{n}_k)$ estimated using QSVT and AA. Due to the fact that all components of the density are updated at each iteration, this method will be referred to as the full-coordinate fixed-point (FCFP) algorithm.

Theorem 4. *For any precision $\epsilon > 0$, a given initial guess $\mathbf{n}_1 \in B_\gamma(\mathbf{n}_*)$ and any failure probability $p_f \in \left(\frac{\|\mathbf{n}_1 - \mathbf{n}_*\|_w^2}{\gamma^2}, 1\right)$, the hybrid algorithm using all coordinates of the electron density can achieve an iterate $\mathbf{n}_t$ for some $t \in [T]$ such that $\|\mathbf{n}_t - \mathbf{n}_*\|_w < \epsilon$ with probability at least $1 - p_f$ with the learning rate*

$$a < \min\left\{a_0, \frac{2r_0}{\frac{\sigma^2}{\epsilon^2} + r_0^2}, \frac{2}{r_0}\right\} \quad (18)$$

and the number of iterations at least

$$T = \mathcal{O}\left(\frac{1}{(1 - c_{\epsilon,a})p_f} \log\left(\frac{\|\mathbf{n}_1 - \mathbf{n}_*\|_w^2}{\epsilon}\right)\right), \quad (19)$$

where $c_{\epsilon,a} := c^2 + \frac{a^2\sigma^2}{\epsilon^2}$. Overall, it involves $\mathcal{O}\left(\frac{N_I}{\epsilon} \log \frac{1}{\epsilon} \left(\log \frac{1}{\delta} + \frac{1}{\zeta}\right)\right)$ queries to O_H .

The proof is provided in the full version [25].

Next, we introduce an alternative to the FCFP method: Rather than updating all components of $F(\mathbf{n})$, we only update the components selectively. The key idea is similar to the randomized coordinate iterative algorithms [35], [48], [38]. The new method will be referred to as the randomized block coordinate fixed-point method (RBCFP). Specifi-

cally, given a fixed-point mapping $\tilde{F}(\mathbf{n})$, RBCFP is defined as

$$\begin{aligned} \tilde{F}_{R,m}(\mathbf{n}) &= \sum_{k \in \{k_{R_j}\}_{j=1}^m} (\mathbf{u}_k, \tilde{F}(\mathbf{n})) \mathbf{u}_k \\ &+ \sum_{k \notin \{k_{R_j}\}_{j=1}^m} (\mathbf{u}_k, \mathbf{n}) \mathbf{u}_k \end{aligned} \quad (20)$$

where $\{k_{R_j}\}_{j=1}^m$ is the set of m indices randomly sampled from the index set $[N_I]$, uniformly without replacement, and the parenthesis $(\cdot, \cdot)$ refers to the standard inner product between vectors. We remark that when $m = N_I$, the method becomes the full coordinate approach. The following theorem shows that despite the partial update of the density, the method still has linear convergence. As proved in the full version [25].

Theorem 5. *For any precision $\epsilon > 0$, a given initial guess $\mathbf{n}_1 \in B_\gamma(\mathbf{n}_*)$, a given $m \in [N_I]$ and any failure probability $p_f \in \left(\frac{\|\mathbf{n}_1 - \mathbf{n}_*\|_w^2}{\gamma^2}, 1\right)$, RBCFP achieves an iterate $\mathbf{n}_t$ for some $t \in [T]$ such that $\|\mathbf{n}_t - \mathbf{n}_*\|_w < \epsilon$ with probability at least $1 - p_f$ with the learning rate*

$$a < \min\left\{a_0, \frac{2r_0}{\frac{N_I\sigma_\infty^2}{m\epsilon^2} + r_0^2}, \frac{2}{r_0}\right\} \quad (21)$$

and the number of iterations is at most

$$T = \mathcal{O}\left(\frac{1}{\frac{m}{N_I}(1 - c_{N_I,\epsilon,a})p_f} \log\left(\frac{\|\mathbf{n}_1 - \mathbf{n}_*\|_w^2}{\epsilon}\right)\right), \quad (22)$$

where $c_{N_I,\epsilon,a} := c^2 + \frac{N_I a^2 \sigma_\infty^2}{\epsilon^2}$. Overall, the algorithm involves $\mathcal{O}\left(\frac{N_I}{\epsilon} \log \frac{1}{\epsilon} \left(\log \frac{1}{\delta} + \frac{1}{\zeta}\right)\right)$ queries to O_H .

