

Quantum algorithm for one-quasiparticle excitations in the thermodynamic limit via cluster-additive block diagonalization

Sumeet ^{*}, M. Hörmann , and K. P. Schmidt 

Department of Physics, *Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)*, 91058 Erlangen, Germany



(Received 17 December 2025; accepted 29 May 2026; published 30 June 2026)

We propose a quantum algorithm for computing one-quasiparticle excitation energies in the thermodynamic limit by combining numerical linked-cluster expansions (NLCEs) and the variational quantum eigensolver (VQE). Our approach uses VQE to block diagonalize the cluster Hamiltonian through a single-unitary transformation. This unitary is then postprocessed using the projective cluster-additive transformation (PCAT) to ensure cluster additivity, a key requirement for NLCE convergence. We benchmark our method on the transverse-field Ising model (TFIM) in one and two dimensions, and with longitudinal field, computing one-quasiparticle dispersions in the high-field polarized phase. We compare two cost function classes, trace minimization and variance based, demonstrating their effectiveness with the Hamiltonian variational ansatz (HVA). For pure TFIM, $\lceil N/2 \rceil$ layers of HVA suffice: NLCE + VQE matches exact diagonalization. For TFIM with longitudinal field, where parity symmetry breaks and PCAT becomes essential, both $\lceil N/2 \rceil$ and N layers of HVA converge with increasing cluster size, with N layers providing improved accuracy. Our results establish PCAT as a cluster-additive framework that extends variational quantum algorithms to excited-state calculations in the thermodynamic limit via NLCE. While demonstrated with VQE, the PCAT postprocessing approach, which requires only low-energy eigenspace information, applies to any quantum eigenstate preparation method.

DOI: [10.1103/sdnp-s42c](https://doi.org/10.1103/sdnp-s42c)

I. INTRODUCTION

Understanding elementary excitations is central to the study of quantum many-body systems. Such excitations govern low-temperature thermodynamics and long-range dynamical behavior, and are directly accessible via spectroscopic techniques such as inelastic neutron or light scattering. Consequently, developing accurate and efficient methods to compute the excitation spectrum is a major objective in condensed matter physics, as well as in quantum chemistry and quantum simulation.

Recent years have witnessed significant progress in hybrid quantum-classical algorithms, particularly in using the variational quantum eigensolver (VQE) [1] to compute ground-state properties of strongly correlated systems. These algorithms are particularly relevant for noisy intermediate-scale quantum devices [2], where limited qubit coherence times and gate fidelities preclude error-corrected quantum computation. Originally designed for quantum chemistry applications [1,3], VQE has been increasingly applied to condensed matter models, including studies of quantum magnets [4,5], spin liquids [6], frustrated systems [7,8], quantum critical points [9,10], and our recent numerical linked-cluster

expansion (NLCE) + VQE approach for ground-state energies [11]. Recent no-go theorems have highlighted trainability challenges for variational quantum algorithms under certain assumptions [12,13]. However, systematic exploration with problem-adapted ansätze remains essential for developing successful strategies, and our work contributes to this by exploring VQE for degenerate subspaces within the NLCE framework using physically motivated initialization.

VQE has proven successful for ground-state calculations, providing a robust method to minimize the energy expectation value on finite quantum systems. Extending these approaches to compute excited states is an important next step, with multiple strategies having been proposed. These include variational quantum deflation [14], quantum equation of motion [15], subspace-search VQE (SSVQE) [16], and other ensemble and penalty-based methods [17,18]. Some of these techniques (such as SSVQE) can handle degenerate subspaces, though they are primarily designed for extracting individual eigenstates rather than constructing effective models for low-energy physics.

In many quantum lattice models, elementary excitations can be interpreted as quasiparticles: stable or long-lived excitations carrying a well-defined momentum $\vec{k}$ and energy $\omega(\vec{k})$. Constructing an effective quasiparticle Hamiltonian provides a compact description of these excitations through a few local coupling constants, reveals the underlying interaction structure, and connects directly to experimental observables.

Classical approaches to construct such effective Hamiltonians rely on block-diagonalization methods in quantum many-body physics, such as the Schrieffer-Wolff transformation [19], continuous unitary transformations [20–22],

^{*}Contact author: sumeet.sumeet@fau.de

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.

and the contractor renormalization group [23], which aim to isolate quasiparticle sectors by transforming the full Hamiltonian into a block-diagonal form where each block conserves the number of quasiparticles. Perturbative linked-cluster expansions [24,25] provide a systematic way to reach the thermodynamic limit by combining such block-diagonalization transformations with cluster expansion techniques. Perturbative approaches are limited to perturbative regimes. However, all of the previously mentioned methods can also be formulated nonperturbatively. Then, NLCEs [26–29], meaning replacing the perturbative with the exact result on a cluster, can be used to overcome the perturbative limitation. However, the exponential scaling still limits accessible cluster sizes.

While NLCE combined with exact diagonalization (ED) has been successful for both ground states and excitations, our recent work has demonstrated NLCE + VQE for ground-state energies [11], a method where VQE is used as a cluster solver instead of ED for each finite cluster in NLCE. However, a quantum or hybrid quantum-classical analog of these NLCE methods for constructing cluster-additive effective Hamiltonians of excited states does not exist yet. Recently, VQE-based approaches have been applied to extract quasiparticle excitations from finite periodic systems [30], but these methods target individual eigenstates on finite periodic clusters rather than constructing cluster-additive transformations suitable for thermodynamic-limit extrapolation via NLCE. Extending the NLCE framework to excitations with VQE as a cluster solver presents unique challenges due to higher dimensionality of excitation subspaces on finite clusters, which requires ensuring cluster additivity, a property not automatically satisfied by standard VQE block-diagonalization unitaries.

In this work, we propose a hybrid quantum-classical method to construct an effective one-quasiparticle (1QP) Hamiltonian in the thermodynamic limit. Our approach combines VQE with NLCEs and generalizes the NLCE + VQE method that we have recently introduced for ground-state energies [11]. Extending this framework to excitations presents significant challenges: Unlike the ground-state energy (a scalar quantity), 1QP excitations form a multi-dimensional subspace on finite clusters, requiring careful treatment to maintain the cluster additivity essential for NLCE convergence.

To ensure the transformation satisfies cluster additivity, which is essential for NLCE convergence, we proceed in two stages. Within each cluster of the NLCE, we employ a VQE-based cost function to approximate a unitary transformation that decouples the ground and 1QP subspaces. From this, we construct the projective cluster-additive transformation (PCAT) [31], which guarantees cluster additivity, a property that the transformation factorizes for disconnected clusters [32]. We use a single unitary to decouple the full low-energy subspace (ground state and 1QP states) from higher excitations, rather than targeting individual eigenstates. As a result, a single VQE optimization simultaneously yields the entire N -dimensional 1QP block on a finite cluster, from which the full dispersion $\omega(\vec{k})$ in the thermodynamic limit is recovered after NLCE embedding, in contrast to approaches that obtain quasiparticle energies one momentum at a time. This block-diagonalization approach exploits the fact that en-

ergy gaps within the low-energy subspace are typically small, while gaps to higher excitations are large, facilitating efficient decoupling. We use the Hamiltonian variational ansatz (HVA) [33], which constructs the quantum circuit from the problem Hamiltonian structure, and find that a circuit depth proportional to the system size is sufficient to achieve accurate results.

As a benchmark, we apply our method to the transverse-field Ising model (TFIM) on both the one-dimensional (1D) chain and the two-dimensional (2D) square lattice, computing the 1QP dispersion in the high-field polarized phase. We demonstrate convergence of the NLCE with increasing cluster size and validate our results by comparing VQE and ED as cluster solvers. Additionally, we investigate the effect of a longitudinal field (LF) on TFIM, which breaks the parity symmetry. In this case, the ground state and 1QP excitations are no longer in different symmetry sectors, making it a particularly stringent test of our approach. The longitudinal field case requires the full machinery of PCAT to ensure cluster additivity.

Our numerical benchmarks reveal that $\lceil N/2 \rceil$ layers of HVA suffice for the pure TFIM, where NLCE + VQE results are indistinguishable from NLCE + ED. For TFIM + LF, convergence is slower but still achieved. We investigate two classes of cost functions (trace based and variance based) and find that both can achieve convergence, though with different initialization requirements and trade-offs discussed in Sec. V B.

The remainder of this paper is organized as follows. In Sec. II, we provide an overview of the NLCE + VQE approach for 1QP excitations and outline the six-step workflow. Section III discusses the linked-cluster expansion for obtaining effective 1QP Hamiltonians, including details on block diagonalization (Sec. III A), the NLCE method (Sec. III B), the details of PCAT (Sec. III C), and the calculation of 1QP dispersions in the thermodynamic limit (Sec. III D). Section IV introduces VQE as the cluster solver, with details on cost functions (Sec. IV A) and state preparation (Sec. IV B). Section V presents numerical results for the 1QP dispersion of the high-field phase of the TFIM, including applications to the one-dimensional chain (Sec. V A 1), the two-dimensional square lattice (Sec. V A 2), and the TFIM + LF (Sec. V B). We conclude in Sec. VI with a summary and outlook.

II. OVERVIEW OF NLCE + VQE APPROACH FOR 1QP EXCITATIONS

In this section, we provide an overview of our NLCE + VQE framework for computing 1QP excitations in the thermodynamic limit, outlining the key components and workflow before presenting the detailed methodology in Secs. III and IV. Here, a cluster C is defined as a finite subset of a lattice consisting of a set of sites connected with bonds. Perturbative linked-cluster expansions (LCEs) are a common tool for calculating high-order series of excitation energies and spectral signatures of elementary excitations in the thermodynamic limit, but are limited by their perturbative nature. Practically, one performs perturbation theory on each cluster and then embeds the result in the thermodynamic limit by a proper




FIG. 1. Workflow of the NLCE + VQE approach for computing 1QP excitation energies in the thermodynamic limit. The six-step procedure combines VQE cluster solving with PCAT postprocessing and NLCE embedding to obtain the dispersion $\omega(\vec{k})$ in the thermodynamic limit.

embedding procedure. NLCEs replace the perturbation theory on finite clusters by numerical cluster solvers, typically ED for both scalar quantities and subspace calculations.

In our recent work [11], we introduced the NLCE + VQE approach for ground-state energies, which uses the VQE as a cluster solver to develop NLCEs for much larger cluster sizes in the future. Here, we extend this to 1QP excitations using VQE and PCAT to block diagonalize cluster Hamiltonians by decoupling the ground-state and 1QP subspace with a unitary U_{PCAT} . Computing 1QP excitations within NLCE is substantially harder than the ground-state case. While the ground-state energy is a single scalar per cluster, the 1QP effective Hamiltonian is an $N \times N$ matrix, where N is the number of sites in the cluster, so $\mathcal{O}(N^2)$ matrix elements must be computed instead of a single number. Moreover, NLCE embedding of nonscalar quantities is nontrivial since cluster additivity, which holds automatically for the ground-state energy, is not guaranteed for the 1QP block and must be enforced through PCAT. The transformed Hamiltonian is $U_{\text{PCAT}}^\dagger H U_{\text{PCAT}} = H_{\text{eff}}$, the effective Hamiltonian. To obtain 1QP excitation energies, we extract the 1QP block of H_{eff} and subtract the ground-state energy.

We benchmark our approach on the TFIM, a paradigmatic quantum spin system that exhibits a quantum phase transition. The Hamiltonian for the TFIM with an optional longitudinal field is given by

$$H = -h \sum_v Z_v - J \sum_{\langle v', v \rangle} X_{v'} X_v - h_l \sum_v X_v, \quad (1)$$

where Z_v and X_v denote Pauli matrices at site v , J is the nearest-neighbor Ising coupling, h is the transverse-field strength, and h_l is the longitudinal field. We set $h = 1$ as our energy unit throughout this work. We investigate three cases: (1) the one-dimensional TFIM chain ($h_l = 0$, Sec. V A 1), (2) the two-dimensional TFIM on the square lattice ($h_l = 0$, Sec. V A 2), and (3) the one-dimensional TFIM with longitudinal field ($h_l \neq 0$, Sec. V B). The longitudinal field breaks the $\mathbb{Z}_2$ parity symmetry, making the third case a particularly stringent test where PCAT becomes essential for ensuring cluster additivity. The expansion is developed in the high-field polarized phase ($J < h$), where the ground state is gapped and 1QP excitations correspond to localized spin flips.

We explain the workflow of our NLCE + VQE approach for 1QP energies in a stepwise fashion (see Fig. 1).

Step 1: Define the model. Choose a gapped quantum lattice model such as the TFIM + LF in Eq. (1).

Step 2: Decompose the lattice for NLCE. Generate topologically distinct, connected finite clusters C . We use rectangular graph expansions for lattice decomposition, which takes the clusters of dimensions $L_m \times L_n \leq N_{\text{max}}$, as detailed in Sec. III B.

Step 3: Solve each cluster using VQE. Block diagonalize each cluster C by minimizing a cost function (Sec. IV A) to find optimal variational parameters θ such that $H_{\text{dec},C} = U_{\text{VQE},C}^\dagger(\theta) H_C U_{\text{VQE},C}(\theta)$ decouples the ground-state and 1QP subspaces from the rest of the spectrum.

Step 4: Construct cluster-additive unitary using PCAT. Use the PCAT scheme (Sec. III C) to construct a unitary $U_{\text{PCAT},C}$ from $U_{\text{VQE},C}$, yielding the effective Hamiltonian $H_{\text{eff},C} = U_{\text{PCAT},C}^\dagger H_C U_{\text{PCAT},C}$. This transformation ensures cluster additivity: Reduced contributions of disconnected clusters vanish.

