

Evaluating a quantum-classical quantum Monte Carlo algorithm with Matchgate shadows

Benchen Huang^{1,2}, Yi-Ting Chen³, Brajesh Gupta¹, Martin Suchara^{4,*}, Anh Tran,¹
Sam McArdle,^{5,†} and Giulia Galli^{2,6,7,‡}¹AWS Worldwide Specialist Organization, Seattle, Washington 98170, USA²Department of Chemistry, University of Chicago, Chicago, Illinois 60637, USA³Amazon Braket, New York, New York 10018, USA⁴Microsoft Azure Quantum, Redmond, Washington 98052, USA⁵AWS Center for Quantum Computing, Pasadena, California, California 91125, USA⁶Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637, USA⁷Materials Science Division and Center for Molecular Engineering, Argonne National Laboratory, Lemont, Illinois 60439, USA

(Received 7 May 2024; accepted 3 September 2024; published 24 October 2024)

Solving the electronic structure problem of molecules and solids to high accuracy is a major challenge in quantum chemistry and condensed matter physics. The rapid emergence and development of quantum computers offer a promising route to systematically tackle this problem. Recent work by [Huggins *et al.*, *Nature (London)* **603**, 416 (2022)] proposed a hybrid quantum-classical quantum Monte Carlo (QC-QMC) algorithm using Clifford shadows to determine the ground state of a Fermionic Hamiltonian. This approach displayed inherent noise resilience and the potential for improved accuracy compared to its purely classical counterpart. Nevertheless, the use of Clifford shadows introduces an exponentially scaling postprocessing cost. In this work, we investigate an improved QC-QMC scheme utilizing the recently developed Matchgate shadows technique [Commun. Math. Phys. **404**, 629 (2023)], which removes the aforementioned exponential bottleneck. We observe from experiments on quantum hardware that the use of Matchgate shadows in QC-QMC is inherently noise robust. We show that this noise resilience has a more subtle origin than in the case of Clifford shadows. Nevertheless, we find that classical postprocessing, while asymptotically efficient, requires hours of runtime on thousands of classical CPUs for even the smallest chemical systems, presenting a major challenge to the scalability of the algorithm.

DOI: [10.1103/PhysRevResearch.6.043063](https://doi.org/10.1103/PhysRevResearch.6.043063)

I. INTRODUCTION

The ability to accurately solve the Schrödinger equation for interacting electrons will help tackle a multitude of problems in physics, chemistry and materials science, relevant to applications ranging from drug discovery [1] to the design of functional materials [2,3]. In the past century, multiple efforts have been devoted to solving the Schrödinger equation on classical computers, either by using suitable approximations and mean-field theories [4] or by employing nearly exact methods such as full configuration interaction (FCI) [5] and quantum Monte Carlo (QMC) [6,7]. However, known algorithms for FCI scale exponentially on classical computers, while scalable solutions using QMC often suffer from the

sign problem [8] and other numerical instabilities. In essence, highly entangled many-body wave functions of interacting electrons are hard to represent and optimize on classical computers.

The advent of quantum computers offers a promising route to tackle the interacting electron problem, either through quantum [9] or hybrid quantum-classical algorithms [10,11]. One popular algorithm to compute Fermionic ground states is the variational quantum eigensolver (VQE) and its variants [12–15], where a parameterized quantum circuit is used to generate the wave function, and its optimization is off-loaded to classical hardware. Nevertheless, current noisy intermediate-scale quantum (NISQ) hardware is limited by the depth of quantum circuits that it can implement. Moreover, VQE faces optimization challenges resulting from ansatz or noise-induced barren plateaus in energy landscapes [16], as well as large measurement overheads.

Recently Huggins *et al.* [17] proposed a hybrid quantum-classical algorithm for quantum Monte Carlo (QC-QMC) and applied it to study the dissociation of diatomic molecules on Google's Sycamore quantum processor. The Monte Carlo algorithm is driven by sampling from a trial state $|\Psi_T\rangle$ prepared on the quantum computer, with the aim of producing smaller bias and better accuracy than its purely classical counterpart. The algorithm uses classical shadows [18] with

*M.S. contributions to this work were made while employed by Amazon.

†Contact author: sammcard@amazon.com

‡Contact author: gagalli@uchicago.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.

random Clifford circuits to avoid iterative communication between classical and quantum hardware, which is desirable on near-term quantum devices as it minimizes latency from quantum-classical communication. Surprisingly, an inherent noise resilience was observed [17], spurring both academic and industrial interest in the technique [19–25]. Nevertheless, currently known techniques for classically postprocessing the Clifford shadows scale exponentially with system size, providing a barrier to achieving quantum advantage.

The inefficiencies of Clifford shadows for QC-QMC were recently addressed in Ref. [26] by replacing the Clifford circuits with Matchgate circuits [27]. To the best of our knowledge, the proposal has not yet been experimentally evaluated in the literature (although we note recent numerical investigations probing the statistical properties of Matchgate shadows applied to chemical systems [25,28]). In particular, it is unclear whether the Matchgate shadows approach exhibits the noise resilience that underpinned the success of the original (unscalable) Clifford shadows approach [17], as the theoretical justification for the noise resilience of Clifford shadows does not immediately extend to the Matchgate case. In addition, while the proposal is formally efficient, its practicality for studying realistic systems has yet to be established.

In this work, we numerically and experimentally study the QC-QMC algorithm, incorporating Matchgate shadows. In particular, we probe the noise resilience of the algorithm on quantum hardware. We observe and theoretically characterize the persistence of the noise resilience that was observed for Clifford shadows, uncovering a more subtle microscopic origin. We also developed improvements to existing robust classical shadow protocols [29–31] that can mitigate state preparation noise, and may be of independent interest. The Matchgate shadows approach reduces the classical postprocessing cost from exponential to formally polynomial. However, its use in QC-QMC has high-degree polynomial scaling in the system size, resulting from computing the local energy of each QMC walker. The algorithm also has large constant factors resulting from the number of samples used in the classical shadow postprocessing, as well as the number of QMC walkers and the number of QMC time steps. These final two factors are properties of the QMC algorithm itself, rather than aspects of the quantum algorithm. We observe hours of postprocessing runtime for even the smallest chemical systems, using thousands of CPUs. This fundamentally challenges the scalability of the approach. We thus validate the persistence of one of the major strengths of QC-QMC, as well as highlight the challenges that must be overcome if the technique is to provide practical quantum advantage.

The rest of the paper is organized as follows. In Sec. II, we discuss how QC-QMC is performed and the basics of classical shadows, which serve as prerequisites for subsequent discussions. In Sec. III, we first discuss prior work on noise robust classical shadows [29,32] (and their recently developed Matchgate variants [30,31]) and give an extension to these methods that mitigate state preparation noise in QC-QMC. Second, we experimentally and analytically confirm the natural noise resilience of the Matchgate shadow protocol as used in QC-QMC. Finally, we apply QC-QMC to simulate the dissociation of hydrogen, and the ground state energy of a solid-state spin defect system, on different quantum hardware

systems. These calculations reach agreement with the reference values even in the absence of error mitigation. In Sec. IV, we discuss the merits and shortcomings of QC-QMC and how it is related to other quantum-classical hybrid algorithms and purely classical QMC methods. Finally, Sec. V concludes our work with a summary.

II. PRELIMINARIES

The workflow of the QC-QMC method is shown in Fig. 1, where the computation on quantum and classical hardware is partitioned by classical shadows. In this section, we first provide a brief overview of the auxiliary-field quantum Monte Carlo (AFQMC) algorithm [33] used in the QC-QMC method. Then we introduce the classical shadows formalism and discuss its use in QC-QMC. Readers familiar with this background material can skip to the summary of our results in Sec. III.

A. Auxiliary-field quantum Monte Carlo

Within the electronic structure community, QMC refers to a family of methods that bypass the explicit optimization of the many-body wave function in an exponentially large Hilbert space. Specifically, probabilistic sampling is applied in a subspace of the full Hilbert space, resulting in a polynomial scaling of the memory required for the evaluation of the ground state energy of a given system. The auxiliary-field quantum Monte Carlo (AFQMC) method is one example of projector QMC methods where a stochastic imaginary time evolution of the wave function is carried out by propagating samples on a manifold of nonorthogonal Slater determinants. Each sample $\{|\phi_l\rangle\}$ is usually called a “walker.” The ground state energy is estimated as

$$E = \frac{\sum_l w_l e^{i\theta_l} E_l^{\text{loc}}}{\sum_l w_l e^{i\theta_l}}, \quad E_l^{\text{loc}} = \frac{\langle \Psi_T | H | \phi_l \rangle}{\langle \Psi_T | \phi_l \rangle}, \quad (1)$$

where E_l^{loc} is the local energy of each walker, w_l and θ_l are its weight and phase, H is the electronic structure Hamiltonian, and $|\Psi_T\rangle$ is a trial state defined as an approximation to the true ground state. The walker wave function and its weight and phase are updated every time step according to functions of $\frac{\langle \Psi_T | \phi_l \rangle}{\langle \Psi_T | \phi_l \rangle}$ that perform the imaginary time evolution, see Appendix A for details.

For a generic Hamiltonian, however, similar to other projector QMC methods, AFQMC suffers from the sign problem, or more precisely, the phase problem where the phases θ_l of the walkers are evenly distributed in $[0, 2\pi)$ [33,34]. It leads to the ground state energy estimator, i.e., Eq. (1) (computed from the ensemble average of all the walkers) experiencing an exponentially fast decay of the signal-to-noise ratio and a large statistical error. To remedy the phase problem, a common solution is the so-called phaseless approximation (ph) [33], where $\theta_l = 0$ is enforced by a modified update rule, thus constraining the evolution of the walkers, see Appendix A for details. Such an approximation, while controlling the phase problem, introduces a bias to the estimated ground state energy. The magnitude of the bias is largely determined by how well the trial state $|\Psi_T\rangle$ represents the exact ground state. On classical computers, $|\Psi_T\rangle$ is usually chosen as either a single

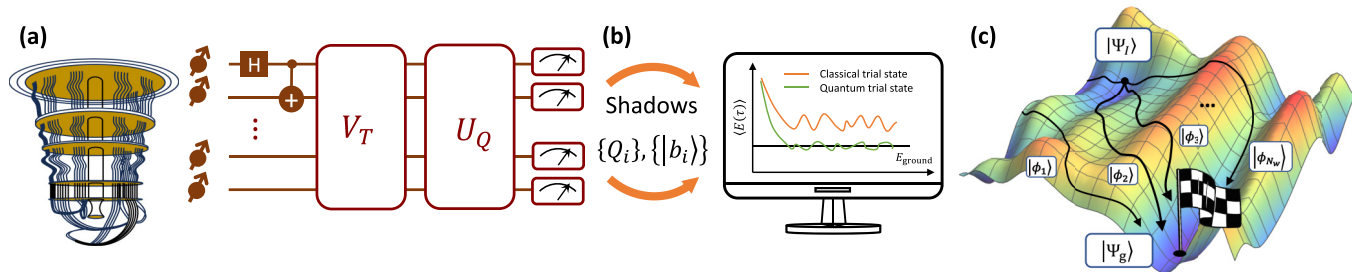


FIG. 1. Workflow for the hybrid quantum-classical quantum Monte Carlo (QC-QMC) algorithm. (a) An equal superposition of the all-zero state $|0\rangle$, and the quantum trial state $|\Psi_T\rangle$ is prepared on a quantum computer, followed by the twirling of random unitary circuits U_Q and measurements. (b) The measurement outcomes $\{|b_i\rangle\}$ and the random circuit unitary U_Q —which together constitute the information required to reconstruct classical shadows—are communicated to the classical computers. (c) The QMC procedure is carried out on a classical computer, where the walker states $|\phi_i\rangle$ evolve on the potential energy surface from the initial state $|\Psi_T\rangle$ towards the target ground state $|\Psi_g\rangle$.

Slater determinant, e.g., the Hartree-Fock state, or a linear combination of Slater determinants, to ensure a polynomial scaling in computational time [35].

The insight of Ref. [17] is that on a quantum computer, one may efficiently compute overlaps between Slater determinant walker states and a much wider range of trial states, e.g., the unitary coupled cluster (UCC) state [10,36]. We refer to implementing AFQMC in this way as QC-AFQMC. Such states are potentially “closer” to the target ground state than those adopted as classical trials, and could lead to more accurate ground state energy estimations. It is currently unclear what are the best AFQMC trial states that can be prepared on a quantum computer, in low circuit depth. This is a similar issue to that of choosing a well-motivated and implementable ansatz circuit in VQE. We will discuss this problem in Sec. IV, and assume for now that a suitable AFQMC trial state can be efficiently prepared on a quantum computer. Then, the central issue for the QC-AFQMC scheme lies in how to evaluate the overlap amplitude $\langle\Psi_T|\phi\rangle$ for each walker, at each time step. Efforts in the literature have branched out in two directions: Xu and Li [19] designed circuits based on the Hadamard test [37] to efficiently compute the overlap between the walker states and the trial state. On the other hand, Huggins *et al.* [17] computed the amplitude by using the classical shadows technique applied to $|\Psi_T\rangle$. The former approach suffers a significant overhead in quantum computation due to the large number of walkers typically used in AFQMC. In addition, it requires iterative communication between the classical and quantum hardware, since $|\Psi_T\rangle$ is queried at every time step during the time evolution. We therefore followed the approach of Huggins *et al.* in this work, and we provide a detailed scaling comparison of the two approaches in Sec. IV.

Finally, as originally pointed out by Ref. [17] and then re-emphasized by Refs. [38,39], the overlap amplitudes $\langle\Psi_T|\phi\rangle$ decay exponentially with system size in typical systems. This decay suggests that to reduce the relative errors of overlap amplitudes within acceptable thresholds, an exponentially large number of measurement shots would be necessary, resulting in the QC-QMC approach becoming non-scalable. This is supported by benchmarks on a transverse-field Ising model with varying sizes and measurement shots (on a noiseless simulator) using a Green’s function Monte Carlo method [38,39] (as the QMC component). This issue is expected to plague both classical and quantum-assisted quantum Monte Carlo algorithms, and can be seen as a manifestation of the

general Quantum Merlin-Arthur hardness of electronic structure calculations [17]. Nevertheless, Ref. [39] explored strategies for postponing or alleviating this exponential bottleneck, including using more sophisticated walker wave functions (e.g. beyond a single Slater determinant) in QMC. It is an open question as to whether these strategies can be made compatible with classical shadows QC-QMC (note that these overlaps could alternatively be measured using the Hadamard test). If the vanishing overlaps issue can be mitigated in sufficiently large system sizes (e.g., ~ 100 orbitals), then it may be possible to achieve practical quantum advantage over classical QMC methods (and potentially other classical quantum chemistry algorithms) for systems of interest [17]. In this study, we only consider small molecular systems which are not at the scale where the exponentially vanishing overlaps become a concern, opting to focus solely on the practical implementation challenges of QC-QMC. We leave investigating methods for testing and mitigating vanishing overlaps for future work.

