Pauli Propagation Enables Fast Classical Simulation of Strongly Correlated Quantum Systems

Chinmay Shrikhande,[1] Arnab Bachhar,[1] José Aarón Rodríguez-Jiménez,[1] Edward F. Valeev,[2] and Nicholas J. Mayhall[1,*]

[1]Department of Chemistry, Indiana University, Bloomington, Indiana 47405, USA
[2]Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA

(Received 31 May 2026; revised 22 September 2026; accepted 30 September 2026; published 2 November 2026)

Ground-state energy estimation for strongly correlated quantum systems remains a central challenge in computational physics and chemistry. While tensor network methods like Density Matrix Renormalization Group (DMRG) provide efficient solutions for one-dimensional systems, higher-dimensional problems remain difficult. Here we present a variational double bracket flow (vDBF) algorithm that leverages Pauli Propagation, a technique originally developed for classical simulation of quantum circuits, to efficiently approximate ground-state energies. By combining greedy operator selection with coefficient-based fluctuation truncation and energy-variance extrapolation, we obtain sub-1% relative errors for Heisenberg and Hubbard lattices of up to 128 qubits. For a 10×10 Heisenberg lattice (100 qubits), vDBF obtains accurate results in approximately 2 min on a single CPU thread, compared to over 50 h on 64 threads for DMRG. Accurate results are also obtained for systems that are difficult for quantum Monte Carlo: on the maximally frustrated 8×8 J_1–J_2 lattice and the hole-doped 8×8 Hubbard model, the tightest thresholds extrapolate to within a few tenths of a percent of the best published variational estimates. We further test vDBF on the 84-qubit π-valence active space of hexabenzocoronene, where the most reliable extrapolated correlation energies agree with DMRG to within 1%. These results demonstrate that classical simulation techniques developed in the context of quantum advantage benchmarking can provide practical tools for many-body physics.

DOI: 10.1103/kq3v-whcq

I. INTRODUCTION

The simulation of quantum many-body systems provides an indispensable tool for studying molecules and materials, often revealing features that are difficult or impossible to access experimentally. Improving our ability to perform such simulations is therefore one of the most important computational challenges in physics and chemistry, with direct implications for problems ranging from high-temperature superconductivity to heterogeneous catalysis. Unfortunately, strongly correlated systems remain extremely challenging, and no general-purpose classical method exists that can provide both reliable and efficient results for arbitrary systems.

Most traditional techniques for approximating low-energy quantum systems fall into one of three categories: perturbative methods (MPn [1], CCn [2–5], GW [6–8], etc.), stochastic methods (Diffusion Monte Carlo [9], Full Configuration Interaction Quantum Monte Carlo [10,11], Auxiliary-Field quantum Monte Carlo (AFQMC) [12–14], etc.), or tensor network states (DMRG [15–21], Projected Entangled Pair States [22–27], Multi-Scale Entanglement Renormalization Ansatz [28,29], etc.). Each framework excels for certain classes of systems but has fundamental limitations. Perturbative methods are effective when mean-field treatments are qualitatively correct but break down when many-body interactions become strong. Stochastic methods can provide benchmark-quality results for ground-state energies but are limited when the sign problem cannot be controlled [30]. Tensor network states provide ideal compactness for low-dimensional systems but become less efficient in higher dimensions. Despite the breadth of problems treatable by these and other methods, systems that combine strong correlation with high dimensionality remain challenging for classical simulation.

In light of rapid advances in quantum information science, quantum computation has become an appealing approach for such classically challenging problems. Quantum hardware continues to improve [31–33], with growing qubit counts and decreasing error rates, while algorithmic developments have reduced the quantum resources required for accurate simulations. However, current quantum hardware and algorithms are not yet capable of outperforming classical methods for many-body simulation, and the realization of an exponential speedup remains uncertain [34].

*Contact author: nmayhall@iu.edu

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

While quantum simulation has not yet demonstrated a definitive advantage over classical methods, its development has driven significant advances in classical simulation techniques. The 2023 IBM Eagle experiment [32], which claimed evidence for quantum utility in simulating kicked Ising dynamics, sparked rapid development of classical methods that ultimately matched or exceeded the quantum results [35–37]. Central to these classical responses were algorithms for approximating Heisenberg evolution of Pauli operators (Sparse Pauli Dynamics (SPD), Pauli Propagation (PP), and related methods), which proved remarkably effective at simulating deep quantum circuits [35,36,38–53]. These tools were developed to benchmark quantum hardware, but their potential for classical ground-state energy calculations remains largely unexplored. In this work, we combine Pauli Propagation with a variational double bracket flow (vDBF) approach for ground-state energy estimation. Double bracket flows have recently emerged as a powerful framework in quantum computing [54], serving as the foundation for algorithms such as Double-Bracket Quantum Imaginary Time Evolution (DB-QITE) [55], optimized initializations for variational ground-state optimization [51], and a variety of other advanced quantum protocols [56–58].

By adapting this quantum-inspired framework for classical execution, we benchmark vDBF on the 10×10 antiferromagnetic Heisenberg model, the frustrated 8×8 $J_1 - J_2$ lattice, the 8×8 Hubbard model at half-filling and at one-eighth hole doping, and the π-valence active space of hexabenzocoronene (HBC).

II. THEORY

To simulate the results of a quantum computation, one generally aims to compute the expectation values of observables for states that have been evolved by a sequence of unitaries,

$$\langle O \rangle = \langle 0 | U_1^\dagger U_2^\dagger \cdots U_n^\dagger O U_n \cdots U_2 U_1 | 0 \rangle. \tag{1}$$

While this can be expressed equivalently in either the Schrödinger or Heisenberg picture, the choice to evolve the state, $\langle \psi_n | O | \psi_n \rangle$, or the operator, $\langle 0 | O_n | 0 \rangle$, or both [36] presents distinct opportunities for approximations.

Due, in large part, to the fact that the wave function is generally a lower-rank object, the Schrödinger picture is typically preferable. However, considerations of sparsity can often lean in favor of the Heisenberg picture. Physically relevant Hamiltonians are generally sparse, while ground states are generally not [59]. Furthermore, when the state is no longer pure (i.e., noisy circuits or finite temperature), a density operator is needed, further reducing the computational distinction between the Heisenberg and Schrödinger pictures.

Choosing Pauli strings, P_i, (defined as the positive Hermitian tensor products of Pauli operators) as an operator basis, we can identify several computational advantages for working in the Heisenberg picture arising from their extremely simple and well-characterized mathematical properties [60–63]. Paulis can be conveniently indexed in binary form, and operator multiplication can be efficiently evaluated through simple binary operations on integers. Because any two Pauli strings either commute or anticommute, these can also be efficiently obtained through binary operations. Due to the involution property, Pauli strings have closed-form exponentiation, making evolution of a state,

$$U_j(\theta)|\psi\rangle = e^{-i\frac{\theta}{2}P_j}|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|\psi\rangle - i\sin\left(\frac{\theta}{2}\right)P_j|\psi\rangle, \tag{2}$$

or Pauli string,

$$U_j(\theta)^\dagger P_k U_j(\theta) = \begin{cases} \cos(\theta)P_k + i\sin(\theta)P_jP_k & \text{if } [P_k, P_j] \neq 0 \\ P_k & \text{if } [P_k, P_j] = 0 \end{cases}, \tag{3}$$

exceptionally convenient. More general considerations of exact exponentiation have recently been explored by Izmaylov and coworkers [64] and Evangelista and Magoulas [65,66]. Furthermore, the Heisenberg-picture Pauli evolution naturally exposes simplifications from Clifford or low-magic unitaries. While a $\pi/2$ rotation about any Pauli string has the potential to create a maximally entangled state in the Schrödinger picture, this angle corresponds to a Clifford gate, which, by definition, maps Pauli strings to Pauli strings, leaving the coefficient structure (and thus sparsity pattern) of any Heisenberg operator unchanged. For example, compare the action of a $\pi/2$ rotation on either a state or an operator:

$$U_i\left(\frac{\pi}{2}\right)|\psi\rangle = \frac{1}{\sqrt{2}}(|\psi\rangle - i|P_i\psi\rangle), \tag{4}$$

$$U_i\left(\frac{\pi}{2}\right)^\dagger P_j U_i\left(\frac{\pi}{2}\right) = P_k, \tag{5}$$

assuming P_i and P_j do not commute. In the Schrödinger picture, this Clifford evolution of a state generally increases in complexity, whereas the Pauli string, P_j, is just converted to a different Pauli, P_k.

A. Pauli evolution truncation

Simulating the sequential rotation of an operator by a large number of Pauli strings, Eq. (1), results in an exponentially increasing number of Paulis. However, Eq. (3) illustrates that this Pauli sum will potentially have a considerable amount of structure. Each rotation splits each noncommuting Pauli in the Pauli sum into two

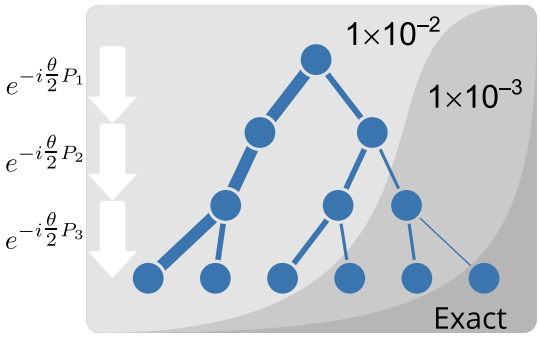


FIG. 1. Schematic depiction of the evolution of a single Pauli string (top) following three sequential small-angle Pauli rotations. Background colors indicate the additional set of Paulis that are included under tighter levels of coefficient truncation.

branches, a cosine and a sine branch, forming a binary tree. If each rotation angle is small, then the cos branches dominate, with the path weights diminishing exponentially in the number of sin branches. In contrast, for near-Clifford rotations ($\theta \approx \pi/2$), the sine branch dominates. The schematic in Fig. 1 illustrates the operator spreading that occurs during sequential evolution.

Pauli evolution, therefore, creates a binary tree, such that the evaluation of the final expectation value requires one to sum the expectation values of all the leaves, weighted by the weight of the branch. While this is clearly exponential to do exactly, the fact that many of these paths will have extremely small weights leads to a number of truncation options.

One could take a depth-first search approach to this, evaluating the sum path-by-path and aborting a given path once the current weight drops below a threshold. This requires essentially zero memory and is embarrassingly parallelizable; however, it performs rather poorly due to the need to go to very tight thresholds for convergence. One reason for this is that each path can visit the same Pauli, resulting in an opportunity for interference between the paths. Instead, we use a breadth-first search approach where the full operator at each time step is stored in memory, providing the opportunity to account for interference between Pauli paths. Keeping track of the full operator stabilizes truncation and facilitates the computation of properties such as four-point correlation functions. In this paper, we perform truncations after each Pauli rotation, deleting all Paulis whose coefficient absolute value is smaller than a given threshold ϵ.

This truncation approach is not new and has been used in several implementations, including Refs. [35,52].

The approximation of Pauli Heisenberg dynamics has been described with a variety of names, `PP` [38,52], Pauli Backpropagation [67], Pauli Path simulation [53], and `SPD` [35], each with different but related goals. Although our current treatment of approximate Pauli evolution is most closely related to `SPD`, since we are currently only approximating the evolution by truncating small-coefficient Paulis, we have also found truncations based on Pauli weight truncation to be effective, and so we will refer to our Pauli evolution treatment by the name of Pauli Propagation, referencing the first detailed presentation of the simulation strategy [38].

B. Double bracket flows

Having access to an efficient algorithm for approximating unitary transformations has the potential to be useful for a variety of applications. The utility of quantum circuit simulation is clear and well documented. In this work, we are interested in applying `PP` to the problem of ground-state energy estimation.

In order to define a unitary transformation that evolves a state into the ground state, we start by adopting a double bracket flow formalism. Double bracket flows are differential equations whose right-hand sides involve nested commutators and which, when integrated, drive the system toward specific targets. This approach was originally introduced as the Similarity Renormalization Group in 1993 by Głazek and Wilson [68,69] and was independently developed as a flow equation by Wegner in 1994 [70].

The basic idea is to define a Hamiltonian transformation along a continuous timelike parameter, s, that moves in a direction that minimizes some cost function. The "canonical commutator" [68,70] is defined using the commutator of the Hamiltonian with the projection onto its own diagonal, H_D:

$$\frac{d}{ds}H(s) = [G(s), H(s)] = [[H_D(s), H(s)], H(s)]. \quad (6)$$

Since $G(s)$ is an anti-Hermitian flow generator operator, this is analogous to a unitary evolution of H by a time-dependent drive $G(s)$, a process that preserves the spectrum of $H(s)$ at all times, and is thus a similarity transformation. This choice of $G(s)$ generates rotations that minimize the off-diagonal components to zero when the Hamiltonian commutes with its own diagonal. This form of $G(s)$ has been widely used, finding applications in physics [71], chemistry [72–74], and quantum algorithms [51,54].

1. State-specific flow generator

While the canonical commutator would certainly suffice for obtaining our target eigenstate because it achieves full diagonalization, it may be unnecessarily powerful, and we thus choose a flow generator that targets a single reference state instead, $G(s) = [H(s), \rho]$.

This generates evolution anytime the Hamiltonian does not have ρ as a stationary state. We choose ρ to be the projector onto a reference computational basis state, which we will take, without loss of generality [75], to be the zero state, $\rho = |0\rangle\langle 0|$.

Since the reference-state expectation value of the right-hand side is just the energy variance, which is nonnegative, the energy is guaranteed to decrease until $|0\rangle$ becomes an eigenstate:

$$\frac{d}{ds}E(s) = -2(\langle H^2\rangle - \langle H\rangle^2). \tag{7}$$

We note that this choice of state-specific double bracket flow generator was rigorously derived in the DB-QITE framework [55], building on earlier Double Bracket Quantum Algorithms work [54]. In particular, Ref. [55] provides the first systematic derivation of Eq. (7) and of the ansatz $[H(s),\rho]$ used here. While this natural choice converges well, the projector expanded in the Pauli basis is highly nonsparse, requiring 2^N diagonal Pauli strings,

$$|0\rangle\langle 0| \propto I + \sum_i Z_i + \sum_{i<j} Z_i Z_j + \cdots. \tag{8}$$

To avoid this exponential cost, we choose to truncate this projector, where we approximate $|0\rangle\langle 0|$ by keeping only Z strings up to a certain Pauli weight. For all results presented here, we keep only the one-body terms. While this might seem like a dramatic approximation, evolving by only the one-body terms will still work, as this generator will only go to zero once $H(s)$ becomes block-diagonal, where Hamiltonian coupling between Hamming weight blocks is driven to zero. However, with this approximation, the integration of this equation takes longer to converge, as demonstrated in Fig. 10 of Appendix A. Similar observations were also noted in Refs. [51,76]. Using this one-body form, our double bracket flow generator becomes:

$$G(s) = \sum_i [H(s), Z_i], \tag{9}$$

leading to the following double bracket flow:

$$\frac{d}{ds}H(s) = \sum_i [[H(s), Z_i], H(s)], \tag{10}$$

which is identical to the form derived in Refs. [51,54].