Step 5: Perform NLCE embeddings. Extract the 1QP effective Hamiltonian $H_{\text{eff},C}^{[1]}$ from $H_{\text{eff},C}$ and subtract the ground-state energy $E_C^{[0]}$, which is obtained from the same VQE procedure (see Sec. IV A). Compute the reduced contributions $\tilde{H}_{\text{eff},C}^{[1]}$ using the inclusion-exclusion principle [Eq. (7)] and embed into the infinite lattice, yielding $H_{\text{eff}}^{[1]}$ in the thermodynamic limit.

Step 6: Compute the 1QP dispersion. Diagonalize the effective 1QP Hamiltonian (obtained in Step 5) through Fourier transform, as described in Sec. III D, to obtain the 1QP dispersion $\omega(\vec{k})$.

Finally, let us stress that PCAT is indeed a necessary ingredient of our approach. NLCEs require observables to be cluster additive, restricting calculations to connected finite clusters with disconnected parts contributing zero. For degenerate subspaces without symmetries, infinitely many unitaries can decouple the Hamiltonian, but most violate cluster additivity. The unitary U_{VQE} from VQE does not satisfy cluster additivity. Furthermore, unitaries must show consistency across clusters for NLCE convergence. VQE efficiently extracts the information needed for PCAT: energy expectation values and overlaps with unperturbed eigenstates from the low-energy subspace. Crucially, PCAT requires only this

low-energy eigenspace information, without needing access to higher excitation sectors, avoiding exponentially costly state tomography. This suffices to construct the PCAT [31], U_{PCAT} , from U_{VQE} , ensuring cluster additivity. Embedding the 1QP block of H_{eff} into the infinite lattice via NLCE yields the effective 1QP Hamiltonian and dispersion $\omega(\vec{k})$ in the thermodynamic limit.

The computational cost is highly nonuniform across the six steps. Steps 1, 2, 5, and 6 are classical and scale polynomially with the number of clusters, while Steps 3 and 4 carry the quantum cost. The VQE optimization in Step 3 is the dominant bottleneck. Step 4 requires a single round of $\mathcal{O}(N^2)$ measurements on the converged circuit to extract the matrices needed for PCAT and avoids the exponential cost of full state tomography.

III. LINKED-CLUSTER EXPANSION FOR EFFECTIVE ONE-QUASIPARTICLE PICTURE

In this section, we establish the mathematical and computational framework for extracting 1QP excitations in the thermodynamic limit using NLCEs. We first introduce the concept of block diagonalization and explain why cluster additivity is essential for NLCE convergence. We then present the NLCE formalism for 1QP excitations and describe how the PCAT ensures that the effective Hamiltonian satisfies the required cluster-additivity property. Finally, we show how the 1QP dispersion relation is obtained from the effective Hamiltonian via Fourier transformation.

A. Block diagonalization

A quasiparticle ansatz typically characterizes the low-energy excitation spectrum by noninteracting dressed excitations. These correspond to block diagonalizing a strongly correlated many-body Hamiltonian into independent sectors labeled by quasiparticle number. Formally, given a Hamiltonian

$$H = H_0 + xV, \quad (2)$$

with unperturbed part H_0 and perturbation V at strength x , we seek to partition the Hilbert space into subspaces based on the number of elementary excitations above the ground state. The unperturbed Hamiltonian H_0 is assumed to be exactly solvable and serves as the reference. Its eigenstates organize the Hilbert space into quasiparticle sectors. Denoting by $E_0^{[n]}$ the unperturbed energy of the n quasiparticle subspace, the Hilbert space decomposes as

$$\mathcal{H} = \bigoplus_n \mathcal{H}_0^{[n]}. \quad (3)$$

Each $\mathcal{H}_0^{[n]}$ carries all states with eigenvalue $E_0^{[n]}$ of H_0 corresponding to n quasiparticles. We call H_0 block diagonal if

$$H_0 = \bigoplus_{n=0}^N H_0^{[nm]}, \quad (4)$$




FIG. 2. Diagrammatic view of a block diagonal Hamiltonian. The ground-state (0QP) block is shown in dark green and the 1QP block is shown in dark blue. The light shaded blocks are off-diagonal terms that are minimized during block diagonalization.

and $H_0^{[nn]}$ keeps $\mathcal{H}_0^{[n]}$ invariant. We call a transformation U block diagonalizing if

$$H_{\text{eff}} = U^\dagger H U, \quad (5)$$

and H_{eff} retains the block structure of H_0 , i.e.,

$$H_{\text{eff}} = \bigoplus_{n=0}^N H_{\text{eff}}^{[nm]}. \quad (6)$$

Here, $H_{\text{eff}}^{[nm]}$ denotes the block corresponding to the n quasi-particle sector. For example, $H_{\text{eff}}^{[11]}$ is the $N \times N$ matrix for the 1QP block, where N is the number of sites.

The subspace $\mathcal{H}_{\text{eff}}^{[n]}$ is defined as containing those eigenstates of H that adiabatically follow the states of $\mathcal{H}_0^{[n]}$ as the perturbation strength x' varies from 0 to x in $H = H_0 + x'V$.

In realistic settings, one can not block diagonalize the full Hamiltonian but only the low-energy part. In our case, we restrict the decomposition to three sectors consisting of the ground-state subspace $\mathcal{H}_{\text{eff}}^{[0]}$, the IQP subspace $\mathcal{H}_{\text{eff}}^{[1]}$, and all higher excitations. The pictorial representation of such a block diagonal Hamiltonian is shown in Fig. 2, where the block diagonalizing U should eliminate the faint blue and green off-diagonal sectors $H_{\text{eff}}^{[nm]}$ with $n \neq m$.

A complete diagonalization of an infinite system is intractable, but LCEs provide a systematic approximation by combining results from finite clusters to recover thermodynamic-limit properties. For LCEs to converge, the transformation must be cluster additive: Disconnected clusters should contribute zero. However, even for the block diagonalizing part of U infinitely many choices exist for a degenerate subspace n as discussed in Ref. [34]. While the ground state is uniquely determined, any unitary rotation within the 1QP subspace yields an equally valid block diagonalization due to the quasidegeneracy. Higher-energy sectors have even more freedom.

This freedom necessitates, but also allows one to tailor the structure of the effective Hamiltonian to ensure cluster additivity. In particular, the PCAT [31] yields a specific unitary U_{PCAT} that preserves cluster additivity. Cluster additivity ensures that only connected clusters need to be calculated, as

disconnected clusters contribute zero, significantly reducing computational cost.

B. Linked-cluster expansions

In this subsection, we first outline the LCE framework assuming a suitable effective Hamiltonian exists, and then specify the requirements it must satisfy.

The general formalism for perturbative and nonperturbative LCEs is reviewed in Refs. [24,27,28]. First introduced for perturbative calculations of ground-state energies in quantum many-body systems, LCEs were subsequently extended to the study of excitation energies [25]. The core idea of LCEs is to decompose the lattice into topologically distinct, finite connected subgraphs called clusters C , and use the inclusion-exclusion principle to sum up their additive contributions to recover thermodynamic-limit properties. Let us for the moment assume that the 1QP effective Hamiltonian $H_{\text{eff},C}^{[1]}$ is such an additive quantity acting as a linear operator in the space of N excitations. Then, the LCE scheme goes as follows.

A cluster is a finite set of sites connected by bonds, where sites carry on-site Hamiltonian terms (transverse and longitudinal fields) and bonds carry the two-body interactions (XX couplings). These clusters are chosen with rectangular graph expansion having the dimensions $L_m \times L_n \leq N_{\text{max}}$, where N_{max} is the maximum number of spins taken in NLCE calculation.

The reduced contribution of a cluster C is defined recursively as

$$\tilde{H}_{\text{eff},C}^{[1]} = H_{\text{eff},C}^{[1]} - \sum_{C' \subsetneq C} w_{C'/C} \tilde{H}_{\text{eff},C'}^{[1]}, \quad (7)$$

where $w_{C'/C}$ counts embeddings of the subgraph C' inside C . Summing the reduced contributions gives the effective 1QP Hamiltonian in the thermodynamic limit,

$$H_{\text{eff}}^{[1]} = \sum_C w_C \tilde{H}_{\text{eff},C}^{[1]}, \quad (8)$$

where weight w_C counts embeddings of cluster C per lattice site.

Often, one uses a full graph decomposition for the NLCE where the calculation is performed on topologically distinct graphs of minimal size. However, the number of graphs increases exponentially with the number of sites, which is challenging for the NLCE + VQE approach. As for the NLCE + VQE for ground-state energy [11], we adopt a rectangular graph expansion [35] for NLCE. In this expansion, only clusters of size $L_m \times L_n$ with $L_m \cdot L_n \leq N_{\text{max}}$ are included, where N_{max} is the maximum number of spins in the largest cluster of the NLCE. These clusters have open boundary conditions. The rectangular graph expansion is applicable for the chain and the square lattice discussed below and is beneficial because the number of clusters increases polynomially with N . For the one-dimensional chain, this reduces to considering open segments of length $L_m \leq N_{\text{max}}$. We will apply this expansion to the TFIM + LF in Sec. V. The rectangular graph expansion for the square lattice is illustrated in Fig. 3.

We now specify the calculation of $H_{\text{eff},C}^{[1]}$ and how it can be made into a cluster-additive quantity. The effective 1QP Hamiltonian on a cluster C with N 1QP states is an $N \times N$




FIG. 3. Rectangular graph expansion for the square lattice, up to $N_{\text{max}} = 9$, showing all clusters of dimensions $L_m \times L_n$ with $L_m \cdot L_n \leq N_{\text{max}}$. For the NLCE embedding, each nonsquare cluster ($L_m \neq L_n$) carries weight $w_C = 2$ from its two lattice orientations (horizontal and vertical), while the 2×2 and 3×3 square clusters have $w_C = 1$.

matrix defined as

$$H_{\text{eff},C}^{[1]} = (U^\dagger H U)^{[11]} - E_C^{[0]}, \quad (9)$$

where $^{[11]}$ denotes the 1QP block and U is the block-diagonalizing transformation. $E_C^{[0]}$ is obtained from the same VQE procedure: either via a separate ground-state minimization (separate-unitary case), or in the single-unitary case via diagonalization of the Hamiltonian projected onto the VQE-prepared subspace (see Sec. IV C 2). The ground-state energy $E_C^{[0]}$ must be subtracted to ensure that $H_{\text{eff},C}^{[1]}$ is intensive [25]. Importantly, since $E_C^{[0]}$ enters only as a diagonal shift of $H_{\text{eff},C}^{[1]}$, any estimation error from VQE affects only the on-site energies and leaves the hopping matrix elements unchanged. After Fourier transformation in the NLCE embedding, this propagates at most to a momentum-independent shift of $\omega(\vec{k})$, leaving its shape and bandwidth unaffected.

The need for subtracting the ground-state energy is best understood through a concrete example. Consider two disconnected clusters A and B (i.e., they share no sites or bonds). Since the Hamiltonians H_A and H_B on the two disconnected clusters commute, $[H_A \otimes \mathbb{1}_B, \mathbb{1}_A \otimes H_B] = 0$, it follows for the eigenfunctions on $A \cup B$ that

$$|\Psi_{A \cup B}^{[n+m]}\rangle = |\Psi_A^{[n]}\rangle \otimes |\Psi_B^{[m]}\rangle. \quad (10)$$

Here, $|\Psi^{[n]}\rangle$ denotes an eigenstate of the full Hamiltonian H with n quasiparticles, and $[n]$ ($[m]$) refers to the number of quasiparticles on cluster A (B). The effective Hamiltonian part on $A \cup B$ that describes 1QP excitations within B would always have a contribution of the ground-state energy on A . Hence, one has to subtract it ($n = 0$ and $m = 1$ in the above formula).

This is linked to the concept of cluster additivity. A transformation U is cluster additive if

$$U^\dagger H_{A \cup B} U = H_{\text{eff},A \cup B} = H_{\text{eff},A} \otimes \mathbb{1}_B + \mathbb{1}_A \otimes H_{\text{eff},B}. \quad (11)$$

If fulfilled,

$$\bar{H}_{\text{eff},A \cup B}^{[1]} = H_{\text{eff},A \cup B}^{[1]} - E_{A \cup B}^{[0]} = \bar{H}_{\text{eff},A}^{[1]} \oplus \bar{H}_{\text{eff},B}^{[1]} \quad (12)$$

will hold. Subtracting the ground-state energy alone is necessary, but not sufficient to ensure cluster additivity.

In the next subsection, we explain how PCAT constructs a transformation U_{PCAT} that guarantees both properties using only limited information from the ground-state and 1QP subspaces.

C. Projective cluster-additive transformation

In LCEs, both perturbative and nonperturbative, it is essential that observables on large systems can be consistently constructed from smaller, disconnected clusters. To understand why PCAT [31] achieves cluster additivity, we first establish the mathematical requirements.

Two clusters are disconnected if the unperturbed Hamiltonian is cluster additive on them (see Sec. III B) [32]:

$$H_{A \cup B} = H_A \otimes \mathbb{1}_B + \mathbb{1}_A \otimes H_B. \quad (13)$$

We demand of a transformation U that when applied to the full Hamiltonian $H = H_0 + xV$, $U^\dagger H U$ remains cluster additive on two disconnected clusters A and B . This ensures that no nonlinked clusters contribute to the expansion. We first illustrate why standard approaches fail to maintain cluster additivity, then show how PCAT resolves this issue through a modified state construction.