VII. APPLICATIONS

To mimic our hybrid algorithm on a classical computer, we conducted numerical tests for the approximation of the density-matrix in Eq. (12) within the MATLAB platform M-SPARC, a real-space DFT code [16]. We chose Barium titanate (BaTiO₃) and an H₂O sheet as our test models from the set of examples in M-SPARC¹. In the models,

¹<https://github.com/SPARC-X/M-SPARC/tree/master/tests>




Fig. 3: Examining the convergence of the hybrid algorithms using the system H_2O -sheet. The density-matrix is approximated by a Chebyshev polynomial with degree $\ell = 500$. The x -axis labels the number of coordinate evaluations. Left: the y -axis is the error of the electron density. Right: the values of the chemical potential μ during the iterations. The FCFP method runs with damping parameter $a = 0.4$; the RBCFP method is run with $a = 0.99$ and block sizes $m = 4, 12, 72$. The damping parameter for updating the chemical potential is $\eta = 0.1$.

temperatures are set to $T = 300K$, and we set up the computation for BaTiO₃ in a cubic supercell and for H₂O in the $x - y$ plane with Dirichlet boundary condition imposed in the z direction.

Within these numerical studies, we compared the performance of FCFP and RBCFP in terms of the number of updated coordinates of the electron density or coordinate evaluations, which in a hy-

brid algorithm, will be directly proportional to the number of quantum measurements. Fig. 3 shows the convergence of the electron density n and chemical potential μ , where Chebyshev approximation of the density-matrix is employed to take into account the QSVT implementation. From the left panel of Fig. 3, we observe that RBCFP exhibits rapid convergence despite the error due to the polynomial approximation. Meanwhile, our hybrid algorithm also updates the chemical potential at each SCF step using the stochastic algorithm [42] (see the full version [25]). The convergence of the chemical potential is shown in the right panel of Fig. 3.

To focus on testing the convergence rate, we generate the density-matrix exactly (no polynomial approximation error). Fig. 4 shows linear convergence of both methods. Similar to the results in Fig. 3, we observe that RBCFP can tolerate a larger learning rate: FCFP becomes unstable when $a > 0.4$, but RBCFP remains stable. The convergence rate of RBCFP is consistently better than that of FCFP, suggesting that its actual performance can be much better than the theoretical bounds: it requires fewer coordinate evaluations than the FCFP method, showing both robustness and efficiency as a hybrid quantum-classical algorithm.

VIII. DISCUSSIONS

We demonstrate the quantum advantage in implementing the density-functional theory. The complexity of our quantum algorithm scales linearly with the dimension of the density update $F(\mathbf{n})$, which typically scales as N_e , the number of electrons. Therefore, this can be considered as linear/sublinear scaling, which compared to the cubic scaling in classical algorithms, is a significant reduction.

The advantages of the algorithm presented in this paper are not limited to the mean-field model of DFT: It can be applied to other mean-field quantum descriptions where self-consistency is central. Another common practice in DFT calculations is to exclude core electrons and incorporate their effects using pseudopotentials [33]. While it is not clear whether this is needed in a quantum algorithm, it is still of theoretical interest to explore how such




Fig. 4: Performance comparison of the FCFP and RBCFP methods. In both panels, the x -axis labels the number of coordinate evaluations. The y -axis labels the error of the electron density on a logarithmic scale. Left: BaTiO₃ system with $N_e = 40$ electrons fixed; Right H₂O with $N_e = 8$ electrons fixed. In the left panel, the FCFP method runs with damping parameter $a = 0.3$, but the RBCFP method with $a = 0.95$. In the right panel, the FCFP runs with $a = 0.4$ which is close to the edge of stability, and the RBCFP with $a = 0.99$.

potentials can be block encoded into U_H . These issues will be explored in separate works.

Quantum algorithms [1], [30] have been designed to compute the ground state energy based on the many-body descriptions, e.g., the second quantization form. In general, those algorithms work with a Hamiltonian that consists of $\mathcal{O}(N_e^4)$ terms, thereby requiring a $\mathcal{O}(N_e^4)$ query complexity. In contrast, our algorithm is based on DFT, which is a mean-field model, and it only involves a $\mathcal{O}(N_e)$ complexity, due to the dimension of a Hamiltonian. Once the ground state density is obtained from our algorithm, we can compute relevant properties by combining it with other quantum algorithms: [11] for the computations of ground-state (GS) energy and energy bands and [40] for the density of states. Table I summarizes these approaches. One assumption in the use of the method [11] is non-zero overlaps with the KS eigenstates in the com-