Integrating Eq. (10) for a small time, δs, we obtain the following recursive transformation:

$$H(s+\delta s) = e^{G(s)\delta s} H(s) e^{-G(s)\delta s}. \tag{11}$$

Since our aim will be to approximate this process using PP, we will want to first expand $G(s)$ in the Pauli basis:

$$G(s) = \sum_i g_i P_i, \tag{12}$$

where g_i depend on s and are purely imaginary. Since δs is small, we can expand this solution in a product formula:

$$H(s+\delta s) = e^{g_n P_n \delta s} \cdots e^{g_1 P_1 \delta s} H(s) e^{-g_1 P_1 \delta s} \cdots e^{-g_n P_n \delta s}. \tag{13}$$

The exact Riemannian-gradient flow is guaranteed to converge to the ground state for almost every initial state with nonzero ground-state overlap: the energy decreases monotonically, the critical points of the flow are exactly the eigenstates of H, and all intermediate eigenstates are saddle points, which first-order methods avoid with probability 1 [77,78]. In practice, however, a trajectory approaching a saddle point can significantly slow convergence [77,79,80]. Approximating this flow via sequential Pauli rotations (Eq. 13) with adaptive step sizes and a greedy selection of operators (detailed in the next section) departs from the exact gradient direction; randomized variants of such updates retain almost-sure convergence guarantees [78,81], and while these guarantees do not transfer as a formal theorem to our deterministic, truncated implementation, the fixed-point structure is preserved: each angle optimization is energy-nonincreasing, and the greedy gradients vanish only when the reference has become an eigenstate of the evolved Hamiltonian.

While recent approaches have sought to implement these Riemannian gradient flows on quantum hardware by approximating the exact gradient via random projections into polynomial-sized Pauli subspaces [77], our deterministic vDBF approach allows for scalable classical simulation of much larger systems. Using the equation above, we *could* proceed by using PP to sequentially evolve $H(s)$ by the series of Paulis in $G(s)$, each time truncating $H(s)$ by deleting Paulis with magnitudes smaller than ϵ. However, keeping δs small (a requirement to explicitly integrate Eq. 10) would require too many rotations for efficient convergence. In the next section, we propose a modification of this strategy that accelerates this convergence.

C. Pauli propagation—vDBF

At this point, we depart from the direct integration of the double bracket flow defined in Eq. (10). Rather than choosing the time steps to be "small," which would inevitably require many Pauli rotations to converge, we will instead choose an *optimal* angle, θ_i^*, for each Pauli, P_i, in the sequence, chosen to minimize the energy of the reference state, $|\psi\rangle$ [82] by analytically minimizing the cost function:

$$\begin{aligned} F(\theta_i) &= \langle\psi| e^{iP_i\theta_i} H e^{-iP_i\theta_i} |\psi\rangle \\ &= \cos^2(\theta_i)\langle\psi|H|\psi\rangle + \sin^2(\theta_i)\langle\psi|P_i H P_i|\psi\rangle \\ &\quad - 2i\cos(\theta_i)\sin(\theta_i)\langle\psi|HP_i|\psi\rangle. \end{aligned} \tag{14}$$

This will ensure that each Pauli rotation minimizes the energy by the maximum amount possible. While this leads to a trajectory that deviates from the specific double

bracket flow defined in Eq. (10), we ultimately only care about the final state. Since we are no longer faithful to the flow of Eq. (10), the s parameter is no longer meaningful, so we track the progress of this sequential transformation with an index, indicating the number of iterations: $H(s) \to H^{(n)}$.

1. Greedy selection of generators

Depending on the number of rotations performed during the vDBF, $H^{(n)}$ can accumulate many terms, resulting in a flow generator for the current iteration, $G^{(n)} = \sum_j [H^{(n)}, Z_j]$, that is even larger. However, not all Paulis in $G^{(n)}$ are equally important. Ideally, we would only rotate by the most "variationally helpful" generators within each iteration to avoid unnecessary computation. To do this, we sort the operators in $G^{(n)}$ by the derivative of the energy with respect to rotation about each operator:

$$\frac{d}{d\theta_i}\left\langle e^{i\theta_i g_i P_i} H^{(n)} e^{-i\theta_i g_i P_i}\right\rangle = g_i \left\langle \left[P_i, H^{(n)}\right]\right\rangle \tag{15}$$

$$= 2 g_i \mathrm{Im}\left\{\langle 0|P_i|\sigma^{(n)}\rangle\right\}, \tag{16}$$

where $|\sigma^{(n)}\rangle = H^{(n)}|0\rangle$ is computed only once per iteration. This form allows the gradient vector to be efficiently computed since we only need to read off the coefficient of the ket in $|\sigma^{(n)}\rangle$ whose on bits match the X string in P_i.

Having the operators in $G^{(n)}$ sorted by their reference-state energy gradient, we then proceed to evolve by the first n_rots angles per iteration (user-specified simulation parameter). Setting n_rots = 1 would therefore result in a single rotation by the operator with the largest gradient in the $G^{(n)}$. This would be a dynamic pool analog to how Adaptive Derivative Assembled Pseudo Trotter Ansatz Variational Quantum Eigensolver (ADAPT-VQE) [83,84] chooses operators out of a fixed operator pool. However, since evaluating the flow generator is not free, we find that it is generally more efficient to perform several rotations for each iteration, and so choosing n_rots to be tens to hundreds (e.g., 50–500) seems to work well. We note that this algorithm for designing a double bracket flow, which uses optimal steps to converge to the ground state, can be viewed as a classical version of the related quantum algorithm, double bracket-diagonal operator iterations [51].

2. Operator truncation

In order to tame the exponential growth of the Hamiltonian, we follow the coefficient truncation schemes used in both SPD [35,36,46] and PP [48–50,52], by which an operator expressed as a linear combination of Paulis, $O = \sum_i c_i P_i$, is modified by deleting Paulis, P_i, that are considered unimportant. However, unlike the generic non-equilibrium dynamics of a quantum circuit, the vDBF approach has a well-defined reference state that allows us to significantly reduce truncation error by truncating only the operator's reference-state fluctuation instead of the full operator,

$$P_i = \delta P_i + \langle 0|P_i|0\rangle I, \tag{17}$$

where, $\delta P_i = P_i - \langle 0|P_i|0\rangle I$. In our simulations, whenever a Pauli's coefficient magnitude drops below a threshold, $|c_i| < \epsilon$, we delete the operator's fluctuation:

$$P_i \to \langle 0|P_i|0\rangle I. \tag{18}$$

Note that this truncation rule preserves the reference-state energy contribution of P_i exactly, since $c_i\langle 0|P_i|0\rangle$ remains in H regardless of whether the operator's fluctuation is retained. The truncation therefore introduces no first-order error in the energy, with errors entering only through subsequent rotations of the truncated operator. We refer to this truncation as clip in Line 19 of the Algorithm 1. While Pauli weight [47,52] (or Majorana weight for fermionic Hamiltonians) [50] can also be used to further decrease computational cost, we delay an in-depth analysis of this for future work (initial calculations seem to suggest it can be helpful) [85].

ALGORITHM 1. Variational double bracket flow

Input: H, ψ, ϵ, n_rots, max_iter, conv_thresh
Output: $H^{(n)}$, *generators*, *angles*

1 *generators* ← [];
2 *angles* ← [];
3 $H^{(1)} \leftarrow H$;
4 **for** $n \leftarrow 1$ **to** max_iter **do**
// Compute generator
5 $G^{(n)} \leftarrow [H^{(n)}, \sum_j Z_j]$;
// Sort G by gradient
6 *gradients* ← [];
7 *paulis* ← [];
8 **for** $g_i, P_i \in G^{(n)}$ **do**
9 Append $g_i \times \langle\psi|[P_i, H^{(n)}]|\psi\rangle$ to *gradients*;
10 Append P_i to *paulis*;
11 *perm* ← sortperm(*gradients*);
12 *paulis* ← *paulis*[*perm*]
// Check for convergence
13 **if** norm(*gradients*) < conv_thresh **then**
14 break;
// Rotate by optimal angles
15 **for** $j \leftarrow 1$ **to** n_rots **do**
16 $P_j \leftarrow$ *paulis*[j];
17 $\theta_j^* \leftarrow$ COMPUTE_OPTIMAL_ANGLE($H^{(n)}, \psi, P_j$);
18 $H^{(n)} \leftarrow$ EVOLVE($H^{(n)}, P_j, \theta_j^*$);
19 $H^{(n)} \leftarrow$ CLIP($H^{(n)}, \epsilon$);
20 Append θ_j^* to *angles*;
21 Append P_j to *generators*;
22 $H^{(n+1)} \leftarrow H^{(n)}$
23 **return** $H^{(n)}$, *generators*, *angles*

3. Zero variance extrapolation

Algorithm 1 lists the steps required to converge a vDBF calculation for a given, user-specified accuracy control, ϵ. Because convergence can be slow, iterating the algorithm until fully converged can become intractable for large systems. However, as one approaches convergence, a linear relationship between energy and variance ($\langle H^2\rangle - \langle H\rangle^2$) is expected. We leverage commonly used zero-variance extrapolations to estimate the energy at convergence, allowing us to stop iterating once a reliably linear relationship is observed.

D. Comparison to existing methods

The vDBF method is related to several other simulation techniques, both quantum and classical.

1. ADAPT-VQE

In ADAPT-VQE [83,84], one performs a variational minimization of an ansatz that has the form of a sequence of unitary transformations. The specific sequence of generators is defined iteratively such that at each iteration, the ground-state energy gradient is used as an importance measure for determining which operator from a static, user-defined operator pool should be added to the ansatz. After adding each new generator, all the operators are allowed to variationally relax. Our current vDBF algorithm can be viewed as a classical realization of a modified ADAPT-VQE, where only the last operator added is minimized (avoiding the challenging nonlinear optimization). This stepwise parameter assignment shares conceptual similarities with the feedback-based quantum optimization algorithm Feedback-based algorithm for quantum optimization (FALQON) [86], which also sets circuit parameters sequentially based on local gradient measurements rather than relying on global variational relaxation. However, unlike both ADAPT-VQE and FALQON, which typically draw from static, predefined operator pools, vDBF constructs its generators from a dynamically evolving operator pool. We also note that shortly after the first preprint of this paper was posted, an interesting related method called Adaptive Derivative-Assembled Pseudo-Trotter ansatz Variational Majorana Propagation Eigensolver was published, which combined ADAPT-VQE with `PP` [87].

2. Iterative qubit coupled cluster

vDBF can be viewed as a classical version of iterative qubit coupled cluster (iQCC) [88,89] where the entanglers are chosen based on the flow generator in Eq. (9) instead of the direct interaction space, and the evolution is approximated using `PP`. Recent advancements also showed ways to reduce operator growth while lowering energy [90], but vDBF relies solely on greedy steepest-descent selection. A recent preprint has also demonstrated a multistate variant, which showed promising results on small systems [91].

3. Double bracket-diagonal operator iteration

In the recently proposed quantum algorithm double bracket-diagonal operator iteration (DB-DOI), the Hamiltonian is diagonalized through a double bracket flow on a quantum computer [51], where each sequential rotation can be determined variationally to minimize the off-diagonal elements of the Hamiltonian. The current vDBF method can be viewed as a classical version of the DB-DOI approach, where the time steps are chosen to minimize the expectation values of the reference state, and the time evolution is approximated using the `PP` technique.

III. NUMERICAL SIMULATIONS

In order to assess the ability of vDBF, we first consider two canonical quantum many-body problems as examples: an antiferromagnetic Heisenberg spin lattice and the Fermi Hubbard model. In both systems, we demonstrate the vDBF method's ability to provide fast estimates of ground-state energies. Finally, we extend our analysis to a complex molecular system, the polycyclic aromatic hydrocarbon HBC.

A. Extrapolation procedure

To obtain stable extrapolations to the exact limit, we leverage the fact that when the trajectory is near convergence, we expect a linear relationship between the energy and the energy variance [92–95]. However, further from convergence, the deviation from linearity might be non-negligible, and a simple linear regression might be less reliable despite having a seemingly respectable R^2 value. In order to obtain a more conservative estimate of errors in our extrapolations, we take the extrapolated result to be the average of a linear (b_1) and a quadratic (b_2) fit, with the uncertainties reported as half the difference. While we find that this extrapolation works surprisingly well, the fact that each calculation often generates thousands of data points means that one must decide how many data points to include in the extrapolation.

In the calculations below, we start with the final step of a given optimization trajectory (which occurs due to either convergence or reaching the maximum number of iterations allowed) and then include all the previous points up to a cutoff which is chosen to minimize the uncertainty ($(b_1 - b_2)/2$) and maximize the R^2 value of the linear fit. This balances precision with goodness of fit. Using this "optimal" cutoff point in our data, our zero-variance extrapolations are reported as the average of the linear and quadratic fits, with the difference as the uncertainty: $E^{\infty}_{(v=0)} = \frac{b_1+b_2}{2} \pm \frac{b_1-b_2}{2}$.

All vDBF calculations were performed with our own software, `DBF.jl` [96] and `PauliOperators.jl` [97], which use standard symplectic Pauli representations for efficient binary manipulations. The DMRG calculations for Heisenberg and Hubbard models were carried out

using the `ITensors.jl` package [98], and the HBC was done using `BLOCK2` [99]. Selected configuration interaction (SCI) computations used the `MPQC` package [100].

B. Heisenberg Hamiltonian

The isotropic Heisenberg Hamiltonian is a standard model not only for describing magnetic interactions in materials but also for evaluating computational methods for strongly correlated quantum systems. While the one-dimensional (1D) lattice admits an analytic solution via the Bethe ansatz, and DMRG provides an optimally efficient numerical treatment, the simulation in higher dimensions has no analytic solution and precludes efficient simulation with a 1D Matrix Product State (MPS). While belief propagation algorithms for general tensor network state contractions have resulted in significant improvements recently [37,101–103], the description of higher-dimensional strongly correlated systems is still challenging.

We start by computing the ground-state energy of the antiferromagnetic Heisenberg spin Hamiltonian:

$$H = -2J\sum_{\langle ij\rangle}\vec{S}_i\cdot\vec{S}_j, \tag{19}$$

on two lattices: (a) a large one-dimensional 1×100 lattice and (b) a large two-dimensional (2D) 10×10 square lattice. For the DMRG calculations, we used the `ITensors.jl` software [98]. For 2D square lattices, we used a snake pattern to define our Matrix Product Operator Hamiltonian and initial MPS and ran symmetry-number-preserved calculations for the half-filled state.

For the vDBF results, we chose the Néel state for our reference state $|0101\cdots\rangle$. Because the form of the projector, ρ, and the resulting approximate flow generator (Eq. (9)) assume the zero state, we simply transform our initial Hamiltonian by a sequence of X gates so that we can work directly with the zero state:

$$E^{(0)} = \langle 0101\cdots|H^{(0)}|0101\cdots\rangle \tag{20}$$

$$= \langle 0|\cdots X_4X_2H^{(0)}X_2X_4\cdots|0\rangle \tag{21}$$

Since each X_i gate is a Clifford, the coefficient structure of $H^{(0)}$ is left unchanged. This feature will allow us to consider entangled reference states in future work.

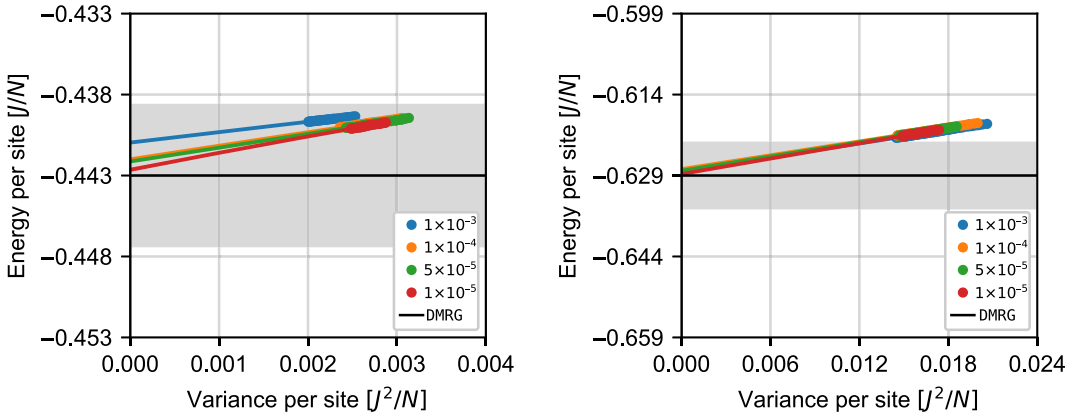
The results for the Heisenberg lattices are presented in Table I, and complementary convergence plots are provided in Fig. 2.