To understand that PCAT fulfills cluster additivity, we have to investigate the product state structure of eigenstates on $A \cup B$. Let $\{|\Phi_i^{[n]}\rangle\}$ denote the eigenstates of H_0 in the $(n+1)$ th subspace, and let $\{|\Psi_i^{[n]}\rangle\}$ denote the corresponding eigenstates of the full Hamiltonian H , where i labels states within the degenerate manifold. $P^{[n]}$ is the corresponding projector, and $\bar{P}^{[n]}$ the adiabatically connected projector of the full Hamiltonian H . On $A \cup B$ eigenstates have the product form as in Eq. (10).

It is instructive to see why a widely used approach to decouple an energy block fails to be cluster additive: The canonical two-block (2b) transformation (Schrieffer-Wolff, des Cloizeaux, Takahashi, etc. [19,24,34,36–38]), which can be written in projector form for the transformation restricted to block n as

$$U_{2b}^{[n]} = \bar{P}^{[n]}(P^{[n]}\bar{P}^{[n]}P^{[n]})^{-1/2}, \quad (14)$$

where $(M)^{-1/2}$ for a rank-deficient matrix M denotes the inverse square root restricted to the range of M , or equivalently, the Moore-Penrose pseudoinverse square root. This convention applies throughout this section whenever inverse operations are applied to projector products. The full transformation $U_{2b} = \sum_n U_{2b}^{[n]}$ minimizes $\|1 - U_{2b}\|$. Let $U^{[n]}$ denote the $2^N \times N_n$ matrix whose columns are the exact eigenvectors of H corresponding to the n th subspace, where N_n is the dimension of that subspace. When embedded in the full $2^N \times 2^N$ space (as needed for sums like $\sum_n U^{[n]}$), this is understood as a partial isometry extended by zeros outside the n -quasiparticle subspace. The transformation can also be given in terms of the overlap matrix $X^{[n]}$, an $N_n \times N_n$ matrix




FIG. 4. Comparison of 2b and PCAT transformations on disconnected clusters. The 2b transformation allows unphysical hopping of a dressed quasiparticle (yellow) between clusters A and B that share no bonds. PCAT forbids such hopping, ensuring cluster additivity.

with elements

$$X_{ij}^{[n]} = \langle \Phi_i^{[n]} | \Psi_j^{[n]} \rangle, \quad (15)$$

where $|\Phi_i^{[n]}\rangle$ are the unperturbed eigenstates and $|\Psi_j^{[n]}\rangle$ are the eigenstates of H in block n . Then,

$$U_{2b}^{[n]} = U^{[n]} X^{[n]\dagger} (X^{[n]} X^{[n]\dagger})^{-1/2}. \quad (16)$$

Written this way, the problem with cluster additivity becomes apparent: For simplicity, consider the 1QP excitations ($n = 1$). Then, $X^{[1]}$ becomes non-block-diagonal for disconnected clusters owing to the product form when the unperturbed ground state overlaps with excited states (i.e., $P^{[1]}|\Psi^{[0]}\rangle \neq 0$). Contributions from $|\Psi_A^{[0]}\rangle \otimes |\Psi_B^{[1]}\rangle$ lead to nonzero elements in $X^{[1]}$ at the same entries as $|\Psi_A^{[1]}\rangle \otimes |\Psi_B^{[0]}\rangle$. This leads to intercluster matrix elements between blocks A and B . It also modifies matrix elements in the A -only or B -only blocks compared to the isolated A and B clusters.

Thus, cluster additivity is violated and dressed quasiparticles can hop between disconnected clusters (see Fig. 4). In the thermodynamic limit, disconnected clusters can be arbitrarily far apart, so such hopping represents unphysical infinite-range coupling. This artifact arises because the 2b transformation creates entanglement between clusters that should be independent.

In order to resolve this, PCAT [31] constructs modified states $\{|\tilde{\Psi}^{[n]}\rangle\}$ from the original eigenstates by eliminating all components projected onto lower-energy subspaces of H_0 , using only subtractions of states from the low-energy spectrum. These modified states lie within the span of low-energy eigenspaces but are not eigenstates themselves. From these states, one constructs a modified overlap $\tilde{X}^{[n]}$, which in contrast to $X^{[n]}$ is now block diagonal again on A and B . In detail, the construction of modified states goes as follows.

The recursion for constructing modified 1QP states is

$$|\tilde{\Psi}^{[0]}\rangle = |\Psi^{[0]}\rangle, \quad |\tilde{\Psi}_i^{[1]}\rangle = |\Psi_i^{[1]}\rangle - \frac{\langle \Phi^{[0]} | \Psi_i^{[1]} \rangle}{\langle \Phi^{[0]} | \Psi^{[0]} \rangle} |\Psi^{[0]}\rangle, \quad (17)$$

which ensures that $P^{[0]}|\tilde{\Psi}_i^{[1]}\rangle = 0$ for all 1QP states. This removes any admixture of the ground state from the 1QP eigenstates. Note that in the special case where $\langle \Phi^{[0]} | \Psi^{[1]} \rangle = 0$ the transformation remains the same for the one-particle subspace. This is the case for the transverse-field Ising model, but as soon as the parity is broken, e.g., by a longitudinal field, the corrections become necessary.

For higher subspaces, the construction generalizes by recursively projecting out all unperturbed lower-energy

subspaces. Using the low-energy projector

$$R^{[n]} = \sum_{m < n} P^{[m]} \quad (18)$$

and its adiabatically connected counterpart $\tilde{R}^{[n]}$, the general subtraction scheme reads

$$|\tilde{\Psi}^{[n]}\rangle = (1 - \tilde{R}^{[n]}(R^{[n]}\tilde{R}^{[n]}R^{[n]})^{-1}R^{[n]})|\Psi^{[n]}\rangle. \quad (19)$$

This ensures that each modified state $|\tilde{\Psi}^{[n]}\rangle$ has no projection onto lower-energy subspaces of H_0 .

The key insight is that a state $|\Psi_{AUB}^{[n+m]}\rangle = |\Psi_A^{[n]}\rangle \otimes |\Psi_B^{[m]}\rangle$ is transformed into

$$|\tilde{\Psi}_{AUB}^{[n+m]}\rangle = |\tilde{\Psi}_A^{[n]}\rangle \otimes |\tilde{\Psi}_B^{[m]}\rangle. \quad (20)$$

This factorization is the essence of why PCAT preserves cluster additivity. Since modified states on $A \cup B$ factor into modified states on A and B separately, the transformation introduces no artificial correlations between disconnected clusters. Each cluster's quasiparticles remain localized on the cluster.

From this factorization property, several important consequences follow. In particular, it yields $P^{[s]}|\tilde{\Psi}_{AUB}^{[n+m]}\rangle = 0$ for $s < n + m$ but also $P^{[s]}|\tilde{\Psi}_A^{[n]}\rangle = 0$ for $s < n$ and $P^{[s]}|\tilde{\Psi}_B^{[m]}\rangle = 0$ for $s < m$. Indeed, the identity for modified states on disconnected clusters A and B is sufficient to guarantee cluster additivity of the transformation:

$$U_{\text{PCAT}}^\dagger H_{AUB} U_{\text{PCAT}} = H_{\text{eff},A} \otimes \mathbb{1}_B + \mathbb{1}_A \otimes H_{\text{eff},B}. \quad (21)$$

We can now write the explicit form of the U_{PCAT} transformation. The modified overlap matrix $\tilde{X}^{[n]}$ is an $N_n \times N_n$ matrix with elements

$$\tilde{X}_{ij}^{[n]} = \langle \Psi_i^{[n]} | \tilde{\Psi}_j^{[n]} \rangle, \quad (22)$$

where $|\Psi_i^{[n]}\rangle$ are the unperturbed eigenstates and $|\tilde{\Psi}_j^{[n]}\rangle$ are the modified eigenstates in block n . The transformation for a single block n is then

$$U_{\text{PCAT}}^{[n]} = U^{[n]} \tilde{X}^{[n]\dagger} (\tilde{X}^{[n]} \tilde{X}^{[n]\dagger})^{-1/2}, \quad (23)$$

and for the full Hilbert space (with each block extended to $2^N \times 2^N$ as partial isometries),

$$U_{\text{PCAT}} = \sum_n U_{\text{PCAT}}^{[n]}. \quad (24)$$

Since the modified states $|\tilde{\Psi}^{[n]}\rangle$ can be written using only projectors (as in the $R^{[n]}$ form above), one can write U_{PCAT} in terms of projectors only, which shows that it has a unique local perturbative expansion when using Kato's perturbative series [39] for projection operators. This property implies that the transformation has the same form on every cluster. It enables a consistent perturbative expansion across clusters and ensures that a given perturbative order can be reached using the smallest possible clusters.

The projector form also establishes that the transformation satisfies

$$U_{\text{PCAT}}^{[n]}(U^{[0]}, U^{[1]}, \dots, U^{[n]}) = U_{\text{PCAT}}^{[n]}(U^{[0]}, \dots, U^{[n]}\mathcal{W}) \quad (25)$$

for arbitrary unitaries $\mathcal{W} \in \mathcal{U}(N_n)$ (the group of $N_n \times N_n$ unitary matrices) acting within subspace n , as rotations within a

subspace do not affect the projection operators. This invariance property is essential for practical implementations.

Thus, PCAT yields a local, cluster-additive effective Hamiltonian built only from the low-energy eigenspaces and is compatible with both perturbative and nonperturbative LCEs.

This highlights the global character of the transformation; i.e., it can be fully constructed from the eigenspace of interest and those of lower energy. This global character is essential for our approach: VQE (Sec. IV) targets only the low-energy eigenspace, and PCAT constructs the transformation from this limited information, namely, energy expectations and state overlaps, without requiring knowledge of higher excitation sectors.

D. 1QP dispersion in the thermodynamic limit

Once we obtain the effective Hamiltonian for finite clusters, the goal is to obtain the effective 1QP Hamiltonian and therefore the 1QP dispersion in the thermodynamic limit. For translationally invariant systems, 1QP eigenstates are labeled by momentum $\vec{k}$. With lattice spacing set to unity and assuming a single-site unit cell, the Brillouin zone is $k \in [-\pi, \pi]$ for the one-dimensional chain and $\vec{k} = (k_x, k_y) \in [-\pi, \pi]^2$ for the square lattice.

For this purpose, the rectangular clusters are embedded into the infinite lattice as explained in Sec. III B using the inclusion-exclusion principle. For the rectangular expansion, the embedding weight w_C for a cluster of size $L_m \times L_n$ on an infinite lattice is simply unity (one embedding per lattice site in the thermodynamic limit). For translationally invariant models, eigenstates are also momentum eigenstates and Fourier transforming the real-space Hamiltonian gives the 1QP energies. The dispersion relation for both the chain and square lattice reads

$$\omega(\vec{k}) = \sum_{L_m, L_n} \sum_{\vec{\mu}, \vec{\nu}} \exp(i\vec{k} \cdot (\vec{\nu} - \vec{\mu})) \tilde{H}_{\text{eff}, L_m \times L_n, \vec{\mu}, \vec{\nu}}^{[1]}, \quad (26)$$

where $\vec{\mu}, \vec{\nu}$ label lattice sites within the cluster. For the one-dimensional chain, $L_n = 1$ and $\vec{k}$ reduces to scalar momentum k . For the square lattice, $\vec{k} = (k_x, k_y)$ spans the two-dimensional Brillouin zone.

The calculation of reduced contribution is particularly simple for the rectangular expansion. It is easy to see that the contribution $\tilde{H}_{\text{eff}, L_m \times L_n, \vec{k}}^{[1]}$ of a cluster in k space can be rewritten in terms of nonreduced contributions in the following way:

$$\begin{aligned} \tilde{H}_{\text{eff}, L_m \times L_n, \vec{k}}^{[1]} &= H_{\text{eff}, L_m \times L_n, \vec{k}}^{[1]} - 2H_{\text{eff}, L_m - 1 \times L_n, \vec{k}}^{[1]} \\ &\quad - 2H_{\text{eff}, L_m \times L_n - 1, \vec{k}}^{[1]} + 4H_{\text{eff}, L_m - 1 \times L_n - 1, \vec{k}}^{[1]} \\ &\quad + H_{\text{eff}, L_m - 2 \times L_n, \vec{k}}^{[1]} + H_{\text{eff}, L_m \times L_n - 2, \vec{k}}^{[1]} \\ &\quad - 2H_{\text{eff}, L_m - 1 \times L_n - 2, \vec{k}}^{[1]} - 2H_{\text{eff}, L_m - 2 \times L_n - 1, \vec{k}}^{[1]} \\ &\quad + H_{\text{eff}, L_m - 2 \times L_n - 2, \vec{k}}^{[1]}. \end{aligned} \quad (27)$$

The coefficients in this expression and the alternating signs follow from the inclusion-exclusion principle applied to rectangular geometry.

Each subtracted term removes overcounting from smaller clusters. This identity reveals an important insight: Only the largest clusters and their boundaries contribute to the rectangular expansion, as interior contributions cancel in the inclusion-exclusion sum. This property reduces error accumulation. Specifically, once clusters of size $L_m \times L_n$ are included in the expansion, all clusters of size $(L_m - 2) \times (L_n - 2)$ or smaller have reduced contributions that vanish identically. Their physics is therefore fully captured by the larger embeddings. This is analogous to the 1D case where only clusters of size N and $N - 1$ contribute to the dispersion at order $N_{\max} = N$.

Having established the NLCE framework, we now describe how VQE is used as the cluster solver.