	QOI	total run time
This work + [11]	GS energy energy bands	$\tilde{\mathcal{O}}\left(\frac{N_e}{\epsilon}\right)$
This work + [40]	Density of States	$\tilde{\mathcal{O}}\left(\frac{N_e + d^2}{\epsilon}\right)$
Method in [30]	GS energy	$\tilde{\mathcal{O}}\left(\frac{N_e^4}{\epsilon}\right)$

TABLE I: Comparison of properties that can be computed and the total run time needed to reach precision ϵ . d denotes the maximal degree of the kernel polynomial method [50]. N_e denotes the number of electrons.

putation of energy bands, otherwise, the problem would be QMA-hard. Meanwhile, the variational quantum eigensolver (VQE) [39] is another hybrid algorithm that is designed for ground-state calculations, perhaps having a comparable complexity to our hybrid approach. However, unlike DFT, VQE's ability to universally represent a ground state for a wide variety of physical systems has not been established [4].

IX. ACKNOWLEDGMENT

The authors thank Dr. Chao Yang for the discussions and references related to this work.

X. APPENDICES

A. Algorithm

The following pseudo-code illustrates the hybrid algorithm for the self-consistent field (SCF) calculations in density functional theory (DFT),

Notice that the full coordinate fixed-point iteration (FCFP) is obtained by setting $m = N_I$. Again, the second step requires quantum computations to estimate m components of the electron density. Once it is done, we update the electron density and the chemical potential on a classical computer.

B. Estimating the Chemical potential

The hybrid algorithm 1 has the flexibility to incorporate an additional step to determine chemical potential μ . This step may be neglected in the case of the grand canonical ensemble, where μ is given. On the other hand, when the number of electrons N_e is fixed (e.g. NVT ensemble), the chemical potential μ is determined on the ground that N_e is preserved

Algorithm 1: Randomized block coordinate fixed-point iteration

Input: initial guess $\mathbf{n}_0$, damping parameters $a, \eta \in (0, 1)$, index parameter $m \in [N_I]$

Output: converged density and chemical potential $(\mathbf{n}_*, \mu_*)$

for $k = 0 : T$

- 1) Sample m different indices $\{k_{R,j}\}_{j=1}^m \subset [N_I]$ uniformly
- 2) Estimate $\tilde{F}_{R,m}(\mathbf{n}_k)$ using the QSVT and the unbiased Amplitude Amplification.
- 3) $\mathbf{n}_{k+1} = (1 - a)\mathbf{n}_k + a\tilde{F}_{R,m}(\mathbf{n}_k)$
- (*) $\mu_{k+1} = \mu_k - \eta(G(\mathbf{n}_{k+1}) - N_e)$ if the chemical potential μ needs to be adjusted
- 4) Update the Hamiltonian

end

during the SCF iteration. At the continuous level, the constraint is formulated as

$$\int n(\mathbf{r}) d\mathbf{r} = N_e, \quad n(\mathbf{r}) = \langle \mathbf{r} | f(H - \mu I) | \mathbf{r} \rangle. \quad (23)$$

Here the first equation can be cast into a nonlinear equation,

$$G(n, \mu) = 0, \quad G(n, \mu) := \int n(\mathbf{r}) d\mathbf{r} - N_e. \quad (24)$$

Notice that given $n(\mathbf{r})$, G is a monotone function of μ .

To incorporate the constraint in 23 into our hybrid algorithm, we should consider a discretized version of the nonlinear function G . Then, we update the chemical potential on classical computers as we do the electron density, (see the algorithm 1). These considerations are summarized as follows

$$\begin{aligned} \mathbf{n}_{k+1} &= (1 - a)\mathbf{n}_k + a\tilde{F}_{R,m}(\mathbf{n}_k), \\ \mu_{k+1} &= \mu_k - \eta(G(\mathbf{n}_{k+1}) - N_e), \end{aligned} \quad (25)$$

where

$$G(\mathbf{n}) := \sum_j |\mathbf{n}(j)|\delta^3,$$

is a discretization of the function in 24. Here δ^3 is the infinitesimal volume from the finite-difference approach with grid size δ ; $\eta \in (0, 1)$ is the damping

parameter for updating μ . This solver for μ is motivated by the stochastic approximation method by Robbins and Monro [42] for solving nonlinear equations.