B. Classical shadows

As mentioned in the previous subsection, the need to evaluate $\langle\Psi_T|\phi\rangle$ for all walkers at each time step is a key consideration for the practicality of the algorithm. In Ref. [17], it was shown that the measurements required in QC-AFQMC can be recast into the following framework. Let ρ denote the density matrix of an n -qubit quantum state that we know how to prepare, and $\{O_i\}$ denote a collection of M observables whose expectation values, i.e., $\text{tr}(O_i\rho)$ we wish to estimate. Classical shadow tomography [18] provides a way to estimate these quantities with a cost that only scales logarithmically with M . Specifically, it estimates each $\text{tr}(O_i\rho)$ up to some error ϵ by the following procedure: i) choose a distribution \mathcal{D} of unitary transformations; ii) sample random unitaries $U \in \mathcal{D}$ from the chosen distribution and iii) measure the state $U\rho U^\dagger$ in the computational basis $\{|b\rangle\}$ (we refer to the tensor products of $\{|0\rangle, |1\rangle\}$ of each qubit as the *computational basis*, i.e., $\{|b\rangle\}_{b \in \{0,1\}^n}$) to obtain the measurement outcome $|b\rangle\langle b|$. Consider the state $U^\dagger|b\rangle\langle b|U$; in expectation, the mapping from ρ to $U^\dagger|b\rangle\langle b|U$ defines a quantum channel,

$$\begin{aligned} \mathcal{M}(\rho) &:= \mathbb{E}_{U \sim \mathcal{D}} [U^\dagger |\hat{b}\rangle\langle \hat{b}| U] \\ &= \mathbb{E}_{U \sim \mathcal{D}} \sum_{b \in \{0,1\}^n} \langle b|U\rho U^\dagger|b\rangle U^\dagger|b\rangle\langle b|U, \end{aligned} \quad (2)$$

where \mathbb{E} denotes the operation of averaging, and the hat represents a statistical estimator.

In the classical shadows framework, we require \mathcal{M} to be invertible, which is true if and only if the collection of measurement operators defined by drawing $U \in \mathcal{D}$ and measuring in the computational basis is tomographically complete. Assuming that these conditions are satisfied, we can apply \mathcal{M}^{-1} to both sides of above equation, yielding

$$\begin{aligned} \rho &= \mathbb{E}_{U \sim \mathcal{D}} [\hat{\rho}] = \mathcal{M}^{-1} \left(\mathbb{E}_{U \sim \mathcal{D}} [U^\dagger |\hat{b}\rangle \langle \hat{b}| U] \right) \\ &= \mathbb{E}_{U \sim \mathcal{D}} [\mathcal{M}^{-1}(U^\dagger |\hat{b}\rangle \langle \hat{b}| U)]. \end{aligned} \quad (3)$$

We call the collection $\{\mathcal{M}^{-1}(U^\dagger |\hat{b}\rangle \langle \hat{b}| U)\}$ the classical shadows of ρ . These shadows can be used to estimate the expectation values $\text{tr}(O_i \rho)$,

$$\langle O_i \rangle = \mathbb{E}_{U \sim \mathcal{D}} \text{tr}[O_i \mathcal{M}^{-1}(U^\dagger |\hat{b}\rangle \langle \hat{b}| U)], \quad (4)$$

each within error ϵ with probability at least $1 - \delta$, with a number of samples that scales as:

$$N_{\text{sample}} = \mathcal{O} \left(\frac{\log(M/\delta)}{\epsilon^2} \max_{1 \leq i \leq M} \text{Var}[\hat{o}_i] \right). \quad (5)$$

In the above equation, we define \hat{o}_i as an estimator of $\text{tr}(O_i \rho)$ and $\text{Var}[\hat{o}_i]$ represents the variance of estimator \hat{o}_i , which could be bounded by the shadow norm of O_i [18]. Importantly for QC-AFQMC, N_{sample} only scales logarithmically with the number of target observables M .¹

Formally, the condition that the measurement channel is invertible is sufficient for performing the classical shadows protocol. In practice, it is desirable that the protocol is efficient both in terms of quantum and classical resources. This means that there should be an efficient procedure to sample unitaries U from \mathcal{D} and implement them on a quantum computer, and in addition the variance of the estimates \hat{o}_i is polynomial in system size. Moreover, it should be efficient to compute the expectation values with respect to the shadows, $\text{tr}[O_i \mathcal{M}^{-1}(U^\dagger |\hat{b}\rangle \langle \hat{b}| U)]$, on a classical computer.

In QC-AFQMC, the overlap amplitude $\langle \Psi_T | \phi \rangle$ is not a physical observable. Nevertheless, it can be measured within the framework of classical shadows by rewriting it as [17]

$$\begin{aligned} \langle \Psi_T | \phi \rangle &= 2 \text{tr}(|\phi\rangle \langle \mathbf{0}| \rho) \\ &= 2 \mathbb{E}_{U \sim \mathcal{D}} \text{tr}[|\phi\rangle \langle \mathbf{0}| \mathcal{M}^{-1}(U^\dagger |\hat{b}\rangle \langle \hat{b}| U)], \end{aligned} \quad (6)$$

where ρ is the density matrix corresponding to the state $\frac{1}{\sqrt{2}}(|\mathbf{0}\rangle + |\Psi_T\rangle)$ and $|\mathbf{0}\rangle = |\mathbf{0}\rangle^{\otimes n}$. Here $|\phi\rangle \langle \mathbf{0}|$ plays the role of the operator O .

Reference [17] used random Clifford circuits to perform classical shadows, where the overlap amplitude is

postprocessed as

$$\langle \Psi_T | \phi \rangle = 2f^{-1} \mathbb{E}_{U \sim \text{Cl}(2^n)} [|\mathbf{0}\rangle \langle U^\dagger | \hat{b}\rangle \langle \hat{b}| U | \phi\rangle]. \quad (7)$$

In the above equation, $f = (2^n + 1)^{-1}$ is the only (nontrivial) eigenvalue of \mathcal{M} , since the Clifford group has only one nontrivial irreducible representation (irrep). The first term $\langle \mathbf{0} | U^\dagger | b \rangle$ can be efficiently computed using the Gottesman-Knill theorem [41], but the second term $\langle b | U | \phi \rangle$, in general, cannot be efficiently estimated due to $|\phi\rangle$ being a random Slater determinant. As noted in Ref. [17], the method of Ref. [42] could be used to efficiently estimate this quantity up to an additive error. Nevertheless, the f^{-1} prefactor exponentially amplifies this error, eliminating the efficiency of the proposed scheme.

Another interesting observation of Ref. [17] is that the evaluation of overlap ratios, which drives the AFQMC algorithm, is resilient to quantum hardware noise. This phenomenon can be understood by first noting that noise will change the prefactor f to a modified prefactor \tilde{f} (as explained in detail in Sec. III A). This prefactor cancels out when taking the ratio, recovering the noiseless result without requiring any additional error mitigation. We refer to this as having an inherent noise resilience.

References [26,27,43,44] investigated the replacement of the random Clifford ensemble with a random Matchgate ensemble (in some cases restricting to Clifford Matchgates and/or number conserving Matchgates). The Matchgate circuit is the qubit representation of the Fermionic Gaussian transformation assuming the Jordan-Wigner (JW) transformation [45] is used. We therefore adopt the JW mapping throughout this work. Unlike the Clifford group, the Matchgate group has $(n + 1)$ (even) irreps, see Appendix D for details. We focus on the results of Ref. [26], which

(i) proved the equivalence between the continuous and discrete Matchgate ensembles for classical shadows;

(ii) explicitly showed how to efficiently compute the overlap between trial states and Slater determinant states using Matchgate shadows, including bounding the worst case variance in the estimate.

We refer the interested reader to Ref. [[26], Sec. III D] for a more detailed comparison between these recent works. The results of Ref. [26] overcame the postprocessing limitations of Clifford shadows, showing that overlap amplitude can be computed as

$$\langle \Psi_T | \phi \rangle = 2 \sum_{l=0}^n f_{2l}^{-1} \mathbb{E}_{Q \sim B(2n)} \text{tr}[|\phi\rangle \langle \mathbf{0}| \Pi_{2l}(U_Q^\dagger |\hat{b}\rangle \langle \hat{b}| U_Q)], \quad (8)$$

where the random Matchgate circuits U_Q are prepared by sampling random signed permutation matrices Q [$Q \in B(2n)$, where B represents the Borel group]. In Eq. (8), Π_{2l} is the projector associated with the l th even irrep, and its eigenvalue f_{2l} is

$$f_{2l} = \binom{2n}{2l}^{-1} \binom{n}{l}. \quad (9)$$

Importantly, Eq. (8) can be efficiently solved using the matrix Pfaffian, with a scaling of $\mathcal{O}((n - \zeta/2)^4)$ by polynomial interpolation [26], where ζ denotes the number of electrons.

¹We note that the scaling in Eq. (5) was rigorously proven in Ref. [18] using a median-of-means estimator. Refs. [27,40] discussed when a mean estimator would suffice for Clifford and Matchgate shadows, respectively. We observed numerically that a mean estimator suffices for shadows in our small system QC-QMC experiments.

As such, the exponential postprocessing bottleneck present in the Clifford shadows approach is eliminated through the use of Matchgate shadows.

Inherent noise resilience, analogous to that observed for overlap ratios computed via Clifford shadows, has not yet been established for overlap ratios computed via Matchgate shadows. One sufficient but not necessary condition, (following from the Clifford case), would be $\tilde{f}_{2l} = \alpha f_{2l}$, $l \in \{0, 1, \dots, n\}$, where α is a proportionality constant. In Sec. III B 2, we investigate the noise resilience of overlap ratios computed via Matchgate shadows. While our experimental results do not display a universal proportionality constant between each \tilde{f}_{2l} and f_{2l} , we nevertheless find that the computed ratios are inherently robust to noise, and provide justification for this observation.

III. RESULTS

In this section, we present our results, divided into three parts. We first introduce prior work on noise robust classical shadows, and provide an improvement that mitigates state preparation noise in QC-AFQMC. Second, we experimentally test for noise resilience of both overlap amplitudes and ratios of overlap amplitudes, evaluated via the Matchgate shadow technique. Notably, we observe that the ratios are resilient to the effects of noise (as was previously observed for Clifford shadows [17]). We provide a theoretical explanation for the observed noise resilience and its limitations, which has more subtle origins than in the Clifford case. Third, we apply the QC-AFQMC algorithm to compute (i) the dissociation curve of the hydrogen molecule on a superconducting qubit quantum computer; and (ii) the ground state of a negatively charged nitrogen-vacancy (NV) center in diamond on a trapped ion quantum computer.

We use the following notation in this section: we denote overlap amplitudes obtained from noiseless shadows with an ordinary bracket, $\langle \Psi_T | \phi \rangle$, quantities subject to quantum noise with a tilde bracket, $\langle \tilde{\Psi}_T | \phi \rangle$, and we use a tilde plus subscript r to denote quantities obtained from a robust shadow protocol, $\langle \tilde{\Psi}_T | \phi \rangle_r$.

A. Robust Matchgate shadow protocol

References [29,32] developed a “robust shadow protocol”, showing that when the quantum noise is gate-independent, time-stationary, and Markovian (GTM), and the state preparation of ρ is perfect, a noise-free expectation value can be obtained. The protocol works by replacing the eigenvalue(s) f of \mathcal{M} with modified value(s) \tilde{f} that compensates for the effect of noise. The robust scheme uses additional shadow-like circuits to estimate the value(s) of \tilde{f} . It can be viewed as passive error mitigation, as the updates are made purely in the classical postprocessing of the shadow expectation value. The framework was originally applied to global and local Clifford shadows. Recently, two separate works extended the robust shadow protocol to the Matchgate setting [30,31]. The latter of these works observed a close connection between the circuit used for the robust Matchgate shadow protocol and prior work on Matchgate randomized benchmarking [46]. We

will closely follow this approach, presenting derivations for completeness in Appendix E.

The robust Matchgate shadows protocol [31] estimates $\{\tilde{f}_{2l}\}$ using the following steps: (i) prepare the all-zero state $|\mathbf{0}\rangle$ on a quantum computer; (ii) sample $Q \in B(2n)$ and apply the Matchgate circuit U_Q to $|\mathbf{0}\rangle$; and (iii) measure in the computational basis and collect measurement outcomes $|b\rangle$. The expressions used to calculate $\{\tilde{f}_{2l}\}$ from these measurement outcomes are given in Appendix E. The sample complexity is asymptotically equivalent to that used in the noiseless case [31].

The established robust shadow scheme focuses on noise in the shadow unitary [U_Q in Fig. 1(a)] and does not mitigate any noise occurring during state preparation [V_T in Fig. 1(a)]. We extend the method to (partly) account for noise that occurs during the state preparation unitary. In QC-AFQMC, we evaluate the overlap amplitudes by measuring classical shadows of an equal superposition state $\frac{1}{\sqrt{2}}(|\Psi_T\rangle + |\mathbf{0}\rangle)$. This state is prepared via a circuit that first generates a superposition of the Hartree-Fock state with $|\mathbf{0}\rangle$ state, and then applies a unitary V_T , such that $V_T|\Psi_{\text{HF}}\rangle = |\Psi_T\rangle$ and $V_T|\mathbf{0}\rangle = |\mathbf{0}\rangle$ [47]. As discussed above, to estimate the values of $\{\tilde{f}_{2l}\}$ required in the robust Matchgate protocol, we apply a Matchgate shadow circuit to the initial state $|\mathbf{0}\rangle$. We can partially account for the noise in the state preparation circuit in the following way. We minimally alter the QC-AFQMC circuit in Fig. 1(a), such that it measures $\{\tilde{f}_{2l}\}$, by removing the initial Hadamard gate. In a noiseless setting, this modified version of V_T would result in $\rho = |\mathbf{0}\rangle\langle\mathbf{0}|$, as required. In the noisy setting, we can treat the noise as originating in the Matchgate circuit, which is then mitigated by the robust protocol. While we do not yet have a full mathematical characterization of the noise models that can be mitigated by our enhanced robust Matchgate shadows technique, we demonstrate the efficacy of this approach through numerical simulations and experiments on quantum hardware, presented in the following section.