IV. VQE AS CLUSTER SOLVER

The VQE is a hybrid quantum-classical algorithm for near-term quantum devices. In this section, we describe how VQE serves as the cluster solver within our NLCE framework, block diagonalizing the Hamiltonian on each finite cluster to decouple the ground state and 1QP subspace from higher excitations.

As discussed in Sec. III C, the unitary transformation U_{VQE} obtained from VQE optimization alone is generally not cluster additive and therefore not suitable for direct use in NLCE. Instead, VQE provides the information, namely, the low-energy eigenspaces and their overlaps with the unperturbed states, needed to construct the cluster-additive transformation U_{PCAT} via the PCAT scheme. From the full effective Hamiltonian $H_{\text{eff}} = U_{\text{PCAT}}^\dagger H U_{\text{PCAT}}$, we extract the 1QP block

$$H_{\text{eff}}^{[1]} = (U_{\text{PCAT}}^\dagger H U_{\text{PCAT}})^{[1]} - E^{[0]}, \quad (28)$$

where the superscript $^{[1]}$ denotes the 1QP block and $E^{[0]}$ is the ground-state energy.

A key advantage of the PCAT approach is that it requires only energy expectation values and overlaps, quantities that VQE can provide, rather than full state tomography, which would be exponentially costly. This efficiency stems from the structure of PCAT itself, as detailed in Sec. III C.

Throughout the paper, we use VQE as cluster solver, though the PCAT framework is algorithm agnostic and applies equally to other eigenstate preparation methods such as quantum phase estimation or adiabatic state preparation.

VQE operates by minimizing a cost function $C(\theta)$ that depends on variational parameters θ characterizing a quantum circuit $U_{\text{VQE}}(\theta)$. For ground-state calculations, the cost function is simply the energy expectation value, and the Rayleigh-Ritz variational principle ensures convergence to the lowest-energy state. For excited states, the situation is more subtle. Several approaches exist within the VQE framework for accessing excited states, including deflation methods [14], purified state approaches [17], and subspace-search methods [16]. Deflation methods sequentially compute excited states by enforcing orthogonality with previously found states. Purified state and subspace-search methods treat entire degenerate subspaces simultaneously, which is more natural for our purpose where we need to decouple a full quasidegenerate 1QP subspace on finite clusters.

The performance of VQE depends on two components: (1) the cost function, which must be designed to achieve block diagonalization rather than simply finding individual eigenstates, and (2) the ansatz for the variational circuit $U_{\text{VQE}}(\theta)$, which must be sufficiently expressive yet trainable. For the ansatz, we adopt the HVA, which we successfully employed for ground-state energy calculations in our previous NLCE + VQE work [11]. The HVA is also well suited for excited states. Since the eigenstates of the full Hamiltonian are superpositions across multiple bare quasiparticle sectors of H_0 , the ansatz must generate intersector mixing. The HVA achieves this naturally, as it is constructed directly from the terms of H .

The remainder of this section is organized as follows. In Sec. IV A, we discuss the design of cost functions for block diagonalization, introducing both variance-based and trace-based approaches and comparing their properties. In Sec. IV B, we describe the HVA ansatz and our initialization strategy. Finally, in Sec. IV C, we detail the quantum measurement protocol required to construct PCAT. This subsection describes how to extract energy expectations and state overlaps from the VQE solution and presents the explicit algorithm for computing U_{PCAT} .

A. Cost functions

For clarity, all cost functions below are evaluated on a finite cluster C within the NLCE, though we suppress the cluster label for notational simplicity. The unitary U appearing in the cost functions defined in this section is generic and can be realized by any ansatz.

VQE aims to find the parameters θ that minimize a cost function $C(\theta)$. For ground states, the ground-state energy $E_C^{[0]}$ can be obtained separately by direct energy minimization via the Rayleigh-Ritz principle,

$$C^{\text{GS}}(\theta) = \langle \Phi^{[0]} | U^\dagger(\theta) H U(\theta) | \Phi^{[0]} \rangle, \quad (29)$$

or jointly as part of the combined GS + 1QP cost functions described below for ground state and excited states.

1. Variance-based cost functions

One effective approach is to minimize the energy variance [40,41], $\langle H^2 \rangle - \langle H \rangle^2$. For a state $|\Psi\rangle = \sum_n c_n |\phi_n\rangle$ expanded in the eigenbasis of H with eigenvalues E_n , the variance equals $\sum_n |c_n|^2 (E_n - \bar{E})^2$, where $\bar{E} = \sum_n |c_n|^2 E_n$ is the energy expectation value and the sum runs over energy eigenstates. The variance vanishes if and only if the state lies within a single energy eigenspace. For nondegenerate spectra, this requires an eigenstate; for degenerate subspaces, any superposition within the degenerate manifold yields zero variance. This quantity measures the extent to which H has been block diagonalized and is related to residual off-diagonal elements in the context of continuous unitary transformations.

To decouple a quasidegenerate 1QP subspace from the rest of the spectrum, we generalize to multiple states. For a set of 1QP basis states $\{|\Phi_i^{[1]}\rangle\}_{i=1}^N$ transformed by unitary $U(\theta)$, the

variance cost function is

$$C_{\text{var}}^{\text{1QP}}(\theta) = \sum_{i=1}^N \langle \Phi_i^{[1]} | U^\dagger(\theta) H^2 U(\theta) | \Phi_i^{[1]} \rangle - \sum_{i,j=1}^N \left| \langle \Phi_i^{[1]} | U^\dagger(\theta) H U(\theta) | \Phi_j^{[1]} \rangle \right|^2. \quad (30)$$

This cost function is unitarily invariant within the 1QP subspace: Any rotation of the basis $\{|\Phi_i^{[1]}\rangle\}$ by a unitary acting only within this subspace leaves $C_{\text{var}}^{\text{1QP}}$ unchanged. Physically, this means the cost function enforces block diagonalization (decoupling the 1QP subspace from other sectors) without requiring full diagonalization within the 1QP block itself; the internal structure of the block remains flexible. In contrast, minimizing the variance for each state individually (as done in Ref. [41]) would not preserve this unitary invariance in the full subspace. The residual variance after optimization directly indicates the quality of decoupling achieved.

We can use variance-based cost functions in two ways. Either optimize $C_{\text{var}}^{\text{1QP}}$ alone for the 1QP sector (with a separate ground-state calculation using energy minimization) or include both sectors simultaneously:

$$C_{\text{var}}^{\text{GS,1QP}}(\theta) = \langle \Phi^{[0]} | U^\dagger(\theta) H^2 U(\theta) | \Phi^{[0]} \rangle - |\langle \Phi^{[0]} | U^\dagger(\theta) H U(\theta) | \Phi^{[0]} \rangle|^2 + C_{\text{var}}^{\text{1QP}}(\theta), \quad (31)$$

where $|\Phi^{[0]}\rangle$ is the ground state of the unperturbed Hamiltonian H_0 . We use equal weights (coefficient 1) for both the ground-state variance and the 1QP variance, treating both sectors symmetrically. Unlike the trace cost function (discussed below), which is unitarily invariant in the combined (GS + 1QP) space, this sum has separate unitary invariances in each subspace. This has the advantage that minimizing variance in each sector separately actively drives orthogonality between the ground-state and 1QP subspace. Perfect optimization (zero variance) guarantees orthogonality; in practice, finite residual variance can reflect approximate orthogonality with the degree of mixing controlled by the residual.

2. Trace-based cost functions

An alternative approach minimizes the sum of weighted energies [16]:

$$C_{\text{tr}}^{\text{GS,1QP}}(\theta) = w_0 \langle \Phi^{[0]} | U^\dagger(\theta) H U(\theta) | \Phi^{[0]} \rangle + \sum_{i=1}^N w_i \langle \Phi_i^{[1]} | U^\dagger(\theta) H U(\theta) | \Phi_i^{[1]} \rangle. \quad (32)$$

We use uniform weights $w_0 = w_i = 1$ for all i , making the cost function basis invariant. It is thus the trace over the combined ground-state and 1QP subspace. This is natural for linked-cluster expansions where we require the full subspace information and treat all sites equivalently. We apply this trace cost function in two ways: $C_{\text{tr}}^{\text{1QP}}$ targets only the 1QP sector [$w_0 = 0$ in Eq. (32) and separate ground-state optimization],

while $C_{\text{tr}}^{\text{GS,1QP}}$ treats both ground-state and 1QP subspace with a single unitary.

The trace minimization is motivated by the fact that block diagonalization preserves the sum of eigenvalues within each block. However, when ground and excited states are included together, minimizing the trace does not guarantee that $U(\theta)|\Phi^{[0]}\rangle$ is the true lowest-energy state. A postprocessing diagonalization step (described in Sec. IV C) is therefore required to identify the actual ground state and extract 1QP energies.

The trace cost function can also be evaluated using a purified approach where the whole subspace is treated as a wavefunction in a larger Hilbert space [17,18].

3. Comparison and choice

We distinguish between two different implementation strategies. With single-unitary approaches ($C_{\text{tr}}^{\text{GS,1QP}}$, $C_{\text{var}}^{\text{GS,1QP}}$), orthogonality is preserved throughout optimization by construction. However, $C_{\text{tr}}^{\text{GS,1QP}}$ requires postprocessing diagonalization to identify the true ground state, while $C_{\text{var}}^{\text{GS,1QP}}$ directly yields approximate eigenstates. With the separate-unitary approach ($C_{\text{var}}^{\text{1QP}}$), orthogonality must be verified and may require a generalized eigenvalue problem when nonorthogonality occurs.

We tested both strategies across all models studied in Sec. V. For the pure TFIM cases, all cost functions achieve similar convergence and stability, successfully decoupling the low-energy block. The trace-based approach requires fewer measurements, as it only needs energy expectation values $\langle H \rangle$, while variance-based approaches additionally require $\langle H^2 \rangle$ measurements. Variance-based approaches provide residual variance as a diagnostic for block-diagonalization quality. However, for TFIM + LF where parity breaking makes optimization more challenging, trace-based cost functions show greater robustness to ground-state initialization, while variance-based approaches require near-zero initialization for larger clusters (see Sec. VB).

The variance-based approach $C_{\text{var}}^{\text{1QP}}$ has the additional flexibility of applying only to the excitation sector with separate ground-state calculation. However, with separate unitaries, orthogonality between ground-state and 1QP subspace is not enforced during optimization and must be verified. The other two cost functions ($C_{\text{tr}}^{\text{GS,1QP}}$ and $C_{\text{var}}^{\text{GS,1QP}}$) maintain orthogonality throughout optimization via their single-unitary structure. The residual value of C_{var} directly quantifies block-diagonalization quality, making it useful for convergence monitoring.

B. State preparation

The performance of VQE depends on the ansatz choice, which must balance expressibility and trainability. The HVA [33] provides this balance by incorporating the problem structure directly into the circuit design. Motivated by the quantum approximate optimisation algorithm [42] and adiabatic quantum computing [43], HVA is formulated based on the structure of the Hamiltonian under investigation and has been successfully applied for ground-state calculations in quantum many-body systems [6,33,44–47].




FIG. 5. Quantum circuit representation of one layer (green box) of the HVA for the one-dimensional TFIM + LF model. The blue, orange, and red blocks implement $e^{i \sum_{(v,\mu)} \theta_{v,\mu}^{XX} X_v X_\mu}$, $e^{i \sum_v \theta_v^X X_v}$, and $e^{i \sum_v \theta_v^Z Z_v}$, respectively, showing the XX - X - Z ordering within each layer.

1. Hamiltonian variational ansatz

Our recent work [11] investigated HVA performance for ground-state energy calculations. Here, we employ HVA to minimize the cost functions in Eqs. (30) and (32) that block diagonalize the Hamiltonian. For a Hamiltonian decomposed as $H = \sum_\mu H_\mu$, the HVA unitary takes the form

$$U_{\text{HVA}}(\theta) = \prod_l \left(\prod_\mu e^{i \theta_{l,\mu} H_\mu} \right), \quad (33)$$

where l indexes the layers. The HVA construction follows the physical intuition of adiabatic evolution: Each layer implements a discrete evolution step alternating between the unperturbed Hamiltonian H_0 and the perturbation V , with variational parameters that can transform the unperturbed eigenstates into eigenstates of the full Hamiltonian $H = H_0 + xV$. Following the bang-bang digitized adiabatic approach [44], we group commuting terms together and separate noncommuting terms to create blocks that are individually classically solvable, thereby minimizing Trotter error.

For the TFIM + LF [Eq. (1)] on any cluster graph C , the HVA unitary is

$$U_{\text{HVA},C}(\theta) = \prod_{l=1}^{N_l} e^{i \sum_v \theta_{v,l}^Z Z_v} \cdot e^{i \sum_{v,l} \theta_{v,l}^X X_v} \cdot e^{i \sum_{(v,\mu)} \theta_{v,\mu}^{XX} X_v X_\mu}. \quad (34)$$

The XX - X - Z ordering exemplifies this strategy: Since $[X_v X_\mu, X_\rho] = 0$, the XX -coupling and X -field blocks commute and are placed adjacently, forming $H_{XX-X} = -\sum_{(v,\mu)} X_v X_\mu - h_X \sum_v X_v$, which is classically solvable (eigenstates are product states in the X basis). The Z -field block $H_Z = -h_Z \sum_v Z_v$ is also diagonal. This creates two classically solvable parts per layer. On a cluster of N spins, each HVA layer requires N qubits, $2N$ single-qubit rotations (N for the Z block and N for the X block), and one two-qubit gate per nearest-neighbor bond from the XX block ($N-1$ for an open one-dimensional chain). This layer structure is illustrated in Fig. 5 for the one-dimensional case. The general

principle of grouping commuting terms to form classically solvable blocks applies to arbitrary cluster geometries.

