

Quantum solver for single-impurity Anderson models with particle-hole symmetry

M. Karabin,^{1,2,*} T. Sohail³, D. Bykov,³ E. A. Coello Pérez^{3,†}, S. Ghosh,³ M. Gopalakrishnan Meena³, S. Kim³,
A. Shehata³, I.-S. Suh³, H. Terletska,^{1,2} and M. Eisenbach³

¹Department of Physics and Astronomy, *Middle Tennessee State University*, Murfreesboro, Tennessee, USA

²Quantum Research Interdisciplinary Science and Education (QRISE) Center, *Middle Tennessee State University*, Murfreesboro, Tennessee, USA

³National Center for Computational Sciences, *Oak Ridge National Laboratory*, Oak Ridge, Tennessee, USA



(Received 16 January 2026; accepted 8 June 2026; published 16 July 2026)

Quantum embedding methods, such as dynamical mean-field theory (DMFT), provide a powerful framework for investigating strongly correlated materials. A central computational bottleneck in DMFT is in solving the Anderson impurity model (AIM), whose exact solution is classically intractable for large bath sizes. In this work, we benchmark a quantum-classical hybrid solver tailored for particle-hole symmetric AIMs, using the variational quantum eigensolver to prepare the ground state of the model with shallow quantum circuits. The solver uses shallow quantum ansätze and one set of variational parameters to prepare the ground state and its particle and hole excitations, enabling the construction of the impurity Green's function through a continued-fraction expansion. We evaluate the performance of this approach across a few bath sizes and interaction strengths under noisy, shot-limited conditions. We compare three optimization routines (COBYLA, Adam, and L-BFGS-B) in terms of convergence and fidelity, assess the benefits of estimating a quantum-computed moment correction to the variational energies, and benchmark the approach by comparing the density of states computed from the impurity Green's function against that obtained using a classical pipeline. Our results demonstrate the feasibility of Green's function construction on near-term devices and establish practical benchmarks for quantum impurity solvers embedded within self-consistent DMFT loops.

DOI: [10.1103/PhysRevResearch.8.033066](https://doi.org/10.1103/PhysRevResearch.8.033066)

I. INTRODUCTION

Chemistry and condensed matter physics are governed by the collective behavior of interacting electrons [1]. Since the direct numerical solution of the Schrödinger equation is computationally intractable for all but the smallest systems, many different approaches have been devised that employ various approximations to reduce the computational complexity of solving the many-electron problem [2]. Among these, density functional theory (DFT) [3,4] is a widely used and popular method that, while theoretically exact, relies on approximations to the exchange-correlation functional. These approximations often fail to capture strong local correlations, especially in systems exhibiting Mott physics or dynamical electron localization. One method to reintroduce electron correlation contributions missed by these approximations augments this approach with dynamical mean-field theory (DMFT) [5–9], which maps the lattice model onto a local Anderson impurity model (AIM) [10–12], providing the

mean field that describes the electrons in a lattice. The impurity model captures local dynamical fluctuations and becomes exact in the limit of infinite spatial dimensions. This mapping has been highly instrumental in describing correlated electron materials ranging from transition metal oxides to heavy fermion systems.

Solving the AIM within DMFT typically requires computationally expensive techniques such as exact diagonalization, continuous-time quantum Monte Carlo, or numerical renormalization group. Although accurate, these classical solvers scale poorly with bath size and may struggle with real-frequency or sign problems [6,13], limiting their application to more complex systems. Furthermore, these methods can also become computationally expensive at low temperatures [14,15], where quantum entanglement and nonperturbative effects like the Kondo phenomenon play a crucial role.

Quantum computing introduces approaches to solving impurity models with the potential to overcome some of these limitations, owing to its ability to efficiently represent large Hilbert spaces and capture strong electron correlations [8,16]. Some of these approaches construct a continued fraction representation of the model's Green's function [17,18] from quantum states that represent the ground state of the AIM and the particle and hole excitations of such a state at the impurity sites. The limitation of circuit depth in current quantum hardware makes the optimization of a shallow parametrized quantum circuit (PQC) via the variational quantum eigensolver (VQE) [19,20] a common strategy to prepare the

*Contact author: mariia.karabin@mtsu.edu

†Contact author: coellopereaa@ornl.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

AIM ground state in the noisy intermediate-scale quantum era [21].

In this work, we evaluate the performance of a quantum solver for AIMs with particle-hole symmetry that incorporates multiple quantum subroutines to construct the impurity Green's function. The focus on the half-filled case in the present work serves as an important benchmark for our approach as this case has been extensively studied in the context of understanding Mott insulators before its generalization for inclusion in a DFT + DMFT workflow. To solve these AIMs, we introduce a minimal yet expressive set of variational ansätze capable of preparing the model's ground state and the particle and hole excitations needed to compute the impurity Green's function using a common set of variational parameters. This approach eliminates the computational overhead associated with the separate optimization of circuits preparing excited states [18], or the additional measurement costs incurred by methods that rely on midcircuit measurement [22]. We benchmark variational results against exact diagonalization to assess the accuracy and robustness of this ground-state preparation method. The corresponding excited states are generated via structured parameter shifts applied to the optimized ground-state circuit, avoiding the need to reoptimize each excited state. The performance of the proposed solver is evaluated by computing the impurity Green's functions for multiple values of the Hubbard interaction and bath sizes. Finally, we analyze the computational overhead associated with different variational optimization strategies, including COBYLA, Adam, and L-BFGS-B, in the presence of finite sampling noise. These resource estimates provide realistic expectations for embedding this quantum solver into DMFT workflows on current quantum hardware.

The remainder of this paper is organized as follows. In Sec. II, we describe the methods used to calculate the continued-fraction representation of the impurity Green's function, present the ansatz architecture employed to approximate the ground and excited states of the AIM, and introduce the tested optimization routines. Section III reports our numerical results, including variational ground-state preparation, the fidelity and cost of optimization, moment-based energy corrections, and Green's function reconstruction. We conclude in Sec. IV with a summary of our findings and a discussion of future directions for integrating quantum impurity solvers within dynamical mean-field theory workflows.