B. Noise resilience of overlap amplitudes and their ratios

As discussed previously, all of the key quantities used in the AFQMC algorithm, e.g., the local energy, can be expressed as linear combinations of the ratio of overlaps $\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi \rangle}$. Hence we can view the propagation in imaginary time as being driven by these overlap ratios. We can estimate the individual overlaps from Matchgate shadows. In this section, we experimentally probe the impact of noise on both the overlap amplitudes, and their ratios, by measuring these quantities via Matchgate shadows implemented on the IBM Hanoi superconducting qubit quantum computer. As a test system, we consider the hydrogen molecule in its minimal STO-3G basis, mapped to four qubits via the Jordan-Wigner transform. We use a trial state of the form $|\Psi_T\rangle = \alpha|1100\rangle + \beta|0011\rangle$, which is a linear combination of the Hartree-Fock configuration and double excited state, respectively.

1. Evaluating the overlap amplitudes

We compute the overlaps $\langle \Psi_T | \phi_i \rangle$ between the trial state and 16 randomly sampled Slater determinants, using Eq. (8), as a function of the number of Matchgate shadows used. We

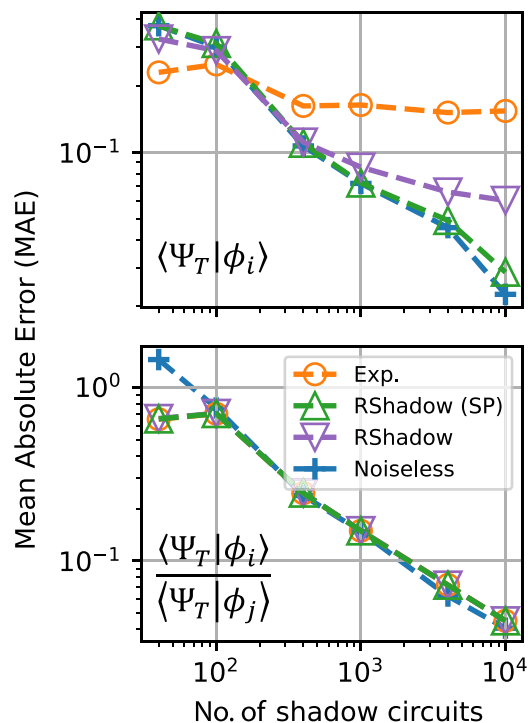


FIG. 2. The mean absolute error (MAE) of overlap amplitudes $\langle \Psi_T | \phi_i \rangle$ and overlap ratios $\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi_j \rangle}$ with respect to the number of Matchgate circuits, varying from 40 to 10 000. A total of 16 000 Matchgate shadow circuits are used in this experiment, each using 1024 measurement shots [48]. The raw experimental results (“Exp.”) and noiseless simulation (“Noiseless”) are plotted in orange circles and blue crosses, respectively. For the two robust protocols (“RShadow” and “RShadow (SP)”), we allocated another 16,000 Matchgate circuits for each to determine the coefficients \hat{f}_{2l} , thus doubling the measurement cost.

consider both noisy Matchgate shadows, and those corrected using the robust Matchgate shadow protocol introduced in Sec. III A.

In Fig. 2 (upper panel), we show the mean absolute error (MAE) of the noisy and noise-robust overlaps, with respect to the ideal noiseless value. We observe that the noisy results differ significantly from the true values, and their accuracy is limited by the effects of noise. The two robust shadow approaches (accounting for state preparation error (SP) and not accounting for it) mitigate the impact of hardware noise. In particular, compensating for the effects of state preparation noise results in a much smaller deviation from the ideal noiseless values than the standard robust approach. This observation signifies the importance of considering the state preparation noise when using robust shadows on quantum hardware.

2. Noise resilience in the evaluation of overlap ratios

Using the results presented in the previous section for $\{\langle \Psi_T | \phi_i \rangle\}$, we can compute the 120 possible ratios of overlaps originating from the 16 Slater determinants sampled. In Fig. 2 (lower panel), we show the MAE of these overlap ratios with respect to the ideal values (and verify in Fig. 9 in Appendix H that this behavior holds for not only the MAE, but the errors

associated with each overlap ratio). We observe the following behavior:

$$\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi_j \rangle} \approx \frac{\widetilde{\langle \Psi_T | \phi_i \rangle}_r}{\widetilde{\langle \Psi_T | \phi_j \rangle}_r} \approx \frac{\widetilde{\langle \Psi_T | \phi_i \rangle}}{\widetilde{\langle \Psi_T | \phi_j \rangle}}, \quad (10)$$

when the number of shadow circuits exceeds 100. This is an indication of noise resilience in evaluating the overlap ratios, suggesting that the robust shadow protocol may not be required for QC-AFQMC. We provide theoretical justification for these observations below.

The first (approximate) equality in Eq. (10) will not hold for all noise models, as the robust Matchgate shadow protocol is only guaranteed to correct the effects of noise if its assumptions (GTM noise, noiseless state preparation) are satisfied. Nevertheless, for cases where the robust Matchgate protocol is able to correct the effects of noise the equality of the ratios follows directly. In our experiments, we found that the robust Matchgate protocol was able to correct the individual overlaps up to a small residual error attributed to residual state preparation noise. This leads to the approximate equality between the ratios.

The second (approximate) equality shows that the ratio of any two uncorrected overlap amplitudes (which are themselves inaccurate, see Fig. 2 (upper panel)) yields the same results as ratios of overlaps computed via the robust Matchgate protocol. To understand this result, we first restate Eq. (8) for computing the overlap in the absence of noise

$$\langle \Psi_T | \phi \rangle = 2 \sum_{l=0}^n f_{2l}^{-1} \mathbb{E}_{Q \sim B(2n)} \text{tr}[\langle \phi | \mathbf{0} | \Pi_{2l}(U_Q^\dagger | \hat{b} \rangle \langle \hat{b} | U_Q)].$$

We make use of the following theorem.

Theorem 1. For an n qubit state $|\phi\rangle = \sum_i c_i |i\rangle$, which is a linear combination of computational basis states $|i\rangle$ with fixed Hamming weight ζ ,

$$\mathcal{P}_{0,\zeta}[\Pi_{2l}(|\phi\rangle\langle \mathbf{0}|)] = b_{2l} |\phi\rangle\langle \mathbf{0}|, \quad (11)$$

with

$$b_{2l} = \begin{cases} 2^{\zeta-n} \binom{n-\zeta}{l-\frac{\zeta}{2}}, & \text{if } \frac{\zeta}{2} \leq l \leq n - \frac{\zeta}{2} \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

such that $\sum_{l=0}^n b_{2l} = 1$. Here $\mathcal{P}_{0,\zeta}$ denotes the projector onto the subspace spanned by the Hamming weight zero state $|\mathbf{0}\rangle$ and computational basis states of Hamming weight ζ .

Proof. A full proof is given in Appendix D4, and we sketch a brief outline here. The operation $\Pi_{2l}(|\phi\rangle\langle \mathbf{0}|)$ can be re-expressed as $2^{-n} \sum_{j_l} \langle \mathbf{0} | \gamma_{j_l}^\dagger | \phi \rangle \gamma_{j_l}$ where γ_{j_l} are strings of $2l$ Majorana operators (defined in Appendix D). The action of the Majorana strings is to create “1’s” in the $|\mathbf{0}\rangle$ state. Only terms with ζ 1’s created contribute to the sum. Projecting the resulting sum into the restricted subspace ensures the “eigenoperator” property. The prefactor is given by counting arguments for the number of Majorana strings that can yield ζ 1’s for a given l value. ■

The justification for restricting to the $\{0, \zeta\}$ -electron subspace is that ρ is an equal superposition of $|\mathbf{0}\rangle$ and $|\Psi_T\rangle$, and therefore the trace of ρ with operators outside this subspace will necessarily be zero. This will still hold approximately true when ρ is reconstructed from a finite number of classical

shadows. This operator only appears in the classical postprocessing step and is therefore not affected by noise. Therefore the noisy overlap amplitude can be expanded as

$$\langle \widetilde{\Psi}_T | \phi \rangle \approx 2 \left(\sum_{l=0}^n f_{2l}^{-1} b_{2l} \right) \mathbb{E}_{Q \sim B(2n)} [\langle \mathbf{0} | U_Q^\dagger | \hat{b} \rangle \langle \hat{b} | U_Q | \phi \rangle], \quad (13)$$

where in deriving this equation, we have used $\text{tr}(AM^{-1}B) = \text{tr}(\mathcal{M}^{-1}(A)B)$, where operators A, B are in the subspace where \mathcal{M} is invertible [26]. This expression can be compared to Eq. (7) for Clifford shadows, noting the common feature of the separation of the expression into a sum over the channel index l , and the expectation over the shadow unitaries acting on the walker state $|\phi\rangle$. Equation (13) can also be understood as the way that the Matchgate channel is twirling the noise in each subspace, in a similar manner to that of the Clifford shadows. In order to mitigate the effects of noise with the robust Matchgate shadows approach, we could replace f_{2l}^{-1} with \widetilde{f}_{2l}^{-1}

$$\langle \widetilde{\Psi}_T | \phi \rangle_r \approx 2 \left(\sum_{l=0}^n \widetilde{f}_{2l}^{-1} b_{2l} \right) \mathbb{E}_{Q \sim B(2n)} [\langle \mathbf{0} | U_Q^\dagger | \hat{b} \rangle \langle \hat{b} | U_Q | \phi \rangle]. \quad (14)$$

If the noise satisfies GTM assumptions (and state preparation is noise-free), Eq. 14 is able to correct the effect of noise that manifests from the deviation of $|\hat{b}\rangle\langle\hat{b}|$ from $|\hat{b}\rangle\langle\hat{b}|$. If these assumptions on the noise are not satisfied, robust Matchgate shadows may not be able to perfectly recover the noiseless value. We can now observe that regardless of whether the robust procedure is used, the terms corresponding to the sum over l in Eqs. (13) and (14) are both independent of $|\phi\rangle$.

As such, when computing the ratios $\frac{\langle \widetilde{\Psi}_T | \phi_i \rangle_r}{\langle \widetilde{\Psi}_T | \phi_j \rangle_r}$ and $\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi_j \rangle}$, the prefactor will cancel between the numerator and denominator, leaving behind a ratio of estimators $\mathbb{E}_{Q \sim B(2n)} [\dots]$ that is identical for the robust case and the nonrobust case. As such, if the ratio of overlaps evaluated via robust Matchgate shadows is able to approximately recover the noiseless results (as observed in our data), we conclude that the ratio of overlaps evaluated via regular Matchgate shadows is also able to approximately recover the noiseless value. This explains the data shown in Fig. 2 (lower panel).

We emphasize that the cancellation of the prefactor found here, i.e., $\sum_{l=0}^n \widetilde{f}_{2l}^{-1} b_{2l}$, should not be mistaken with that of the Clifford case, although they have a very similar form. This prefactor of Matchgate shadows, as a sum, contains the noise information of all $(n+1)$ even subspaces of the Matchgate channel. From the Clifford case, one might naturally think that the cancellation for the Matchgate shadow is due to $f_{2l} = \alpha f_{2l}$ where α is independent of l . We verified that this is not the case by experimentally determining f_{2l} . We observed that α_l could differ by up to 20% for different l . These data are recorded in Appendix H. Therefore this cancellation should be credited to the structure of $|\phi\rangle\langle\mathbf{0}|$, namely, being an eigenoperator of Π_{2l} in the restricted subspace, as proven in Theorem 1. If another observable, without this property, were chosen, any observed noise resilience would have an alternative origin. For example, Ref. [28] shows a similar contractive property in evaluating the ratios, but only for noise models that ensure

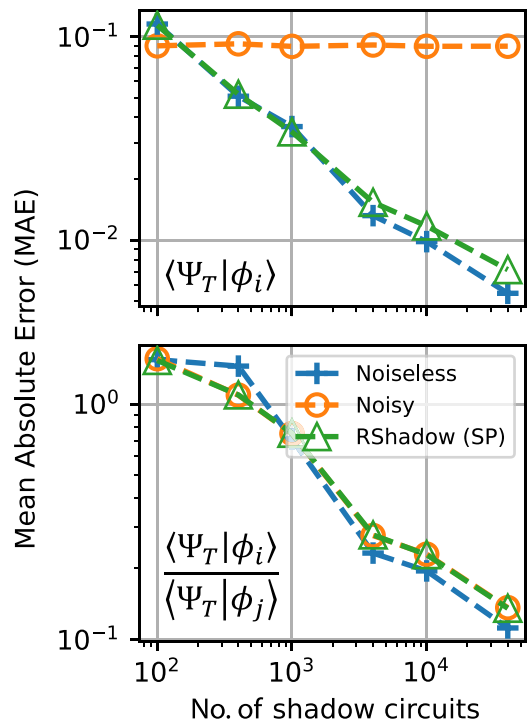


FIG. 3. The MAE of estimating overlap amplitudes and overlap ratios with respect to the number of Matchgate shadow circuits for the water molecule (8 qubits). These results were obtained through numerical simulations under single-qubit asymmetric Pauli noise, where the Pauli X, Z , and Y error rates are set to 1%, 3%, and 2%, respectively.

$\langle \mathbf{0} | \rho_{\text{noise}} | \phi \rangle$ is zero, where ρ_{noise} is the noise corrupted part of the density matrix.