We emphasize that the presence of symmetries allows for parameter reduction. Symmetries of the Hamiltonian translate into symmetries of the optimal unitary, constraining the variational parameters. If $[H, \mathcal{S}] = 0$ for symmetry operation $\mathcal{S}$, then any block-diagonalizing unitary U can be chosen to commute with $\mathcal{S}$, imposing parameter constraints. For the one-dimensional chain, reflection symmetry about the center implies $\theta_i^q = \theta_{N+1-i}^q$ (where q labels the parameter type), reducing parameters by roughly half. For NLCE, most clusters have only discrete symmetries (reflection, rotation); periodic chains additionally have translational invariance allowing further reduction.

2. Initialization strategy

Good initialization is crucial for successful optimization. We employ a two-stage protocol. All qubits are initialized in the computational basis state $|0\rangle$, denoted as $|\text{ref}\rangle$. For the high-field polarized phase of the TFIM, this coincides with the unperturbed ground state $|\Phi^{[0]}\rangle$ of H_0 . The initial parameters θ_{initial} determine the starting configuration via $U_{\text{HVA}}(\theta_{\text{initial}})|\text{ref}\rangle$.

Stage 1: We first optimize $U_{\text{HVA}}(\theta)$ for the ground state alone by minimizing energy expectation value (not variance), ensuring convergence to the lowest-energy state via the Rayleigh-Ritz variational principle. Initializing all parameters to small values close to zero ensures reliable convergence.

Stage 2: We use the optimized ground-state parameters θ_{GS} as initialization for optimizing the full subspace cost functions:

$$U_{\text{HVA}}(\theta) X_i |\text{ref}\rangle = U_{\text{HVA}}(\theta) X_i U_{\text{HVA}}(\theta_{\text{GS}})^\dagger U_{\text{HVA}}(\theta_{\text{GS}}) |\text{ref}\rangle, \quad (35)$$

where $X_i |\text{ref}\rangle$ form the 1QP basis of H_0 . If the ground-state unitary provides a good approximation to the true ground state, then $U_{\text{HVA}}(\theta_{\text{GS}}) X_i U_{\text{HVA}}(\theta_{\text{GS}})^\dagger$ acts as an approximate excitation operator. For the one-dimensional TFIM, this initialization already provides nearly exact 1QP states, requiring minimal further optimization. For two-dimensional TFIM and models with longitudinal field, the initialization provides a good starting point that improves convergence compared to random parametrization, though substantial optimization is still required to achieve accurate decoupling.

The ground-state solution provides a physics-informed initialization strategy that leverages the already-optimized ground-state parameters. This approach reflects warm-start principles, beginning optimization with lower initial cost function values compared to random initialization, and proves effective in most cases. However, as we show in Secs. V A 2 and V B, it can lead to local minima for variance-based cost functions when optimizing larger clusters, while trace-based approaches remain robust. In cases where convergence issues arise, initializing parameters near zero provides more stable convergence despite starting with higher initial cost function values.

C. Quantum circuit implementation and measurement requirements

Once VQE optimization is complete and we have obtained the optimized parameters θ_{opt} , we must perform additional quantum measurements to extract the information needed for constructing the PCAT unitary U_{PCAT} and the effective 1QP Hamiltonian $H_{\text{eff}}^{[1]}$. As discussed in Sec. III C, PCAT requires constructing modified states that eliminate projections onto lower-energy subspaces of H_0 . Since VQE provides a unitary transformation U_{VQE} rather than explicit state vectors, we must extract the necessary information through quantum measurements.

We organize this procedure into two parts: Sec. IV C 1 describes the quantum measurements required to obtain three essential $(N+1) \times (N+1)$ matrices ($\mathcal{O}$, S , $\tilde{H}$), and Sec. IV C 2 describes the classical postprocessing steps to construct U_{PCAT} and extract $H_{\text{eff}}^{[1]}$.

1. Quantum measurements

After VQE optimization yields optimal parameters θ_{opt} , we have a set of $N+1$ VQE-prepared states that approximately span the ground-state and 1QP subspace. We label the VQE-prepared ground state as $|\chi^{[0]}\rangle$ and the VQE-prepared 1QP states as $|\chi_i^{[1]}\rangle$ for $i = 1, \dots, N$. We use χ to denote VQE-prepared states, reserving Ψ for exact eigenstates. Here, $U_{\text{VQE}}^{[0]} = U_{\text{HVA}}(\theta_{\text{GS}})$ and $U_{\text{VQE}}^{[1]} = U_{\text{HVA}}(\theta_{1\text{QP}})$, where θ_{GS} and $\theta_{1\text{QP}}$ are the optimized parameters for ground state and excited states. $U_{\text{VQE}}^{[0,1]} = U_{\text{HVA}}(\theta_{\text{GS},1\text{QP}})$ is used for the common unitary obtained after the optimization of the cost functions in Eqs. (31) and (32).

How these states are prepared depends on the cost function used:

- (1) For $C_{\text{tr}}^{\text{GS},1\text{QP}}$ or $C_{\text{var}}^{\text{GS},1\text{QP}}$ (single unitary): All states use the same optimized unitary: $|\chi^{[0]}\rangle = U_{\text{VQE}}^{[0,1]}|\Phi^{[0]}\rangle$, $|\chi_i^{[1]}\rangle = U_{\text{VQE}}^{[0,1]}|\Phi_i^{[1]}\rangle$, where $|\Phi^{[0]}\rangle$ is the unperturbed ground state (which equals $|\text{ref}\rangle$ for our models) and $|\Phi_i^{[1]}\rangle = X_i|\text{ref}\rangle$ are the unperturbed 1QP states.
- (2) For $C_{\text{var}}^{1\text{QP}}$ with separate ground-state optimization: We use different unitaries: $|\chi^{[0]}\rangle = U_{\text{VQE}}^{[0]}|\Phi^{[0]}\rangle$, $|\chi_i^{[1]}\rangle = U_{\text{VQE}}^{[1]}|\Phi_i^{[1]}\rangle$.

The three essential $(N+1) \times (N+1)$ matrices we must measure are as follows:

Matrix $\mathcal{O}$ (overlap with unperturbed states): This $(N+1) \times (N+1)$ matrix contains overlaps between VQE-prepared states (rows) and unperturbed eigenstates $|\Phi_j^{[0]}\rangle$ (columns, $j \in \{0, 1, \dots, N\}$). The ground-state row is $\mathcal{O}_{0j} = \langle\chi^{[0]}|\Phi_j^{[0]}\rangle$, and the 1QP rows are $\mathcal{O}_{ij} = \langle\chi_i^{[1]}|\Phi_j^{[0]}\rangle$ for $i = 1, \dots, N$.

Matrix S (overlap between VQE-prepared states): This matrix contains overlaps between the VQE-prepared states themselves. Elements are $S_{00} = \langle\chi^{[0]}|\chi^{[0]}\rangle$, $S_{0i} = \langle\chi^{[0]}|\chi_i^{[1]}\rangle$, and $S_{ij} = \langle\chi_i^{[1]}|\chi_j^{[1]}\rangle$.

For single-unitary cost functions, $S = \mathbb{1}_{N+1}$ exactly (diagonal by construction). For separate-unitary approaches, off-diagonal elements $S_{0i} = \langle\chi^{[0]}|\chi_i^{[1]}\rangle$ will be nonzero (unless there is an additional symmetry as in the pure TFIM), reflecting imperfect orthogonality.

Matrix $\tilde{H}$ (Hamiltonian expectation values): This matrix contains Hamiltonian matrix elements in the basis of VQE-prepared states. Elements are $\tilde{H}_{00} = \langle\chi^{[0]}|H|\chi^{[0]}\rangle$, $\tilde{H}_{0i} = \langle\chi^{[0]}|H|\chi_i^{[1]}\rangle$, and $\tilde{H}_{ij} = \langle\chi_i^{[1]}|H|\chi_j^{[1]}\rangle$.

The overlap matrix $\mathcal{O}$ and (when needed) the overlap matrix S are measured using the Hadamard test [48]. The Hamiltonian matrix $\tilde{H}$ can be measured more directly: Diagonal elements $\tilde{H}_{00}$ and $\tilde{H}_{ii}$ are expectation values obtained via standard Pauli measurements, while off-diagonal elements can be measured as expectation values on superposition states when both states are prepared using the same unitary. Specifically, by preparing superposition states such as $|\chi_{\pm}\rangle = (|\chi^{[0]}\rangle \pm |\chi_i^{[1]}\rangle)/\sqrt{2}$ or $(|\chi_i^{[1]}\rangle \pm |\chi_j^{[1]}\rangle)/\sqrt{2}$ and measuring $\langle H \rangle$, one obtains the real part of the off-diagonal elements; the imaginary part is obtained similarly using complex superpositions.

For separate-unitary approaches, overlaps and Hamiltonian matrix elements that couple different unitaries require the Hadamard test with composite unitaries. For example, $S_{0i} = \langle\chi^{[0]}|\chi_i^{[1]}\rangle = \langle\Phi^{[0]}|U_{\text{VQE}}^{[0]\dagger}U_{\text{VQE}}^{[1]}|\Phi_i^{[1]}\rangle$ requires Hadamard test. Similarly, Hamiltonian expectation values coupling different unitaries, such as $\tilde{H}_{0i} = \langle\chi^{[0]}|H|\chi_i^{[1]}\rangle = \langle\Phi^{[0]}|U_{\text{VQE}}^{[0]\dagger}H U_{\text{VQE}}^{[1]}|\Phi_i^{[1]}\rangle$, also require Hadamard tests with composite unitaries. Within the 1QP block using the same unitary, $\tilde{H}_{ij} = \langle\chi_i^{[1]}|H|\chi_j^{[1]}\rangle$ can use superposition state measurements.

The total number of measurements scales as $O(N^2)$ for each matrix. However, these measurements are performed only once after VQE optimization is complete, so they do not represent the computational bottleneck; the iterative VQE optimization itself is far more resource demanding.

2. Classical postprocessing

Having obtained the three matrices $\mathcal{O}$, S , and $\tilde{H}$ from quantum measurements, we now proceed entirely classically to construct $U_{\text{PCAT}}^{[1]}$ and extract $H_{\text{eff}}^{[1]}$. This construction proceeds in four steps.

Step 1: Identify the true ground state and handle nonorthogonality. When using $C_{\text{tr}}^{\text{GS},1\text{QP}}$, the optimization does not guarantee that the nominal ground state $|\chi^{[0]}\rangle$ is the true lowest-energy eigenstate. When using separate unitaries ($C_{\text{var}}^{1\text{QP}}$), there may be nonzero overlaps $S_{0i} \neq 0$ between ground and 1QP states. In both cases, we must solve an eigenvalue problem to find orthogonal eigenstates.

For single-unitary cases where $S = \mathbb{1}_{N+1}$, we diagonalize $\tilde{H}$:

$$\tilde{H} = W D W^\dagger, \quad (36)$$

where W is the unitary matrix of eigenvectors and D is the diagonal matrix of eigenvalues.

For separate-unitary cases where $S \neq \mathbb{1}_{N+1}$, we solve the generalized eigenvalue problem:

$$\tilde{H}v = \lambda S v, \quad (37)$$

which can be transformed to standard form and yields eigenvectors that form the columns of matrix W . The lowest (most negative) eigenvalue identifies the true ground state.

This transformation redefines our basis to the true orthonormal eigenstates. The true ground state is

$$|\chi_{\text{true}}^{[0]}\rangle = W_{00}|\chi^{[0]}\rangle + \sum_{i=1}^N W_{0i}|\chi_i^{[1]}\rangle, \quad (38)$$

and the true 1QP eigenstates are

$$|\chi_{i,\text{true}}^{[1]}\rangle = W_{i0}|\chi^{[0]}\rangle + \sum_{j=1}^N W_{ij}|\chi_j^{[1]}\rangle, \quad \text{for } i = 1, \dots, N. \quad (39)$$

The matrices in this new basis are

$$\begin{aligned} \mathcal{O}_{\text{true}} &= W^\dagger \mathcal{O}, \quad \tilde{H}_{\text{true}} = W^\dagger \tilde{H} W = D, \\ S_{\text{true}} &= W^\dagger S W = \mathbb{1}_{N+1}, \end{aligned} \quad (40)$$

where $S_{\text{true}} = \mathbb{1}_{N+1}$ confirms that the transformed states are orthonormal. For the generalized eigenvalue problem with nonorthogonal initial states ($S \neq \mathbb{1}_{N+1}$), the matrix W obtained from the eigenvalue solver is not unitary. The transformation can be equivalently understood as a two-step process: First, preorthogonalize the basis using the symmetric orthogonalization $(S^\dagger S)^{-1/2}$, and then apply a unitary transformation to diagonalize the Hamiltonian in the orthogonal basis. The generalized eigenvalue solver performs both steps simultaneously, directly producing orthonormal eigenvectors with respect to the metric S .

For single-unitary variance cost functions ($C_{\text{var}}^{\text{GS},1\text{QP}}$) with successful optimization, $S = \mathbb{1}_{N+1}$ already, but diagonalizing $\tilde{H}$ can still improve state quality by finding exact eigenstates within the VQE-prepared subspace.

We proceed to Step 2 using the transformed matrices. For notational simplicity, we drop the true subscript and continue to denote the VQE-prepared states as $|\chi^{[0]}\rangle$ and $|\chi_i^{[1]}\rangle$, understanding these now refer to the orthonormalized eigenstates. The matrices $\mathcal{O}$, $\tilde{H}$, and S refer to these corrected, orthonormalized versions.