1. 1×100 *Heisenberg*

Looking first at the one-dimensional 1×100 Heisenberg lattice, we find that all four thresholds achieve

TABLE I. Energy and site values for Heisenberg models for DMRG and vDBF. vDBF calculations are each run with the threshold, ϵ, defined in the top row, and `max_iter` = 100 and `n_rots` = 100. Energy denotes variational energies, while "Extrap. E" denotes extrapolated energies; note that while vDBF is variational in principle, coefficient truncation removes the strict-upper bound in practice. Uncertainties in the extrapolated vDBF energies, shown in parentheses, are given by half the difference between linear and quadratic extrapolation, as described in Sec. III A. Discarded W. is the discarded weight from Sec. II C 3. 1×100 lattice has periodic boundary conditions, and 10×10 has open boundary conditions.

Heisenberg Lattice	DMRG		vDBF				
			ϵ:	1×10^{-3}	1×10^{-4}	5×10^{-5}	1×10^{-5}
1×100	*Energy*:	−0.443230	*Energy*:	−0.439909	−0.440248	−0.440287	−0.440336
			Extrap. E:	−0.44120(3)	−0.44223(4)	−0.442363(2)	−0.4429(1)
			Error:	0.459%	0.226%	0.196%	0.079%
			Variance:	0.0020	0.0024	0.0024	0.0025
			Discarded W.:	0.054012	0.004627	0.002277	0.000376
			Iterations:	100	100	100	100
	χ:	800	len (H):	29 272	249 696	483163	2 340 704
	Time (min):	3.9	Time (min):	1.1	2.2	2.7	10.8
	Threads:	64	Threads:	1	1	1	1
10×10	*Energy*:	−0.628455	*Energy*:	−0.621682	−0.621285	−0.621338	−0.621367
	Extrap E:	−0.628693	*Extrap. E*:	−0.6277(2)	−0.627506(6)	−0.62780(3)	−0.6284(2)
			Error:	0.158%	0.189%	0.141%	0.045%
			Variance:	0.0145	0.0146	0.0147	0.0149
			Discarded W.:	0.39230	0.06304	0.03624	0.00944
			Iterations:	100	100	100	100
	χ:	3500	len (H):	189 593	3 023 903	7 159 792	48 357 051
	Time (min):	3213.9	Time (min):	2.5	18.2	39.4	249.8
	Threads:	64	Threads:	1	1	1	1



(a) 1×100 Heisenberg Lattice (b) 10×10 Heisenberg Lattice

FIG. 2. Energy vs variance extrapolations for Heisenberg lattices using a series of vDBF thresholds: $\epsilon = 1 \times 10^{-3}$ (blue), 1×10^{-4} (orange), 5×10^{-5} (green), and 1×10^{-5} (red). Extrapolations are shown as solid lines. DMRG best estimates (solid black line). 1% error region with respect to the DMRG best estimate (gray-shaded area).

less than 1% error, with $\epsilon = 1 \times 10^{-3}$ doing so in just 1.1 min on a single thread. We observe systematic accuracy improvement upon decreasing ϵ; each tighter calculation in the sequence decreases the deviations from DMRG, up to the $\epsilon = 1 \times 10^{-5}$ result, which has approximately 0.08% error. While both the energy errors and discarded weight decrease as ϵ is decreased, the cost also increases.

This data already illustrates a more general feature that will be seen with the following examples as well: the accuracy obtained at a given threshold is controlled by a balance of two distinct error sources: (i) truncation error, which is directly reduced by decreasing ϵ (as seen in the discarded weight), and (ii) convergence error, which generally grows as ϵ is tightened, both because tighter thresholds empirically slow the convergence of the flow and because the increased per-iteration cost reduces the number of iterations affordable at a fixed compute budget. However, since the energy-variance extrapolation lets us estimate the converged energy from a partially converged trajectory, we still obtain accurate results from relatively few iterations.

In Panel (a) in Fig. 2, the convergence and extrapolations are shown for each of the four threshold values for the 1×100 lattice. As described above, we choose our data cutoff point so as to minimize the uncertainty in our linear fit, and the shaded regions in the curve show the difference between linear and quadratic extrapolations. For this system, the differences between linear and quadratic fits are undetectable in the plots, indicating a good fit. We find that even the raw energy, $\langle 0|H^{(n_{\text{final}})}|0\rangle$ without extrapolation, is already less than 1% error from the ground-state energy for all four truncation thresholds.

To develop better physical insight into the nature of our approximations, in Fig. 3, we have computed the spin-spin correlation function vs distance, $C(r) = \langle S_1 S_r\rangle - \langle S_1\rangle\langle S_r\rangle$, by simply transforming the target observable by the same sequence of unitaries that was constructed by

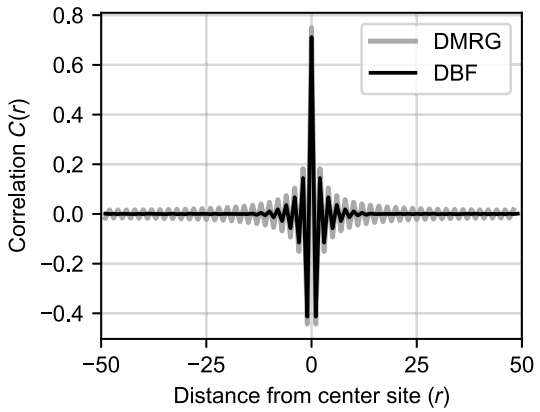


FIG. 3. Spin-spin correlation functions for 1×100 site Heisenberg lattice. $C(r) = \langle S_1 S_r\rangle - \langle S_1\rangle\langle S_r\rangle$. Blue lines are for converged DMRG results. Black line is vDBF (1×10^{-4}).

the vDBF. This gives us raw expectation values, which are not expected to have the same accuracy as the energy, since we do not currently have a way to extrapolate observables to improve their accuracy, as we do with energy. However, this provides direct insight into the origin of the computational speedups. DMRG shows the expected polynomial decay of $C(r)$, with nontrivial alternating correlations extending across the entire 100-site chain. In contrast, the vDBF correlations decay exponentially. Since the sequence of Pauli rotations is chosen to greedily minimize the energy, nonlocal correlations that do not directly affect the energy are naturally discarded. This also indicates that for gapped systems away from critical points (e.g., introducing anisotropy in J), we would expect vDBF to perform even better, since there the correlations also decay exponentially. Future work will need to focus on developing techniques for improving the accuracy of properties beyond just the ground-state energy.

2. 10×10 *Heisenberg*

Compared to the previous example, the 10×10 lattice is both large and high-dimensional, making direct Tensor Network States treatment challenging. In Table I, we obtained relatively good DMRG convergence with a bond dimension of 3500, which led to a very small extrapolation correction of around $0.0002J$/site.

All thresholds used for the vDBF calculations were able to obtain less than 1% error with respect to the extrapolated DMRG energies while requiring significantly fewer resources than DMRG. For example, the $\epsilon = 1 \times 10^{-3}$ calculation took only 2.5 min on a single thread, with an error of 0.158%, whereas the DMRG calculation took more than 2 days on 64 threads.

3. $J_1 - J_2$ *Heisenberg: Frustrated* 8×8 *lattice*

The previous two examples focus on bipartite lattices, which although they remain challenging to many simulation methods, are unique in that they do not suffer from a sign problem (since the rotation of one sublattice by local σ_z makes all the off-diagonal elements nonpositive). Consequently, the previous examples are able to be efficiently sampled using Monte Carlo techniques. In order to test whether or not the performance of vDBF is somehow benefiting from this lack of sign problem, we consider a frustrated (nonbipartite) Heisenberg lattice by introducing next-nearest neighbor (diagonal) couplings, J_2.

$$H = J_1 \sum_{\langle ij \rangle} \vec{S}_i \cdot \vec{S}_j + J_2 \sum_{\langle\langle ij \rangle\rangle} \vec{S}_i \cdot \vec{S}_j, \tag{22}$$

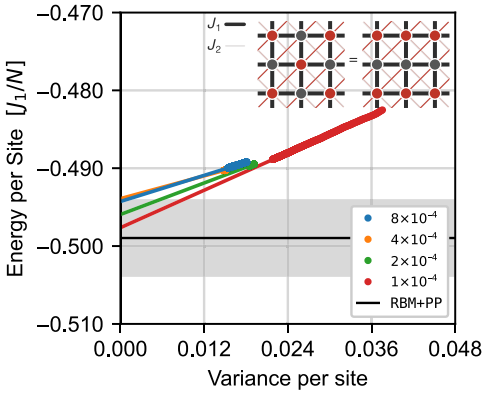
with periodic boundary conditions. When $J_2 = 0$, the model reduces to simple Heisenberg model, and the Néel state is the lowest energy single configuration. When J_2 becomes large, a striped configuration becomes the lowest energy configuration, and the point at which these two phases cross ($J_2 = J_1/2$) represents the maximal frustration point, where the two configurations are degenerate; see Fig. 4(a) (inset).

While this degeneracy prevents one from selecting a "good" reference state for vDBF, we can accelerate convergence by warm starting the vDBF by first optimizing a sequence of rotations to minimize the reference-state energy without truncation. In order to keep this untruncated warm-starting optimization from creating excessive operator growth, we choose the sequence of rotations to form a set of qubit-disjoint Pauli transformations. This sequence was defined by first coupling all horizontal edges with XY Paulis, then vertical edges with XY Paulis, then a 2×2 tiling of $XXXY$ plaquette Paulis, e.g.,

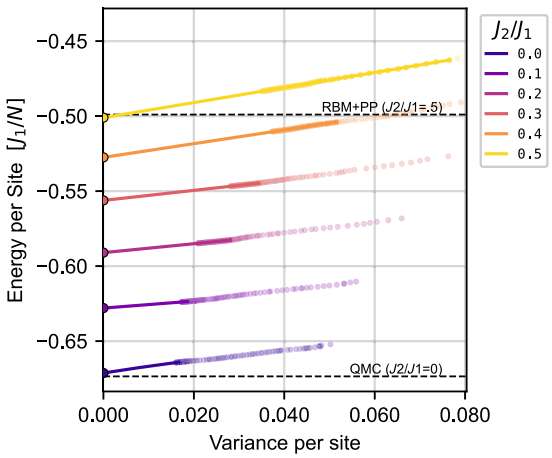
$$U(\theta) = e^{i\theta_5 X_1 X_2 X_3 Y_4} e^{i\theta_4 Y_2 X_4} e^{i\theta_3 Y_1 X_3} e^{i\theta_2 Y_3 X_4} e^{i\theta_1 Y_1 X_2}. \tag{23}$$

Since each 2×2 plaquette is disjoint [equivalent to choosing a cluster mean field (Refs. [108–112])], the operator spreading is minimal (H went from 768 terms to 15 088 terms), making this optimization efficient such that no truncation was needed. This improved reference state lowered the initial energy from $-0.25\ J_1$/site (Néel) to $-0.4375\ J_1$/site (plaquette) [113].

In order to evaluate the performance of our vDBF extrapolations, we compare directly to the Variational Benchmark dataset provided in Ref. [105]. From this



(a) Energy Extrapolations at classical critical point: $J_2/J_1 = \frac{1}{2}$



(b) Energy Extrapolations for range of J_2/J_1 ratios: $[0, \frac{1}{2}]$, $\epsilon = 8 \times 10^{-4}$, `max_iter` $= 100$

FIG. 4. Frustrated Heisenberg 8×8 lattice. (a) Energy vs variance extrapolations for $J_2/J_1 = \frac{1}{2}$ using a series of vDBF thresholds: $\epsilon = 8 \times 10^{-4}$ (blue), 4×10^{-4} (orange), 2×10^{-4} (green), and 1×10^{-4} (red). (b) Loose extrapolations ($\epsilon = 8 \times 10^{-4}$, `maxiter` $= 100$) for range of J_2/J_1 ratios. Extrapolations are shown as solid lines. Best estimate (black line) for $J_2/J_1 = \frac{1}{2}$ is Restricted Boltzmann Machine, RBM + PP obtained from Ref. [105]. 1% error region with respect to best estimate (gray-shaded area). Best estimate for $J_2/J_1 = 0$ from QMC. obtained from Refs. [106,107]. Periodic boundary conditions.

TABLE II. Energy and Ssite values for the frustrated 8×8 J_1–J_2 model at $J_2/J_1 = 0.5$ (periodic boundary conditions) for vDBF compared against the RBM + PP reference energy [104], the lowest V-score entry in the VarBench dataset [105]. vDBF calculations are each run with the threshold, ϵ, defined in the top row and `n_rots` = 50. Energy denotes variational energies, while Extrap. E denotes extrapolated energies; uncertainties, shown in parentheses, are given by half the difference between linear and quadratic extrapolation, as described in Sec. III A. Error is measured on the extrapolated energy relative to the reference. Discarded W. is the discarded weight from Sec. II C 3.

J_1–J_2 8×8 Lattice	RBM + PP		vDBF			
		ϵ:	8×10^{-4}	4×10^{-4}	2×10^{-4}	1×10^{-4}
$J_2/J_1 = 0.5$	*Energy*: −0.498964	*Energy*:	−0.487724	−0.487900	−0.487763	−0.487671
		Extrap. E:	−0.494272(1)	−0.493903(2)	−0.495942(7)	−0.497638(3)
		Error:	0.940%	1.014%	0.606%	0.266%
		Variance:	0.0155	0.0149	0.0185	0.0218
		Discarded W.:	15.341	11.327	7.152	4.382
		Iterations:	7278	5238	1632	798
		len (H):	4 679 162	14 494 361	40 485 851	106 469 502
		Time (min):	17.5	49.9	236.1	482.0
		Threads:	1	1	1	1

dataset, we took −0.49896348, computed with RBM + PP, to be our best estimate reference number, as it had the lowest V-score for this system (0.00098).

Considering the zero-variance extrapolations in Fig. 4(a), we again observe similar behavior in that tightening the truncation threshold generally improves the extrapolated energy estimates, quickly bringing them to within the target 1% of the best variational estimate from RBM + PP calculations [104,105], with our tightest (1×10^{-4}) calculation providing an energy and site of −0.4976, as shown in Table II. Although the bare vDBF calculations are far from converged (e.g., the calculation still had a variance and site of 0.02 after 798 iterations), the extrapolations are quite accurate. In fact, without extrapolation, the bare $\epsilon = 1 \times 10^{-4}$ vDBF energy was comparable to the RBM($\alpha = 1$) energy from the same dataset, but with twice the variance.

While computing the absolute ground-state energy accurately is clearly desirable, often times, one is more interested in relative energies. In Fig. 4(b), we show extrapolations from comparatively cheap calculations (100 iterations at $\epsilon = 8 \times 10^{-4}$ taking around 20 min total for all systems) on a series of J_2/J_1 ratios, ranging from 0 (the Heisenberg limit) to $\frac{1}{2}$ (the classical critical point). While we only have reference values for the end points, we observe quite systematic behavior allowing us to compare the extrapolated energies. The computed energy difference between $J_2/J_1 = 0$ and $\frac{1}{2}$ is only in error of about 2%.

While the vDBF results in Table II show that we are indeed able to obtain sub-1% accuracy for this spin-frustrated problem, these calculations were considerably more demanding than for the bipartite lattices considered in the previous sections. In Table III, we show how both the number of terms in the Hamiltonian, as well as the final variance, increases as J_2 is increased from 0 to the maximally frustrated point $J_2 = J_1/2$. We observe that after 100 iterations, increasing J_2 results in a monotonic increase in the size of the resulting transformed Hamiltonian, with the critical point containing over 5× the number of Paulis than in the Heisenberg case. Further, the rate of convergence is also slowed for frustrated systems, with the variance at 100 iterations at $J_2 = .5J_1$ over twice that of the Heisenberg case. Nonetheless, vDBF still attains sub-1% accuracy at the maximally frustrated point despite this nearly order-of-magnitude increase in operator count and the roughly twofold larger residual variance. This indicates that the accuracy of the method does not rely on the sign-problem-free structure of the bipartite examples, but rather on the energy-variance extrapolation, which seems to remain reliable even when the underlying trajectory is far from converged.