As discussed above, robust shadows have been shown to correct for noise in the shadow circuit, assuming noiseless state preparation [29], or state preparation affected by global depolarizing noise [17,28]. In a realistic setting, the state preparation step may suffer from more complicated noise which may prevent the technique from recovering the noiseless result. We performed numerical simulations to compute overlap ratios for the water molecule in a $(4e, 4o)$ active space (8 qubits). We employed asymmetric Pauli errors as the noise model in this simulation (including in the state preparation circuit), and we refer to Appendix H for more details. The results, shown in Fig. 3 corroborate the analysis provided above. Specifically, we observe from the upper panel that when computing $\langle \Psi_T | \phi_i \rangle$ the results differ significantly from the noiseless value. In contrast, when computing $\langle \Psi_T | \phi_i \rangle_r$ using the robust Matchgate protocol (with compensation for state preparation error) we are able to almost recover the noise-free value. We attribute the deviation from the noiseless value to the presence of some residual state preparation error, which violates the assumptions of the robust shadow protocol. In contrast, from the lower panel we observe that the ratios of overlaps evaluated via the robust Matchgate shadow protocol are no more accurate than those evaluated using the regular Matchgate shadow protocol. Moreover, both of these values deviate from the noiseless value by a small error due to residual (uncorrectable) state preparation noise. In the

case of coherent errors during state preparation, we effectively prepare $|\tilde{\Psi}_T\rangle = U|\Psi_T\rangle$ for some unitary U . As such, while we would still observe that $\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi_j \rangle} \approx \frac{\langle \Psi_T | \phi_i \rangle_r}{\langle \Psi_T | \phi_j \rangle_r}$, neither ratio would likely recover the noiseless value.

Overall, the results of this subsection highlight a limitation of the robust Matchgate protocol, and its use in QC-AFQMC. Nevertheless, these results should be seen as an attractive feature of QC-AFQMC, as they suggest that the algorithm, driven by overlap ratios, is naturally resilient to the impact of many types of noise.

C. Computing ground states with QC-AFQMC

Having established the inherent noise resilience of the overlap ratios that drive the AFQMC algorithm, we demonstrate that this property contributes to noise resilience in the overall algorithm. We first compute the dissociation curve of the four-qubit hydrogen molecule studied in the previous subsections using the QC-AFQMC algorithm implemented on the IBM Hanoi superconducting qubit quantum processor. We then apply the QC-AFQMC algorithm to compute the ground state of an NV center in diamond, using the Aria trapped ion quantum processor from IonQ. In both settings, we emulate practical applications of QC-AFQMC by assuming a trial state generated by an imperfect VQE calculation. In both cases, we use a tailored UCCSD ansatz, see Appendix H, which was optimized until the energy was approximately 30–80 mHa higher than the FCI reference. We emphasize that these quantum hardware calculations do not seek to demonstrate any advantage over the classical QMC algorithm, as using an unrestricted Hartree-Fock or a multi-Slater trial state in classical AFQMC would also guarantee accurate ground state energies for these small system sizes.

1. Hydrogen molecule dissociation

In Fig. 4 (upper panel), we show the ground state energies obtained from QC-AFQMC at five different molecular geometries. If we target an error $\epsilon = 10^{-2}$ on any overlap, with $\geq 99.99\%$ confidence, the rigorous sample complexity bounds presented in Eq. (5) imply the need to use $\sim 1.1 \times 10^5$ shadow circuits (with one measurement shot per circuit). As shown in Fig. 2, it was sufficient to use only 16 000 circuits for each bond length, suggesting the bounds of Eq. (5) are overly pessimistic for the small system studied here. Given the noise resilience observed in evaluating the overlap ratios, the use of robust shadows was not necessary. In our AFQMC calculations, we used 4800 walkers for all numerical data presented here to ensure a negligible variance, and a time step $\Delta\tau = 0.005$ H.a.⁻¹ resulting in a negligible Trotter error. We chose the initial walker state to be the Hartree-Fock state, for each of the five distances studied here. The calculation was parallelized across 4800 CPU cores (1 core per AFQMC walker), as discussed in more detail below.

In the interest of conserving computational resources, only two of the five data points (solid points) were obtained using the scalable Matchgate shadows approach outlined above. The remaining three data points were obtained using an exponentially scaling approach used in Ref. [17] and detailed in Appendix H 1 d, that is ultimately more efficient than the




FIG. 4. Estimations of the ground state energy of the hydrogen molecule at five different bond distances (upper). The QC-AFQMC calculation run on the IBM Hanoi quantum computer (solid and hollow orange circles labeled “exp.”) converges within computational accuracy for all distances (see inset), with a difference of $\sim 10^{-4}$ H.a. from its simulated noiseless counterpart (blue crosses). As discussed in the main text, filled circles denote results obtained with the scalable Matchgate shadows approach, while the empty circles denote results obtained with the non-scalable approach outlined in Appendix H 1 d. The quantum trial state is obtained from a noisy VQE simulation (purple diamond). An imaginary time evolution at 0.75 Å is plotted in the lower panel, where the auxiliary fields sampled at each time step are synchronized between the raw noisy experiment and noiseless reference.

scalable Matchgate approach for small system sizes. We verified that the two schemes give the same results for overlap amplitudes. Thus the two Matchgate-obtained points are used as a complete evaluation of the practicality of the algorithm, while the remaining three points are only presented to confirm the accuracy of the algorithm in the presence of hardware noise.

We observe that all five QC-AFQMC calculations converge to within computational accuracy of the FCI reference value, even without the use of the robust shadow protocol. In Fig. 4 (lower panel), we show the variation in energy as a function of imaginary time for the 0.75 Å geometry. There is little observable difference (~ 0.1 mHa) between the noisy experiment and its noiseless classical emulation.

We further note that even when the state preparation suffers from coherent errors, the QC-AFQMC algorithm is still able to recover accurate ground state energies, providing the effective quantum trial state $|\tilde{\Psi}_T\rangle$, reconstructed from the noisy shadows, does not appreciably deviate from $|\Psi_T\rangle$. This is a result of the dissipative nature of the imaginary time process emulated by AFQMC, which will drive any initial state towards the ground state. This further enhances the inherent noise resilience of the QC-AFQMC algorithm.

We emphasize that even for the hydrogen molecule in a minimal basis set, executing the QC-AFQMC algorithm using Matchgate shadows requires a substantial amount of resources, with classical computation dominating over the quantum subroutine. It was necessary to perfectly parallelize the classical AFQMC propagation, with one walker per CPU core² (this is only possible when not using population control³). For each walker, it took approximately 1 minute to postprocess the shadows from 16 000 circuits in each time step. As each step of the AFQMC algorithm must be performed sequentially, it could take up to hours or even days to complete the full evolution, see Appendix H for details. This total runtime partly comes from the large number of shadows, which plays the role of a prefactor and could be further improved by implementing additional parallelization among the shadows, for example, processing the shadows corresponding to a single walker in parallel across a number of cores, rather than sequentially on the same core. The total runtime also stems from a scaling as high as $\mathcal{O}(n^8)$ to evaluate the local energy of each walker, which is discussed in detail in Appendix B. This high cost presents a practical concern for scaling this algorithm to larger systems of practical interest.

2. Nitrogen-vacancy center ground state

We applied the QC-AFQMC algorithm to study the electronic states of a point defect in a solid. We chose the NV center in diamond, which is widely considered a promising candidate for quantum sensing and has recently been applied to imaging high-pressure phase transitions [49] and superconducting systems [50]. To obtain multireference states we employed a quantum defect embedding theory (QDET) [51], which allows us to derive an effective Hamiltonian H_{eff} , by defining an active space, see Appendix G for details. This effective Hamiltonian is used as the starting point of the QC-AFQMC calculation.

Although it has been shown [52] that a rather large active space is needed to fully converge the computed neutral excitation energies, here we only chose a minimum model of (4e, 3o) to carry out the QC-AFQMC calculations, as a proof of principle. The ground state has a 3A_2 irrep due to the C_{3v} symmetry of the defect and its wave function can be written as

$$|{}^3A_2\rangle = \frac{1}{\sqrt{2}}(|a_1\bar{a}_1e_x\bar{e}_y\rangle + |a_1\bar{a}_1e_xe_y\rangle), \quad (15)$$

where a_1, e denotes the irrep of the single-particle orbital and the bar symbol represents the spin-down channel. In Ref. [53], a subset of the current authors investigated this system using VQE and found that error mitigation techniques were necessary to obtain an accurate ground state energy. Here, we use

²Intel Xeon Platinum Processor.

³In QMC, population control refers to a technique to retain sampling efficiency by adjusting the walkers so that the algorithm avoids spending a disproportionate amount of time keeping track of walkers that contribute little to the energy estimate [33]. It makes the algorithm more robust against sampling noise but necessitates communication between all the walkers and therefore renders the parallelization less effective.




FIG. 5. The QC-AFQMC calculation of an NV center in diamond using a noiseless quantum simulator (blue crosses), and the IonQ Aria quantum computer (orange circles). Converged results of both have an error on the order of 0.1 mHa compared to classical reference. The inset shows an atomistic model of the defect center.

the tailored UCCSD ansatz in Appendix H as the quantum trial state for our QC-AFQMC calculations.

In QC-AFQMC, we used 4000 shadow circuits (with 100 shots each), 4800 walkers, and an imaginary time step of $\Delta\tau = 0.4 \text{ Ha}^{-1}$. As can be seen from Fig. 5, the results obtained on quantum hardware (orange curve) agree with the noiseless reference (blue curve), which is within our expectation, given the noise resilience discussed in Sec. III B 2. Both curves converge to the classical reference FCI limit, which has been renormalized to zero.

The computation was again fully parallelized, with one walker per physical CPU core (4800 physical cores total). The total runtime required for 400 time steps was approximately 1.5 hours.

IV. DISCUSSION

As we discussed in Sec. I, variational algorithms like VQE have been the methods of choice for quantum chemistry problems on near-term quantum computers. The QC-AFQMC scheme discussed here is complementary to VQE. Specifically, QC-AFQMC could be viewed as an error mitigation technique for VQE, where both the inaccuracy of the ansatz used and the noise effects would in principle be corrected by imaginary time evolution. In turn, the quantum trial state used in the QC-AFQMC algorithm could come from VQE calculations (in a smaller active space, or for a reduced problem size). However, one notable difference from the VQE ansatz is that the quantum trial preparation circuit needs to satisfy $V_T|\mathbf{0}\rangle = |\mathbf{0}\rangle$, if following the approach used in this work. This additional constraint prevents the use of some popular ansatz circuits in VQE, such as qubit coupled cluster [54], and circuit simplification techniques [55]. For some ansatz circuits it may be possible to circumvent this limitation by introducing another ancilla qubit, and controlling the implementation of the ansatz (which can often be done by controlling select gates, rather than every gate), following the procedure in Appendix E of Ref. [31].

TABLE I. Comparison of different measurement schemes for QC-AFQMC.

Method for computing $\langle \Psi_T \phi \rangle$ Only offline QC access?	Vacuum reference [19]	Clifford shadows [17]	Matchgate shadows [26]
	\times	\checkmark	\checkmark
Circuit sampling complexity ^a	$\mathcal{O}\left(\frac{Mn^4}{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2}\sqrt{n}\log n\right)$
Circuit depth complexity ^b		$\mathcal{O}(N_T + n)$	
Noise-resilient overlap ratios	\times	\checkmark	\checkmark
$\{ \langle \Psi_T \phi \rangle \}$ postprocessing complexity	$\mathcal{O}(M)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2} M \exp(n)\right)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2} Mn^{4.5} \log n\right)$
$\{ \langle \Psi_T H \phi \rangle \}$ postprocessing complexity ^c	$\mathcal{O}(Mn^4)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2} M \exp(n)\right)$	$\mathcal{O}\left(\frac{\log(Mn^4/\delta)}{\epsilon^2} Mn^{8.5} \log n\right)$

^a M in the scaling should be interpreted as the total number of walker states involved in the algorithm. The logarithmic factor entering the complexities for the shadow-based methods arises from the use of a median-of-means estimator, and is not present in the heuristic mean estimator used in this work.

^b $\mathcal{O}(N_T)$ represents the quantum trial state preparation circuit depth complexity. And the $|\phi\rangle$ preparation, random Clifford, and Matchgate circuit for the three schemes all scale as $\mathcal{O}(n)$ in depth.

^cThe local energy numerator can be computed by decomposing it into a linear combination of $\mathcal{O}(n^4)$ overlap amplitudes, due to the complexity of the Hamiltonian. The number of shadow samples required also contributes to the postprocessing cost.

The possibility of achieving exponential quantum advantage in solving quantum chemistry problems remains controversial [56]. This is arguably not the direct aim of QC-AFQMC, given that phaseless-AFQMC is already a polynomially scaling method. Instead, the motivation of ph-QC-AFQMC is that the trial state used may provide energies with smaller biases than the purely classical algorithm. A full comparison requires an improved understanding of the advantages offered by quantum trial states over classical trial states and experimenting with larger systems, which we leave for future investigations. As a proxy for this, we could compare the cost of classical methods using state-of-the-art techniques [57–60], such as multi-Slater trial states. Two popular implementations using Wick’s theorem have the following scaling for evaluating the local energy, $\mathcal{O}(MXn^3 + MXN_c)$ [35] and $\mathcal{O}(MXn^4 + MN_c)$ [61], where M, X, N_c are the number of walkers, Cholesky vectors, and Slater determinants in the trial, respectively. This can be compared against the complexities of the quantum algorithm, which is summarized in Table. I for the three different proposed approaches for measuring the overlap amplitudes. We see that the classical algorithms are more efficient in n than the classical shadows-based quantum algorithms, which scale as $\mathcal{O}(n^{8.5} \log^2 n)$. Following this initial posting of this paper, Ref. [62] proposed using algorithmic differentiation to compute the local quantities, which may improve the high polynomial scaling of classical postprocessing. Integrating this approach with the quantum algorithm is an important direction for future research. Faced with the steep overhead of classical postprocessing, the quantum trial state would need to lead to a much smaller bias than a classical multi-Slater trial state with a large number of determinants to be practically advantageous. It is currently an open question as to what kind of quantum trial state is best suited for QC-AFQMC. The only criterion adopted so far is the trial state energy $\langle \Psi_T | H | \Psi_T \rangle$, but there might exist other, better-suited criteria.