Step 2: Construct modified overlap matrix using PCAT recursion. According to the PCAT prescription (Sec. III C), modified states eliminate projections onto lower-energy subspaces of H_0 . For the 1QP sector, the recursion from Eq. (17) is

$$|\tilde{\chi}_i^{[1]}\rangle = |\chi_i^{[1]}\rangle - \frac{\langle \Phi^{[0]} | \chi_i^{[1]} \rangle}{\langle \Phi^{[0]} | \chi^{[0]} \rangle} |\chi^{[0]}\rangle \quad (41)$$

for $i = 1, \dots, N$. Taking the overlap with an unperturbed state $|\Phi^j\rangle$ yields

$$\begin{aligned} \langle \Phi_j^{[0]} | \tilde{\chi}_i^{[1]} \rangle &= \langle \Phi_j^{[0]} | \chi_i^{[1]} \rangle - \frac{\langle \Phi^{[0]} | \chi_i^{[1]} \rangle}{\langle \Phi^{[0]} | \chi^{[0]} \rangle} \langle \Phi_j^{[0]} | \chi^{[0]} \rangle \\ &= \mathcal{O}_{ij} - \frac{\mathcal{O}_{i0}}{\mathcal{O}_{00}} \mathcal{O}_{0j}. \end{aligned} \quad (42)$$

The modified overlap matrix $\tilde{\mathcal{O}}$ for the 1QP sector is an $N \times (N+1)$ matrix:

$$\tilde{\mathcal{O}}_{ij} = \mathcal{O}_{ij} - \frac{\mathcal{O}_{i0}}{\mathcal{O}_{00}} \mathcal{O}_{0j} \quad (43)$$

for rows $i = 1, \dots, N$ (1QP states) and columns $j = 0, 1, \dots, N$ (all unperturbed states). Restricting to the 1QP subspace of H_0 (columns $j = 1, \dots, N$), we obtain the $N \times N$ modified overlap matrix:

$$\tilde{\chi}_{ik}^{[1]} = \mathcal{O}_{i,k} - \frac{\mathcal{O}_{i0}}{\mathcal{O}_{00}} \mathcal{O}_{0,k} \quad (44)$$

for $i, k = 1, \dots, N$. This ensures that $P^{[0]}|\tilde{\chi}_i^{[1]}\rangle = 0$ for all 1QP states.

Step 3: Construct the PCAT correction unitary. The PCAT correction within the 1QP subspace is the $N \times N$ unitary matrix:

$$V^{[1]} = \tilde{\chi}^{[1]\dagger} (\tilde{\chi}^{[1]} \tilde{\chi}^{[1]\dagger})^{-1/2}. \quad (45)$$

This corresponds to the transformation described in Sec. III C, where the full PCAT unitary for the 1QP block is $U_{\text{PCAT}}^{[1]} = U^{[1]} V^{[1]}$, with $U^{[1]}$ being the matrix whose columns are the VQE-approximated 1QP eigenstates. Here, $V^{[1]}$ represents the additional correction needed beyond the VQE unitary to enforce cluster additivity.

Step 4: Extract the effective 1QP Hamiltonian. The effective 1QP Hamiltonian is obtained by restricting the Hamiltonian matrix (from Step 1) to the 1QP sector and applying the PCAT correction. We extract the $N \times N$ subblock corresponding to the 1QP sector (rows and columns with indices $1, \dots, N$) from the $(N+1) \times (N+1)$ transformed Hamiltonian. This restriction removes the ground state (index 0) and isolates the 1QP subspace:

$$\tilde{H}_{ij}^{[1]} = \tilde{H}_{i,j}, \quad \text{for } i, j \in \{1, \dots, N\}. \quad (46)$$

The effective 1QP Hamiltonian is then

$$H_{\text{eff}}^{[1]} = V^{[1]\dagger} \tilde{H}^{[1]} V^{[1]} - E^{[0]}. \quad (47)$$

This polynomial-time classical computation transforms the VQE output, encapsulated in the three matrices $\mathcal{O}$, S , and $\tilde{H}$, into the cluster-additive effective Hamiltonian required for NLCE. The resulting $H_{\text{eff}}^{[1]}$ is an $N \times N$ matrix in the 1QP basis, which after subtracting the ground-state energy $E^{[0]}$ gives the effective 1QP Hamiltonian for the cluster as defined in Sec. III:

$$H_{\text{eff},C}^{[1]} = (U_{\text{PCAT},C}^\dagger H U_{\text{PCAT},C})^{[11]} - E_C^{[0]}. \quad (48)$$

This quantity is then used in the NLCE embedding procedure [Eq. (7)] to obtain the 1QP dispersion in the thermodynamic limit.

V. APPLICATIONS

We now apply the NLCE + VQE approach of the previous sections to compute 1QP excitation energies in the TFIM. We investigate the one-dimensional chain at the critical point, the two-dimensional square lattice at its quantum critical point, and additionally the one-dimensional TFIM with longitudinal field (TFIM + LF) in the disordered phase. The longitudinal field fundamentally changes the model by breaking parity symmetry, coupling ground state to excitations, and making the one-dimensional model nonintegrable.

To benchmark NLCE + VQE, we compare its performance with NLCE + ED, where ED is used to compute properties on finite clusters. Furthermore, we investigate whether the

variance- and trace-based cost functions yield similar results or not. In all three cases, we use the conjugate gradient method for the classical optimization in the VQE process.

A. Transverse-field Ising model

The 1D TFIM is one of the paradigmatic models for studying quantum phase transitions driven by quantum fluctuations. The Hamiltonian of the TFIM is given by

$$H = -h \sum_v Z_v - J \sum_{\langle v', v \rangle} X_{v'} X_v, \quad (49)$$

where Z_v and X_v denote the Pauli matrices acting on site v , J represents the nearest-neighbor Ising exchange coupling, and h the strength of the transverse field. Throughout this work, we set $h = 1$ as our energy unit, such that all energies are measured in units of the transverse-field strength, and we consider ferromagnetic $J > 0$. For $h = 0$, the ground state is doubly degenerate, corresponding to ferromagnetic order along the x axis and spontaneous breaking of the global $\mathbb{Z}_2$ symmetry generated by $\prod_v X_v$. In contrast, for $J = 0$, the ground state is a product state of spins polarized along the z direction, forming a paramagnetic phase that preserves $\mathbb{Z}_2$ symmetry. The competition between these two noncommuting terms leads to a quantum critical point at J_c separating the ordered and disordered phases.

The TFIM in one dimension is exactly solvable via Jordan-Wigner transformation [49], making it an ideal starting benchmark for validating our NLCE + VQE framework. In higher dimensions, the TFIM becomes nonintegrable but retains the same qualitative features. We perform expansions from the high-field polarized phase where the ground state is gapped and 1QP excitations correspond to localized spin flips, and study the quantum critical point where the gap closes.

1. One-dimensional chain

The one-dimensional TFIM undergoes a second-order phase transition at $J_c = 1$. The system is disordered when $J < 1$ and ordered (Ising ferromagnetic) when $J > 1$. We perform an expansion from the disordered phase as the unperturbed limit and present NLCE + VQE results at the critical point.

The exact 1QP dispersion in the thermodynamic limit is given by

$$\omega_{\text{exact}}(k) = 2\sqrt{1 + J^2 - 2J \cos(k)}, \quad (50)$$

where k labels the momentum quantum number of the 1QP state, $k \in [-\pi, \pi]$. At the critical point $J_c = 1$, the excitation gap closes at $k = 0$, resulting in a linear dispersion $\omega(k) \propto |k|$ characteristic of the $(1+1)$ -dimensional Ising universality class. This vanishing gap at criticality provides a stringent test of NLCE convergence.

For the one-dimensional chain, rectangular graphs have dimensions $L_m \times 1 \leq N_{\text{max}}$. The embedding factors in NLCE are equal to 1. Hence, the expansion reduces to the contributions from only two clusters corresponding to N and $N-1$ and the terms corresponding to smaller clusters cancel out. The dispersion $\omega(k)$ in the momentum states k for




FIG. 6. 1QP dispersion $\omega(k)$ for the 1D TFIM at the critical point $J_c = 1$ in the thermodynamic limit. Gradient of red lines with plus signs shows NLCE + VQE results for different maximum cluster sizes, yellow dotted line with circles shows NLCE + ED for $N_{\text{max}} = 16$, and blue dotted line with diamonds shows the exact solution [49]. The inset shows an enlarged view of the region around $k = 0$, where the gap closes. Results use $[N/2]$ HVA layers for finite clusters in NLCE + VQE.

one-dimensional chain can thus be written as

$$\omega(k) = \sum_{v,\mu}^N H_{\text{eff},N,v,\mu}^{[1]} e^{ik(v-\mu)} - \sum_{v',\mu'}^{N-1} H_{\text{eff},N-1,v',\mu'}^{[1]} e^{ik(v'-\mu')}, \quad (51)$$

where $H_{\text{eff},N,v,\mu}^{[1]}$ are the matrix elements of H_{eff} obtained with VQE for a hopping from site v to site μ for a cluster N given by $\langle v | H_{\text{eff},N}^{[1]} | \mu \rangle$. The matrix $H_{\text{eff},N}^{[1]}$ for a cluster of size N is given in Eq. (48).

Remarkably, for the one-dimensional TFIM, the VQE unitary optimized solely to minimize ground-state energy automatically decouples all excitation sectors simultaneously. This is a special property of this model that warrants explanation. After the Jordan-Wigner transformation, the 1D TFIM maps to a free-fermion model described by a quadratic Hamiltonian. In such systems, a single Bogoliubov transformation diagonalizes the entire Hamiltonian across all quasiparticle sectors. The VQE unitary U_{VQE} effectively learns this Bogoliubov transformation when minimizing the ground-state energy. Importantly, while any block-diagonalization method must decouple the ground state from excitations, for a generic nonintegrable system it would generally not decouple all higher quasiparticle sectors from each other. That the ground-state-optimized VQE unitary achieves complete block diagonalization is thus a remarkable feature of the free-fermion structure and integrability of this model, not a general expectation. This could explain why $[N/2]$ HVA layers suffice for convergence: The target transformation may be efficiently expressible within the HVA ansatz structure.

In practice, we compute the ground-state parameters with NLCE + VQE using $[N/2]$ layers of HVA and then construct the 1QP subspace with this unitary. We performed state-vector simulations for 1D TFIM. Figure 6 shows the resulting




FIG. 7. Convergence of NLCE with maximum cluster size for 1D TFIM at the critical point $J_c = 1$ in the thermodynamic limit. The plot shows the dispersion $\omega(0)$ at momentum $k = 0$ obtained using NLCE + VQE (red dotted line with diamonds) and NLCE + ED (blue solid line with circles). Both exhibit identical algebraic convergence, demonstrating that $[N/2]$ HVA layers suffice for the critical point.

dispersion $\omega(k)$ for cluster sizes up to $N_{\max} = 16$, where $N_{\max}$ denotes the largest cluster included in the NLCE expansion. As shown in the figure, one can observe that as the cluster size increases, the NLCE + VQE results converge rapidly toward both the NLCE + ED (for $N_{\max} = 16$) and the exact dispersion, confirming the validity of the convergence of NLCE with implementation of variational approach for finite clusters. The inset of Fig. 6 highlights the region near $k = 0$, the critical momentum where the gap vanishes at $J_c = 1$. Despite the vanishing gap making NLCE convergence more demanding, the progressive approach of NLCE + VQE data toward the exact curve with increasing cluster size demonstrates that the method correctly captures the critical behavior and the linear dispersion $\omega(k) \propto |k|$ near the gapless point.

Quantitative convergence is illustrated in Fig. 7, which plots the dispersion $\omega(0)$ in the thermodynamic limit at momentum $k = 0$ as a function of maximum cluster size $N_{\max}$. Both NLCE + VQE and NLCE + ED converge algebraically toward the exact solution, with VQE-ED differences negligible compared to NLCE truncation errors. The excellent agreement between NLCE + VQE and NLCE + ED confirms that $[N/2]$ HVA layers provide sufficient expressivity for this integrable model. This establishes a proof of principle for the NLCE + VQE approach on an exactly solvable benchmark before proceeding to more challenging cases.

2. Square lattice

The two-dimensional TFIM on a square lattice provides a more stringent test of our framework. In contrast to the one-dimensional case, the 2D TFIM is nonintegrable and lacks an exact analytical solution. Nevertheless, extensive quantum Monte Carlo and series-expansion studies have accurately located the quantum critical point at $J_c = 0.328$ [50,51]. The transition separates a ferromagnetically ordered phase at large J from a quantum paramagnet polarized along the transverse field at small J .

Near the critical point, the system exhibits universal scaling behavior belonging to the $(2+1)$ -dimensional Ising universality class. At the quantum critical point, the excitation gap $\Delta(\vec{k})$ closes at momentum $(k_x, k_y) = (0, 0)$, resulting in critical fluctuations and a linear dispersion $\omega(\vec{k}) \propto |\vec{k}|$. The vanishing gap at criticality presents a challenging test for NLCE convergence, as larger cluster sizes are required to capture the diverging correlation length.

Using the NLCE + VQE framework, we compute the 1QP dispersion for the square lattice in the thermodynamic limit at the quantum critical point $J_c = 0.328$, as shown in Fig. 8. The figure presents a comparison between dispersions obtained from NLCE + VQE and NLCE + ED for different NLCE orders $N_{\max}$. For these calculations, we employ the rectangular graph expansion defined in Eq. (27) and use $[N/2]$ layers of the HVA. The energy variance cost function given in Eq. (30) is minimized to obtain the lowest 1QP excitation energy for each cluster. The resulting NLCE + VQE data are compared with ED-based NLCE results and with a high-order series expansion up to 12th order [50], which we obtained using perturbative continuous unitary transformations [22], serving as our benchmark.