REFERENCES

- [1] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe. Chemical basis of trotter-suzuki errors in quantum chemistry simulation. *Physical Review A*, 91(2):022311, 2015.
- [2] Thomas L. Beck. Real-space mesh techniques in density-functional theory. *Reviews of Modern Physics*, 72(4):1041–1080, October 2000.
- [3] Constantine Bekas, Effrosini Kokiopoulou, and Yousef Saad. Computation of large invariant subspaces using polynomial filtered lanczos iterations with applications in density functional theory. *SIAM Journal on Matrix Analysis and Applications*, 30(1):397–418, 2008.
- [4] Jacob Biamonte. Universal variational quantum computation. *Physical Review A*, 103(3):L030401, 2021.
- [5] DR Bowler and MJ Gillan. An efficient and robust technique for achieving self consistency in electronic structure calculations. *Chemical Physics Letters*, 325(4):473–476, 2000.
- [6] Daan Camps, Lin Lin, Roel Van Beeumen, and Chao Yang. Explicit quantum circuits for block encodings of certain sparse matrices. *arXiv preprint arXiv:2203.10236*, 2022.
- [7] Eric Cancès, Gaspard Kemlin, and Antoine Levitt. Convergence analysis of direct minimization and self-consistent iterations. *SIAM Journal on Matrix Analysis and Applications*, 42(1):243–274, 2021.
- [8] Shuai Chen, Zachary H Aitken, Subrahmanyam Pattamatta, Zhaoxuan Wu, Zhi Gen Yu, David J Srolovitz, Peter K Liaw, and Yong-Wei Zhang. Simultaneously enhancing the ultimate strength and ductility of high-entropy alloys via short-range ordering. *Nature communications*, 12(1):4953, 2021.
- [9] Fabrizio Cleri and Vittorio Rosato. Tight-binding potentials for transition metals and alloys. *Physical Review B*, 48(1):22, 1993.
- [10] Yael Cytter, Eran Rabani, Daniel Neuhauser, and Roi Baer. Stochastic density functional theory at finite temperatures. *Physical Review B*, 97(11):115207, 2018.
- [11] Zhiyan Ding and Lin Lin. Simultaneous estimation of multiple eigenvalues with short-depth quantum circuit on early fault-tolerant quantum computers. *arXiv preprint arXiv:2303.05714*, 2023.
- [12] Marcus Elstner, Th Frauenheim, E Kaxiras, G Seifert, and S Suhai. A self-consistent charge density-functional based tight-binding scheme for large biomolecules. *physica status solidi (b)*, 217(1):357–376, 2000.
- [13] Julian Gale. Siesta: A linear-scaling method for density functional calculations. In *Computational Methods for Large Systems-Electronic Structure Approaches for Biotechnology and Nanotechnology*, pages 45–75. Wiley & Sons Inc., 2011.
- [14] C. J. García-Cervera, J. Lu, and W. E. A sub-linear scaling algorithm for computing the electronic structure of materials. *Communications in Mathematical Sciences*, 5(4):999–1026, 2007.