While Matchgate classical shadows are formally efficient, they introduce a large postprocessing cost. This presents a practical challenge for scaling up the QC-AFQMC algorithm to larger system sizes. To highlight this, in Fig. 6, we present the classical runtime for postprocessing a single time step of the QC-AFQMC algorithm, assuming parallelization with 1

CPU core per walker. Numerical results were obtained for hydrogen and then extrapolated to larger system sizes using the scaling of the QC-AFQMC algorithm. The results show that our implementation of the algorithm is only realistic up to eight spin-orbitals (water). Within each CPU, computation associated with postprocessing the results of each of the shadow circuits can be further parallelized. While not investigated in this paper, parallelization between shadow circuits can potentially speed up the classical runtime. For example, using one million CPU cores [63] would drive the classical postprocessing of benzene toward a realistic regime, of around 6 hours per time step. Nevertheless, for larger systems such as FeMoco, further optimizations in the algorithm and/or classical postprocessing are still required to become practical.




FIG. 6. Prediction of the runtime estimation of classical postprocessing for a single time step in the QC-AFQMC algorithm performed on a single CPU core. These results were extrapolated from the runtime of postprocessing results from hydrogen, based on a scaling of classical postprocessing of $\mathcal{O}(n^{8.5} \log^2 n)$. We make the optimistic assumptions that the error thresholds ϵ, δ in classical shadows do not need to decrease as the system size increases. The scaling is dominated by the local energy estimation. The system sizes (plotted on the x axis with a log scale) are quantified using active spaces, taken from the literature [54,65–68].

Compared to the classical shadows approach, the Hadamard-test based approach of Ref. [19] reduces the cost of classical postprocessing. However, this comes at a price of an increased number of quantum circuit repetitions (scaling linearly with the number of walkers), leading to similar measurement bottlenecks as observed for VQE methods [64]. In addition, this approach requires iterative communication between quantum and classical processors, which may be subject to latencies. As such, despite the higher classical postprocessing costs of the classical shadow based approach, we currently view it as the more promising approach to QC-AFQMC.

Beyond consideration of computational efficiency, it is important to consider how noise affects the final energy estimate from QC-AFQMC. We have investigated above the effects of hardware noise on the evaluation of overlaps $\langle \Psi_T | \phi_i \rangle$ and their ratios $\frac{\langle \Psi_T | \phi_i \rangle}{\langle \Psi_T | \phi_j \rangle}$, and we have found natural noise resilience for the latter. While this noise resilience will not persist in all settings (e.g., coherent state preparation noise, or noise that violates the GTM assumptions), our results show that there are no benefits to be gained from the robust classical shadows schemes. Since the AFQMC algorithm is driven by overlap ratios, it is important to understand how AFQMC performs with unbiased estimates of overlap ratios, subject to shot noise from finite measurement statistics. This question has recently been addressed by Ref. [25], which observed that the error in the AFQMC energy estimate can be upper bounded by the 2-norm of the Hamiltonian matrix elements. Hence, we can conclude that if the noise obeys the conditions specified above, QC-AFQMC can still achieve accurate results, despite the presence of noise.

Finally, it has been acknowledged [17] and highlighted [38] that the overlap amplitudes $\langle \Psi_T | \phi \rangle$ are expected to decay exponentially with system size, as discussed in Sec. II A. In the worst case scenario, this decay implies an exponentially large number of measurements may still be required to control the uncertainty in estimating these overlaps. Establishing the practical viability of QC-AFQMC will require developing an improved understanding of these challenges, and their possible solutions.

V. CONCLUSIONS

In this work, we carried out the first end-to-end experimental evaluation of the recently proposed Matchgate shadows [26] powered QC-AFQMC algorithm for quantum chemistry [17]. We observed that the algorithm is inherently noise robust, which we find to be a consequence of the natural noise resilience of evaluating overlap ratios via Matchgate shadows. We provided a theoretical explanation for this observed phenomenon, which also elucidates limitations in recently proposed robust Matchgate shadow protocols. We also developed improvements to those protocols which can mitigate state preparation noise, and may be of independent interest. Nevertheless, despite the tantalizing noise resilience of the algorithm, our optimized practical implementations have uncovered a number of challenges to the scalability of the algorithm. Most prominently, while the Matchgate shadows protocol is asymptotically efficient, its high-degree polynomial scaling for evaluating the local energy necessitates significant parallel compute resources for classical

postprocessing. We estimate that for larger molecules, we would require significant amounts of classical postprocessing, as shown in Fig. 6. As such, future work should focus on developing new methods for efficiently computing the local energy in QC-AFQMC, with lower postprocessing costs. Many open questions remain about the best trial states to use for QC-AFQMC, and how to ensure nonvanishing overlaps between the trial state and walker states. The merits of QC-AFQMC found in this work, together with these open research questions, motivate the importance of further study of this algorithm applied to larger system sizes, on real quantum devices.

ACKNOWLEDGMENTS

We thank Senrui Chen, Steve Flammia, Liang Jiang, Andrew Zhao, Kianna Wan, William Huggins, Siyuan Chen, Yu Jin, Guan Wang, and Ji Liu for fruitful discussions. We thank Eric Kessler and Alexander Dalzell for comments on this manuscript. This work was in part supported by the Next Generation Quantum Science and Engineering (Q-NEXT) center that develops quantum science and engineering technologies. Q-NEXT is supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. We thank the Amazon Braket team for facilitating the access to IonQ quantum computers through Amazon Web Service (AWS), and the access to HPC services on AWS. We acknowledge the use of IBM Quantum services for this work and to advanced services provided by the IBM Quantum Researchers Program. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

APPENDIX A: AUXILIARY-FIELD QUANTUM MONTE CARLO

The auxiliary-field quantum Monte Carlo (AFQMC) algorithm is a stochastic implementation of the imaginary time evolution (ITE) process:

$$e^{-\tau(H-E_0)}|\Psi_I\rangle = [e^{-\Delta\tau(H-E_0)}]^N|\Psi_I\rangle, \quad \Delta\tau = \frac{\tau}{N}, \quad (\text{A1})$$

where E_0 is some approximated ground state energy typically from mean-field calculations. ITE projects out the excited state components in the initial state $|\Psi_I\rangle$, and the true ground state $|\Psi_g\rangle$ is ultimately obtained as long as $|\Psi_I\rangle$ is not orthogonal to it. In general, it is often as difficult to realize the ITE in Eq. (A1) as solving the time-independent Schrödinger equation. To remedy this challenge, the Hubbard-Stratonovich transformation [69,70] can be employed, where the two-body interaction in the following form is recast into one-body coupled to an auxiliary field:

$$e^{\frac{\Delta\tau}{2}v_\gamma^2} = \int \frac{dx_\gamma}{\sqrt{2\pi}} e^{-\frac{x_\gamma^2}{2}} e^{\sqrt{\Delta\tau}x_\gamma v_\gamma}, \quad (\text{A2})$$

where v_γ is a matrix that represents a one-body interaction, and x_γ is the auxiliary field. To achieve that, the Hamiltonian,

typically written in second quantized form, has to be rewritten into a Cholesky decomposed format

$$\begin{aligned} H &= H_0 + \sum_{i,j} h_{ij} a_i^\dagger a_j + \frac{1}{2} \sum_{i,j,k,l} V_{ijkl} a_i^\dagger a_j^\dagger a_l a_k \\ &= H_0 + v_0 - \frac{1}{2} \sum_{\gamma} v_{\gamma}^2, \end{aligned} \quad (\text{A3})$$

where H_0 is a constant, h, V are the one- and two-electron integrals, v_0 is the modified one-body interaction, $v_{\gamma} = i\mathcal{L}_{\gamma}$, and \mathcal{L}_{γ} is the Cholesky vector of the two-body interactions satisfying

$$\mathcal{L}_{\gamma} = \sum_{pq} L_{pq}^{\gamma} a_p^\dagger a_q, \quad V_{ijkl} = \sum_{\gamma} L_{ik}^{\gamma} L_{lj}^{\gamma*}. \quad (\text{A4})$$

Note that we have assumed a mean-field subtraction [33] in the above Hamiltonian. With it, and the second-order Trotter formula [71,72]: $e^{A+B} \approx e^{A/2} e^B e^{A/2}$, the short-time propagator can be approximated as

$$\begin{aligned} e^{-\Delta\tau(H-E_0)} &\approx e^{-\Delta\tau(H_0-E_0)} e^{-\frac{\Delta\tau}{2} v_0} \prod_{\gamma} e^{\frac{\Delta\tau}{2} v_{\gamma}^2} e^{-\frac{\Delta\tau}{2} v_0} \\ &= \int d\mathbf{x} p(\mathbf{x}) B(\mathbf{x}) + O(\Delta\tau^2), \end{aligned} \quad (\text{A5})$$

where we have defined

$$B(\mathbf{x}) = e^{-\frac{\Delta\tau}{2} v_0} \prod_{\gamma} e^{\sqrt{\Delta\tau} x_{\gamma} v_{\gamma}} e^{-\frac{\Delta\tau}{2} v_0}. \quad (\text{A6})$$

The constant prefactor is usually left out and $p(\mathbf{x})$ is the multivariable standard normal distribution. After this transformation, we have only one-body terms left, and the original interacting problem is now mapped onto an ensemble of noninteracting problems coupled to a set of auxiliary fields. The probability distribution $p(\mathbf{x})$ in the above expression is realized by an ensemble of walkers $|\phi_l\rangle$, which are evolved as

$$|\phi^{(i+1)}\rangle = B(\mathbf{x})|\phi^{(i)}\rangle, \quad (\text{A7})$$

where the superscript represents the time step. The walkers in AFQMC are chosen as nonorthogonal Slater determinants and stay in the manifold of Slater determinants \mathcal{S} during evolution, due to the Thouless theorem [73,74]. The ground state is then estimated by a mixed energy estimator which is exact and can lead to considerable simplifications in practice

$$E = \frac{\langle \Psi_T | H e^{-N\Delta\tau(H-E_0)} | \phi^{(0)} \rangle}{\langle \Psi_T | e^{-N\Delta\tau(H-E_0)} | \phi^{(0)} \rangle}, \quad (\text{A8})$$

where $|\Psi_T\rangle$, $|\phi^{(0)}\rangle$ are the trial state and initial Slater determinant, respectively.

So far we have not imposed any constraints on the evolution of walkers, which is known as free-projection AFQMC. The only potential systematic error comes from the Trotterization error. However, free-projection AFQMC is prone to large fluctuations in the estimated energy because of an asymptotic instability in τ , also known as the phase problem. The phase problem manifests as the phase θ_l of walkers evolves into the whole $[0, 2\pi)$ range during the random walk (e.g., see Fig. 3 in Ref. [34] for visualization), and results in the denominator in Eq. (1), i.e., $\sum_l w_l e^{i\theta_l}$ approaching 0 exponentially quickly

with τ . This problem could be viewed as a generalization of the sign problem, see Refs. [33,34] for a detailed discussion.

Importance sampling is a variance-reduction technique widely used in QMC methods. It's also employed to control the phase problem in AFQMC. To do that, we first note that the Eq. (A5) still holds up to a complex-valued shift $\bar{\mathbf{x}}$. Therefore it can be modified as

$$p(\mathbf{x}) \rightarrow p(\mathbf{x} - \bar{\mathbf{x}}), \quad B(\mathbf{x}) \rightarrow B(\mathbf{x} - \bar{\mathbf{x}}). \quad (\text{A9})$$

To find the best $\bar{\mathbf{x}}$, we first expand Eq. (A8) as

$$\begin{aligned} E &\approx \frac{\int \prod_{k=0}^{N-1} d\mathbf{x}_k p(\mathbf{x}_k - \bar{\mathbf{x}}_k) \langle \Psi_T | H | \phi^{(N)} \rangle}{\int \prod_{k=0}^{N-1} d\mathbf{x}_k p(\mathbf{x}_k - \bar{\mathbf{x}}_k) \langle \Psi_T | \phi^{(N)} \rangle} \\ &= \frac{\int \prod_{k=0}^{N-1} d\mathbf{x}_k p(\mathbf{x}_k) I(\mathbf{x}_k, \bar{\mathbf{x}}_k, \phi^{(k)}) \frac{\langle \Psi_T | H | \phi^{(N)} \rangle}{\langle \Psi_T | \phi^{(N)} \rangle}}{\int \prod_{k=0}^{N-1} d\mathbf{x}_k p(\mathbf{x}_k) I(\mathbf{x}_k, \bar{\mathbf{x}}_k, \phi^{(k)})}, \end{aligned} \quad (\text{A10})$$

where we have defined

$$|\phi^{(N)}\rangle = B(\mathbf{x}_{N-1} - \bar{\mathbf{x}}_{N-1}) \dots B(\mathbf{x}_0 - \bar{\mathbf{x}}_0) |\phi^{(0)}\rangle, \quad (\text{A11})$$

and an importance function

$$I(\mathbf{x}, \bar{\mathbf{x}}, \phi) = \frac{\langle \Psi_T | B(\mathbf{x} - \bar{\mathbf{x}}) | \phi \rangle}{\langle \Psi_T | \phi \rangle} e^{\mathbf{x} \cdot \bar{\mathbf{x}} - \frac{\bar{\mathbf{x}}^2}{2}}. \quad (\text{A12})$$

By choosing

$$\bar{x}_{\gamma} = -\sqrt{\Delta\tau} \langle v_{\gamma} \rangle, \quad \langle v_{\gamma} \rangle = \frac{\langle \Psi_T | v_{\gamma} | \phi \rangle}{\langle \Psi_T | \phi \rangle}, \quad (\text{A13})$$

fluctuations in the importance function to first order in $\sqrt{\Delta\tau}$ is canceled [33], leading to a more stable random walk. The dynamic shift, usually referred to as a force bias, modifies the sampling of the auxiliary fields by shifting the center of the Gaussian distribution $p(\mathbf{x})$ according to the overlap $\langle \Psi_T | \phi \rangle$. The importance function can therefore be further approximated as $I(\mathbf{x}, \bar{\mathbf{x}}, \phi) \approx \exp[-\Delta\tau(E^{\text{loc}} - E_0)]$, known as the local energy formalism. Introducing the force bias into the importance sampling would solve the instability if $B(\mathbf{x})$ were real, which could be true in special cases. For a general phase problem, however, it does not lead to complete control of the phase problem.