Overall, the NLCE + VQE dispersions show very good agreement with both NLCE + ED and the series-expansion benchmark, demonstrating that the VQE-based evaluation of cluster properties can accurately capture the low-energy excitations of the model. The inset highlights the behavior near $(k_x, k_y) = (0, 0)$, which corresponds to the point of gap closing at the critical point. In this region, the NLCE convergence becomes increasingly demanding because the excitation gap approaches zero, and larger cluster sizes are required to obtain quantitatively reliable results. Nevertheless, the systematic improvement with increasing $N_{\max}$ indicates that the NLCE + VQE approach is capable of reproducing the correct critical behavior (just as NLCE + ED) when sufficient cluster sizes are included. At $N_{\max} = 14$, both NLCE + VQE and NLCE + ED agree to within $\Delta\omega \lesssim 10^{-2}$ at the critical momentum, demonstrating that $[N/2]$ HVA layers remain sufficient even at criticality for the 2D TFIM. Let us note that scaling of the results in the cluster size is needed to obtain the gapless behavior of the quantum critical point, which is expected due to the diverging correlation length.

Figure 9 quantifies this convergence behavior, showing that the dispersion $\omega(0, 0)$ decreases systematically with increasing $N_{\max}$ for all cost functions tested, confirming both the NLCE convergence and the equivalence of cost function choices at the critical point.

Cost function comparison. For the 2D TFIM, we tested four cost functions for the 1QP sector: variance based ($C_{\text{var}}^{\text{1QP}}$, $C_{\text{var}}^{\text{GS,1QP}}$) and trace based ($C_{\text{tr}}^{\text{1QP}}$, $C_{\text{tr}}^{\text{GS,1QP}}$). Figure 9 shows that all cost functions yield essentially identical NLCE convergence, with variations $\lesssim 10^{-3}$ in the thermodynamic-limit dispersion at $(k_x, k_y) = (0, 0)$. Importantly, these cost-function-induced variations are significantly smaller than the NLCE truncation errors (differences between successive $N_{\max}$ values, typically $\sim 5 \times 10^{-2}$) and comparable to or smaller than the VQE-ED differences. However, for the 5×2 cluster, variance-based cost functions with ground-state initialization encountered convergence difficulties, while trace-based




FIG. 8. 1QP dispersion $\omega(k_x, k_y)$ for the 2D TFIM on a square lattice at the quantum critical point $J_c = 0.328$ [50,51]. Gradient of solid blue lines shows NLCE + ED results and gradient of dashed red lines shows NLCE + VQE results for different maximum cluster sizes $N_{\max}$. The inset shows an enlarged view of the critical region near $(k_x, k_y) = (0, 0)$, where the gap closes. These results use $[N/2]$ HVA layers for finite clusters in NLCE + VQE. The solid black line refers to the bare series expansion of order 12 about the high-field limit.

approaches remained robust. Near-zero initialization resolved this issue. The underlying causes of this initialization sensitivity are analyzed in detail for TFIM + LF (Sec. VB).

For the TFIM + LF case (Sec. VB), where both parity symmetry breaking and nonintegrability make optimization more challenging, we perform a more detailed analysis of initialization sensitivity and its impact on convergence.

B. Transverse-field Ising chain in longitudinal field

The addition of a longitudinal magnetic field h_l to the TFIM fundamentally alters the model's physics. The




FIG. 9. Cost function comparison for the 2D TFIM at the critical point $J_c = 0.328$ [50,51]. The plot shows $\omega(0, 0)$ for four cost functions: (1QP) variance, (GS + 1QP) variance, (1QP) trace, and (GS + 1QP) trace. All functions converge identically, with NLCE + ED (blue) shown for reference. The x axis shows the highest cluster size $N_{\max}$ used in the NLCE calculation.

Hamiltonian becomes

$$H = -h \sum_v Z_v - J \sum_{(v',v)} X_{v'} X_v - h_l \sum_v X_v, \quad (52)$$

where the longitudinal field h_l explicitly breaks the $\mathbb{Z}_2$ parity symmetry. This symmetry breaking has important consequences: It couples the ground state to excited states, removes the free-fermion structure that made the pure TFIM chain exactly solvable, and renders the model nonintegrable. We investigate the disordered phase at $J = 0.5$, $h_l = 0.5$ (with $h = 1$).

PCAT becomes crucial for this model, as the 2b transformation alone becomes insufficient due to its allowance of hopping processes between disconnected clusters, as illustrated in Fig. 4. The longitudinal field induces large overlaps between the unperturbed ground state and the 1QP eigenstates, necessitating careful treatment to not violate cluster additivity. This is explained in detail in Sec. IIIC. The aforementioned issue is further explained in Fig. 10, which shows the 1QP dispersion for $N_{\max} = 10$ at $J = 0.5$ and $h_l = 0.5$, obtained using NLCE + ED. The blue solid line in this figure corresponds to the effective Hamiltonian H_{eff} derived from the 2b transformation, while the orange solid line represents the PCAT result. The PCAT visibly corrects the 2b data, yielding a much smoother dispersion. Under the 2b transformation, a dressed quasiparticle can nonphysically hop between clusters that do not share a bond, violating cluster additivity. PCAT systematically removes these spurious couplings and restores the required additivity condition for NLCE. Such nonphysical oscillations become increasingly pronounced with larger cluster sizes and are already significant at $N = 9$. Critically, the 2b transformation without PCAT correction fails to produce convergent NLCE results in this regime: The NLCE sum diverges as cluster size increases, making thermodynamic-limit extrapolation impossible. Only after PCAT correction does the




FIG. 10. Dispersion $\omega(k)$ for the TFIM + LF obtained with NLCE + ED, where the effective Hamiltonian is constructed using the 2b transformation, compared against the result obtained using PCAT. Results are shown for $N_{\max} = 10$ spins, $J = 0.5$, and $h_l = 0.5$. Blue line: Effective Hamiltonian from 2b transformation alone shows unphysical oscillations due to non-cluster-additive (nonlocal) hopping. Orange line: PCAT-corrected result yields smooth dispersion restoring cluster additivity.

NLCE converge, demonstrating that enforcing cluster additivity is essential for obtaining meaningful thermodynamic-limit results.

1. Convergence results and cost function analysis

Since we consider the TFIM + LF on a one-dimensional chain, the NLCE again reduces to the expression in Eq. (51). We now perform a detailed comparison of cost functions and initialization strategies for the 1D TFIM + LF. The combination of broken parity symmetry and nonintegrability makes this a particularly challenging optimization problem. We investigate convergence speed, robustness to local minima, and layer requirements for different cost functions.

Figures 11 and 12 show NLCE convergence comparing $[N/2]$ and N layers at both $k = 0$ and $k = \pi$ for all three cost functions: variance of 1QP states alone ($C_{\text{var}}^{1\text{QP}}$), variance of combined ground and 1QP states ($C_{\text{var}}^{\text{GS},1\text{QP}}$), and trace of the combined subspace ($C_{\text{tr}}^{\text{GS},1\text{QP}}$). Regarding layer requirements, both $[N/2]$ and N layers achieve NLCE convergence with increasing cluster size. At $k = 0$ (Fig. 11), N layers provide improved accuracy for intermediate cluster sizes ($N_{\max} \approx 7-9$), though both layer choices converge to comparable accuracy by $N_{\max} = 10$. The transient advantage of N layers contrasts with the pure TFIM (Sec. V A 2) where $[N/2]$ layers proved sufficient throughout. At $k = \pi$ (Fig. 12), both layer choices show comparable accuracy throughout, with results converging to NLCE + ED at the largest cluster sizes.

The effective Hamiltonian construction depends on accurately identifying the low-energy eigenspace, which all cost functions accomplish when optimization converges. However, for the most challenging case ($N = 10$ spins, N layers), we observed that the initialization strategy using ground-state parameters led to convergence failures for variance-based cost functions, while the trace-based cost function remained robust. Similar behavior was also observed for the 5×2




FIG. 11. Dispersion at momentum $k = 0$, denoted as $\omega(0)$, in the thermodynamic limit for the TFIM + LF at $J = 0.5$ and $h_l = 0.5$ as a function of the maximum cluster size $N_{\max}$. NLCE + VQE results are shown for all three cost functions using HVA with $[N/2]$ layers (blue lines with diamonds) and N layers (red lines with circles). The black solid line with stars represents NLCE + ED results.

cluster in the 2D TFIM (Sec. V A 2). For context, successful convergence typically achieves variance values $C_{\text{var}} \sim 0.01$, corresponding to subspace infidelities $1 - \mathcal{F} \sim 10^{-4}$ relative to the target eigenspace, while optimizations trapped in local minima exhibited $C_{\text{var}} \sim 0.5$, indicating substantial contamination from higher excitation sectors. Specifically, $C_{\text{var}}^{1\text{QP}}$ and $C_{\text{var}}^{\text{GS},1\text{QP}}$ became trapped in local minima with cost function values more than 10 times higher than successful optimizations. Initializing all variational parameters near zero resolved this issue for all cost functions, though requiring more optimization iterations due to the less favorable starting point. The initialization sensitivity reveals important structure in the optimization landscape. Figure 13 provides insight for $N = 10$ by tracking where the VQE state resides during optimization.




FIG. 12. Dispersion at momentum $k = \pi$, denoted as $\omega(\pi)$, in the thermodynamic limit for the TFIM + LF at $J = 0.5$ and $h_l = 0.5$ as a function of the maximum cluster size $N_{\max}$. NLCE + VQE results are shown for all three cost functions using HVA with $[N/2]$ layers (blue lines with diamonds) and N layers (red lines with circles). The black solid line with stars represents NLCE + ED results.




FIG. 13. Energy decomposition during VQE optimization for TFIM + LF for $N_{\max} = 10$ spins, $J = 0.5$, and $h_l = 0.5$ using N layers of HVA with conjugate gradient optimization. For selected optimization iterations, we plot the norm $\|P_{\text{VQE}}|\Psi_n\rangle\|$ of the projection of the optimized VQE subspace onto exact energy eigenstates $|\Psi_n\rangle$ as a function of eigenstate energy E_n , where P_{VQE} is the projector onto the optimized subspace. Dashed lines: after 2 conjugate gradient iterations (initial state); solid lines: after 1000 iterations (attempted convergence). Blue lines: Trace cost function $C_{\text{tr}}^{\text{GS,1QP}}$ shows successful convergence with norm values concentrating in ground-state and 1QP sectors, with negligible projection onto higher excitations. Green lines: variance cost functions $C_{\text{var}}^{\text{1QP}}$ and $C_{\text{var}}^{\text{GS,1QP}}$ show failed convergence with norm values $\|P_{\text{VQE}}|\Psi_n\rangle\| \approx 0.5$ to 0.6 persisting in 2QP and 3QP sectors (higher-energy region around $E \approx -8$ to -6), barely evolving from 2 to 1000 iterations. This indicates the optimizer is trapped in a local minimum.

We decompose the optimized VQE subspace at each iteration into exact eigenstates (obtained via ED) and plot the norm $\|P_{\text{VQE}}|\Psi_n\rangle\|$ versus eigenstate energy E_n , where P_{VQE} is the projector onto the optimized subspace.

The trace cost function (successful case) shows clean behavior: The norm values concentrate progressively in the ground-state and 1QP sectors, with negligible projection ($\|P_{\text{VQE}}|\Psi_n\rangle\| < 0.1$, corresponding to $< 1\%$ probability weight) onto higher-energy excitations by convergence. The variance cost functions (failed cases) tell a different story: Norm values of $\|P_{\text{VQE}}|\Psi_n\rangle\| \approx 0.5$ to 0.6 (corresponding to approximately 25% to 35% probability weight) persist in two-quasiparticle (2QP) and three-quasiparticle (3QP) sectors throughout optimization, barely changing from initial to final iterations. The optimizer cannot decouple these higher sectors despite targeting only ground and 1QP states.

This reveals the underlying problem: The two cost functions have different optimization landscapes. Starting from ground-state parameters places the variance optimization in a local minimum where the VQE states are not proper eigenstates but remain superpositions across ground, 1QP, 2QP, and 3QP sectors. While variance vanishes for any eigenstate regardless of which sector it belongs to, here the optimizer fails to find eigenstates altogether; the states remain mixed rather than collapsing to eigenstates of the full Hamiltonian. The trace cost function avoids this trap, likely




FIG. 14. Dispersion $\omega(k)$ in the thermodynamic limit for TFIM + LF at $J = 0.5$ and $h_l = 0.5$. Blue lines show NLCE + VQE results for different maximum cluster sizes up to $N_{\max} = 10$, while the red dotted line shows NLCE + ED for $N_{\max} = 12$. VQE uses N layers of HVA with variance $C_{\text{var}}^{\text{GS,1QP}}$ cost function.

because direct energy minimization provides stronger gradients away from such mixed configurations than variance minimization.