TABLE III. Effect of frustration on vDBF efficiency after 100 iterations with $\epsilon = 8 \times 10^{-4}$. J_1 and J_2 are defined in Eq. (22). Both calculations run on 8×8 lattice with periodic boundary conditions.

J_2	len (H)	Variance/site	Discarded W.
0	650 873	0.016	0.679
0.1	863 750	0.017	1.004
0.2	1 305 895	0.021	1.678
0.3	2 044 238	0.028	3.089
0.4	3 151 350	0.037	5.869
0.5 (critical point)	3 537 249	0.035	6.746

C. Hubbard Hamiltonian

The Fermi–Hubbard Hamiltonian:

$$H = -t \sum_{\langle ij \rangle, \sigma} \hat{a}^{\dagger}_{i\sigma} \hat{a}_{j\sigma} + U \sum_{i} \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} - \mu \sum_{i} (\hat{n}_{i\uparrow} + \hat{n}_{i\downarrow}), \quad (24)$$

TABLE IV. Energy and site values for Hubbard models with 128 qubits. All lattices have open boundary conditions. Results compare DMRG, SCI, and vDBF methodologies. vDBF calculations are each run with the threshold, ϵ, defined in the top row, and `max_iter` = 1000 and `n_rots` = 50. Uncertainties in the vDBF energies, shown in parentheses, are given by half the difference between linear and quadratic extrapolation, as described in Sec. III A. The time reported are in minutes. [†]The $\epsilon_1 = 0$ SCI energy estimated via linear energy vs variance relationship (similar to vDBF) using SCI energies with $\epsilon_1 = \{10^{-3}, 5 \times 10^{-4}, 2.5 \times 10^{-4}\}$.

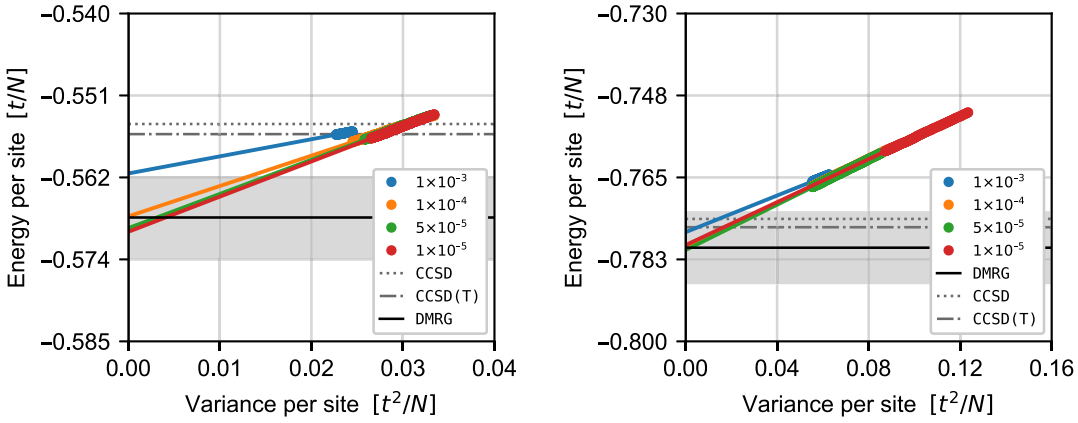
Hubbard Lattice	DMRG		SCI ϵ_1:	2.5×10^{-4}	vDBF ϵ:	1×10^{-3}	1×10^{-4}	5×10^{-5}	1×10^{-5}
1×64	*Energy*:	−0.567997	*Energy*:	−0.54671	*Energy*:	−0.556624	−0.557581	−0.557266	−0.557107
			Extrap:	−0.56426[†]	*Extrap*:	−0.561953(1)	−0.56784(1)	−0.56945(1)	−0.5699547(7)
					Error:	1.064%	0.028%	−0.256%	−0.345%
					Variance:	0.0227	0.0245	0.0258	0.0266
					Discarded W.:	0.7893	0.1294	0.0728	0.0183
					Iterations:	1000	1000	1000	1000
	χ:	211	N_{det}:	13 429 409	len (H):	304 797	4 620 991	10 418 275	69 891 297
	Time:	6.6	Time:	26	Time:	38.1	227.8	451.5	3056.7
	Threads:	64	Threads:	128	Threads:	1	1	1	1
8×8	*Energy*:	−0.77272	*Energy*:	−0.76867	*Energy*:	−0.766327	−0.767390	−0.767522	−0.759833
	Extrap:	−0.7805(2)	*Extrap*:	−0.77589[†]	*Extrap*:	−0.777186(1)	−0.780655(3)	−0.78102(3)	−0.77995(4)
					Error:	0.424%	−0.020%	−0.066%	0.070%
					Variance:	0.0555	0.0555	0.0555	0.0873
					Discarded W.:	1.6133	0.3268	0.2014	0.0277
					Iterations:	1000	1000	1000	307
	χ:	6000	N_{det}:	9 090 855	len (H):	584 905	10 622 036	26 995 422	131 216 940
	Time:	5200.0	Time:	71	Time:	40.8	592.5	1462.9	1693.4
	Threads:	64	Threads:	128	Threads:	1	1	1	1

is a canonical model for describing interactions in electronic systems across a wide range of phenomenological settings. Changes in the balance of kinetic and potential energy (t/U) or connectivity ($\langle ij \rangle$) give rise to diverse physics, including metal-insulator transitions, magnetism, Mott insulating behavior, and high-temperature superconductivity. Furthermore, it serves as a fundamental benchmark for evaluating new computational methods for many-body physics [114].

In this section, we demonstrate vDBF's ability to estimate the ground-state energy of the Hubbard model. Having previously investigated the Heisenberg model, which governs the strong-correlation limit of the half-filled Hubbard model, we now focus on the intermediately correlated regime ($U = 4t$) at half-filling, with open boundary conditions. To target the half-filled ground state, we set the chemical potential within the Mott plateau ($\mu = U/2$) [115]. Analogous to the Heisenberg results, we first compute a large one-dimensional 1×64 lattice (128 qubits), then study a two-dimensional 8×8 lattice (128 qubits) for which DMRG results were challenging to obtain, and validate these energies against SCI calculations. We then turn to the hole-doped 8×8 lattice at $U = 4t$ with periodic boundaries, for which published reference energies exist [116]. Because the half-filled clusters have open boundaries, their energies cannot be compared directly to the periodic lattice and thermodynamic-limit values of Ref. [114]; we compare instead to DMRG and SCI on the open boundary condition and use the doped periodic lattice for the comparison against published benchmarks.

In order to simulate this fermionic system with this Pauli operator representation, we use a Jordan–Wigner transformation to map the fermionic sites onto distinguishable lattice sites. For the vDBF calculations, we chose the Néel state as our reference state (alternating spin up and down). For the DMRG calculations, similar to the 2D Heisenberg model, we used a snakelike mapping to define our site ordering. The vDBF, DMRG and SCI energies and convergence data are presented in Table IV, and the extrapolation plots are shown in Fig. 5.

For the 1×64 case, DMRG is well converged at $\chi = 211$, and no extrapolation is needed; vDBF errors are reported relative to the variational DMRG energy. For the 8×8 case, DMRG requires bond-dimension extrapolation, and vDBF errors are reported relative to the extrapolated DMRG energy.



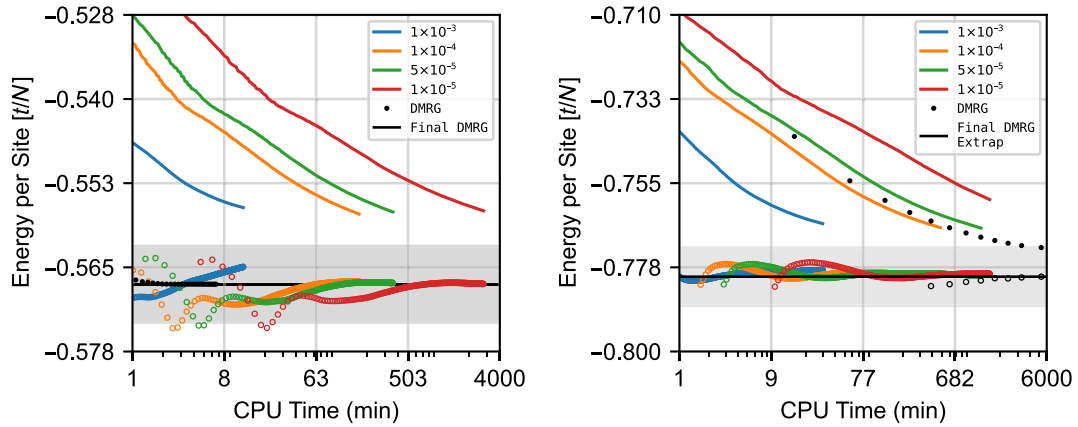
(a) 1×64 Hubbard lattice (b) 8×8 Hubbard lattice

FIG. 5. Energy vs variance extrapolations for Hubbard lattice using a series of vDBF thresholds: $\epsilon = 1 \times 10^{-3}$ (blue), 1×10^{-4} (orange), 5×10^{-5} (green), and 1×10^{-5} (red). Extrapolations are shown as solid lines. DMRG best estimates (solid black line). 1% error region with respect to DMRG best estimate (gray-shaded area).

1. 1×64 *Hubbard*

For the linear 1×64 site Hubbard model, the extrapolated energy from 1×10^{-4}, 5×10^{-5}, and 1×10^{-5} thresholds has an error of less than 1% relative to the DMRG benchmark. As shown in Fig. 5(a), our variational results surpass the accuracy of standard methods like coupled cluster singles and doubles (CCSD) and coupled cluster singles and doubles with triples correction CCSD(T). Using zero-variance extrapolations, the SCI calculations agree with the DMRG baseline, similar to vDBF, as shown in Table IV. Consistent with the Heisenberg model, increasing accuracy by tightening the truncation thresholds naturally increases computational cost, as seen for 1×10^{-5}, which took 3056.7 min, compared to 1×10^{-3}, which took 38.1 min.

However, zero-variance extrapolation dramatically lowers this computational cost. This can be seen from the Fig. 6(a); rather than targeting full convergence, the corresponding extrapolated energies (hollow circles) for the tighter thresholds drop into the sub-1% error regime almost immediately, demonstrating the effectiveness of energy-variance extrapolations. In fact, it took only 10 min to calculate a vDBF trajectory before the extrapolated error dropped to within 1%.



(a) 1×64 Hubbard lattice (b) 8×8 Hubbard lattice

FIG. 6. Energy estimates are plotted on a logarithmic time scale for the Hubbard lattices across four vDBF truncation thresholds: $\epsilon = 1 \times 10^{-3}$ (blue), 1×10^{-4} (orange), 5×10^{-5} (green), and 1×10^{-5} (red). Solid lines indicate the raw variational trajectories, while hollow circles represent the corresponding linear energy-variance extrapolations computed using all prior iterations. The best extrapolated DMRG estimates are indicated by the solid black line, and the 1% error threshold with the gray-shaded region.

2. 8 × 8 *Hubbard*

Shifting our focus now to the two-dimensional Hubbard model, we immediately note that in Table IV, the DMRG results contain both a variational energy and an extrapolated energy, since we were not able to go beyond the bond dimension of 6000 while using a single compute node with 64 threads due to memory and compute time limitations. As such, our DMRG extrapolation yields a substantial correction, amounting to approximately 1% of the final energy estimate. Compared to DMRG, our SCI calculations with around nine million determinants achieve a competitive extrapolated energy at a lower CPU time.

As we see in Fig. 5(b), all the thresholds obtain 99% accuracy compared to the extrapolated DMRG results. We also obtain a better energy estimate compared to CCSD and CCSD(T). Consistent with the Heisenberg model, the computational cost increases substantially as the truncation threshold is lowered. This is explicitly demonstrated by the evolved Hamiltonian at $\epsilon = 1 \times 10^{-5}$, which grows to contain approximately 131 million Pauli terms while only reaching 307 iterations. We become memory-bound at this point, as the number of terms in the commutation calculation also reaches 600 million Pauli strings.

As illustrated in Fig. 6(b), the raw variational trajectories for both vDBF (solid lines) and DMRG (black points) converge slowly. In contrast, the variance-extrapolated vDBF energies (hollow circles) reach the sub-1% error regime (gray-shaded area) more quickly than the DMRG extrapolations, yielding highly accurate ground-state estimates after fewer iterations or less time. To compare them head-to-head, DMRG extrapolations took around 500 min to reach the 1% region, whereas vDBF extrapolations fell to the 1% error regime almost immediately.

3. 8 × 8 *doped Hubbard*

The doped two-dimensional Hubbard model is particularly difficult because it has the sign problem absent at half-filling. We therefore consider a periodic 8 × 8 lattice with 56 electrons, equally divided between the two spin sectors, giving a density of $n = 0.875$ and one-eighth hole doping. We use $U/t = 4$, an intermediate-coupling regime where kinetic and interaction energies compete on equal footing, so a single-determinant description is unreliable. This provides a challenging test of vDBF convergence and energy-variance extrapolation.

For this example, we first converged a unrestricted Hartree-Fock solution from a checkerboard initial guess and used the UCCSD one-particle density matrix to construct natural orbitals. The occupied and virtual spaces were localized separately, producing the localized natural orbital basis, which was used in the vDBF calculation after doing a Jordan–Wigner transformation. Although this orbital transformation increases the number of terms we start with, the benefit from having a well-defined occupied-virtual separation ultimately provides faster convergence. To evaluate the doped Hubbard results, we compare the vDBF energies with four published reference values, as compiled in Table I of Ref. [116]: $-1.01153t$ per site from fixed-node (FN) [114], $-1.01391t$ from VAFQMC [116], $-1.00579t$ from Variational Monte Carlo (VMC) [105], and $-1.01923t$ from AFQMC [116,117]. We use the VAFQMC result as our main reference because it gives the lowest variational energy among these calculations.

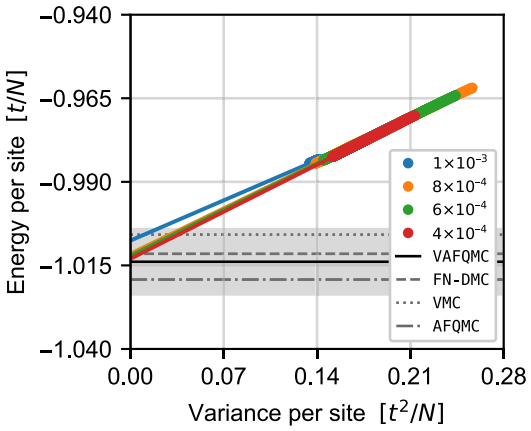


FIG. 7. Energy vs variance extrapolations for Doped 8 × 8 Hubbard lattice using a series of vDBF thresholds: $\epsilon = 1 \times 10^{-3}$ (blue), 8×10^{-4} (orange), 6×10^{-4} (green), and 4×10^{-4} (red). Extrapolations are shown as solid lines. VAFQMC best estimates (solid black line). 1% error region with respect to VAFQMC best estimate (gray-shaded area).

The results are shown in Fig. 7 and Table V. The reference values are more scattered here than for the half-filled lattices; they differ from one another by up to 1.3% (0.8% among the variational methods alone). All four vDBF trajectories ran the full 1000 iterations. As ϵ is tightened, the extrapolated energy moves steadily toward the VAFQMC value, from 0.63% above it at $\epsilon = 1 \times 10^{-3}$ to 0.18% at $\epsilon = 6 \times 10^{-4}$ and 0.09% at $\epsilon = 4 \times 10^{-4}$. Every extrapolation lies between the VMC and VAFQMC values, i.e., within the spread of the reference methods.