II. METHODS

In this section, we outline the workflow developed to compute the impurity Green's function. We first define the Hamiltonian used to model the interacting impurity coupled to a noninteracting bath, followed by a description of the ansatz used to approximate its ground state via VQE and the methods used to optimize it. To improve the variational energy estimates, we discuss how to incorporate leading corrections derived from a cumulant expansion based on low-order Hamiltonian moments. Finally, we use our ansätze to approximate particle and hole excitations and the parameters of the optimized ground-state circuit, and describe the Krylov subspace procedure used to construct the continued-fraction representation of the impurity Green's function.

A. The impurity model

This work focuses on single-impurity AIMs defined by the Hamiltonian [10]

$$\begin{aligned} H_{\text{AIM}} &= H_{\text{imp}} + H_{\text{bath}} + H_{\text{coup}}, \\ H_{\text{imp}} &= \sum_{\sigma} \epsilon_0 n_{0\sigma} + U n_{0\uparrow} n_{0\downarrow}, \\ H_{\text{bath}} &= \sum_{k=1}^{N_b} \sum_{\sigma} \epsilon_k n_{k\sigma}, \\ H_{\text{coup}} &= \sum_{k=1}^{N_b} \sum_{\sigma} V_k (c_{0\sigma}^{\dagger} c_{k\sigma} + c_{0\sigma} c_{k\sigma}^{\dagger}), \end{aligned} \quad (1)$$

where ϵ_0 and ϵ_k are the single-particle energies of the impurity and the k th bath site, respectively, U is the Hubbard interaction, V_k is the hybridization between the impurity and the k th bath sites, $N_b \in \{1, 3, 5, 7\}$ is the number of bath sites, and $\sigma \in \{\uparrow, \downarrow\}$ is the projection of the electron's spin. The operators $c_{i\sigma}^{\dagger}$ and $c_{i\sigma}$ create and annihilate an electron with spin projection σ in site i . The number operators $n_{i\sigma}$ count the number of σ electrons in orbital i .

In this work, we impose particle-hole symmetry on the Hamiltonian by setting $\epsilon_0 = -U/2$, ensuring that the system is half filled, and its ground state has spin projection $S_z = 0$ in the absence of an external magnetic field, and set $V_k = V$ for all k .

These impurity models form the foundation of DMFT, which maps the original interacting lattice problem onto a self-consistently defined quantum impurity coupled to a noninteracting bath. This mapping retains all local interaction effects while approximating nonlocal correlations. The interaction between the impurity and the bath is described by the hybridization function, which reflects the dynamical coupling to the environment and must be updated iteratively through the DMFT loop. In practical implementations, this continuous bath is discretized into a finite number of orbitals, making the bath size a crucial control parameter. As the bath size increases, the model captures finer dynamical features, but also becomes increasingly challenging to solve. These challenges motivate the development of quantum impurity solvers [23,24] capable of efficiently handling large Hilbert spaces and providing direct access to real-frequency observables.

Solving the impurity problem in DMFT involves computing the impurity Green's function [5,6], which as a function of the energy ω takes the form

$$\begin{aligned} G_{ij}(\omega) &= \langle \psi_0 | c_i \frac{1}{\omega - H_{\text{AIM}} + E_0} c_j^{\dagger} | \psi_0 \rangle \\ &+ \langle \psi_0 | c_j^{\dagger} \frac{1}{\omega + H_{\text{AIM}} - E_0} c_i | \psi_0 \rangle. \end{aligned} \quad (2)$$

Here, the ground state ψ_0 of the AIM Hamiltonian (1) solves the Schrödinger equation

$$H_{\text{AIM}} |\psi_0\rangle = E_0 |\psi_0\rangle. \quad (3)$$

Therefore, computing the impurity Green's function involves finding the ground state of the AIM (1), calculating its energy, and evaluating the matrix elements on the right-hand side of Eq. (2).




FIG. 1. Spin-block ordering used to map the model's spin-up and spin-down sites onto a linear chain of qubits.

B. Variational ground-state approximation

Our quantum solver employs VQE to optimize a PQC minimizing the variational energy

$$E_{\text{VQE}}(\theta) = \min_{\theta} \langle \psi(\theta) | H_{\text{AIM}} | \psi(\theta) \rangle, \quad (4)$$

to find an approximation $|\theta\rangle$ to the model's ground state $|\psi_0\rangle$. To this end, the Hamiltonian (1) is recast in terms of spin operators using the Jordan-Wigner transformation [25,26],

$$\begin{aligned} c_i^\dagger &= \prod_{j=0}^{i-1} Z_j \otimes \frac{1}{2}(X_i - iY_i) \otimes \cdots, \\ c_i &= \prod_{j=0}^{i-1} Z_j \otimes \frac{1}{2}(X_i + iY_i) \otimes \cdots, \end{aligned} \quad (5)$$

where X_i , Y_i , and Z_i are Pauli matrices acting on the space of qubit i , which encodes the occupation of a single spin orbital in the model.

Using the orbital-to-qubit mapping depicted in Fig. 1 to assign qubits along a linear chain to the model's spin orbitals leads to the expression

$$\begin{aligned} H_{\text{AIM}} = C &- \sum_{k=1}^{N_b} \frac{\epsilon_k}{2} (Z_{N_b-k} + Z_{N_b+1+k}) + \frac{U}{4} Z_{N_b} Z_{N_b+1} \\ &+ \frac{V}{2} \sum_{k=1}^{N_b} X_{N_b-k} \left(\prod_{j=N_b+1-k}^{N_b-1} Z_j \right) X_{N_b} \\ &+ \frac{V}{2} \sum_{k=1}^{N_b} Y_{N_b-k} \left(\prod_{j=N_b+1-k}^{N_b-1} Z_j \right) Y_{N_b} \\ &+ \frac{V}{2} \sum_{k=1}^{N_b} X_{N_b+1+k} \left(\prod_{j=N_b+2}^{N_b+k} Z_j \right) X_{N_b+1+k} \\ &+ \frac{V}{2} \sum_{k=1}^{N_b} Y_{N_b+1+k} \left(\prod_{j=N_b+2}^{N_b+k} Z_j \right) Y_{N_b+1+k}, \end{aligned} \quad (6)$$

where $C = -U/4 + \sum_{k=1}^{N_b} \epsilon_k$.