To further control the phase problem, the local energy in the importance function has to be replaced by its real part, leading to $\theta_l = 0$ and $w_l > 0$ in Eq. (1). A phaseless approximation is also adopted as

$$\begin{aligned} I(\mathbf{x}, \bar{\mathbf{x}}, \phi_l) &\approx \exp[-\Delta\tau(\Re E_l^{\text{loc}} - E_0)] \\ &\times \max\left(0, \cos\left[\arg\left(\frac{\langle \Psi_T | B(\mathbf{x} - \bar{\mathbf{x}}) | \phi_l \rangle}{\langle \Psi_T | \phi_l \rangle}\right)\right]\right), \end{aligned} \quad (\text{A14})$$

where abrupt phase changes during the random walk are now forbidden. The weight is updated as $w_l^{(k)} \leftarrow w_l^{(k-1)} \times I(\mathbf{x}, \bar{\mathbf{x}}, \phi_l^{(k-1)})$. Combining these techniques finally resolves the phase problem, and the whole algorithm is usually referred to as the phaseless AFQMC (ph-AFQMC). Note that the phaseless approximation also introduces a bias to the mixed energy estimator, which we seek to systematically improve by using quantum computation in QC-AFQMC.

The whole process of the AFQMC algorithm is summarized in Algorithm 1, following the presentation in Ref. [19].

ALGORITHM 1. Auxiliary-field quantum Monte Carlo.

- 1: Input $H, E_0, |\Psi_T\rangle, |\Psi_I\rangle$, number of walkers M , number of time steps N , and step size $\Delta\tau$. (H will be decomposed into L Cholesky vectors.)
 - 2: **for** $l = 1$ to M **do**
 - 3: $|\phi_l^{(0)}\rangle \leftarrow |\Psi_I\rangle, w_l^{(0)} \leftarrow 1$ \triangleright Initialize the walker.
 - 4: **for** $k = 1$ to N **do**
 - 5: **for** $\gamma = 1$ to L **do**
 - 6: $\bar{x}_\gamma = -\sqrt{\Delta\tau} \frac{\langle \Psi_T | \delta_\gamma | \phi_l^{(k-1)} \rangle}{\langle \Psi_T | \phi_l^{(k-1)} \rangle}$
 - 7: Sample \mathbf{x} according to distribution $p(\mathbf{x})$
 - 8: $|\phi_l^{(k)}\rangle \leftarrow B(\mathbf{x} - \bar{\mathbf{x}}) |\phi_l^{(k-1)}\rangle$. \triangleright Update the walker
 - 9: $E_l^{(k-1)} \leftarrow \frac{\langle \Psi_T | H | \phi_l^{(k-1)} \rangle}{\langle \Psi_T | \phi_l^{(k-1)} \rangle}$ \triangleright Compute local energy
 - 10: $\theta \leftarrow \arg\left(\frac{\langle \Psi_T | \phi_l^{(k)} \rangle}{\langle \Psi_T | \phi_l^{(k-1)} \rangle}\right)$ \triangleright Compute the phase
 - 11: $w_l^{(k)} \leftarrow w_l^{(k-1)} \times I(\mathbf{x}, \bar{\mathbf{x}}, \phi_l^{(k-1)})$ \triangleright Update weight
 - 12: $E_l^{(N)} \leftarrow \frac{\langle \Psi_T | H | \phi_l^{(N)} \rangle}{\langle \Psi_T | \phi_l^{(N)} \rangle}$ \triangleright Compute the local energy
 - 13: Output the energy $E \leftarrow \frac{\sum_l^M w_l^{(N)} E_l^{(N)}}{\sum_l^M w_l^{(N)}}$
-

APPENDIX B: LOCAL ENERGY EVALUATION WITH CLASSICAL SHADOWS

The most time-consuming step in QC-AFQMC is the evaluation of the local energy. Here we discuss how this task is performed in our implementation compatible with classical shadows. We first note that the local energy numerator can be rewritten as

$$\begin{aligned}
\langle \Psi_T | H | \phi \rangle &= \langle \Psi_T | U_\phi U_\phi^\dagger H U_\phi | \Phi_0 \rangle \\
&= \sum_{pr} \langle \Psi_T | U_\phi | \Phi_p^r \rangle \langle \Phi_p^r | \bar{H} | \Phi_0 \rangle \\
&\quad + \sum_{pqsr} \langle \Psi_T | U_\phi | \Phi_{pq}^{rs} \rangle \langle \Phi_{pq}^{rs} | \bar{H} | \Phi_0 \rangle \\
&= \sum_{pr} \langle \Psi_T | \phi_p^r \rangle \langle \Phi_p^r | \bar{H} | \Phi_0 \rangle \\
&\quad + \sum_{pqsr} \langle \Psi_T | \phi_{pq}^{rs} \rangle \langle \Phi_{pq}^{rs} | \bar{H} | \Phi_0 \rangle, \quad (\text{B1})
\end{aligned}$$

where $\bar{H} = U_\phi^\dagger H U_\phi$ represents the rotated Hamiltonian and $|\Phi_0\rangle, |\Phi_p^r\rangle (|\Phi_{pq}^{rs}\rangle)$ denotes the Hartree-Fock state and its single (double) excited counterpart. We also have $|\phi_p^r\rangle = U_\phi |\Phi_p^r\rangle$. The second line in the above equation utilizes the resolution of identity and the fact that Hamiltonian has only up to two-body interactions.

The Hamiltonian rotation can be done with $\mathcal{O}(n^5)$ operations. In the above equation, the overlap amplitudes $\langle \Psi_T | \phi_p^r \rangle, \langle \Psi_T | \phi_{pq}^{rs} \rangle$ can be calculated with Matchgate shadows each with complexity $\mathcal{O}(n^4)$, as already discussed in the main

text. Each matrix element can be computed using the Slater-Condon rules with complexity $\mathcal{O}(1)$. Due to the number of possible excited Slater determinants, which scales as $\mathcal{O}(n^4)$, the scaling of the evaluation of the local energy for each walker would therefore scale as $\mathcal{O}(n^8 + n^5 + n^4) = \mathcal{O}(n^8)$ using this approach.

We also note that there exist other ways to evaluate the local energy with Matchgate shadows. Sec.V C of Ref. [26] proposed an alternative way to evaluate the local energy using the Grassmann algebra. In this approach, the Hamiltonian is first rewritten into the Majorana formalism. Then the mixed estimator of each Majorana operator $\langle \Psi_T | \gamma_S | \phi \rangle$ can be estimated in a similar fashion to the overlap amplitude. This approach leads to a scaling between $\mathcal{O}(n^9)$, due to the relevant matrices being noninvertible, making it slightly less favorable than the previous approach.

APPENDIX C: CLASSICAL SHADOWS REVISITED

In this section, we will revisit classical shadows from a group representation perspective, which will be useful when we include noise effects. We closely follow the presentation in Ref. [29].

We start by introducing the notations and conventions used throughout these Appendixes. We consider a n -qubit system with Hilbert space \mathcal{H}_n . Its dimension is denoted by $d \equiv 2^n$. All the n -qubit operators (super-operators) live in a vector space defined as $\mathcal{L}(\mathcal{H}_n)$ ($\mathcal{L}(\mathcal{L}(\mathcal{H}_n))$). For better clarity, we make use of the Liouville representation of operators and super-operators.

(1) Operators are notated using the double kets $|A\rangle\rangle = \frac{1}{\sqrt{\text{tr}(A^\dagger A)}} A$ for its length d^2 vectorization for $A \in \mathcal{L}(\mathcal{H}_n)$ and $\langle\langle B|A\rangle\rangle := \frac{\text{tr}(B^\dagger A)}{\sqrt{\text{tr}(B^\dagger B)\text{tr}(A^\dagger A)}}$;

(2) Superoperators, written as cursive letters, are mapped to $d^2 \times d^2$ matrices: any $\mathcal{E} \in \mathcal{L}(\mathcal{L}(\mathcal{H}_n))$ can be specified by its matrix elements $\mathcal{E}_{ij} := \langle\langle B_i | \mathcal{E} | B_j \rangle\rangle$, where $\{|B_i\rangle\rangle\}$ is an orthonormal basis for $\mathcal{L}(\mathcal{L}(\mathcal{H}_n))$.

For any unitary U , its corresponding channel is denoted by $\mathcal{U}(\cdot) := U(\cdot)U^\dagger$. For any $|\phi\rangle \in \mathcal{H}_n, |\phi\rangle\rangle$ is the vectorization of $|\phi\rangle\langle\phi|$. And hats indicate statistical estimators, e.g., \hat{o} denotes an estimate for $o = \text{tr}(O\rho)$.

Classical shadows are based on a simple measurement primitive: for the quantum state ρ , apply a unitary U randomly drawn from a distribution of unitaries \mathcal{D} and measure in the computational basis. This produces measurement outcomes $b \in \{0, 1\}^{\otimes n}$ with probability $\langle b | U \rho U^\dagger | b \rangle$. One then inverts the unitary on the outcome $|b\rangle$ in postprocessing, which amounts to storing a classical representation of $U^\dagger |b\rangle$. The distribution \mathcal{D} , from which the random unitaries U are drawn, usually can be associated with a group G , e.g., random Clifford unitaries and the Clifford group. \mathcal{M} can be viewed as a twirl of the measurement channel \mathcal{M}_Z , giving:

$$\mathcal{M} := \mathbb{E}_{g \sim G} U_g^\dagger \mathcal{M}_Z U_g, \quad \mathcal{M}_Z = \sum_{b \in \{0,1\}^n} |b\rangle\rangle \langle\langle b|, \quad (\text{C1})$$

where g is an element in G , and the random unitaries \mathcal{U} is a unitary representation of G . Applying Schur's lemma, and

assuming no multiplicities in \mathcal{U} lead to

$$\mathcal{M} = \sum_{\lambda \in R_G} f_\lambda \Pi_\lambda, \quad f_\lambda = \frac{\text{tr}(\mathcal{M}_Z \Pi_\lambda)}{\text{tr}(\Pi_\lambda)}, \quad (\text{C2})$$

where R_G are the irreducible representations (irreps) of G . The super-operators $\Pi_\lambda \in \mathcal{L}(\mathcal{L}(\mathcal{H}_n))$ are orthogonal projectors onto the irreducible subspaces Γ_λ , and f_λ is the eigenvalue associated with each orthogonal projector.

APPENDIX D: MATCHGATE SHADOWS

In this section, we outline the basics of Matchgate shadows, closely following the presentation in Ref. [26]. The Matchgate unitaries are based on the Majorana formalism. For a system of n orbitals, $2n$ Majorana operators can be defined as

$$\gamma_{2j-1} = a_j + a_j^\dagger, \quad \gamma_{2j} = -i(a_j - a_j^\dagger), \quad (\text{D1})$$

for $j \in [n] := \{1, \dots, n\}$. These have qubit representation

$$\gamma_{2j-1} = (\Pi_{i=1}^{j-1} Z_i) X_j, \quad \gamma_{2j} = (\Pi_{i=1}^{j-1} Z_i) Y_j, \quad (\text{D2})$$

under the Jordan-Wigner (JW) transformation [45], where X_i, Y_i, Z_i are the Pauli operators on qubit i . For a subset of indices $S \subseteq [2n]$, we denote by γ_S the product of the Majorana operators indexed by the elements in S in ascending order. That is,

$$\gamma_S := \gamma_{\mu_1} \cdots \gamma_{\mu_{|S|}}, \quad (\text{D3})$$

for $S = \{\mu_1, \dots, \mu_{|S|}\} \subseteq [2n]$ with $\mu_1 < \dots < \mu_k$.

1. Fermionic Gaussian states

Matchgate circuits are qubit representations of fermionic Gaussian unitaries under the JW transformation. The fermionic Gaussian unitaries transform between valid sets of Majorana operators $\{\gamma_{2j-1}, \gamma_{2j}\}, j \in [n]$. Fermionic Gaussian states are the ground states and thermal states of noninteracting fermionic Hamiltonians. An n -mode Gaussian state is any state whose density operator ϱ can be written as

$$\varrho = \prod_{j=1}^n \frac{1}{2} (I - i\lambda_j \gamma_{2j-1} \gamma_{2j}), \quad (\text{D4})$$

for some coefficients $\lambda_j \in [-1, 1]$. If $\lambda_j \in \{-1, 1\}$ for all $j \in [n]$, then ϱ is a pure Gaussian state; otherwise, ϱ is a mixed state. For any computational basis state $|b\rangle$, for example, we have

$$|b\rangle\langle b| = \prod_{j=1}^n \frac{1}{2} [I - i(-1)^{b_j} \gamma_{2j-1} \gamma_{2j}]. \quad (\text{D5})$$

A Gaussian state can also be defined by its two-point correlations $\text{tr}(\varrho \gamma_\mu \gamma_\nu)$, which is also known as its covariance matrix. An n -qubit Gaussian state ρ has its covariance matrix defined as:

$$(C_\rho)_{\mu\nu} := -\frac{i}{2} \text{tr}([\gamma_\mu, \gamma_\nu] \rho) \quad (\text{D6})$$

for $\mu, \nu \in [2n]$, which is a $2n \times 2n$ antisymmetric matrix. For any computational basis state $|b\rangle\langle b|$ as an example, its

covariance matrix is

$$C_{|b\rangle} := \bigoplus_{j=1}^n \begin{pmatrix} 0 & (-1)^{b_j} \\ (-1)^{b_{j+1}} & 0 \end{pmatrix}. \quad (\text{D7})$$

A Gaussian state which is also an eigenstate of the number operator $\sum_{j=1}^n a_j^\dagger a_j$ is known as Slater determinant. Any ζ -fermion Slater determinant $|\phi\rangle$ can also be written as

$$|\phi\rangle = a_1^{\dagger\zeta} \dots a_\zeta^{\dagger\zeta} |0\rangle, \quad a_j^{\dagger\zeta} = \sum_{k=1}^n V_{kj} a_k^\dagger = U_V a_j^\dagger U_V^\dagger, \quad (\text{D8})$$

for some $n \times n$ unitary matrix V . Hence a Slater determinant can also be specified by the first ζ columns of V . Finally the Majorana operators $\{\gamma'_\mu\}_{\mu \in [2n]}$ have the following relation:

$$\gamma'_{2j-1} = \sum_k [\Re(V_{jk}) \gamma_{2j-1} - \Im(V_{jk}) \gamma_{2j}], \quad (\text{D9})$$

$$\gamma'_{2j} = \sum_k [\Im(V_{jk}) \gamma_{2j-1} + \Re(V_{jk}) \gamma_{2j}]. \quad (\text{D10})$$