D. Molecular Hamiltonian—HBC π-system

To evaluate vDBF's performance on strongly correlated chemical systems, we now target the ground-state energy of HBC, $C_{42}H_{18}$. As a large, symmetric polycyclic aromatic hydrocarbon, HBC exhibits nontrivial electron correlation, providing a rigorous molecular benchmark for our approach.

The molecular Hamiltonian in second-quantized form is given as:

$$H = \sum_{p,q} h_{pq} a_p^\dagger a_q + \sum_{p,q,r,s} g_{pqrs} a_p^\dagger a_q^\dagger a_s a_r,$$

TABLE V. Doped Hubbard model on 8×8 lattice. Energy per site at $U/t = 4$, $N_e = 56$ ($n = 0.875$), with periodic boundary conditions. vDBF calculations use the thresholds shown and $n_{\text{rots}} = 50$. Extrap. E is obtained from the energy vs variance fit; uncertainties are half the difference between the linear and quadratic intercepts. Error is measured relative to the variational auxiliary-field quantum Monte Carlo (VAFQMC) reference energy.

Doped	Reference energies				vDBF			
8×8	VAFQMC	VMC	FN	AFQMC	ϵ:	1×10^{-3}	8×10^{-4}	6×10^{-4}	4×10^{-4}
$U/t = 4$	−1.013910	−1.00579	−1.01153	−1.01923	*Energy*:	−0.984450	−0.984399	−0.983505	−0.982522
					Extrap. E:	−1.007577(6)	−1.011519(2)	−1.012138(4)	−1.013027(3)
					Error:	0.625%	0.236%	0.175%	0.087%
					Variance:	0.1344	0.1392	0.1451	0.1517
					Discarded W.:	7.242	6.489	5.576	4.474
					Iterations:	1000	1000	1000	1000
					len(H):	1 852 073	2 556 110	3 873 056	6 961 382
					Time:	165.5	226.1	404.9	750.3
					Threads:	1	1	1	1

where h_{pq} and g_{pqrs} are one-body and two-body electronic integrals. We first perform a restricted Hartree-Fock calculation using the Correlation-Consistent Polarized Valence Double-Zeta basis set, which yields a mean-field energy of −1601.26918095 Ha across a total of 678 spatial orbitals. To make the problem computationally tractable, we restrict the system to an active space of 42 electrons in 42 spatial orbitals, corresponding to the full π-valence space of the molecule. The active-space Hamiltonian is subsequently transformed into a localized natural orbital basis obtained from a CCSD calculation.

The fermionic creation and annihilation operators of this active space are then mapped to a Pauli-string representation using the Jordan–Wigner transformation (84 qubits). To assess the accuracy of our vDBF correlation energies, along with DMRG, we also compare with the TPSCI [118–121] and SCI calculations.

In Table VI and Fig. 8, we summarize the ground-state correlation-energy calculations for the HBC molecule. As a benchmark, we performed DMRG calculations with a maximum bond dimension of $\chi = 1750$, which was the largest value accessible within our node memory limits. The calculations are well converged at this limit, as discarded weight was around 3×10^{-5} and the extrapolated energy was only around 1.5 mH. To verify convergence, we also performed a high-level, computationally intensive TPSCI calculation, which was in close agreement with the DMRG estimate, suggesting that the DMRG and TPSCI values provide a reliable reference window for this active space. The SCI variance extrapolation overshoots the DMRG reference, recovering more correlation energy than

TABLE VI. Correlation energy($E - E_{\text{RHF}}$) for HBC molecule. vDBF calculations are each run with the threshold, ϵ, defined in the top row, and `max_iter` = 1000 and `n_rots` = 50. Uncertainties in the vDBF energies, shown in parentheses, are given by half the difference between linear and quadratic extrapolation, as described in Sec. III A. The percentage errors are calculated with respect to the DMRG correlation energy. Timings for the tensor product selected configuration interaction (TPSCI) results are not compared, as the setup involved a complicated multistep process, but the CPU cost is significantly higher than either DMRG or vDBF. †The $\epsilon_1 = 0$ SCI energy is estimated via a linear energy vs variance relationship (similar to vDBF) using SCI energies with $\epsilon_1 = \{2.5 \times 10^{-5}, 1 \times 10^{-5}\}$.

DMRG		TPSCI		SCI		vDBF				
				ϵ_1:	2.5×10^{-5}	ϵ:	1.6×10^{-4}	8×10^{-5}	4×10^{-5}	2×10^{-5}
Energy:	−0.424931	*Energy*:	−0.415978	*Energy*:	−0.394326	*Energy*:	−0.405024	−0.4058708	−0.404422	−0.401284
Extrap:	−0.42653(2)	*Extrap*:	−0.42618(7)	*Extrap*:	−0.432597†	*Extrap*:	−0.429585(5)	−0.42866(3)	−0.434184(6)	−0.437876(7)
						Error:	−0.716%	−0.499%	−1.794%	−2.660%
						Variance:	0.0201	0.0188	0.0213	0.0248
						DW:	0.1120	0.0693	0.0440	0.0280
						Iterations:	1000	1000	1000	1000
χ:	1750			N_{det}:	5 965 230	len (H):	1 717 121	4 717 996	11 120 138	25 380 248
Time:	1010.7			Time:	37	Time:	108.8	283.6	667.0	1608.3
Threads:	64			Threads:	128	Threads:	1	1	1	1

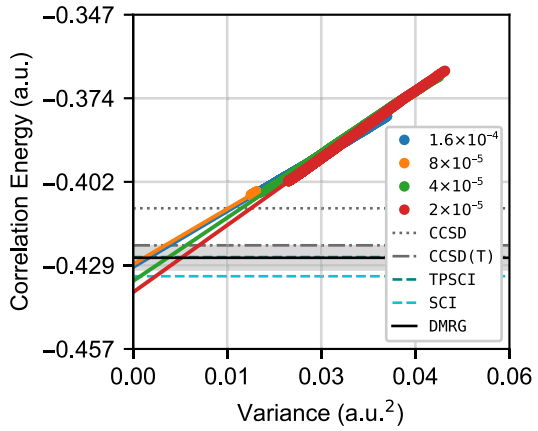


FIG. 8. Correlation energy vs variance extrapolations for the HBC molecule using a series of vDBF thresholds: $\epsilon = 1.6 \times 10^{-4}$ (blue), 8×10^{-5} (orange), 4×10^{-5} (green), and 2×10^{-5} (red). Extrapolations are shown as solid lines, with uncertainties (the difference between linear and quadratic fits) represented by the shaded regions surrounding them. The DMRG best estimate is indicated by the solid black line, while the 1% correlation-energy error region relative to DMRG is shown as a gray-shaded area. Reference energies are included for TPSCI (dashed teal line), SCI (dashed purple line), CCSD (dotted gray line), and CCSD(T) (dash-dot gray line).

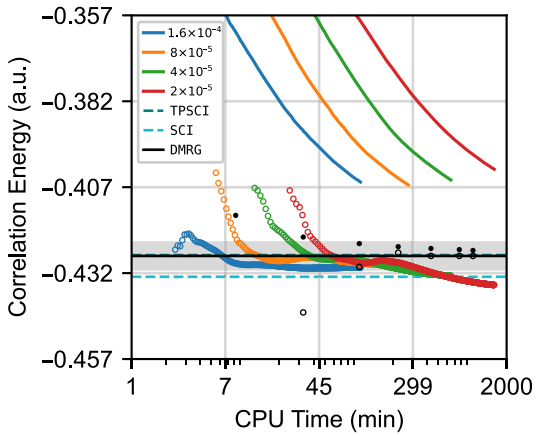


FIG. 9. Correlation energy vs log time for the HBC molecule using a series of vDBF thresholds: $\epsilon = 1.6 \times 10^{-4}$ (blue), 8×10^{-5} (orange), 4×10^{-5} (green), and 2×10^{-5} (red). vDBF variational energies are shown as solid curves, while their corresponding linear extrapolations over time are shown as hollow circles. DMRG variational energies at different sweeps are marked by black dots, with the final DMRG extrapolated estimate shown as the solid black line. The 1% correlation-energy error region relative to the DMRG estimate is denoted by the gray-shaded area. Reference limits for TPSCI and SCI are shown as dashed teal and dashed purple lines, respectively.

the converged DMRG estimate, illustrating that variance extrapolation can fail differently for different methods.

The vDBF results for HBC follow a similar pattern as the lattice models, with one difference worth examining closely. The two loosest thresholds, 1.6×10^{-4} and 8×10^{-5}, extrapolate to within 1% of the DMRG reference, at −0.72% and −0.50%, respectively, while the two tightest, 4×10^{-5} and 2×10^{-5}, overshoot it by 1.79% and 2.66%. This is initially surprising, as one would ideally expect to see systematic improvement as one tightens the ϵ thresholds. Figure 9 shows how this differs from the Hubbard lattices: as in Fig. 6, the extrapolated energies for all thresholds drop into the 1% region early on, but here the tighter thresholds do not remain there, drifting away as the calculation proceeds.

Investigating this nonmonotonic convergence, we first note that the fidelity measure that directly depends on the truncation threshold, the discarded weight, does indeed behave as expected, systematically decreasing as we decrease ϵ. However, the measure that depends on convergence of the flow (or optimization) directly tracks the accuracy here, decreasing from 16×10^{-5} to 8×10^{-5}, then increasing as ϵ is tightened. So this calculation seems to be limited, not by truncation, but rather flow convergence. However, in principle, distance from convergence is not alone a good metric for determining the accuracy of the extrapolated energy. Once the energy-variance plot becomes truly linear, then the extrapolated energy is converged (although it is natural that the smaller the final variance is, the more linear we would expect the fit to be). From this result, we see that the tighter thresholds take longer to reach the linear regime and thus require many more iterations to converge the extrapolated energy. As such, we interpret the 8×10^{-5} calculation to be the most robust HBC energy computed with vDBF.

To better understand the convergence behavior of the extrapolation (not just for HBC, but for all systems in this paper), in Appendix C, we explore how the extrapolations behave as more iterations are included. Figure 12 plots both the energy and the extrapolated energy as a function of variance. Looking at the HBC curve [Fig. 12(g)], we see that while the loose threshold extrapolated energies are quite converged, the tighter calculations are still converging. From this plot is appears as though the tighter calculations 4×10^{-5} and 2×10^{-5} would approach the DMRG result with enough iterations.

IV. CONCLUSIONS

Classical simulation techniques developed to challenge claims of quantum utility, Sparse Pauli Dynamics or Pauli propagation, were originally designed to benchmark quantum hardware. Here, we have shown that the same ideas can also serve as efficient computational kernels for estimating ground-state energies of strongly correlated systems. Our variational double bracket flow method, vDBF, combines Pauli propagation with greedy, gradient-based selection of Pauli generators, coefficient-based fluctuation truncation, and energy-variance extrapolation.

This provides a compact operator-based route to ground-state energy estimation without explicitly storing a many-body wave function.

For the 10×10 antiferromagnetic Heisenberg model (100 qubits), vDBF reaches sub-1% accuracy in approximately 2 min on a single CPU thread; the comparable DMRG calculation required over 50 h on 64 threads. For the 8×8 Hubbard model at half-filling (128 qubits), our DMRG calculations required $\chi = 6000$ and more than 80 h over 64 threads, whereas vDBF reaches this DMRG extrapolated energy to about 0.4% in 40 min and 0.07% or better at tighter thresholds.

We also tested more challenging problems that exhibit a sign problem. The maximally spin-frustrated $J_1 - J_2$ model on an 8×8 lattice extrapolates to within 0.3% of the best variational reference, although the presence of spin-frustration results in a significantly greater number of Paulis that must be tracked. For the hole-doped 8×8 lattice at $U/t = 4$, every vDBF extrapolation considered fell within the target 1% error, and the tightest calculation had an error of only 0.09% compared to the reference VAFQMC value. These results demonstrate that Pauli propagation can provide competitive ground-state energy estimates at substantially reduced computational cost for the lattice benchmarks considered here.

For a molecular example, we considered the full π-space of hexabenzocoronene (84 qubits). Here we found that the two loosest thresholds agree with our reference values to within 1%, the best being -0.50% at $\epsilon = 8 \times 10^{-5}$, while the two tightest overestimate the correlation energy due to the fact that the convergence is slower for tighter thresholds. Overall, these results show that vDBF achieves competitive accuracy across both lattice and molecular systems.

The proposed vDBF method is primarily meant to serve as an energy-estimation approach, as it can recover accurate ground-state energies even when the underlying truncated operator representation does not faithfully reproduce all ground-state observables. This limitation is visible in Fig. 3, where the unextrapolated spin-spin correlations for the one-dimensional Heisenberg chain decay approximately exponentially rather than with the expected algebraic form.

Several technical developments follow directly from these limitations. Further characterizing the convergence behavior as a function of thresholds will be important for understanding the balance between operator fidelity and convergence criterion (as highlighted in the HBC results). A second priority is enforcing physical symmetries more naturally. The current implementation uses individual Pauli-string rotations, which make it difficult to impose particle-number, spin, and other symmetry constraints exactly. While Appendix D demonstrates that the resulting drift in both quantities over the course of a run is relatively modest, strict conservation is desirable for interpretation and, likely, computational efficiency. Efforts in this direction will follow the techniques developed for leveraging symmetry in the Pauli basis detailed in Refs. [122,123], we expect to see important property improvements using symmetry-preserving generators, grouped fermionic excitations, or explicit penalty terms, which may improve both convergence and interpretability. A third priority will focus on developing a scalable parallel implementation following the work of Ref. [39], which will give us access to larger Pauli sums that can fit on a single node. Finally, convergence at late iterations may be accelerated by limited nonlocal reoptimization, such as variationally relaxing the last k selected operators, or by coupling the Heisenberg-picture vDBF flow to a Schrödinger-picture correction inspired by Ref. [36].

More broadly, this work shows that techniques developed to probe the boundary between classical and quantum computation can be redirected toward long-standing classical problems in strongly correlated matter. With improved extrapolation stability, symmetry control, and observable-specific extensions, Pauli-propagation-based methods may become useful tools for difficult problems ranging from frustrated magnetism and Hubbard-like models of correlated materials to extended π-electron systems and molecular active spaces relevant to chemical bonding and catalysis.

ACKNOWLEDGMENTS

We thank Marek Gluza, Zoë Holmes, and Manuel Rudolph for helpful comments and discussion. The authors are grateful for the generous support from the National Science Foundation (Award No. 2414574). The work by EFV was supported by the U.S. Department of Energy via award DE-SC0022327. This research was supported in part by Lilly Endowment, Inc. through its support for the Indiana University Pervasive Technology Institute.

DATA AVAILABILITY

The software used to generate the vDBF results presented in this manuscript is publicly available. The vDBF algorithm is implemented in DBF.jl (Ref. [124]), which depends on PauliOperators.jl (Ref. [125]) for symplectic Pauli representations. DMRG calculations for the Heisenberg and Hubbard models were performed using the ITensors.jl package; DMRG calculations for hexabenzocoronene used BLOCK2, and selected configuration interaction calculations used MPQC. Input files, scripts, and numerical data needed to reproduce the figures and tables in this manuscript are available from the corresponding author upon reasonable request.