The expectation value (4) is calculated adding the expectation values of each term in the Pauli-strings' decomposition of

the Hamiltonian, which can be estimated on a quantum chip by measuring the state $\psi(\theta)$ in the N_{group} bases corresponding to the commuting groups into which the Pauli strings are organized. The absolute error associated with the finite sampling of the state is estimated as

$$\varepsilon \sim \mathcal{O}(\sqrt{N_{\text{Pauli}}/N_{\text{shots}}}), \quad (7)$$

where N_{Pauli} and N_{shots} denote the number of Pauli strings defining the Hamiltonian and the number of measurements taken on each basis [27,28].

The symmetries of the ground state arising from that imposed on the Hamiltonian allow us to propose a highly constrained PQC as an ansatz for the ground state. The first component of this PQC, shown at the top left of Fig. 2, acts on the impurity qubits initializing them to

$$\begin{aligned} |\psi(\theta)\rangle &= \sqrt{\frac{1}{2}} \left\{ \cos\left(\frac{\theta}{2}\right) (|00\rangle + |11\rangle) \right. \\ &\quad \left. + \sin\left(\frac{\theta}{2}\right) (|01\rangle + |10\rangle) \right\}, \end{aligned} \quad (8)$$

a state that is symmetric under spin exchange.

Next, the half-filling gate at the top center of Fig. 2, which “creates” an electron in a spin-orbital pair, is used to prepare an initial state with the proper number of electrons. The most general state with this electron number can be prepared by acting the right number of Givens rotation gates, shown at the top right of Fig. 2, on the initial state [29]. In this work, however, we restrict the number of Givens rotations to those that can be implemented within two layers after the initial state preparation (see, for example, the three-bath ansatz at the bottom right of Fig. 2), thus constraining the PQC's local depth in an attempt to mitigate barren plateaus [30]. More expressive symmetry-preserving ansätze including additional entangling structures or repetitions of the variational layer have been proposed and used to accurately prepare AIM ground states [22]. Here, we adopt a more constrained variant to maintain shallow depths that could help with its optimization on present-day noisy hardware, and allow for the preparation of particle and hole excitations from the same parameter set, eliminating additional optimization costs.

To enforce spin-exchange symmetry, we apply identical unitaries to the top and bottom registers. The state prepared by the PQC is

$$\begin{aligned} |\psi(\theta)\rangle &= \sum_a c_{a00a}(\theta) |a\rangle \otimes |00\rangle \otimes |a\rangle \\ &+ \sum_a c_{a11a}(\theta) |a\rangle \otimes |11\rangle \otimes |a\rangle \\ &+ \sum_a \sum_b c_{a01b}(\theta) |a\rangle \otimes |01\rangle \otimes |b\rangle \\ &+ \sum_a \sum_b c_{b10a}(\theta) |b\rangle \otimes |01\rangle \otimes |a\rangle, \end{aligned} \quad (9)$$

where a, b are bitstrings representing basis states of the bath registers, and $c_{aijb} = c_{bjia}$. In particular, for the one-bath case




FIG. 2. Building blocks used to construct symmetry-preserving ansätze (top), and ansätze for the ground states of particle-hole symmetric AIMs with one, three, five, and seven bath sites (bottom). These ansätze first prepare a spin-exchange symmetric impurity state (with the top left circuit), half fill the rest of the orbitals (with the top center one), and implements Givens rotation (via the top right circuit) to prepare states with the proper number of electrons.

we have

$$|\psi(\theta_0)\rangle = \sqrt{\frac{1}{2}} \left\{ \cos\left(\frac{\theta_0}{2}\right) (|1001\rangle + |0110\rangle) + \sin\left(\frac{\theta_0}{2}\right) (|1010\rangle + |0101\rangle) \right\}. \quad (10)$$

The ansätze for the ground states of single-impurity AIMs with one, three, five, and seven bath sites are shown at the bottom of Fig. 2. Notice that the ansatz for the one-bath case is embedded within those for the ground states of AIMs with more than one bath site.

C. Optimization methods

The estimation of the variational energy depends on finding the parameters that best approximate the ground state, making the choice of optimizer critical for the efficient and accurate computation of this observable. Different optimization strategies navigate the parameter landscape in distinct ways, each presenting trade-offs between convergence speed, sensitivity to sampling noise, and computational cost. To identify the most effective approach, we examined three representative optimization strategies employing gradient-free, first-order, and quasi-second-order methods.

The constrained optimization by linear approximation (COBYLA) gradient-free method [31] constructs a local linear model of the objective function by evaluating it at a dynamical number of parameter points located within a hypersphere of radius ρ_0 , set in this work to π , centered at the initial point. The parameters that minimize this model are then chosen as a new center, and the procedure is repeated for a new radius, $\rho_1 < \rho_0$. Hence, gradient-free optimizers have the

lowest measurement cost per evaluation, requiring only the N_{group} circuit measurements needed to estimate the energy.

Adam was selected as a representative first-order optimizer [32]. Gradient-based methods update the parameters along the direction of the negative gradient, estimated on a quantum computer using the parameter-shift rule,

$$\frac{\partial E_{\text{VQE}}}{\partial \theta_i} \bigg|_{\theta} = \frac{1}{2} [E_{\text{VQE}}(\theta_i^+) - E_{\text{VQE}}(\theta_i^-)], \quad (11)$$

with $\theta_i^\pm$ a set of parameters in which the i th component of θ is set to $\theta_i \pm \pi/2$. Although the number of measurements per evaluation, $N_{\text{group}}(1 + 2N_{\text{params}})$, is substantially higher than that for gradient-free methods, first-derivative information may help these optimizers converge the VQE in fewer overall evaluations. Our Adam optimizations used an initial learning rate $\eta = 1$, standard hyperparameters $\{\beta_1, \beta_2, \epsilon\} = \{0.9, 0.999, 10^{-8}\}$, and initialized all momentum terms to zero.