So the fermionic Gaussian unitary $U_{Q'}$ that implements this transformation is given by the orthogonal matrix

$$Q' = \begin{pmatrix} R_{11} & \cdots & R_{1n} \\ \vdots & \ddots & \vdots \\ R_{n1} & \cdots & R_{nn} \end{pmatrix}, \quad R_{jk} := \begin{pmatrix} \Re(V_{jk}) & -\Im(V_{jk}) \\ \Im(V_{jk}) & \Re(V_{jk}) \end{pmatrix}. \quad (\text{D11})$$

2. Matchgate 3-design

The adjoint action of Fermionic Gaussian unitaries U_Q on the Majorana operators obeys

$$U_Q \gamma_\mu U_Q^\dagger = \sum_{\nu \in [2n]} Q_{\nu\mu} \gamma_\nu, \quad \mu \in [2n], \quad (\text{D12})$$

where Q belongs to the orthogonal group $O(2n)$. Matchgate circuits form a continuous group M_n which is in one-to-one correspondence with the orthogonal group $O(2n)$, up to a global phase:

$$M_n = \{U_Q : Q \in O(2n)\}. \quad (\text{D13})$$

Ref. [26] considered two distributions: (i) ‘‘uniform’’ distribution over M_n , where uniformity is more precisely given by the normalized Haar measure μ on $O(2n)$; (ii) uniform distribution over the discrete Borel group $B(2n)$, which is the intersection of M_n and the n -qubit Clifford group Cl_n . The second consists of $2n \times 2n$ signed permutation matrices:

$$M_n \cap \text{Cl}_n = \{U_Q : Q \in B(2n)\}. \quad (\text{D14})$$

For $j \in \mathbb{Z}_{>0}$, we use $\mathcal{E}_{M_n}^{(j)}$ and $\mathcal{E}_{M_n \cap \text{Cl}_n}^{(j)}$ to denote the j -fold twirl channels corresponding to the distributions over M_n and $M_n \cap \text{Cl}_n$, respectively:

$$\mathcal{E}_{M_n}^{(j)} := \int_{O(2n)} d\mu(Q) \mathcal{U}_Q^{\otimes j}, \quad (\text{D15})$$

$$\mathcal{E}_{M_n \cap \text{Cl}_n}^{(j)} := \frac{1}{|B(2n)|} \sum_{Q \in B(2n)} \mathcal{U}_Q^{\otimes j}, \quad (\text{D16})$$

where we have used the fact that twofold Matchgate twirl is hermitian [26]. Since the measurement channel \mathcal{M} in the classical shadows procedure and the variance of the estimates

obtained from the classical shadows are determined by the two- and threefold twirls, it is necessary to compare $\mathcal{E}_{M_n}^{(j)}$ and $\mathcal{E}_{M_n \cap Cl_n}^{(j)}$ up to $j = 3$.

Reference [26] proved that the j -fold twirl channels of these two distributions are equivalent up to $j = 3$. Thus the discrete ensemble of Clifford Matchgate circuits is a 3-design for the continuous Haar-uniform distribution over all Matchgate circuits, in the same way that the Clifford group is a unitary 3-design [75]. Ref. [26] stated this result informally as ‘‘The group of Clifford Matchgate circuits forms a Matchgate 3-design.’’ In the implemented Matchgate shadow protocol, we adopt the latter since it requires less memory to store a random signed permutation matrix compared to an orthogonal matrix. An additional virtue of sampling in $M_n \cap Cl_n$ is that Clifford circuits can be randomly compiled to convert coherent errors into incoherent errors [76] (within the shadow circuit), leading to better satisfaction of the GTM assumption of noise. Our implementation of Matchgate circuits is given in Appendix F.

3. Matchgate channel

The n -qubit Matchgate group M_n has $(2n + 1)$ irreps, so we can define correspondingly $(2n + 1)$ subspaces

$$\Gamma_k := \text{span} \left\{ \gamma_S : S \in \binom{[2n]}{k} \right\}, \quad k \in \{0, \dots, 2n\}, \quad (\text{D17})$$

where $\binom{[2n]}{k}$ denotes the set of subsets of $[2n]$ of cardinality k . Therefore we have $\mathcal{L}(\mathcal{H}_n) = \bigoplus_{k=0}^{2n} \Gamma_k$. For fermionic simulations, we usually deal with even operators so we only look at the even subspaces $\Gamma_{\text{even}} = \bigoplus_{l=0}^n \Gamma_{2l}$.

Now we can derive the details of the Matchgate channel, which will serve as the basis for calculations involving noise presented in the following section. As we discussed above, the measurement channel for any shadow scheme can be derived by specifying the twofold twirl $\mathcal{E}^{(2)}$

$$\mathcal{M}(\rho) = \text{tr}_1 \left[\sum_{b \in \{0,1\}^n} \mathcal{E}_{M_n \cap Cl_n}^{(2)}(|b\rangle\langle b|^{\otimes 2})(\rho \otimes I) \right], \quad (\text{D18})$$

where tr_1 denotes the partial trace over the first tensor component. The twofold twirl for the Matchgate group has the following form:

$$\mathcal{E}_{M_n \cap Cl_n}^{(2)} = \sum_{k=0}^{2n} |\Upsilon_k^{(2)}\rangle\langle \Upsilon_k^{(2)}|, \quad (\text{D19})$$

$$|\Upsilon_k^{(2)}\rangle = \binom{2n}{k}^{-1/2} \sum_{S \in \binom{[2n]}{k}} |\gamma_S\rangle |\gamma_S\rangle. \quad (\text{D20})$$

The standard derivation can be found in Sec. IV B in Ref. [26]. For our purpose, we will take a shortcut by assuming all the projectors $\Pi_{2l} \in \mathcal{L}(\mathcal{L}(\mathcal{H}_n))$ are already known

$$\Pi_{2l} := \sum_{S \in \binom{[2n]}{2l}} |\gamma_S\rangle\langle \gamma_S|, \quad (\text{D21})$$

where we have defined $|\gamma_S\rangle = \gamma_S / \sqrt{2^n}$, and $\langle \gamma_S | \gamma_{S'} \rangle = \delta_{SS'}$. We directly start from Eq. (C2), and the eigenvalues are hence

computed as

$$\begin{aligned} f_{2l} &= \frac{\text{Tr}[\mathcal{M}_Z \Pi_{2l}]}{\text{Tr}[\Pi_{2l}]} \\ &= \frac{\sum_{S \in \binom{[2n]}{2l}} \sum_{b \in \{0,1\}^n} |\langle \gamma_S | b \rangle|^2}{\sum_{S \in \binom{[2n]}{2l}} \langle \gamma_S | \gamma_S \rangle} \\ &= \binom{2n}{2l}^{-1} \sum_{b \in \{0,1\}^n} \sum_{T \in \binom{[n]}{l}} \left| \frac{(-i)^{|T|}}{\sqrt{2^n}} (-1)^{\sum_{j \in T} b_j} \right|^2 \\ &= \binom{2n}{2l}^{-1} \binom{n}{l}, \end{aligned} \quad (\text{D22})$$

where we have used

$$|b\rangle = \frac{1}{\sqrt{2^n}} \sum_{T \in \binom{[n]}{l}} (-i)^{|T|} (-1)^{\sum_{j \in T} b_j} |\gamma_{\text{pairs}(T)}\rangle, \quad (\text{D23})$$

and $\text{pairs}(T)$ is defined as $\bigcup_{j \in T} \{2j - 1, 2j\}$. We only consider the even subspaces for fermionic simulations. Therefore the Matchgate channel is

$$\mathcal{M} = \sum_{l=0}^n \binom{2n}{2l}^{-1} \binom{n}{l} \Pi_{2l}. \quad (\text{D24})$$

With the above definitions and properties of fermionic Gaussian transformation, we can finally discuss how to use the Matchgate shadow to compute the desired overlap integrals. The overlap can first be reconstructed as

$$\begin{aligned} \langle \Psi_T | \phi \rangle &= 2 \sum_{l=0}^n \binom{2n}{2l} \binom{n}{l}^{-1} \\ &\quad \times \mathbb{E}_{Q \sim B(2n)} \text{tr}[|\phi\rangle\langle \mathbf{0} | \Pi_{2l} (U_Q^\dagger \hat{b} \langle \hat{b} | U_Q)], \end{aligned} \quad (\text{D25})$$

where the trace involving Π_{2l} correspond to the coefficient of z^l in the following polynomial [26]

$$q_{|\phi\rangle, |b\rangle}(z) = \frac{i^{\zeta/2}}{2^{n-\zeta/2}} \text{pf}[(C_{\mathbf{0}}) + z W^* Q^T Q^T C_{|b\rangle} Q Q^T W^\dagger] |_{\bar{s}_\zeta}.$$

The pf in the above equation denotes the matrix Pfaffian, and $M|_S$ represents matrix M restricted to rows and columns in set S . ζ is the particle number, $C_{\mathbf{0}}$ denotes the covariance matrix of the vacuum state, and Q' is an orthogonal matrix defined from the Slater determinant $|\varphi\rangle$ according to Eq. (D11). We have also defined the W matrix

$$W = \bigoplus_{j=1}^{\zeta} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix} \bigoplus_{j=\zeta+1}^n \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (\text{D26})$$

and $\bar{s}_\zeta := [2n] \setminus \{1, 3, \dots, 2\zeta - 1\}$. All these coefficients can be computed using polynomial interpolation in $\mathcal{O}((n - \zeta/2)^4)$ time, where ζ is the number of electrons.

4. Noise resilience

In this subsection, we prove Theorem 1. This result serves as the foundation of noise resilience observed in experiments performed on quantum hardware.

Theorem 1. For an n qubit state $|\phi\rangle = \sum_i c_i |i\rangle$, which is a linear combination of computational basis states $|i\rangle$ with fixed Hamming weight ζ ,

$$\mathcal{P}_{0,\zeta}[\Pi_{2l}(|\phi\rangle\langle\mathbf{0}|)] = b_{2l}|\phi\rangle\langle\mathbf{0}|, \quad (11)$$

with

$$b_{2l} = \begin{cases} 2^{\zeta-n} \binom{n-\zeta}{l-\frac{\zeta}{2}}, & \text{if } \frac{\zeta}{2} \leq l \leq n - \frac{\zeta}{2} \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

such that $\sum_{l=0}^n b_{2l} = 1$. Here $\mathcal{P}_{0,\zeta}$ denotes the projector onto the subspace spanned by the Hamming weight zero state $|\mathbf{0}\rangle$ and computational basis states of Hamming weight ζ .

Proof. We expand

$$\mathcal{P}_{0,\zeta}[\Pi_{2l}(|\phi\rangle\langle\mathbf{0}|)] = \sum_i c_i \mathcal{P}_{0,\zeta}[\Pi_{2l}(|i\rangle\langle\mathbf{0}|)]. \quad (D27)$$

By the results of Lemma 1,

$$\begin{aligned} \mathcal{P}_{0,\zeta}[\Pi_{2l}(|\phi\rangle\langle\mathbf{0}|)] &= \sum_i c_i b_{2l} |i\rangle\langle\mathbf{0}| \\ &= b_{2l} |\phi\rangle\langle\mathbf{0}| \end{aligned} \quad (D28)$$

as required. \blacksquare

The above result is proved using the following lemma. The general proof idea is to write

$$\Pi_\zeta(|i\rangle\langle\mathbf{0}|) = 2^{-n} \sum_{s \in \binom{[2n]}{2l}} \langle\mathbf{0}|\gamma_s^\dagger |i\rangle \gamma_s. \quad (D29)$$

The action of the Majorana strings is to “create electrons” in the $|\mathbf{0}\rangle$ state. Terms in the sum evaluate to zero unless ζ electrons are created. For small/large numbers of Majorana operators in each string, insufficient electrons are created, and the whole sum evaluates to zero. In the intermediate regime, each Majorana string can be divided into two parts; a part that creates ζ electrons in the correct positions, and a part that contributes an ultimately unimportant phase and increases the multiplicity of a given term. The more interesting part is the part that creates the electrons, which acts as a projector from an x -electron subspace to an $x + \zeta$ electron subspace. The action of projecting into the subspace spanned by the vacuum state and computational basis states with Hamming weight ζ retains only terms that project from zero electrons to ζ electrons. Combinatorics give the desired values of b_{2l} .

Lemma 1. Given an n qubit computational basis state $|i\rangle$ with even Hamming weight ζ , then $\mathcal{P}_{0,\zeta}[\Pi_{2l}(|i\rangle\langle\mathbf{0}|)] = b_{2l} |i\rangle\langle\mathbf{0}|$ with

$$b_{2l} = \begin{cases} 2^{\zeta-n} \binom{n-\zeta}{l-\frac{\zeta}{2}}, & \text{if } \frac{\zeta}{2} \leq l \leq n - \frac{\zeta}{2} \\ 0, & \text{otherwise} \end{cases}, \quad (D30)$$

where $\mathcal{P}_{0,\zeta}$ is the projection from onto the subspace spanned by the vacuum state and computational basis states with Hamming weight ζ .

Proof. First expand

$$\Pi_\zeta(|i\rangle\langle\mathbf{0}|) = 2^{-n} \sum_{s \in \binom{[2n]}{2l}} \langle\mathbf{0}|\gamma_s^\dagger |i\rangle \gamma_s \quad (D31)$$

using the definition of super-operators from Appendix C. The set S contains the $\binom{2n}{2l}$ possible Majorana strings formed by

choosing $2l$ Majorana operators acting on n qubits. Recall the qubit representation of Majorana operators under the Jordan-Wigner transform

$$\gamma_{2j-1} = (\Pi_{i=1}^{j-1} Z_i) X_j, \quad \gamma_{2j} = (\Pi_{i=1}^{j-1} Z_i) Y_j,$$

We refer to “connected” Majorana operators as pairs of the form $\gamma_k \gamma_{k+1} = i Z_{\frac{k+1}{2}}$ for k odd. We refer to “unconnected” Majorana operators as pairs of the form $\gamma_i \gamma_j$ such that $j \neq i + 1$ if i is odd. Unconnected Majorana operators take the form $\gamma_i \gamma_j = p_{i,j} O_i \bar{Z}_{i+1}^{j-1} O_j$ where $p_{i,j} \in \{\pm 1, \pm i\}$ and $O \in \{X, Y\}$, and the notation $\bar{Z}_a^b := \bigotimes_{k=a}^b Z_k$. Hence, observe that connected Majorana operators act as the identity on $|\mathbf{0}\rangle$, while unconnected Majorana operators act as $\gamma_i \gamma_j |\mathbf{0}\rangle = |0 \dots 0 1_i 0 \dots 0 1_j 0 \dots 0\rangle$ (up to a complex phase). Because each unconnected pair increases the Hamming weight by 2, we require γ_S to contain exactly $\zeta/2$ unconnected pairs, and all other pairs to be connected pairs. If this requirement is not satisfied, $\langle\mathbf{0}|\gamma_S^\dagger |i\rangle = 0$.