APPENDIX A: COMPARISON TO EXACT PROJECTOR

In the theory section above, we motivated the use of the simple sum of single-qubit Z operators in the flow

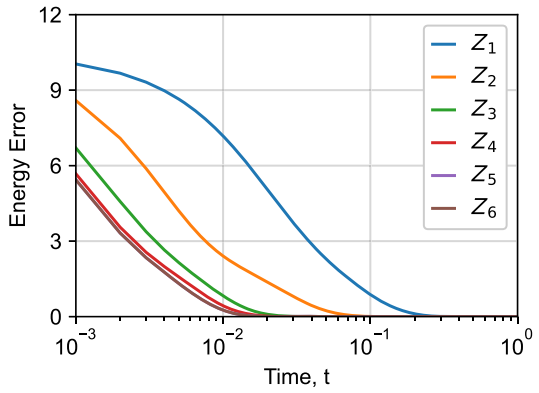


FIG. 10. Convergence of the exact continuous double bracket flow for different approximations to the projector for a 1×6 Heisenberg lattice. Z_i indicates that only terms up to i are kept in the projector.

generator, $[H(s), \sum_i Z_i]$ as an approximation to the more complicated projector-based generator, $[H(s), |\psi\rangle\langle\psi|]$. While this approximation was made for computational considerations, we show in Fig. 10 that this seemingly aggressive approximation still works to drive the system to the ground state. We plot the difference between the exact ground-state energy and the time-dependent expectation value of $H(s)$, for various approximations to the exact projection-based generator for a small system of a one-dimensional six-site Heisenberg model. We see that while the full projector (Z_6) exhibits the fastest convergence (Z_5 is actually identical here, since Z_6 simply adds the $ZZZZZZ$ operator, which is a symmetry), all the other n-body approximations still eventually converge.

While the Z_1 approximation is the slowest, the implementation is much more efficient. Further, in our final algorithm, the final step-size is not determined by directly integrating the double bracket flow, but rather chosen to variationally minimize the energy.

APPENDIX B: TRUNCATION-CORRECTED VARIANCE

In the plots above, we rely on zero-variance extrapolations to obtain accurate ground-state energy estimates. As discussed in Sec. II C 2, we explain how truncating only the energy fluctuation instead of the full operator provides a significant improvement to the accuracy of the energy. To see this, consider the Hamiltonian at iteration n just before truncation: $\tilde{H}^{(n)} = U_n^\dagger H^{(n-1)} U_n$, where $U_n = e^{-i\theta_n P_n}$. Truncation is performed by defining a projection (super)operator, $\mathcal{P}^{(n)}$, onto the Paulis in $\tilde{H}^{(n)}$ which are determined to be important, i.e., the Paulis which have coefficients $|c_i| > \epsilon$. The truncated operator, $H^{(n)}$, is then written in terms of this projector and its orthogonal complement, $\mathcal{Q}^{(n)} = 1 - \mathcal{P}^{(n)}$:

$$\tilde{H}^{(n)} = \mathcal{P}^{(n)}\tilde{H}^{(n)} + \mathcal{Q}^{(n)}\tilde{H}^{(n)}, \tag{B1}$$

$$H^{(n)} = \mathcal{P}^{(n)}\tilde{H}^{(n)} + \left\langle \mathcal{Q}^{(n)}\tilde{H}^{(n)} \right\rangle. \tag{B2}$$

Keeping the expectation value instead of just deleting the $\mathcal{Q}$ projection ensures that the operators expectation is unaffected by truncation, $\langle H^{(n)} \rangle = \langle \tilde{H}^{(n)} \rangle$, and the only approximation is then in the downstream dynamics.

While the interpretation of this as a fluctuation truncation is physically insightful, one could equivalently obtain this same result in a more mechanistic approach by using "bare truncation" (where we do not add in the Q space expectation value) and simply computing the "truncation error" associated with projecting onto the $\mathcal{P}^{(n)}$ space and accumulating this error over the sequence of rotations. For example,

$$\Delta_E^{(n)} = \left\langle \mathcal{P}^{(n)}\tilde{H}^{(n)} \right\rangle - \left\langle \tilde{H}^{(n)} \right\rangle. \tag{B3}$$

After N rotations, the final energy would then be expressed as:

$$E^{(N)} = \left\langle \mathcal{P}^{(N)}\tilde{H}^{(N)} \right\rangle - \sum_{n=1}^{N} \Delta_E^{(n)} \tag{B4}$$

$$= \left\langle H^{(N)} \right\rangle. \tag{B5}$$

We similarly compute the final energy variance as:

$$V^{(N)} = \mathrm{var}(\mathcal{P}^{(n)}\tilde{H}^{(N)}) - \sum_{n=1}^{N} \Delta_{\mathrm{var}}^{(n)}, \tag{B6}$$

where, $\Delta_{\mathrm{var}}^{(n)} = \mathrm{var}(\mathcal{P}^{(n)}\tilde{H}^{(n)}) - \mathrm{var}(\tilde{H}^{(n)})$. This variance correction provides an estimate of how much of the variance reduction resulted from truncation rather than convergence.

All the reported variances in this paper include this truncation correction.

APPENDIX C: CONVERGENCE OF THE EXTRAPOLATED ENERGY

Because vDBF relies on extrapolation for quantitative accuracy, it is important to inspect the behavior and convergence of this extrapolation. In Figs. 11 and 12, we plot the convergence of both the raw energies and the extrapolated energies as a function of iteration and variance, respectively. In order to ensure a more interpretable extrapolation convergence plot, we choose the number of points in a slightly different manner than in the main text. Here, we perform the extrapolations with a rolling window such that $E_{\mathrm{extrap},i}$ is computed from the data points which have variance up to 2 times the variance at

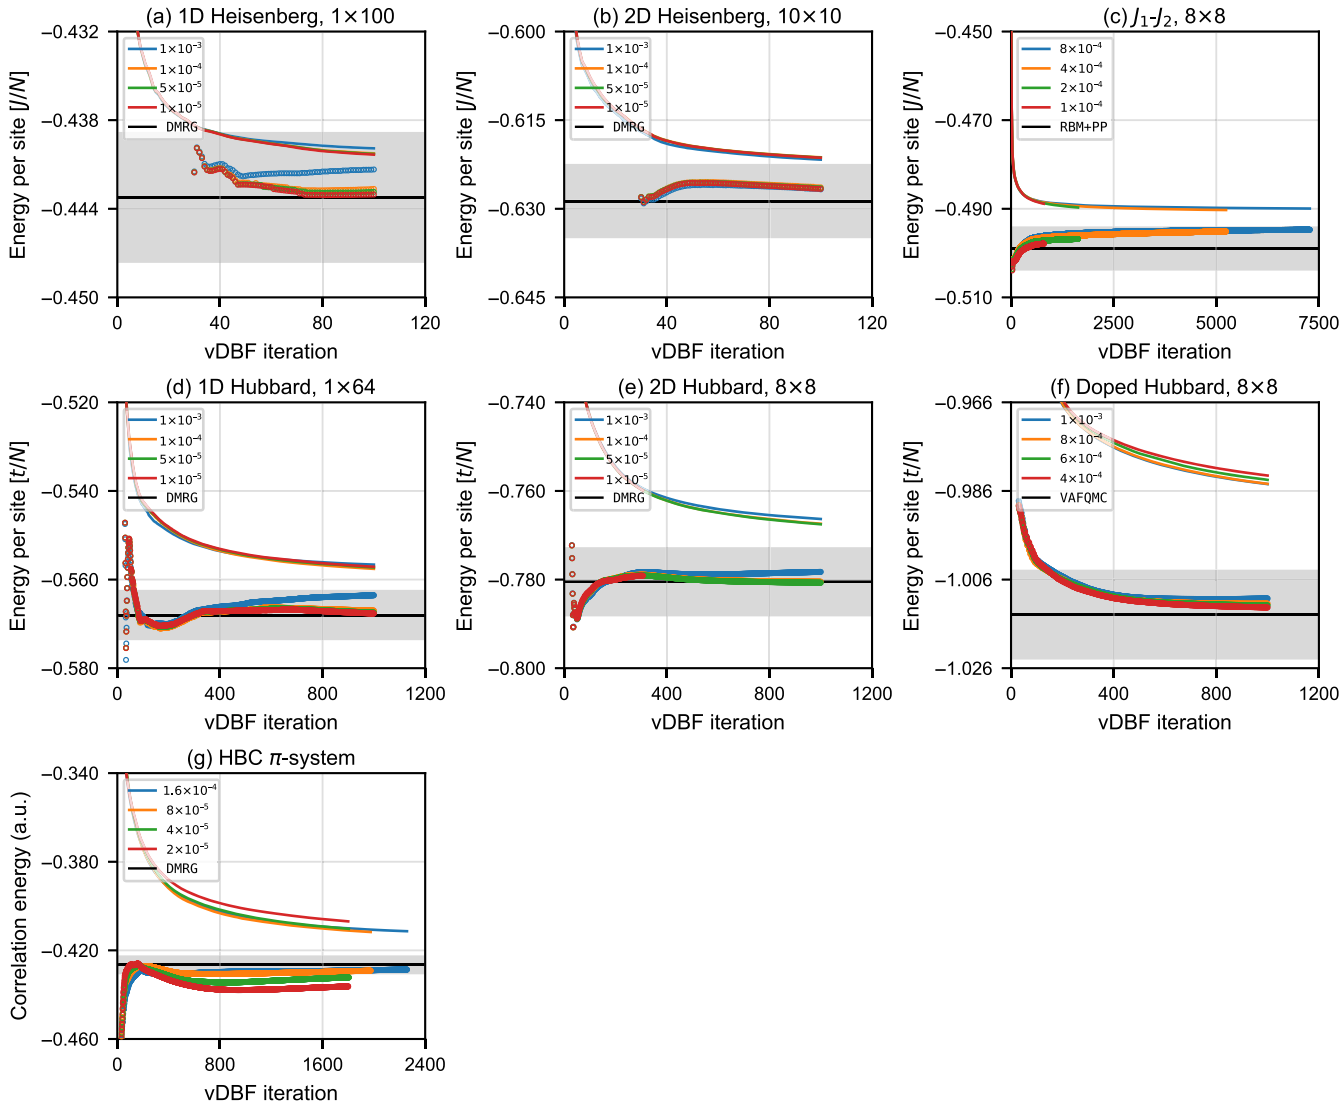


FIG. 11. Convergence of the zero-variance extrapolation across the benchmark systems. Solid curves are the raw vDBF trajectories; open markers are the running intercepts for each iteration. For each raw data point (E_i, V_i), the extrapolation is performed using all the previous data points up to $2V_i$, ensuring that the extrapolation distance does not exceed the data range. Reference energies are the solid black lines, with the 1% window shaded.

iteration i. Using this rolling $[V, 2V]$ window ensures that the distance of extrapolation never exceeds the variance range of the data.

Considering each system in turn, we find that overall, the lattice systems generally display extrapolations that quickly converge within the 1% error window. A couple of exceptions can be noted. For both 1D Hubbard and Heisenberg, we see that the loosest threshold, 1×10^{-3}, seems to exhibit significant truncation error, which then results in an extrapolation trajectory that deviates from the tighter calculations. For the J_1–J_2 model, it appears as though the the extrapolations have not yet sufficiently converged and will possibly increase in error with more iterations, although the tighter thresholds seem to be converging (with variance) a bit faster here. However, for the molecular HBC example, we notice extrapolation behavior that is much further from convergence for the tighter thresholds. While the looser 16×10^{-5} and 8×10^{-5} extrapolations appear to have flatlined within the 1% region (consistent with Table VI), the tighter calculations are converging more slowly. It appears as though 4×10^{-5} and 2×10^{-5} would need significantly more iterations to converge their extrapolated energies to within the target error window.

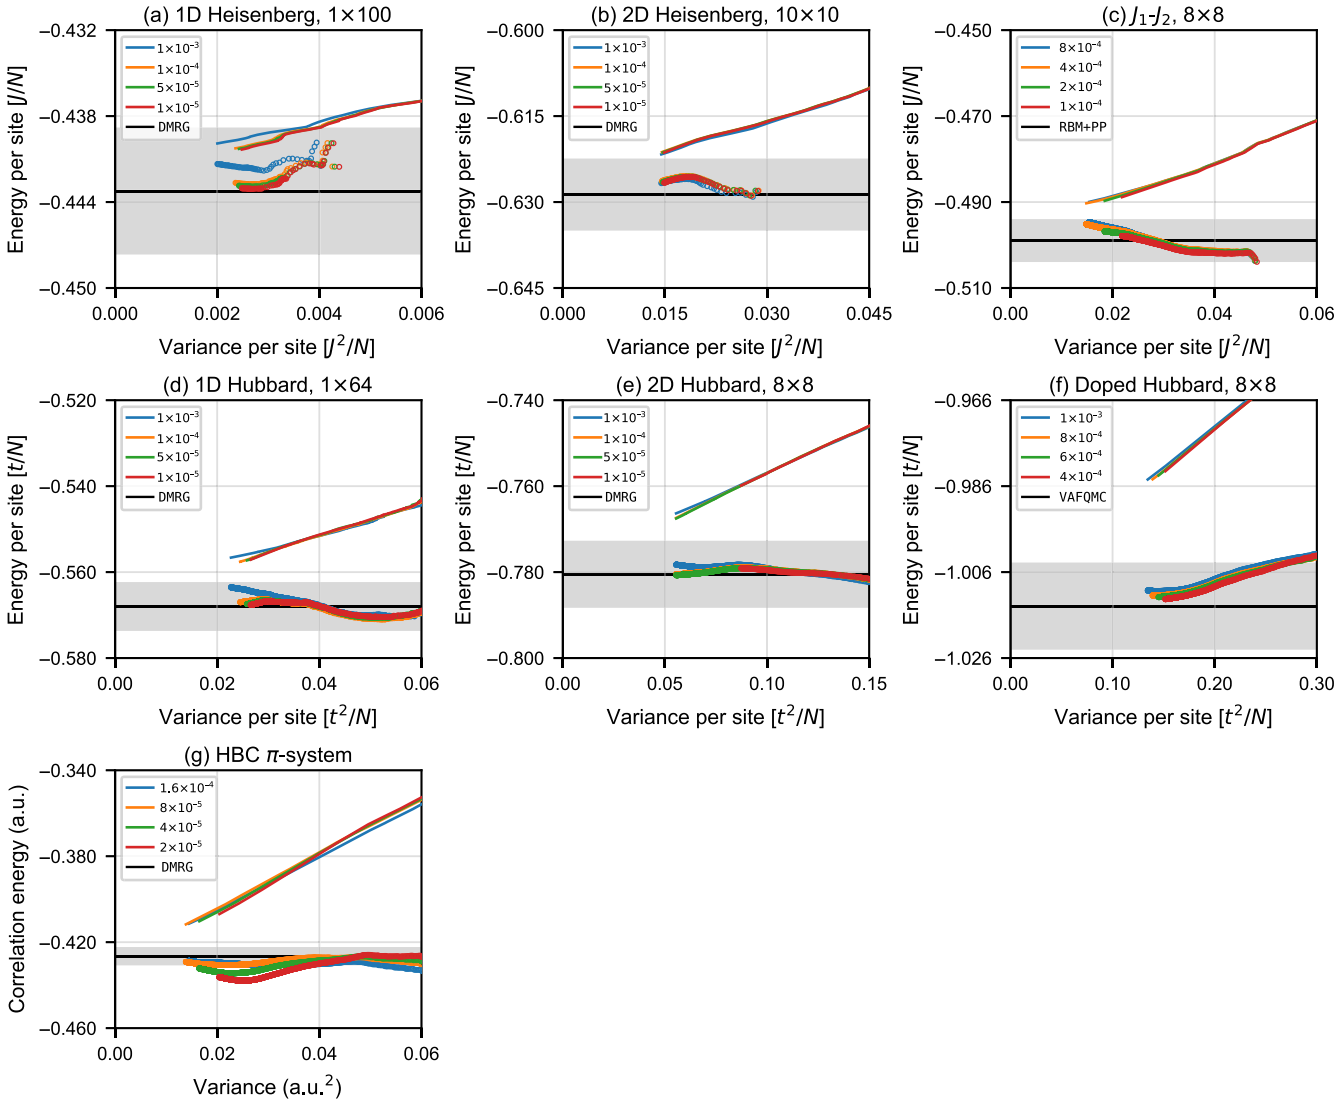


FIG. 12. Convergence of the zero-variance extrapolation across the benchmark systems. Solid curves are the raw vDBF trajectories; open markers are the running intercepts, each drawn at the variance of its window's final iteration. For each raw data point (E_i, V_i), the extrapolation is performed using all the previous data points up to $2V_i$, ensuring that the extrapolation distance does not exceed the data range. Reference energies are the solid black lines, with the 1% window shaded. Same plot as Fig. 11 but x axis is variance instead of iteration.