For comparison, we also tested the quasi-second-order limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS-B) algorithm [33], which approximates the inverse Hessian from a limited number of past gradient estimates. In this context, the number of measurements per evaluation is that of a first-order gradient-based method.

Convergence for all optimizers was defined as either an absolute estimate change $|E_{\text{VQE}}^{(n)} - E_{\text{VQE}}^{(n-1)}| < 10^{-6}$, with $E_{\text{VQE}}^{(n)}$ the energy estimated at iteration n , or a gradient norm $\|\nabla_{\theta} E_{\text{VQE}}^{(n)}\| < 10^{-9}$.

D. Quantum computed moments approach

A key limitation of the VQE is that the extent to which the state $\psi(\theta)$ and the estimated energy E_{VQE} approximate the true ground state and its energy strongly depends on the

expressive power of the chosen ansatz and the convergence of the optimization routine. The latter is severely affected by errors in the expectation value and the presence of barren plateaus.

The quantum computed moments (QCM) approach proposed in Ref. [34] addresses this limitation by estimating a correction to the variational energy from higher-order Hamiltonian moments $\mu_m = \langle \psi(\theta) | H^m | \psi(\theta) \rangle$. In terms of a set of cumulants c_m , defined as

$$c_m = \mu_m - \sum_{p=0}^{m-2} \binom{m-1}{p} c_{p+1} \mu_{m-1-p}, \quad (12)$$

the Lanczos coefficients [35] defining the tridiagonal Hamiltonian in the Krylov basis built on $\psi(\theta)$,

$$H_{\text{AIM}} = \begin{pmatrix} a_0 & b_1 & 0 & \dots \\ b_1 & a_1 & b_2 & \\ 0 & b_2 & a_3 & \\ \vdots & & & \ddots \end{pmatrix}, \quad (13)$$

can be written as expansions in powers of $1/N$ [36],

$$a_\ell = c_1 + \ell \left(\frac{c_3}{c_2} \right) \frac{1}{N} + \mathcal{O}(N^{-2}),$$

$$b_\ell^2 = \ell c_2 + \frac{1}{2} \ell(\ell-1) \left(\frac{c_2 c_4 - c_3^2}{2c_2^2} \right) \frac{1}{N} + \mathcal{O}(N^{-2}), \quad (14)$$

with N the size of the system. For large ℓ and N , these expansions become series in powers of $z = \ell/N$, in terms of which the ground-state energy can be calculated as the infimum of $a(z) - 2b(z)$ for $z > 0$ [37,38]. At first order in z , this infimum takes the approximate form

$$E_{\text{INF}} \approx c_1 - \frac{c_2^2}{c_3^2 - c_2 c_4} \left(\sqrt{3c_3^2 - 2c_2 c_4} - c_3 \right). \quad (15)$$

Notice that this approximation, although no longer a strict upper bound, incorporates corrections from higher-order Hamiltonian moments in its second term, yielding a more accurate estimate of the true ground-state energy than the variational result alone. From a computational perspective, this approach shifts part of the problem complexity away from deep, highly expressible circuits to the measurement of higher-order observables.

Corrections to the ground-state energy are particularly relevant within the context of DMFT, as the accurate determination of this quantity is important to stabilize the self-consistency loop. By incorporating higher Hamiltonian moments, one can potentially mitigate errors due to shallow ansatz, sampling noise, and barren plateaus.

E. Green's function as continued fractions

In order to extract physically meaningful observables from the quantum-computed ground and excited states of the AIM, we reconstruct the impurity Green's function in the real-frequency domain using a Krylov subspace-based continued-fraction expansion.

Since the model lacks coupling between spin-up and spin-down orbitals, the off-diagonal elements of the impurity Green's function (2) vanish identically. The diagonal matrix




FIG. 3. Ansätze for particle (left) and hole (right) excitations of the ground state of AIMs with one bath site, preparing the states $\mathcal{N} c_{i\sigma}^\dagger |\text{GS}\rangle$ and $\mathcal{N} c_{i\sigma} |\text{GS}\rangle$ with spin projections $\sigma = \uparrow$ (top) and $\sigma = \downarrow$ (bottom).

elements can be expressed as continued fractions [39]

$$G_{ii}(\omega) = \frac{1}{\omega + E_0 - a_0^p - \frac{(b_1^p)^2}{\omega + E_0 - a_1^p - \dots}} + \frac{1}{\omega - E_0 + a_0^h - \frac{(b_1^h)^2}{\omega - E_0 + a_1^h - \dots}}, \quad (16)$$

where a_n^p and b_n^p are the Lanczos coefficients of H_{AIM} on the Krylov basis built on top of $c_{i\downarrow}^\dagger |\psi_0\rangle$, and a_n^h and b_n^h the corresponding coefficients on the basis built on top of $c_i |\psi_0\rangle$.

The ansätze in Fig. 3 can be used to prepare normalized versions of these particle and hole excitations. The PQC at the top of the figure prepare the states

$$|\phi(\theta_0)\rangle = \cos\left(\frac{\theta_0}{2}\right) |0111\rangle + \sin\left(\frac{\theta_0}{2}\right) |1011\rangle,$$

$$|\chi(\theta_0)\rangle = \cos\left(\frac{\theta_0}{2}\right) |1000\rangle + \sin\left(\frac{\theta_0}{2}\right) |0100\rangle, \quad (17)$$

with one and three electrons. For the parameter $\pi - \theta_0$,

$$|\phi(\pi - \theta_0)\rangle = \mathcal{N} c_{i\uparrow}^\dagger |\psi(\theta_0)\rangle,$$

$$|\chi(\pi - \theta_0)\rangle = \mathcal{N} c_{i\uparrow} |\psi(\theta_0)\rangle, \quad (18)$$

with $|\psi(\theta_0)\rangle$ the one-bath state in Eq. (10) and $\mathcal{N}$ a normalization constant. Similarly, the PQC at the bottom of Fig. 3 prepare states proportional to $c_{i\downarrow}^\dagger |\psi(\theta_0)\rangle$ and $c_{i\downarrow} |\psi(\theta_0)\rangle$.