- [15] Vikram Gavini, Kaushik Bhattacharya, and Michael Ortiz. Quasi-continuum orbital-free density-functional theory: A route to multi-million atom non-periodic dft calculation. *Journal of the Mechanics and Physics of Solids*, 55(4):697–718, 2007.
- [16] Swarnava Ghosh and Phanish Suryanarayana. Sparc: Accurate and efficient finite-difference formulation and parallel implementation of density functional theory: Isolated clusters. *Computer Physics Communications*, 212:189–204, 2017.
- [17] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204. ACM, 2019.
- [18] S. Goedecker. Linear scaling electronic structure methods. *Reviews of Modern Physics*, 71(4):1085, 1999.
- [19] Jürgen Hafner. Ab-initio simulations of materials using vasp: Density-functional theory and beyond. *Journal of computational chemistry*, 29(13):2044–2078, 2008.
- [20] Geoffroy Hautier, Christopher C Fischer, Anubhav Jain, Tim Mueller, and Gerbrand Ceder. Finding nature’s missing ternary oxide compounds using machine learning and density functional theory. *Chemistry of Materials*, 22(12):3762–3767, 2010.
- [21] R Haydock, Volker Heine, and MJ Kelly. Electronic structure based on the local atomic environment for tight-binding bands. *Journal of Physics C: Solid State Physics*, 5(20):2845, 1972.
- [22] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864, 1964.
- [23] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.
- [24] Taehee Ko and Xiantao Li. Stochastic algorithms for self-consistent calculations of electronic structures. *Mathematics of Computation*, 92(342):1693–1728, 2023.
- [25] Taehee Ko, Xiantao Li, and Chunhao Wang. Implementation of the density-functional theory on quantum computers with linear scaling with respect to the number of atoms. *arXiv preprint arXiv:2307.07067*, 2023.
- [26] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133–A1138, 1965.
- [27] Ruipeng Li, Yuanzhe Xi, Eugene Vecharynski, Chao Yang, and Yousef Saad. A thick-restart lanczos algorithm with polynomial filtering for hermitian eigenvalue problems. *SIAM Journal on Scientific Computing*, 38(4):A2512–A2534, 2016.
- [28] Lin Lin and Jianfeng Lu. *A mathematical introduction to electronic structure theory*. SIAM, 2019.
- [29] Lin Lin, Jianfeng Lu, and Lexing Ying. Numerical methods for kohn–sham density functional theory. *Acta Numerica*, 28:405–539, 2019.
- [30] Lin Lin and Yu Tong. Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers. *PRX Quantum*, 3(1):010318, 2022.
- [31] Lin Lin and Chao Yang. Elliptic preconditioner for accelerating the self-consistent field iteration in kohn–sham density functional theory. *SIAM Journal on Scientific Computing*, 35(5):S277–S298, 2013.
- [32] Kai-Hsin Liou, Chao Yang, and James R Chelikowsky. Scalable implementation of polynomial filtering for density functional theory calculation in parsec. *Computer Physics Communications*, 254:107330, 2020.
- [33] R. M. Martin. *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, 2011.
- [34] Phani Motamarri, Sambit Das, Shiva Rudraraju, Krishnendu Ghosh, Denis Davydov, and Vikram Gavini. Dft-fe—a massively parallel adaptive finite-element code for large-scale density functional theory calculations. *Computer Physics Communications*, 246:106853, 2020.
- [35] Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [36] Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2011.
- [37] R. G. Parr and W. Yang. *Density-functional theory of atoms and molecules*. Oxford University Press, 1995.
- [38] Zhimin Peng, Yangyang Xu, Ming Yan, and Wotao Yin. Arock: an algorithmic framework for asynchronous parallel coordinate updates. *SIAM Journal on Scientific Computing*, 38(5):A2851–A2879, 2016.
- [39] 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.
- [40] Patrick Rall. Quantum algorithms for estimating physical quantities using block encodings. *Physical Review A*, 102(2):022408, 2020.
- [41] Patrick Rall and Bryce Fuller. Amplitude estimation from quantum signal processing. *Quantum*, 7:937, 2023.
- [42] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [43] Grady Schofield, James R Chelikowsky, and Yousef Saad. A spectrum slicing method for the kohn–sham problem. *Computer Physics Communications*, 183(3):497–505, 2012.
- [44] Gotthard Seifert and Jan-Ole Joswig. Density-functional tight binding—an approximate density-functional theory method. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2(3):456–465, 2012.
- [45] Abhiraj Sharma and Phanish Suryanarayana. On the calculation of the stress tensor in real-space kohn–sham density functional theory. *The Journal of chemical physics*, 149(19):194104, 2018.
- [46] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal. The SIESTA method for ab initio order-N materials simulation. *Journal of Physics: Condensed Matter*, 14(11):2745, 2002.
- [47] A. Toth and C. T. Kelley. Convergence Analysis for Anderson Acceleration. *SIAM Journal on Numerical Analysis*, 53(2):805–819, January 2015.
- [48] John N Tsitsiklis. Asynchronous stochastic approximation and q-learning. *Machine learning*, 16:185–202, 1994.

- [49] Eugene Vecharynski, Chao Yang, and John E Pask. A projected preconditioned conjugate gradient algorithm for computing many extreme eigenpairs of a hermitian matrix. *Journal of Computational Physics*, 290:73–89, 2015.
- [50] Alexander Weiß, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. The kernel polynomial method. *Reviews of modern physics*, 78(1):275, 2006.
- [51] Zaiwen Wen, Chao Yang, Xin Liu, and Yin Zhang. Trace-penalty minimization for large-scale eigenspace computation. *Journal of Scientific Computing*, 66:1175–1203, 2016.
- [52] Chao Yang, Juan C Meza, Byounghak Lee, and Lin-Wang Wang. Kssolv—a matlab toolbox for solving the kohn-sham equations. *ACM Transactions on Mathematical Software (TOMS)*, 36(2):1–35, 2009.
- [53] Hyobin Yoo, Rebecca Engelke, Stephen Carr, Shiang Fang, Kuan Zhang, Paul Cazeaux, Suk Hyun Sung, Robert Hovden, Adam W Tsen, Takashi Taniguchi, et al. Atomic and electronic reconstruction at the van der waals interface in twisted bilayer graphene. *Nature materials*, 18(5):448–453, 2019.
- [54] Yunkai Zhou, Yousef Saad, Murilo L Tiago, and James R Chelikowsky. Self-consistent-field calculations using chebyshev-filtered subspace iteration. *Journal of Computational Physics*, 219(1):172–184, 2006.