APPENDIX D: PARTICLE NUMBER AND SPIN-PROJECTION CONSERVATION

In Sec. III B 1, we observed limitations in the ability of vDBF to capture long-distance spin-spin correlations in the Heisenberg example. For quantum chemistry applications, the ability to obtain accurate estimates of observables beyond just the energy is also valuable. In Figs. 13 and 14, we plot the expectation values of the number and spin-projection operators as a function of iteration. Because individual Pauli rotations do not generally commute with either the number operator or the spin-projection operator, the average number of electrons and the spin projection are expected to fluctuate slightly over the course of a vDBF optimization. From these plots, we do indeed see that neither the particle number nor the spin projection is constant, with deviations from the exact values (42 electrons and $S_z = 0$, respectively) given on the y axis. Although particle number and spin projection are not exactly preserved, we do see that the deviations are quite small, generally in the fourth decimal place aside from brief early-iteration transients. The particle-number deviation is also biased toward a small net overall negative charge. In future work, we will explore adding either

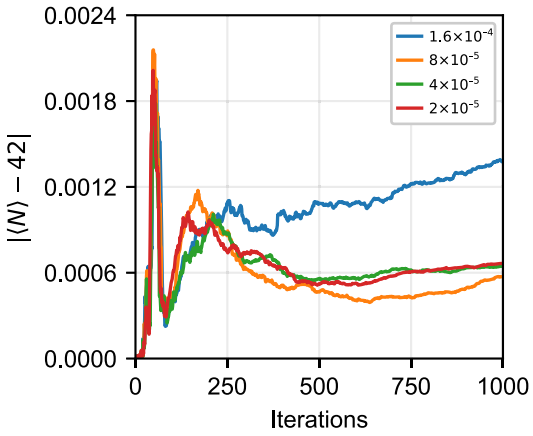


FIG. 13. Absolute error in the corrected expectation value of the evolved number operator with the circuits generated using different thresholds for the HBC molecule.

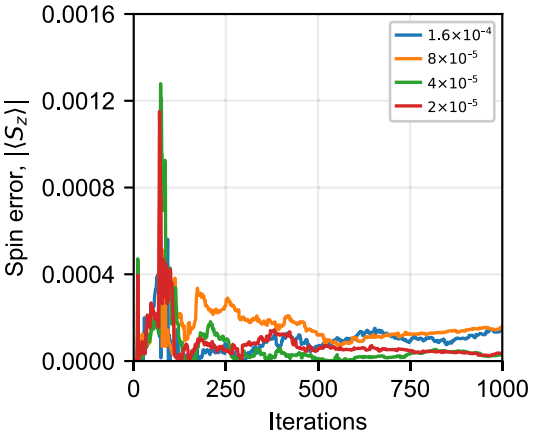


FIG. 14. Absolute error in the corrected expectation value of the evolved spin-projection operator with the circuits generated using different thresholds for the HBC molecule.

symmetry penalties or symmetry-constrained rotations to make the computation of properties more robust.

APPENDIX E: OPERATOR WEIGHT UNDER THE JORDAN–WIGNER MAPPING

The vDBF truncation acts on Pauli coefficients, and the Pauli weight depends on the fermion-to-qubit mapping. Under Jordan–Wigner, an operator acting on a few orbitals picks up a string of Z operators spanning the intervening qubits, so its Pauli weight reflects orbital indices as much as physical locality. The Majorana weight, twice the fermionic excitation rank, is independent of the mapping, so comparing the two separates the encoding from the physics. Figure 15 shows both for the HBC trajectory at $\epsilon = 2 \times 10^{-5}$. In the Pauli basis, panel (a), terms spread across weights 0–84 within the first hundred iterations, with most of the count between 14 and 56. In the Majorana basis, panel (b), the same operator has support only at weights 4, 6, and 8: the initial Hamiltonian is dominated by weight 4, a weight-6 component takes over during the first hundred iterations, and nothing above eight acquires appreciable weight for the rest of the run. The apparent spreading is therefore a property of the encoding, not of the fermionic operator. What Jordan–Wigner renders as a weight-30 Pauli string is a two- or three-body fermionic operator.

The norm tells the same story. Panel (c) places almost all of the centered L_2 weight at very low Pauli weight, and panel (d) at Majorana weights 2–6, with only a small contribution at 8 and above. Both distributions are stationary after the first hundred iterations, so the flow does not progressively drive the operator toward higher fermionic rank. The truncation, therefore, does not discard high-body physics, and the results do not depend on the Jordan–Wigner ordering.

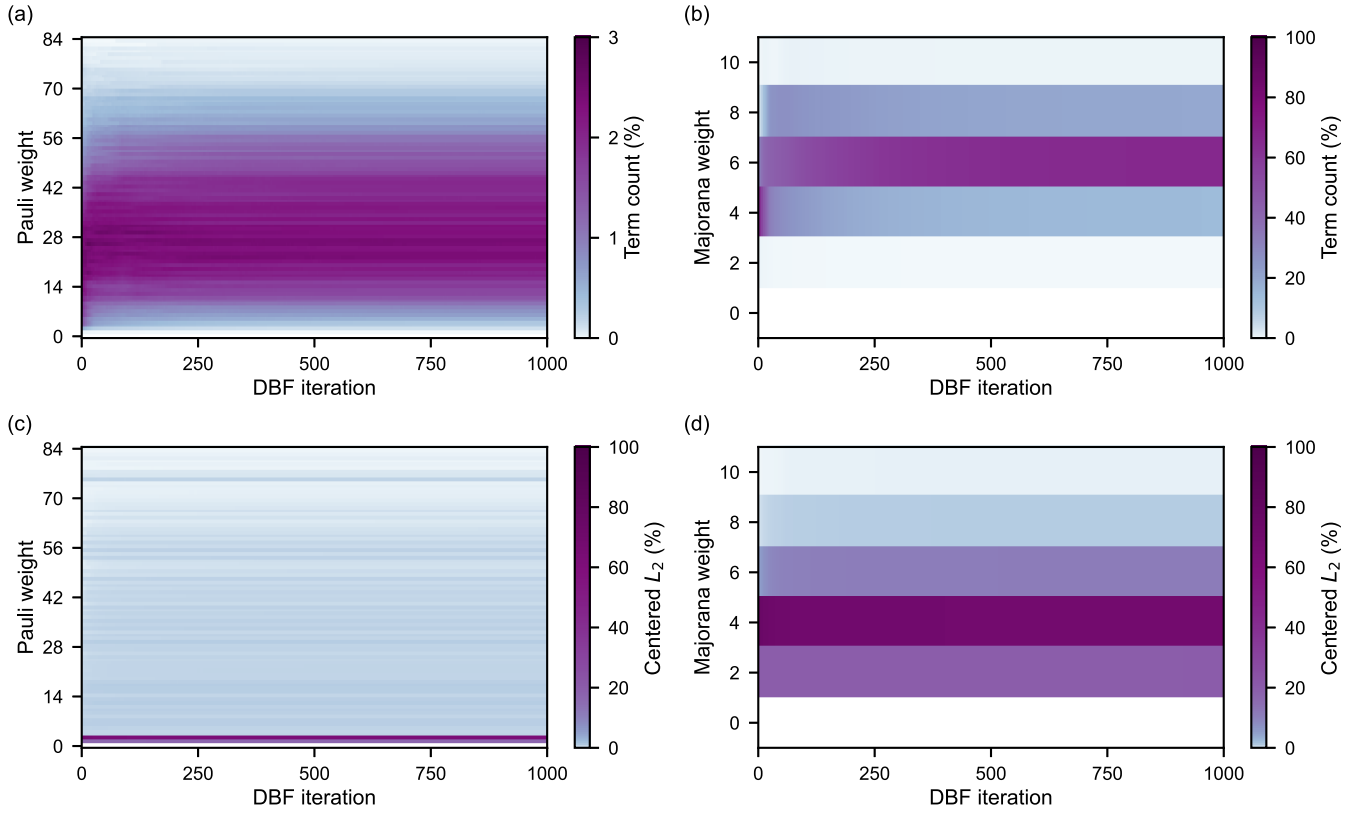


FIG. 15. Operator weight distribution along the HBC trajectory at $\epsilon = 2 \times 10^{-5}$ (84 qubits, 1000 iterations). Left panels are in the Pauli basis; right panels are in the Majorana basis. The top panels show the fraction of terms at each weight; the bottom panels show the fraction of the centered L_2 norm. The identity contribution is removed from weight 0 in the lower panels.

[1] C. K. Møller and M. S. Plesset, *Note on an approximation treatment for many-electron systems*, Phys. Rev. **46**, 618 (1934).

[2] R. J. Bartlett and G. D. Purvis, *Many-body perturbation theory, coupled-pair many-electron theory, and the importance of quadruple excitations for the correlation problem*, Int. J. Quantum Chem. **14**, 561 (1978).

[3] I. Shavitt and R. J. Bartlett, *Many-Body Methods in Chemistry and Physics: MBPT and Coupled-Cluster Theory*, 1st ed. (Cambridge University Press, Cambridge, England, 2009).

[4] R. J. Bartlett and M. Musiał, *Coupled-cluster theory in quantum chemistry*, Rev. Mod. Phys. **79**, 291 (2007).

[5] R. J. Bartlett, *Many-body perturbation theory and coupled cluster theory for electron correlation in molecules*, Ann. Rev. Phys. Chem. **32**, 359 (1981).

[6] L. Hedin, *New method for calculating the one-particle Green's function with application to the electron-gas problem*, Phys. Rev. **139**, A796 (1965).

[7] A. Marie, A. Ammar, and P.-F. Loos, *Advances in Quantum Chemistry* (Elsevier, New York, 2024), Vol. 90, pp. 157–184.

[8] F. Aryasetiawan and O. Gunnarsson, *The GW method*, Rep. Prog. Phys. **61**, 237 (1998).

[9] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Quantum Monte Carlo simulations of solids*, Rev. Mod. Phys. **73**, 33 (2001).

[10] G. H. Booth, A. J. W. Thom, and A. Alavi, *Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space*, J. Chem. Phys. **131**, 054106 (2009).

[11] D. Cleland, G. H. Booth, and A. Alavi, *Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum Monte Carlo*, J. Chem. Phys. **132**, 041103 (2010).

[12] M. Motta and S. Zhang, *Ab Initio computations of molecular systems by the auxiliary-field quantum Monte Carlo method*, Comput. Mol. Sci. **8**, e1364 (2018).

[13] S. Zhang and H. Krakauer, *Quantum Monte Carlo method using phase-free random walks with slater determinants*, Phys. Rev. Lett. **90**, 136401 (2003).

[14] S. Zhang, J. Carlson, and J. E. Gubernatis, *Constrained path Monte Carlo method for fermion ground states*, Phys. Rev. B **55**, 7464 (1997).

[15] G. K.-L. Chan and S. Sharma, *The density matrix renormalization group in quantum chemistry*, Ann. Rev. Phys. Chem. **62**, 465 (2011).

[16] K. H. Marti and M. Reiher, *New electron correlation theories for transition metal chemistry*, Phys. Chem. Chem. Phys. **13**, 6750 (2011).

[17] S. Wouters and D. Van Neck, *The density matrix renormalization group for ab initio quantum chemistry*, Eur. Phys. J. D **68**, 272 (2014).

[18] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69**, 2863 (1992).

[19] S. R. White and R. L. Martin, *Ab initio quantum chemistry using the density matrix renormalization group*, J. Chem. Phys. **110**, 4127 (1999).

[20] S. R. White, *Density matrix renormalization group algorithms with a single center site*, Phys. Rev. B **72**, 180403 (2005).

[21] Y. Ma, J. Wen, and H. Ma, *Density-matrix renormalization group algorithm with multi-level active space*, J. Chem. Phys. **143**, 034105 (2015).

[22] K. Hyatt and E. M. Stoudenmire, *DMRG approach to optimizing two-dimensional tensor networks*, arXiv:1908.08833.

[23] M. T. Fishman, L. Vanderstraeten, V. Zauner-Stauber, J. Haegeman, and F. Verstraete, *Faster methods for contracting infinite 2D tensor networks*, arXiv:1711.05881.

[24] L. Ma, M. Fishman, E. M. Stoudenmire, and E. Solomonik, *Approximate contraction of arbitrary tensor networks with a flexible and efficient density matrix algorithm*, Quantum **8**, 1580 (2024).

[25] V. Zauner-Stauber, L. Vanderstraeten, M. T. Fishman, F. Verstraete, and J. Haegeman, *Variational optimization algorithms for uniform matrix product states*, Phys. Rev. B **97**, 045145 (2018).

[26] G. Vidal, *Classical simulation of infinite-size quantum lattice systems in one spatial dimension*, Phys. Rev. Lett. **98**, 070201 (2007).

[27] F. Verstraete and J. I. Cirac, *Renormalization algorithms for quantum-many body systems in two and higher dimensions*, arXiv:cond-mat/0407066.

[28] E. M. Stoudenmire, *Learning relevant features of data with multi-scale tensor networks*, Quantum Sci. Technol. **3**, 034003 (2018).

[29] G. Vidal, *Entanglement renormalization*, Phys. Rev. Lett. **99**, 220405 (2007).

[30] M. Troyer and U.-J. Wiese, *Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations*, Phys. Rev. Lett. **94**, 170201 (2005).

[31] Google Quantum AI and Collaborators, *Quantum error correction below the surface code threshold*, Nature (London) **638**, 920 (2025).

[32] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, *Evidence for the utility of quantum computing before fault tolerance*, Nature (London) **618**, 500 (2023).

[33] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, *Logical quantum processor based on reconfigurable atom arrays*, Nature (London) **626**, 58 (2024).

[34] S. Lee *et al.*, *Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry*, Nat. Commun. **14**, 1952 (2023).

[35] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, *Evidence for the utility of quantum computing before fault tolerance*, Nature (London) **618**, 500 (2023).

[36] T. Begušić, J. Gray, and G. K. Chan, *Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance*, Sci. Adv. **10**, eadk4321 (2024).

[37] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, *Efficient tensor network simulation of IBM's eagle kicked Ising experiment*, PRX quantum **5**, 010308 (2024).

[38] P. Rall, D. Liang, J. Cook, and W. Kretschmer, *Simulation of qubit quantum circuits via Pauli propagation*, Phys. Rev. A **99**, 062337 (2019).

[39] L. Broers, R.-Y. Sun, and S. Yunoki, *Scalable simulation of quantum many-body dynamics with OR-represented quantum algebra*, arXiv:2506.13241.

[40] M. S. Rudolph, E. Fontana, Z. Holmes, and L. Cincio, *Classical surrogate simulation of quantum systems with LOWESA* (2023).

[41] N. Loizeau, J. C. Peacock, and D. Sels, *Quantum many-body simulations with PauliStrings.jl*, SciPost Phys. Codebases 054 (2025).

[42] L. Müller, A. Bärligea, A. Knapp, and J. S. Kottmann, *PauliEngine: High-performant symbolic arithmetic for quantum operations*, arXiv:2601.02233.

[43] T. Schuster, C. Yin, X. Gao, and N. Y. Yao, *A polynomial-time classical algorithm for noisy quantum circuits* (2024).

[44] C. Cirstoiu, *A Fourier analysis framework for approximate classical simulations of quantum circuits* (2024).

[45] E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, and C. Cîrstoiu, *Classical simulations of noisy variational quantum circuits*, npj Quantum Inf. **11**, 84 (2025).

[46] T. Begušić, K. Hejazi, and G. K. Chan, *Simulating quantum circuit expectation values by Clifford perturbation theory*, J. Chem. Phys. **162**, 154110 (2025).