Substituting the one-bath ground-state ansatz embedded in the ansätze for the ground states of AIMs with more than one bath site, with the PQC shown in Fig. 3, and replacing θ with a new parameter set in which θ_0 is replaced by $\pi - \theta_0$, yields the corresponding particle and hole excitations $\mathcal{N} c_{i\sigma}^\dagger |\psi(\theta)\rangle$ and $\mathcal{N} c_{i\sigma} |\psi(\theta)\rangle$, with $\sigma \in \{\uparrow, \downarrow\}$.

Using the ansätze for the particle and hole excitations, we compute the Lanczos coefficients in Eq. (16) as [40,41]

$$a_n = \frac{\Delta'_{n-2}\Delta_n}{\Delta'_{n-1}\Delta_{n-1}} + \frac{\Delta'_n\Delta_{n-1} - 1}{\Delta'_{n-1}\Delta_n},$$

$$b_n^2 = \frac{\Delta_n\Delta_{n-2}}{\Delta_{n-1}^2}, \quad (19)$$

with

$$\Delta_n = \begin{vmatrix} \mu_0 & \mu_1 & \dots & \mu_n \\ \mu_1 & \mu_2 & & \\ \vdots & & \ddots & \\ \mu_n & & & \mu_{2n} \end{vmatrix},$$

$$\Delta'_n = \begin{vmatrix} \mu_1 & \mu_2 & \dots & \mu_{n+1} \\ \mu_2 & \mu_3 & & \\ \vdots & & \ddots & \\ \mu_{n+1} & & & \mu_{2n+1} \end{vmatrix} \quad (20)$$

Hankel determinants. Since the number of Pauli strings, and therefore the number of quantum circuits, involved in the estimation of μ_m increases rapidly with m , we estimate Lanczos coefficients up to order $n = 2$, corresponding to the computation of moments up to μ_5 .

Preparing particle and hole excitations enables the direct evaluation of the moments μ_m , rather than the indirect ground-state-based approach in which powers of the Hamiltonian are dressed with creation and annihilation operators [17]. This construction reduces the number of Pauli strings entering the moments' estimation and the number of shots required to achieve a target absolute error ϵ .

F. Classical reference calculations

To validate our quantum ansatz, we also perform classical calculations of the Green's function. For the model sizes investigated in the present paper, exact diagonalization of the Hamiltonian is easily accomplished without the need for high-performance computing resources. We follow the procedure outlined in Sec. II E to represent the Green's function as a continued fraction with its coefficients being determined using Lanczos iteration. For this, we first construct a basis for the ground-state calculations, which consists of all sets of creation operators that preserve the total charge and spin for the half-filled system considered here, i.e., a total of $2^{\binom{N_b+1}{(N_b+1)/2}}$ of basis states. The Hamiltonian for the AIM can be represented as a sparse matrix and we first find the lowest eigenvalue state $|\psi_0\rangle$ with the ground-state energy E_0 . Using this, we can apply creation or annihilation operators to this vector and follow Eq. (2) to calculate the many-body Green's function using the continued fraction in Eq. (16), where the coefficients can be obtained using Lanczos iteration [39,42].

III. RESULTS

The results presented in this section, used to assess the performance of the proposed quantum solver, were obtained from classical simulations of the quantum subroutines involved in the computation of the impurity Green's function (2). These simulations solved the AIM described in Sec. II A, with




FIG. 4. Best absolute relative error between true ground-state energies and variational estimates for them obtained from noiseless simulations. For models with seven bath sites, the proposed ansatz can yield energy estimates accurate to within about 5% of the true ground-state energy.

Hubbard interactions $U \in \{2, 4, 6, 8\}$. To account for the potential for the optimizer to become trapped in local minima, the VQE results for AIMs with one, three, and five bath sites reported here are averaged over simulations initiated from ten distinct starting points. Moreover, each optimization was repeated five times to mitigate statistical errors arising from the finite number of shots used to evaluate the variational estimates. For the more computationally demanding seven-bath case, simulations were performed using a single initialization, with two optimization repeats. The scalability of the solver with respect to the number of bath sites is assessed by averaging results over the different values of U , initializations, and optimization repetitions. As discussed in Sec. II C, the ansatz was optimized using a gradient-free method (COBYLA), a first-order gradient-based method (Adam), and a quasi-second-order gradient-based method (L-BFGS-B) to compare the performance of different optimization strategies.

A. Ground-state optimization

The ansatz proposed for the ground state does not span the full n -qubit Hilbert space and therefore may struggle to approximate the true ground state of the AIM as the number of bath sites increases. To assess the performance of our ansatz, we compared ground-state energies obtained via classical diagonalization with those produced by noiseless VQE simulations. Figure 4 shows the absolute relative errors $\delta E \equiv \text{abs}([E_{\text{ideal}} - E_{\text{exact}}]/E_{\text{exact}})$ produced by these simulations. For AIMs with five and seven bath sites, the ansatz was capable of yielding variational energies that deviate from the exact ground-state energy by approximately 5%.

In practice, the variational energy is estimated from a finite number of measurements, causing sampling noise that may impede convergence to the best achievable approximation to the ground state. To quantify this effect, we performed VQE simulations with a varying number of shots for AIMs with one and three bath sites and Hubbard interaction $U = 2$. The resulting absolute relative errors, $\delta E \equiv \text{abs}([E_{\text{samp}} - E_{\text{ideal}}]/E_{\text{ideal}})$, are shown in Fig. 5.




FIG. 5. Averaged absolute relative error between converged variational energies in the presence and absence of sampling noise. Simulations solving AIMs with on-site energy $U = 2$ and one (top) or three (bottom) bath sites show a decrease of the relative error with the number of shots. Means and standard deviations result from simulations initiated from ten distinct starting points, each repeated five times to account for optimization difference due to sampling noise.