In Fig. 14, we show the 1QP dispersion $\omega(k)$ in the thermodynamic limit obtained from NLCE + VQE (using N layers and variance $C_{\text{var}}^{\text{GS,1QP}}$ with near-zero initialization) compared to NLCE + ED. NLCE + VQE converges systematically toward NLCE + ED with increasing maximum cluster size $N_{\max}$. Convergence is slower than for both the pure 1D TFIM and the 2D TFIM cases. This reflects the increased optimization difficulty: The longitudinal field breaks integrability and parity symmetry, coupling odd and even quasiparticle sectors and creating a more complex optimization landscape than in the pure TFIM cases. The interplay between optimization difficulty and physical correlation structure requires further investigation: The longitudinal field may increase entanglement within clusters, or the slower NLCE convergence may stem primarily from the harder optimization problem rather than from a genuinely longer correlation length. Determining the relative contributions of these effects remains an open question.

VI. CONCLUSIONS

In this work, we have developed a hybrid quantum-classical method for computing quasiparticle excitation energies in the thermodynamic limit, introducing a quantum algorithm capable of constructing cluster-additive effective Hamiltonians for excited states. The key challenge is ensuring cluster additivity for degenerate subspaces, a property not automatically satisfied by variational optimization. This is resolved by integrating the PCAT [31] with VQE. This integration is nontrivial: VQE provides the essential information (energy expectations and state overlaps) needed to construct PCAT without exponentially costly state tomography, enabling NLCE convergence for excited states with quantum algorithms as cluster solvers.

Our benchmarks on the TFIM demonstrate that VQE with the HVA achieves convergence using $\lceil N/2 \rceil$ layers for the pure TFIM in both one and two dimensions. For the TFIM with longitudinal field, where parity symmetry breaking makes PCAT essential, the situation becomes more nuanced: While $\lceil N/2 \rceil$ layers show convergence with increasing cluster size, using N layers provides improved accuracy.

For the variance-based cost function, we observe that near-zero initialization outperforms the ground-state-initialized strategy, while the trace cost function remains robust to ground-state initialization. This reveals that the two cost functions have qualitatively different optimization landscapes: The variance landscape contains local minima near ground-state parameters that trap the optimizer, while the trace landscape does not. Cluster sizes beyond $N \sim 10$ become challenging for this combination of variance cost function with suboptimal initialization.

The PCAT framework provides several advantages beyond ensuring cluster additivity. By constructing the transformation from measured energy expectations and state overlaps, it requires $\mathcal{O}(N^2)$ measurements per cluster, polynomial overhead that scales favorably compared to exponential state space growth.

Beyond the specific application to quasiparticle excitations, this work demonstrates that the PCAT framework extends naturally to thermodynamic-limit calculations through appropriate enforcement of cluster additivity. Importantly, while we have focused on VQE implementation, the PCAT construction is not intrinsically tied to variational optimization. The framework relies on measuring energy expectations and state overlaps, as detailed in Sec. IV C 1, making it immediately ready for application when fault-tolerant quantum protocols achieve quantum advantage. The same measurement protocol applies equally to quantum phase estimation, adiabatic state preparation, and quantum annealing. This algorithmic flexibility positions the approach to scale with advancing quantum hardware and ensures it remains relevant even if variational methods face fundamental limitations such as barren plateaus in certain regimes. The measurement protocol thus provides a practical pathway for implementation across different quantum algorithmic paradigms.

Implementation of NLCE + VQE with PCAT on real quantum hardware is currently underway to assess performance under realistic noise and finite measurement sampling. We will investigate how gate errors and shot noise propagate through the PCAT construction and subsequent NLCE summations. To mitigate error accumulation, the hybrid classical-quantum nature suggests a strategy where the largest clusters, those contributing most to NLCE convergence but being classically intractable, are treated with VQE, while smaller clusters use exact diagonalization to minimize accumulated errors.

Several extensions present themselves naturally. Computing two-quasiparticle sectors would enable investigation of bound states and scattering continua, though the quadratically

growing subspace dimensions make this significantly more challenging. Dynamical structure factors, the natural extension of static structure factors to finite-frequency response, are directly accessible through our framework by evaluating local observables between the ground state and quasiparticle excitation states constructed via NLCE. These spectral functions are directly measurable in inelastic neutron or light scattering experiments and provide detailed information about the elementary excitation spectrum. The framework's basis-independent formulation makes it applicable to various quantum lattice models, though we note that gapped systems where quasiparticle excitations are well defined represent the most natural domain of applicability given the current formalism.

For conventional NLCE with classical algorithms, the logarithm $S = \ln T$ of the transformation possesses cluster-additive structure when properly constructed [31]. Whether analogous principles can be established for constructing approximate unitaries for larger clusters from solutions on smaller constituent subgraphs in variational circuits, potentially extending beyond HVA to other ansätze, remains an open question that could improve robustness to optimization errors and help circumvent local minima.

The challenges encountered already with TFIM with longitudinal field (where optimization becomes difficult) underscore the importance of systematic investigation across diverse lattice models. Future applications to frustrated quantum magnets represent a long-term vision where quantum approaches might prove particularly valuable: Frustrated systems suffer from severe sign problems that make quantum Monte Carlo exponentially harder with system size, rendering many frustrated systems practically inaccessible. While NLCE converges more slowly for frustrated systems and requires larger clusters, quantum algorithms can potentially access these larger cluster sizes where classical exact diagonalization fails. The combination of NLCE with quantum algorithms could thus provide access to parameter regimes where classical methods cannot reach, offering a domain where quantum advantages emerge naturally. Establishing the approach's general applicability and practical utility across different model classes remains essential.

ACKNOWLEDGMENTS

Sumeet and M.H. acknowledge Lucas Marti for discussions on quantum hardware application. We acknowledge the support by the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

DATA AVAILABILITY

The data that support the findings of this article are openly available in Zenodo [52]

[1] 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).

- [2] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [3] P. J. J. O'Malley *et al.*, Scalable quantum simulation of molecular energies, *Phys. Rev. X* **6**, 031007 (2016).
- [4] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [5] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum chemistry calculations on a trapped-ion quantum simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [6] R. Wiersema, C. Zhou, Y. De Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the Hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [7] J. Kattemölle and J. van Wezel, Variational quantum eigensolver for the Heisenberg antiferromagnet on the kagome lattice, *Phys. Rev. B* **106**, 214429 (2022).
- [8] A. C. Y. Li, M. S. Alam, T. Iadecola, A. Jahin, J. Job, D. M. Kurkcuoglu, R. Li, P. P. Orth, A. B. Özgüler, G. N. Perdue, and N. M. Tubman, Benchmarking variational quantum eigensolvers for the square-octagon-lattice Kitaev model, *Phys. Rev. Res.* **5**, 033071 (2023).
- [9] R. Watanabe, K. Fujii, and H. Ueda, Variational quantum eigensolver with embedded entanglement using a tensor-network ansatz, *Phys. Rev. Res.* **6**, 023009 (2024).
- [10] J. Wang and R. Jaiswal, Scalable quantum ground state preparation of the Heisenberg model: A variational quantum eigensolver approach, *arXiv:2308.12020*.
- [11] Sumeet, M. Hörmann, and K. P. Schmidt, Hybrid quantum-classical algorithm for the transverse-field Ising model in the thermodynamic limit, *Phys. Rev. B* **110**, 155128 (2024).
- [12] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [13] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [14] O. Higgott, D. Wang, and S. Brierley, Variational quantum computation of excited states, *Quantum* **3**, 156 (2019).
- [15] P. J. Ollitrault, A. Kandala, C. F. Chen, P. K. Barkoutsos, A. Mezzacapo, M. Pistoia, S. Sheldon, S. Woerner, J. M. Gambetta, and I. Tavernelli, Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor, *Phys. Rev. Res.* **2**, 043140 (2020).
- [16] K. M. Nakanishi, K. Mitarai, and K. Fujii, Subspace-search variational quantum eigensolver for excited states, *Phys. Rev. Res.* **1**, 033062 (2019).
- [17] C. L. Benavides-Riveros, L. Chen, C. Schilling, S. Mantilla, and S. Pittalis, Excitations of quantum many-body systems via purified ensembles: A unitary-coupled-cluster-based approach, *Phys. Rev. Lett.* **129**, 066401 (2022).
- [18] Y. Guo, T. Angelides, K. Jansen, and S. Kühn, Concurrent VQE for simulating excited states of the Schwinger model, *arXiv:2407.15629*.
- [19] S. Bravyi, D. P. DiVincenzo, and D. Loss, Schrieffer-Wolff transformation for quantum many-body systems, *Ann. Phys.* **326**, 2793 (2011).
- [20] F. Wegner, Flow-equations for Hamiltonians, *Ann. Phys.* **506**, 77 (1994).
- [21] S. D. Glazek and K. G. Wilson, Renormalization of Hamiltonians, *Phys. Rev. D* **48**, 5863 (1993).
- [22] C. Knetter and G. S. Uhrig, Perturbation theory by flow equations: Dimerized and frustrated $S = 1/2$ chain, *Eur. Phys. J. B* **13**, 209 (2000).
- [23] C. J. Morningstar and M. Weinstein, Contractor renormalization group technology and exact Hamiltonian real-space renormalization group transformations, *Phys. Rev. D* **54**, 4131 (1996).
- [24] J. Oitmaa, C. Hamer, and W. Zheng, *Series Expansion Methods for Strongly Interacting Lattice Models* (Cambridge University Press, Cambridge, England, 2006).
- [25] M. P. Gelfand, Series expansions for excited states of quantum lattice models, *Solid State Commun.* **98**, 11 (1996).
- [26] A. C. Irving and C. J. Hamer, Linked cluster expansions for U(1) lattice gauge theory in $2 + 1$ and $3 + 1$ dimensions, *Nucl. Phys. B* **235**, 358 (1984).
- [27] B. Tang, E. Khatami, and M. Rigol, A short introduction to numerical linked-cluster expansions, *Comput. Phys. Commun.* **184**, 557 (2013).
- [28] M. Rigol, T. Bryant, and R. R. P. Singh, Numerical linked-cluster approach to quantum lattice models, *Phys. Rev. Lett.* **97**, 187202 (2006).
- [29] K. Coester, S. Clever, F. Herbst, S. Capponi, and K. P. Schmidt, A generalized perspective on non-perturbative linked-cluster expansions, *Europhys. Lett.* **110**, 20006 (2015).
- [30] R. Jaiswal, I. Lovas, and L. Balents, Simulating a quasiparticle on a quantum device, *Phys. Rev. A* **111**, 012629 (2025).
- [31] M. Hörmann and K. P. Schmidt, Projective cluster-additive transformation for quantum lattice models, *SciPost Phys.* **15**, 097 (2023).
- [32] C. Knetter, K. P. Schmidt, and G. S. Uhrig, The structure of operators in effective particle-conserving models, *J. Phys. A: Math. Gen.* **36**, 7889 (2003).
- [33] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
- [34] L. S. Cederbaum, J. Schirmer, and H. D. Meyer, Block diagonalisation of Hermitian matrices, *J. Phys. A: Math. Gen.* **22**, 2427 (1989).
- [35] I. G. Enting, Series expansions from the finite lattice method, *Nucl. Phys. B Proc. Suppl.* **47**, 180 (1996).
- [36] J. des Cloizeaux, Extension d'une formule de Lagrange à des problèmes de valeurs propres, *Nucl. Phys.* **20**, 321 (1960).
- [37] M. Takahashi, Half-filled Hubbard model at low temperature, *J. Phys. C* **10**, 1289 (1977).
- [38] I. Shavitt and L. T. Redmon, Quasidegenerate perturbation theories: A canonical van Vleck formalism and its relationship to other approaches, *J. Chem. Phys.* **73**, 5711 (1980).
- [39] T. Kato, On the convergence of the perturbation method. I, *Prog. Theor. Phys.* **4**, 514 (1949).
- [40] A. Kardashin, A. Uvarov, D. Yudin, and J. Biamonte, Certified variational quantum algorithms for eigenstate preparation, *Phys. Rev. A* **102**, 052610 (2020).

- [41] D. B. Zhang, B. L. Chen, Z. H. Yuan, and T. Yin, Variational quantum eigensolvers by variance minimization, *Chin. Phys. B* **31**, 120301 (2022).
- [42] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [43] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [44] R. Wang, T. H. Hsieh, and G. Vidal, Bang-bang algorithms for quantum many-body ground states: A tensor network exploration, *Phys. Rev. B* **106**, 195133 (2022).
- [45] B. Anselme Martin, P. Simon, and M. J. Rančić, Simulating strongly interacting Hubbard chains with the variational Hamiltonian ansatz on a quantum computer, *Phys. Rev. Res.* **4**, 023190 (2022).
- [46] A. A. Mele, G. B. Mbeng, G. E. Santoro, M. Collura, and P. Torta, Avoiding barren plateaus via transferability of smooth solutions in a Hamiltonian variational ansatz, *Phys. Rev. A* **106**, L060401 (2022).
- [47] C. Y. Park and N. Killoran, Hamiltonian variational ansatz without barren plateaus, *Quantum* **8**, 1239 (2024).
- [48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary ed. (Cambridge University Press, Cambridge, England, 2010).
- [49] P. Pfeuty, The one-dimensional Ising model with a transverse field, *Ann. Phys.* **57**, 79 (1970).
- [50] H. X. He, C. J. Hamer, and J. Oitmaa, High-temperature series expansions for the (2+1)-dimensional Ising model, *J. Phys. A: Math. Gen.* **23**, 1775 (1990).
- [51] S. Hesselmann and S. Wessel, Thermal Ising transitions in the vicinity of two-dimensional quantum critical points, *Phys. Rev. B* **93**, 155157 (2016).
- [52] Sumeet, M. Hörmann, and K. P. Schmidt, Data for Quantum algorithm for one-quasiparticle excitations in the thermodynamic limit via cluster-additive block diagonalization, Zenodo, 2026, doi:10.5281/zenodo.20775920.