[47] T. Schuster, C. Yin, X. Gao, and N. Y. Yao, *A polynomial-time classical algorithm for noisy quantum circuits*, Phys. Rev. X **15**, 041018 (2025).

[48] A. Angrisani, A. Schmidhuber, M. S. Rudolph, M. Cerezo, Z. Holmes, and H. Y. Huang, *Classically estimating observables of noiseless quantum circuits*, Phys. Rev. Lett. **135**, 170602 (2025).

[49] A. Angrisani, A. A. Mele, M. S. Rudolph, M. Cerezo, and Z. Holmes, *Simulating quantum circuits with arbitrary local noise using pauli propagation*, PRX Quantum **7**, 020313 (2026).

[50] A. Miller, Z. Holmes, Ö. Salehi, R. Chakraborty, A. Nykänen, Z. Zimborás, A. Glos, and G. García-Pérez, *Simulation of fermionic circuits using majorana propagation*, arXiv:2503.18939.

[51] M. Robbiati, E. Pedicillo, A. Pasquale, X. Li, O. Kiss, A. Wright, R. M. Farias, K. U. Giang, J. Son, J. Knörzer, and S. T. Goh, *Double-bracket quantum algorithms for high-fidelity ground-state preparation*, Phys. Rev. Res. **8**, 033078 (2026).
[52] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, *Pauli propagation: A computational framework for simulating quantum systems*, PRX Quantum **7**, 032001 (2026).
[53] G. González-García, J. I. Cirac, and R. Trivedi, *Pauli path simulations of noisy quantum circuits beyond average case*, Quantum **9**, 1730 (2025).
[54] M. Gluza, *Double-bracket quantum algorithms for diagonalization*, Quantum **8**, 1316 (2024).
[55] M. Gluza, J. Son, B. H. Tiang, R. Zander, R. Seidel, Y. Suzuki, Z. Holmes, and N. H. Y. Ng, *Double-bracket quantum algorithms for quantum imaginary-time evolution*, Phys. Rev. Lett. **136**, 020601 (2026).
[56] Y. Suzuki, B. H. Tiang, J. Son, N. H. Y. Ng, Z. Holmes, and M. Gluza, *Double-bracket algorithm for quantum signal processing without post-selection*, Quantum **9**, 1954 (2025).
[57] Y. Suzuki, M. Gluza, J. Son, B. H. Tiang, N. H. Ng, and Z. Holmes, *Grover's algorithm is an approximation of imaginary-time evolution*, arXiv:2507.15065.
[58] M. Alghadeer, K. U. Giang, S. Cao, S. D. Fasciati, M. Piscitelli, N. Ng, P. J. Leek, M. Gluza, and M. Bakr, *Double-bracket algorithmic cooling*, arXiv:2510.00302.
[59] While one can always find a change of basis that makes the ground-state sparse (e.g., the eigenbasis), this is computationally intractable, whereas Hamiltonians are generally sparse in computationally tractable bases (e.g., Pauli, or fermionic operator bases).
[60] G. Aguilar, S. Cichy, J. Eisert, and L. Bittel, *Full classification of Pauli lie algebras*, arXiv:2408.00081v1.
[61] E. Kökcü, T. Steckmann, Y. Wang, J. K. Freericks, E. F. Dumitrescu, and A. F. Kemper, *Fixed depth Hamiltonian simulation via cartan decomposition*, Phys. Rev. Lett. **129**, 070501 (2022).
[62] R. Wiersema, E. Kökcü, A. F. Kemper, and B. N. Bakalov, *Classification of dynamical lie algebras of 2-local spin systems on linear, circular and fully connected topologies*, npj Quantum Inf. **10**, 110 (2024).
[63] I. Magoulas and F. A. Evangelista, *Clifford transformations for fermionic quantum systems: From Paulis to Majoranas to Fermions*, arXiv:2510.23923.
[64] P. Jayakumar, T. Zeng, and A. F. Izmaylov, *On the feasibility of exact unitary transformations for many-body Hamiltonians*, J. Chem. Theory Comput. **22**, 3916 (2026).
[65] F. A. Evangelista and I. Magoulas, *Exact closed-form unitary transformations of fermionic operators*, Phys. Rev. A **111**, 042825 (2025).
[66] I. Magoulas and F. A. Evangelista, *Spin-adapted fermionic unitaries: From lie algebras to compact quantum circuits*, arXiv:2511.13485.
[67] V. Martinez, A. Angrisani, E. Pankovets, O. Fawzi, and D. S. França, *Efficient simulation of parametrized quantum circuits under non-unital noise through Pauli backpropagation*, Phys. Rev. Lett. **134**, 250602 (2025).
[68] S. D. Glazek and K. G. Wilson, *Perturbative renormalization group for Hamiltonians*, Phys. Rev. D **49**, 4214 (1994).
[69] S. D. Głazek and K. G. Wilson, *Renormalization of Hamiltonians*, Phys. Rev. D **48**, 5863 (1993).
[70] F. Wegner, *Flow-equations for Hamiltonians*, Ann. Phys. **506**, 77 (1994).
[71] H. Hergert, S. K. Bogner, J. G. Lietz, T. D. Morris, S. J. Novario, N. M. Parzuchowski, and F. Yuan, in *An Advanced Course in Computational Nuclear Physics: Bridging the Scales from Quarks to Neutron Stars*, edited by M. Hjorth-Jensen, M. P. Lombardo, and Kolck van (Springer International Publishing, Cham, 2017), p. 570.
[72] F. A. Evangelista, *Driven similarity renormalization group approach to quantum many-body problems*, J. Chem. Phys. **141**, 054109 (2014).
[73] A. Marie and P.-F. Loos, *A Similarity renormalization group approach to Green's function methods*, J. Chem. Theory Comput. **19**, 3943 (2023).
[74] S. Wang, C. Li, and F. A. Evangelista, *Analytic energy gradients for the driven similarity renormalization group multireference second-order perturbation theory*, J. Chem. Theory Comput. **17**, 7666 (2021).
[75] Any computational basis state can be rotated into $|0\rangle$, folding the resulting unitary into $H(0)$.
[76] L. Xiaoyue, M. Robbiati, A. Pasquale, E. Pedicillo, A. Wright, S. Carrazza, and M. Gluza, *Strategies for optimizing double-bracket quantum algorithms*, arXiv:2408.07431.
[77] M. Pervez, A. Haqq, N. A. McMahon, and C. Arenz, *Riemannian gradient descent-based quantum algorithms for ground state preparation with guarantees*, arXiv:2512.13401.
[78] E. Malvetti, C. Arenz, G. Dirr, and T. Schulte-Herbrüggen, *Randomized gradient descents on riemannian manifolds: Almost sure convergence to global minima in and beyond quantum optimization*, arXiv:2405.12039.
[79] R. Zander, R. Seidel, L. Xiaoyue, and M. Gluza, *Role of Riemannian geometry in double-bracket quantum imaginary-time evolution*, in *Proceedings of the International Conference on Geometric Science of Information.* (2025), pp. 105–114.
[80] R. Wiersema and N. Killoran, *Optimizing quantum circuits with Riemannian gradient flow*, Phys. Rev. A **107**, 062421 (2023).
[81] N. A. McMahon, M. Pervez, and C. Arenz, *Equating quantum imaginary time evolution, Riemannian gradient flows, and stochastic implementations*, arXiv:2504.06123.
[82] While this approach is valid for any computational basis reference state, the choice of Eq. (8) means that we must work with $|0\rangle$ as reference state by first transforming H by a sequence of X gates that transform $|\psi\rangle$ to $|0\rangle$, which is essentially a particle-hole transformation.
[83] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, *An adaptive variational algorithm for exact molecular simulations on a quantum computer*, Nat. Commun. **10**, 3007 (2019).
[84] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou,

Qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient Ansätze on a quantum processor, PRX Quantum **2**, 020310 (2021).

[85] We note that additional applications of the clip function are applied at various points to improve efficiency but are not included in Algorithm for clarity. For example, after forming the flow generator, G, this operator is clipped. However, this is clipped with a tight threshold (default 10^{-6}) so as to not affect the accuracy.

[86] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, *Feedback-based quantum optimization*, Phys. Rev. Lett. **129**, 250502 (2022).

[87] R. Chakraborty, A. Miller, A. Nykänen, S. Özlem, F. Tarocco, F. Pavošević, P. A. B. Haase, M. Stella, and A. Glos, *Scalable quantum circuit generation for iterative ground state approximation using Majorana Propagation*, arXiv:2603.23444.

[88] I. G. Ryabinkin, R. A. Lang, S. N. Genin, and A. F. Izmaylov, *Iterative qubit coupled cluster approach with efficient screening of generators*, J. Chem. Theory Comput. **16**, 1055 (2020).

[89] I. G. Ryabinkin, T.-C. Yen, S. N. Genin, and A. F. Izmaylov, *Qubit coupled-cluster method: A systematic approach to quantum chemistry on a quantum computer*, J. Chem. Theory Comput. **14**, 6317 (2018).

[90] R. A. Lang, A. Ganeshram, and A. F. Izmaylov, *Growth reduction of similarity-transformed electronic Hamiltonians in qubit space*, J. Chem. Theory Comput. **19**, 6656 (2023).

[91] R. A. Lang, S. G. Mehendale, I. G. Ryabinkin, and A. F. Izmaylov, *Multistate iterative qubit coupled cluster (MS-iQCC): A quantum-inspired, state-averaged approach to ground- and excited-state energies* (2025).

[92] T. Mizusaki and M. Imada, *Precise estimation of shell model energy by second-order extrapolation method*, Phys. Rev. C **67**, 041301 (2003).

[93] N. Shimizu, Y. Utsuno, T. Mizusaki, T. Otsuka, T. Abe, and M. Honma, *Novel extrapolation method in the Monte Carlo shell model*, Phys. Rev. C **82**, 061305 (2010).

[94] N. Shimizu, Y. Utsuno, T. Mizusaki, T. Otsuka, T. Abe, and M. Honma, *Extrapolation method in the Monte Carlo shell model and its applications*, AIP Conf. Proc. **1355**, 138 (2011).

[95] T. Mizusaki and N. Shimizu, *New variational Monte Carlo method with energy variance extrapolation for large-scale shell-model calculations*, Phys. Rev. C **85**, 021301 (2012).

[96] https://github.com/nmayhall/DBF.jl (2025).

[97] https://github.com/nmayhall/PauliOperators.jl (2025).

[98] M. Fishman, S. R. White, and E. M. Stoudenmire, *The ITensor software library for tensor network calculations*, arXiv:2007.14822.

[99] H. Zhai, H. R. Larsson, S. Lee, Z.-H. Cui, T. Zhu, C. Sun, L. Peng, R. Peng, K. Liao, J. Tölle, J. Yang, S. Li, and G. K..-L. Chan, *Block2: A comprehensive open source framework to develop and apply state-of-the-art DMRG algorithms in electronic structure and beyond*, J. Chem. Phys. **234801**, 159 (2023).

[100] C. Peng, C. A. Lewis, X. Wang, M. C. Clement, K. Pierce, V. Rishi, F. Pavošević, S. Slattery, J. Zhang, N. Teke, A. Kumar, C. Masteran, A. Asadchev, J. A. Calvin, and E. F. Valeev, *Massively parallel quantum chemistry: A high-performance research platform for electronic structure*, J. Chem. Phys. **153**, 044120 (2020).

[101] R. Alkabetz and I. Arad, *Tensor networks contraction and the belief propagation algorithm*, Phys. Rev. Res. **3**, 023073 (2021).

[102] G. Evenbly, N. Pancotti, A. Milsted, J. Gray, and G. K.-L. Chan, *Loop series expansions for tensor networks*, arXiv:2409.03108.

[103] J. Tindall and M. T. Fishman, *Gauging tensor networks with belief propagation*, arXiv:2306.17837.

[104] Y. Nomura and M. Imada, *Dirac-type nodal spin liquid revealed by refined quantum many-body solver using neural-network wave function, correlation ratio, and level spectroscopy*, Phys. Rev. X **11**, 031034 (2021).

[105] D. Wu *et al.*, *Variational benchmarks for quantum many-bodyproblems*, Science **386**, 296 (2024).

[106] A. W. Sandvik, *Finite-size scaling of the ground-state parameters of the two-dimensional Heisenberg model*, Phys. Rev. B **56**, 11678 (1997).

[107] A. W. Sandvik, *High-precision ground state parameters of the two-dimensional spin-1/2 Heisenberg model on the square lattice*, J. Stat. Mech. 2026, 043101 (2026).

[108] C. A. Jiménez-Hoyos and G. E. Scuseria, *Cluster-based mean-field and perturbative description of strongly correlated fermion systems: Application to the one- and two-dimensional Hubbard model*, Phys. Rev. B **92**, 085101 (2015).

[109] V. Abraham and N. J. Mayhall, *Selected configuration interaction in a basis of cluster state tensor products*, J. Chem. Theory Comput. **16**, 6098 (2020).

[110] N. J. Mayhall, *Using higher-order singular value decomposition to define weakly coupled and strongly correlated clusters: The n-body tucker approximation*, J. Chem. Theory Comput. **13**, 4818 (2017).

[111] V. Abraham and N. J. Mayhall, *Cluster many-body expansion: A many-body expansion of the electron correlation energy about a cluster mean field reference*, J. Chem. Phys. **155**, 054101 (2021).

[112] A. Bachhar and N. J. Mayhall, *Restricted open-shell cluster mean-field theory for strongly correlated systems*, J. Phys. Chem. A **128**, 9015 (2024).

[113] While this improved reference state made our calculations more efficient, we were still able to obtain sufficiently accurate results when starting from the original Néel state.

[114] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, E. Gull,, *Solutions of the two-dimensional Hubbard model: Benchmarks and results from a wide range of numerical algorithms*, Phys. Rev. **5**, 041041 (2015).

[115] E. Pavarini, E. Koch, J. van den Brink, and G. Sawatzky (eds.), *Quantum Materials: Experiments and Theory, Modeling and Simulation* 6, 1 (Verlag des Forschungszentrum Jülich, Jülich, 2016).

[116] R. Levy, M. A. Morales, and S. Zhang, *Automatic order detection and restoration through systematically*

improvable variational wave functions, Phys. Rev. Res. **6**, 013237 (2024).

[117] H. Shi and S. Zhang, *Symmetry in auxiliary-field quantum Monte Carlo calculations*, Phys. Rev. B **88**, 125132 (2013).

[118] V. Abraham and N. J. Mayhall, *Selected configuration interaction in a basis of cluster state tensor products*, J. Chem. Theory Comput. **16**, 6098 (2020).

[119] N. M. Braunscheidel, A. Bachhar, and N. J. Mayhall, *Accurate and interpretable representation of correlated electronic structure via Tensor Product Selected CI*, Faraday Discuss. **254**, 130 (2024).

[120] N. M. Braunscheidel, V. Abraham, and N. J. Mayhall, *Generalization of the tensor product selected CI method for molecular excited states*, J. Phys. Chem. A **127**, 8179 (2023).

[121] A. Bachhar and N. J. Mayhall, *Computing exchange coupling constants in transition metal complexes with tensor product selected configuration Interaction*, J. Chem. Theory Comput. **22**, 1699 (2026).

[122] R. Haghshenas, E. Chertkov, M. Mills, W. Kadow, S. H. Lin, Y. H. Chen, C. Cade, I. Niesen, T. Begušić, M. S. Rudolph, and C. Cirstoiu, *Digital quantum magnetism on a trapped-ion quantum computer*, Nature (London) **653**, 56 (2026).

[123] Y. Teng, S. Y. Chang, M. S. Rudolph, and Z. Holmes, *Leveraging symmetry merging in Pauli propagation*, arXiv:2512.12094.

[124] https://github.com/nmayhall/DBF.jl.

[125] https://github.com/nmayhall/PauliOperators.jl.