These values show that for all three optimizers, the relative energy error decreases as the number of shots per commuting group increases. The rate and stability of this improvement depend strongly on whether gradients are used in the optimization. Across the range of shots, COBYLA exhibits large errors, indicating that sampling noise can cause the gradient-free search to get trapped in local minima or terminate early. In contrast, gradient-based methods can achieve higher accuracy, but doing so requires information about second derivatives. For the best-performing optimization method (L-BFGS-B), relative differences of roughly 2% are achieved in the one- and three-bath cases from 100 and 2500 measurements per quantum circuit, respectively. This behavior highlights that reliable gradient information and accurate Hessian approximations are crucial to avoid noise-induced convergence to suboptimal states.

Using relation (7), we estimate that absolute errors $\epsilon \sim 0.1$ on each energy estimate can be achieved with 700, 2500, 5000, and 7500 shots per commuting group for AIMs with one, three, five, and seven bath sites, respectively.




FIG. 6. Comparison between averaged total number of shots required by different optimization methods to converge VQEs solving single-impurity AIMs with $N_b \in \{1, 3, 5, 7\}$.

B. Optimization performance, fidelity, and cost

We use the number of shots required to converge the VQE as a metric of the subroutine's quantum cost. Figure 6 shows the scaling of this cost with the number of bath sites for the three different optimization methods. The costs presented there, which are averages over all runs that solved AIMs with the corresponding number of bath sites, highlight the differences in computational expense between the tested optimization methods.

The COBYLA optimizer measures the lowest number of shots per evaluation and thus exhibits the best scaling with N_b . Unfortunately, without the computation of a gradient pushing in the direction of maximum change, the objective function tends to get stuck in local minima, as reflected by the larger δE values for this method in Fig. 5. The gradient-based Adam and L-BFGS-B optimizers substantially increase the number of circuit measurements per evaluation, leading to worse scaling than that of COBYLA. Furthermore, results in Fig. 5 from runs using the Adam optimizer suggest that sampling errors in the gradient estimate can push the optimization toward local minima, yielding larger δE values. This effect seems to be mitigated via the estimation of the inverse Hessian carried on by the L-BFGS-B optimizer.

These cost differences are considered together with the accuracy results discussed in Sec. III A. In ideal simulations, the proposed ansatz produces ground-state energies that differ from the exact values by at most a few percent for $N_b \leq 7$, as shown in Fig. 4, indicating that the ansatz retains a reasonable ground-state fidelity across the range of bath sizes studied in this work. In the presence of sampling noise, COBYLA is more prone to being trapped in shallow local minima, especially at low shot counts. This leads to larger energy errors despite its lower circuit cost. The gradient-based optimizers make more efficient use of the measurement statistics. For a fixed target accuracy, they achieve smaller relative energy errors at moderate shot numbers, with L-BFGS-B providing the most favorable balance between accuracy and quantum cost.

These results highlight a trade-off between circuit depth, optimizer robustness, and measurement overhead. Gradient-

TABLE I. Number of Pauli strings in H^n up to $n = 4$, number of commuting groups in H^n with $n \in \{4, 5\}$, and number of shots per group required to achieve absolute errors $\epsilon \sim 10^{-1}$ on estimates for μ_n with $n \in \{4, 5\}$.

N_b	N_{Pauli}			N_{group}		N_{shots}	
	H	H^4	H^5	H^4	H^5	μ_4	μ_5
1	6	23	24	2	2	2K	2K
3	18	1190	2146	60	63	110K	185K
5	30	10 995	33 672	278	597	1M	3M
7	42	47 000	212 222	859	2143	4M	18M

free methods minimize the raw number of circuits but can incur a significant accuracy penalty, particularly when measurements are noisy. Gradient-based methods require more circuits but deliver higher-fidelity approximations to the impurity ground state.

C. Energy correction estimates

As discussed in Sec. II B, the variational energy can be refined by incorporating higher-order Hamiltonian moments measured in the optimal state. Starting from the moments μ_n up to $n = 4$, we used the approximation to the infimum (15). This approximation corrects the raw variational estimate c_1 with fluctuations encoded in c_2 , c_3 , and c_4 , which capture increasingly nontrivial dynamical correlations beyond those described by the shallow ansatz alone.

For consistency, the number of shots used to estimate μ_4 is increased to achieve an absolute error of $\epsilon \sim 10^{-1}$. These numbers are listed in Table I. From those, the total number of shots required to compute the correction to the variational energy are approximately 4.0×10^3 , 6.6×10^6 , and 2.8×10^8 for the one-, three-, and five-bath cases.

The AIMs considered here provide a controlled setting to assess the effectiveness of this correction, since the exact ground-state energies are available from classical diagonalization. For each VQE run, we compute the absolute relative error, both for the bare variational estimate E_{VQE} and for the infimum approximation E_{INF} . As shown in Fig. 7, the correction estimate systematically reduces the absolute relative error introduced by both ansatz limitations and sampling noise. For the gradient-free optimal states, the correction yields a substantial reduction for both the three- and five-bath cases, indicating that higher-order moments can partially recover correlation energy that the variational state fails to capture. Gradient-based optimizations already yield smaller variational errors, particularly L-BFGS-B, yet, they also benefit from the correction. The largest relative gains appear in the larger-bath regimes. Computation of the QCM correction requires additional measurements to estimate higher-order Hamiltonian moments. For the three-bath case, this overhead is comparable to the VQE optimization cost for the L-BFGS-B optimizer, as shown in Fig. 6. In the five-bath case, computing this correction is more expensive than solving the variational problem.

Since the QCM energy is not constrained to the upper bound, it can overshoot the energy for individual runs. However, the cumulant-based estimate provides, on average, a closer approximation to the exact ground-state energy. In the DMFT context, this improvement in the impurity ground-state energy is particularly valuable, as it can improve the stability of the self-consistency loop without requiring deeper ansatz for the ground state.

D. Green's function construction

To assess the accuracy with which the proposed solver can reconstruct the impurity Green's function, we estimated particle and hole Lanczos coefficients (14) from Hamiltonian




FIG. 7. Average absolute relative error δE between the exact ground-state energy and its variational estimate for AIMs with three (left), five (center), and seven (right) bath sites. Bars compare the variational energies obtained with the COBYLA, Adam, and L-BFGS-B methods to the corresponding QCM-corrected estimates (“+QCM” labels). Each bar height is averaged over all interaction strengths and random initializations.

moments computed on the states prepared by the particle and hole excitations. These estimates were computed for every parameter set found by simulated VQE runs optimized via the L-BFGS-B algorithm, and averaged to obtain mean values for the Lanczos coefficients for a given number of bath sites and Hubbard interaction.

The resulting Lanczos coefficients were used to construct an averaged continued-fraction representation of the Green's function for each AIM. From this representation, we then computed the corresponding density of states (DOS),

$$\text{DOS}(\omega) = -\frac{1}{\pi} \text{Im}\{\text{Tr}[G(\omega)]\}, \quad (21)$$

and compared it with that obtained from the exact Lanczos coefficients computed classically from the particle and hole excitations of the ground state determined via classical diagonalization.

This comparison is shown in Fig. 8. The peaks in the DOS for one-bath AIMs, shown at the top of the figure, are reproduced with very good accuracy for all values of the Hubbard interaction, as shown in the top row of the figure. As the interaction strength increases, the spectrum progressively splits into lower- and upper-Hubbard-like features, and the average DOS tracks the relative weight of these peaks across all interaction values.

Results for three-bath AIMs, shown in the middle row of Fig. 8, capture the main features of the exact DOS over the full frequency window. This shows that, once a reasonably accurate variational ground state has been obtained, the combination of excited-state ansatz and Lanczos-based continued fractions yields a reasonable and compact representation of the impurity Green's function.

Due to the significantly higher simulation cost, the five-bath case averages were evaluated using a reduced number of configurations. Specifically, the optimization was repeated only twice per initialization to gather statistical information. These results provide valuable insight into the applicability of our ansatz and VQE-based reconstruction to larger systems. Results for this case, shown at the bottom of Fig. 8, highlight how deviations from the true ground state lead to poor agreement with the true DOS for some values of the Hubbard interaction. Improving the construction of the Green's function for larger bath cases will prepare states closer to the true ground state for the corresponding AIMs.

The additional measurements needed to estimate the Lanczos coefficients and evaluate the continued fractions result in a quantum cost comparable to that of the VQE optimization in the three-bath case. For larger AIMs, however, this subroutine becomes the dominant cost, as the number of Pauli strings involved in the estimation of μ_5 grows rapidly with the number of bath sites.

IV. OUTLOOK

In this work, we evaluated the performance of a hybrid quantum-classical solver computing the impurity Green's function of particle-hole symmetric AIMs. The solver uses shallow symmetry-preserving ansätze to prepare the ground state of the model and its particle and hole excitations from

the same set of variational parameters, enabling the construction of the continued-fraction representation of the Green's function via the estimation of Hamiltonian moments on the excited states.

Our comparison of different optimization strategies highlights practical cost-performance trade-offs, that will become increasingly relevant as the complexity of the AIM grows. Simulations show that a “good” variational approximation to the model's true ground state leads to a DOS, computed from the constructed Green's function, in good agreement with that obtained from a classical pipeline. As the approximation to the true ground state deteriorates with the number of bath sites due to constraints imposed on the ansatz, the fidelity of the computed DOS decreases with this number accordingly.

Therefore, improving the accuracy with which the true ground state is approximated via refinement of the ansatz and implementation of more efficient optimization strategies is paramount to further enhance the fidelity of the constructed Green's function, particularly for larger and more complex models for which the impurity site represents highly correlated orbitals. The proposed ansatz can be modified to prepare states with other electron numbers, enabling the treatment of systems beyond the half-filled case. Moreover, increasing the number of Givens rotation layers will allow the ansatz to prepare more general states, albeit at the cost of increased circuit depth, which may exacerbate the barren plateau problem. An additional route to improved fidelity is adopting a more expressive ansatz in which the impurity block includes additional entangling operators (e.g., ZZ-type interactions), as explored in symmetry-preserving hardware-adaptable circuits for AIMs [22], or the bath block includes double-excitation operators. In this work, we focus on a constrained, shot-limited benchmark, and we treat the accuracy-cost trade-off of such augmented ansätze as a natural follow-on study.

A relevant direction for future work is the integration of the quantum impurity solver into a one-shot DFT + DMFT framework, followed by its extension to a fully self-consistent scheme. Realizing this program will require improved scalability of the solver's subroutines, together with enhanced robustness via error-mitigation strategies suited to noisy intermediate-scale quantum devices. Progress along these directions would push forward toward simulations of correlated real materials with a quantum-accurate treatment of local electron interactions.

ACKNOWLEDGMENTS

This work has been partially supported by U.S. DOE Grant No. DE-FG02-13ER41967. ORNL is managed by UT-Battelle, LLC, under Contract No. DE-AC05-00OR22725 for the U.S. Department of Energy.

The U.S. government retains, and the publisher, by accepting the article for publication, acknowledges that the U.S. government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. government purposes. DOE will provide public access to these




FIG. 8. Comparison between DOS constructed using the proposed quantum solver (green solid lines) and the classical pipeline (black dashed lines). The shown density of states corresponds to AIMs with one (top), three (middle), and five (bottom) bath sites, with Hubbard interactions $U = 2$ (left) and $U = 8$ (right).

results of federally sponsored research in accordance with the DOE Public Access Plan [43].

DATA AVAILABILITY

The data that support the findings of this article are openly available [44].

- [1] P. Fulde, *Electron Correlations in Molecules and Solids*, 3rd ed. (Springer-Verlag, Berlin, 1995), p. 483.
- [2] M. Rossmannek, F. Pavošević, A. Rubio, and I. Tavernelli, Quantum embedding method for the simulation of strongly correlated systems on quantum computers, *J. Phys. Chem. Lett.* **14**, 3491 (2023).
- [3] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [4] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
- [5] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Electronic structure calculations with dynamical mean-field theory, *Rev. Mod. Phys.* **78**, 865 (2006).
- [6] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* **68**, 13 (1996).
- [7] D. Zgid and G. K.-L. Chan, Dynamical mean-field theory from a quantum chemical perspective, *J. Chem. Phys.* **134**, 094115 (2011).
- [8] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, Hybrid quantum-classical approach to correlated materials, *Phys. Rev. X* **6**, 031045 (2016).
- [9] I. Rungger, N. Fitzpatrick, H. Chen, C. Alderete, H. Apel, A. Cowtan, A. Patterson, D. Muñoz Ramo, Y. Zhu, N. Nguyen, E. Grant, S. Chretien, L. Wossnig, N. Linke, and R. Duncan, Dynamical mean field theory algorithm and experiment on quantum computers, [arXiv:1910.04735](https://arxiv.org/abs/1910.04735).
- [10] P. W. Anderson, Localized magnetic states in metals, *Phys. Rev.* **124**, 41 (1961).
- [11] T. Keen, T. Maier, S. Johnston, and P. Lougovski, Quantum-classical simulation of two-site dynamical mean-field theory on noisy quantum hardware, *Quantum Sci. Technol.* **5**, 035001 (2020).
- [12] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Comput. Mater.* **6**, 85 (2020).
- [13] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Continuous-time Monte Carlo methods for quantum impurity models, *Rev. Mod. Phys.* **83**, 349 (2011).
- [14] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [15] R. Bulla, T. A. Costi, and T. Pruschke, Numerical renormalization group method for quantum impurity systems, *Rev. Mod. Phys.* **80**, 395 (2008).
- [16] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [17] G. Greene-Diniz, D. Z. Manrique, K. Yamamoto, E. Plekhanov, N. Fitzpatrick, M. Krompiec, R. Sakuma, and D. M. Ramo, Quantum computed Green's functions using a cumulant expansion of the Lanczos method, *Quantum* **8**, 1383 (2024).
- [18] F. Jamet, A. Agarwal, C. Lupo, D. E. Browne, C. Weber, and I. Rungger, Krylov variational quantum algorithm for first principles materials simulations, [arXiv:2105.13298](https://arxiv.org/abs/2105.13298).
- [19] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [20] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [21] B. Jaderberg, A. Agarwal, K. Leonhardt, M. Kiffner, and D. Jaksch, Minimum hardware requirements for hybrid quantum-classical dmft, *Quantum Sci. Technol.* **5**, 034015 (2020).
- [22] E. B. Jones, C. J. Winkleblack, C. Campbell, C. Rotello, E. D. Dahl, M. Reynolds, P. Graf, and W. Jones, Dynamic, symmetry-preserving, and hardware-adaptable circuits for quantum computing many-body states and correlators of the Anderson impurity model, *Phys. Rev. Res.* **7**, 023186 (2025).
- [23] R. Sakurai, W. Mizukami, and H. Shinaoka, Hybrid quantum-classical algorithm for computing imaginary-time correlation functions, *Phys. Rev. Res.* **4**, 023219 (2022).
- [24] R. Sakurai, O. J. Backhouse, G. H. Booth, W. Mizukami, and H. Shinaoka, Comparative study on compact quantum circuits of hybrid quantum-classical algorithms for quantum impurity models, *Phys. Rev. Res.* **6**, 023110 (2024).
- [25] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [26] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Ann. Phys.* **298**, 210 (2002).
- [27] O. Crawford, B. v. Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [28] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient estimation of Pauli observables by derandomization, *Phys. Rev. Lett.* **127**, 030503 (2021).
- [29] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm, *npj Quantum Inf.* **6**, 10 (2020).
- [30] H.-K. Zhang, S. Liu, and S.-X. Zhang, Absence of barren plateaus in finite local-depth circuits with long-range entanglement, *Phys. Rev. Lett.* **132**, 150603 (2024).
- [31] M. J. D. Powell, A direct search optimization method that models the objective and constraint functions by linear interpolation, in *Advances in Optimization and Numerical Analysis*, edited by S. Gomez and J.-P. Hennart (Springer Netherlands, Dordrecht, 1994), pp. 51–67.
- [32] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [33] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [34] H. J. Vallury, M. A. Jones, C. D. Hill, and L. C. L. Hollenberg, Quantum computed moments correction to variational estimates, *Quantum* **4**, 373 (2020).
- [35] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *J. Res. Natl. Bur. Stand.* **45**, 255 (1950).
- [36] L. C. L. Hollenberg, Plaquette expansion in lattice Hamiltonian models, *Phys. Rev. D* **47**, 1640 (1993).

- [37] L. C. L. Hollenberg and N. S. Witte, Analytic solution for the ground-state energy of the extensive many-body problem, *Phys. Rev. B* **54**, 16309 (1996).
- [38] L. C. L. Hollenberg and N. S. Witte, General nonperturbative estimate of the energy density of lattice Hamiltonians, *Phys. Rev. D* **50**, 3382 (1994).
- [39] R. Haydock, V. Heine, and M. J. Kelly, Electronic structure based on the local atomic environment for tight-binding bands, *J. Phys. C: Solid State Phys.* **5**, 2845 (1972).
- [40] N. S. Witte and L. C. L. Hollenberg, Plaquette expansion proof and interpretation, *Z. Phys. B: Condens. Matter* **95**, 531 (1994).
- [41] W. B. Jones and W. J. Thron, *Continued Fractions: Analytic Theory and Applications*, *Encyclopedia of Mathematics and Its Applications* (Addison-Wesley Publishing Company, Reading, MA, 1980), Vol. 11.
- [42] R. Haydock, *The Recursive Solution of the Schrodinger Equation* (Academic Press, 1980), pp. 215–294.
- [43] U.S. Department of Energy, Department of Energy Public Access Plan (U.S. Department of Energy, Washington, DC, 2023), available at <https://www.energy.gov/doe-public-access-plan>
- [44] M. Karabin E. A. Coello Perez, T. Sohail, S. Ghosh, M. Eisenbach, Optimization performance, fidelity, and cost SIAM VQE, doi: [10.13139/OLCF/3013386](https://doi.org/10.13139/OLCF/3013386)