First consider the case $2l < \zeta$. In this case, γ_S contains $2l < \zeta$ Majorana operators, and it is impossible to form $\zeta/2$ unconnected pairs. Hence, every term in the sum evaluates to zero, and $b_{2l} = 0$, as required.

For $2l > 2n - \zeta$, observe that because the Hamming weight of $|i\rangle$ is ζ , there are $(n - \zeta)$ 0’s in $|i\rangle$. After placing a connected Majorana pair on each of these 0 qubits, the remaining number of Majorana operators to distribute onto the ζ 1’s in $|i\rangle$ is given by $2l - 2(n - \zeta) > 2n - \zeta - 2(n - \zeta) = \zeta$. It is necessary to place greater than ζ Majorana operators on the ζ qubits. As a result, at least two of the Majorana’s will form a connected pair. This leaves fewer than $\zeta/2$ unconnected pairs, so $\langle\mathbf{0}|\gamma_S^\dagger |i\rangle = 0$. Hence, every term in the sum evaluates to zero, and $b_{2l} = 0$, as required.

Proving the result for $\frac{\zeta}{2} \leq l \leq n - \frac{\zeta}{2}$ is more involved, and proceeds via explicit calculation. As discussed above, terms in the sum are only nonzero if γ_S contains exactly $\zeta/2$ unconnected pairs. Denote the ζ qubits with value 1 in $|i\rangle$ as $[i] = [i_1, \dots, i_\zeta]$, with $i_\alpha < i_\beta$ if $\alpha < \beta$. Denote the remaining $(n - \zeta)$ qubits with value 0 in $|i\rangle$ as $[j] = [j_1, \dots, j_{n-\zeta}]$. Observe that $[i] \cap [j] = \emptyset$.

The only nonzero terms in the sum correspond to strings with ζ Majorana operators forming $\zeta/2$ unconnected pairs on $[i]$ and $(2l - \zeta)$ Majorana operators forming $l - \frac{\zeta}{2}$ connected pairs on $[j]$. First consider the part of the string acting on $[i]$. The string $\gamma_{i_1} \gamma_{i_2} \dots \gamma_{i_\zeta}$ can be written as a tensor product of nonoverlapping (commuting) Pauli strings:

$$\bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} p_{i_k, i_{k+1}} O_{i_k} \bar{Z}_{(i_k+1)}^{(i_{k+1}-1)} O_{i_{k+1}}. \quad (D32)$$

Note that the qubits acted on by Z are not in the set $[i]$. Next, consider the part of the string acting on $[j]$. There are $\binom{n-\zeta}{l-\frac{\zeta}{2}}$ (paired) locations to place the connected pairs, denote any such choice as $[\bar{j}]$. The ‘connected part’ of the string can be written as $\bigotimes_{k=1}^{l-\frac{\zeta}{2}} \iota Z_{j_k}$, where here we have used $\iota = \sqrt{-1}$ to avoid confusion with the index i .

Then explicitly compute the sum over these Majorana strings (qubits that are not explicitly acted on in the following expressions are implicitly acted on with the identity operator):

$$\begin{aligned}
 \Pi_{2l}(|i\rangle\langle\mathbf{0}|) &= 2^{-n} \sum_{S \in \binom{[2n]}{2l}} \langle\mathbf{0}|\gamma_S^\dagger|i\rangle\gamma_S \\
 &= 2^{-n} \sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} \sum_{\substack{O_{i_1} \in \{X_{i_1}, Y_{i_1}\} \\ O_{i_2} \in \{X_{i_2}, Y_{i_2}\} \\ \dots \\ O_{i_\zeta} \in \{X_{i_\zeta}, Y_{i_\zeta}\}}} \langle\mathbf{0}|\left(\bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} P_{i_k, i_{k+1}} O_{i_k} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} O_{i_{k+1}} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} i Z_{\bar{j}_\eta}\right)^\dagger |i\rangle \left(\bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} P_{i_k, i_{k+1}} O_{i_k} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} O_{i_{k+1}} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} i Z_{\bar{j}_\eta}\right) \\
 &= 2^{-n} \sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} \sum_{\substack{O_{i_1} \in \{X_{i_1}, Y_{i_1}\} \\ O_{i_2} \in \{X_{i_2}, Y_{i_2}\} \\ \dots \\ O_{i_\zeta} \in \{X_{i_\zeta}, Y_{i_\zeta}\}}} \bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} \langle 0_{i_k} \dots 0_{i_{k+1}} \dots 0_{\bar{j}_\eta} | O_{i_k} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} O_{i_{k+1}} Z_{\bar{j}_\eta} | 1_{i_k} 0 \dots 0_{i_{k+1}} 0 \dots 0_{\bar{j}_\eta} \rangle \\
 &\quad \times (O_{i_k} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} O_{i_{k+1}} Z_{\bar{j}_\eta}) \\
 &= 2^{-n} \left(\bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} \sum_{\substack{O_{i_k} \in \{X_{i_k}, Y_{i_k}\} \\ O_{i_{k+1}} \in \{X_{i_{k+1}}, Y_{i_{k+1}}\}}} \langle 0_{i_k} 0_{i_{k+1}} | O_{i_k} O_{i_{k+1}} | 1_{i_k} 1_{i_{k+1}} \rangle O_{i_k} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} O_{i_{k+1}} \right) \left(\sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} Z_{\bar{j}_\eta} \right) \tag{D33}
 \end{aligned}$$

where in the first term of the final line we have swapped the order of the sum and tensor product, such that we are essentially considering $\zeta/2$ copies of states with Hamming weight two. Using the results of Lemma 2, the first term in the final line above evaluates to

$$\bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} 2^2 |1\rangle\langle 0|_{i_k} \otimes |1\rangle\langle 0|_{i_{k+1}} \otimes \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} = 2^\zeta |i\rangle_{[i]} \langle\mathbf{0}|_{[i]} \bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)}, \tag{D34}$$

where we use $|x\rangle_{[i]}$ to denote the qubits of computational basis state $|x\rangle$ corresponding to the set $[i]$. Hence, denoting $|y\rangle$ as the Hamming weight of a computational basis state $|y\rangle$,

$$\begin{aligned}
 \mathcal{P}_{0,\zeta}[\Pi_{2l}(|i\rangle\langle\mathbf{0}|)] &= 2^{\zeta-n} \sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} \mathcal{P}_{0,\zeta} \left[|i\rangle_{[i]} \langle\mathbf{0}|_{[i]} \bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} Z_{\bar{j}_\eta} \right] \\
 &= 2^{\zeta-n} \sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} \sum_{|x\rangle, |y\rangle = |0, \zeta\rangle} |x\rangle\langle x| \left(|i\rangle_{[i]} \langle\mathbf{0}|_{[i]} \bigotimes_{\substack{k=1 \\ k \text{ odd}}}^{\zeta-1} \bar{Z}_{(i_{k+1})}^{(i_{k+1}-1)} \bigotimes_{\eta=1}^{l-\frac{\zeta}{2}} Z_{\bar{j}_\eta} \right) |y\rangle\langle y| \\
 &= 2^{\zeta-n} \sum_{[\bar{j}] \in \binom{[n-\zeta]}{l-\frac{\zeta}{2}}} |i\rangle\langle\mathbf{0}| \\
 &= 2^{\zeta-n} \binom{n-\zeta}{l-\frac{\zeta}{2}} |i\rangle\langle\mathbf{0}|. \tag{D35}
 \end{aligned}$$

Thus for $\frac{\zeta}{2} \leq l \leq n - \frac{\zeta}{2}$, $\mathcal{P}_{0,\zeta}[\Pi_{2l}(|i\rangle\langle\mathbf{0}|)] = 2^{\zeta-n} \binom{n-\zeta}{l-\frac{\zeta}{2}} |i\rangle\langle\mathbf{0}|$, as required. ■

Lemma 2. The sum

$$\sum_{\substack{O_i \in \{X_i, Y_i\} \\ O_j \in \{X_j, Y_j\}}} \langle 0_i 0_j | O_i O_j | 1_i 1_j \rangle O_i \bar{Z}_{i+1}^{j-1} O_j \tag{D36}$$

evaluates to

$$2^2 |1\rangle\langle 0|_i \otimes |1\rangle\langle 0|_j \otimes \bar{Z}_{i+1}^{j-1}, \tag{D37}$$

where $\bar{Z}_a^b := \bigotimes_{k=a}^b Z_k$.

Proof. Explicitly computing the terms in the sum yields

$$(X_i X_j - Y_i Y_j - \iota(X_i Y_j + Y_i X_j)) \bar{Z}_{i+1}^{j-1}. \quad (\text{D38})$$

Using the decomposition $X = (|0\rangle\langle 1| + |1\rangle\langle 0|)$ and $Y = (-\iota|0\rangle\langle 1| + \iota|1\rangle\langle 0|)$ yields

$$2^2 |1\rangle\langle 0|_i \otimes |1\rangle\langle 0|_j \otimes \bar{Z}_{i+1}^{j-1} \quad (\text{D39})$$

as required. \blacksquare

APPENDIX E: ROBUST MATCHGATE SHADOWS

In this section, we formulate the robust Matchgate shadow protocol, providing an alternative derivation from those in Refs. [30,31].

1. Robust shadow protocol

In this part, we review the robust shadow estimation [29], which serves as a basis for robust Matchgate shadows. Equation (C2) is particularly useful when it comes to noise, where the channel is modified by noise as

$$\tilde{\mathcal{M}} = \sum_{\lambda \in R_G} \tilde{f}_\lambda \Pi_\lambda, \quad \tilde{f}_\lambda = \frac{\text{tr}(\mathcal{M}_Z \Lambda \Pi_\lambda)}{\text{tr}(\Pi_\lambda)}, \quad (\text{E1})$$

$$\begin{aligned} \tilde{f}_{2l} &= \frac{\text{Tr}[\mathcal{M}_Z \Lambda \Pi_{2l}]}{\text{Tr}[\Pi_{2l}]} = \binom{2n}{2l}^{-1} \sum_{S \in \binom{[2n]}{2l}} \sum_{b \in \{0,1\}^n} \langle\langle b | \Lambda | \gamma_S \rangle\rangle \langle\langle \gamma_S | b \rangle\rangle \\ &= \binom{2n}{2l}^{-1} \sum_{b \in \{0,1\}^n} \sum_{T \in \binom{[n]}{l}} \left| \frac{(-i)^{|T|}}{\sqrt{2^n}} (-1)^{\sum_{j \in T} b_j} \right|^2 \langle\langle \gamma_{\text{pairs}(T)} | \Lambda | \gamma_{\text{pairs}(T)} \rangle\rangle \\ &= \binom{2n}{2l}^{-1} \sum_{T \in \binom{[n]}{l}} \langle\langle \gamma_{\text{pairs}(T)} | \Lambda | \gamma_{\text{pairs}(T)} \rangle\rangle. \end{aligned} \quad (\text{E2})$$

To get \tilde{f}_{2l} , we need to determine $\langle\langle \gamma_{\text{pairs}(T)} | \Lambda | \gamma_{\text{pairs}(T)} \rangle\rangle$, which corresponds to the diagonal elements of the Liouville representation of noise channel Λ in the $|\gamma_S\rangle\rangle$ basis.

Interestingly, a recent work on randomized benchmarking [46] proposed an efficient way to characterize the noise in the quantum hardware by defining and computing a set of Majorana fidelities λ_k using random Matchgate circuits. Here we reformulate some of their derivations using our established conventions and get more insights into the interpretation and connection between λ_k and \tilde{f}_{2l} .

The protocol consists of multiple rounds with varying parameters $k \in [2n]$ and circuit sequence lengths m . Due to our focus on fermionic quantum simulations, we will only discuss the case when k is even. Each round starts with the preparation of the all-zero state $|0\rangle$. This is followed by m Matchgate unitaries U_{Q_1}, \dots, U_{Q_m} , chosen uniformly and independently

where \tilde{f}_λ is the noisy eigenvalue, Λ denotes the noise channel, and we have assumed the noise is gate-independent, time-stationary, and Markovian (GTM) [29]. Within this assumption, the noise-free expectation values of $\text{tr}(O_i \rho)$ could be retrieved by inverting this noisy channel, if ρ is perfectly prepared on quantum computers. To achieve that, $\{\tilde{f}_\lambda\}$ need to be determined first, and different measurement schemes have been designed for that according to the specific group G used. Such a protocol is known as robust shadow estimation [29].

2. Noisy Matchgate channel

In this section, we provide derivations for the recently introduced robust Matchgate shadows approach [30,31]. Our derivations closely follow previous work on Matchgate randomized benchmarking [46], highlighting the close connection between these two topics.

The robust shadow protocol assumes the quantum noise satisfies the GTM assumption, in which case the measurement channel is only modified in the eigenvalues associated with each projector. The idea is therefore to determine those noisy coefficients \tilde{f}_λ and apply the noisy inverse channel $\tilde{\mathcal{M}}^{-1}$ when computing the expectation values. For the noisy case, we can generalize the derivation for eigenvalues in the noiseless case into

at random. Finally, all qubits are measured in the Z basis. By averaging over many random sequences, we obtain an estimate $\hat{c}_{2l}(m)$ of the weighted average

$$c_{2l}(m) = \frac{1}{|B(2n)|^m} \sum_{\{Q\} \in B(2n)} \sum_{b \in \{0,1\}^n} \alpha_{2l}(b, Q) p(b|Q, m), \quad (\text{E3})$$

where $Q = Q_m \dots Q_1$ and $p(b|Q, m)$ represents the probability of measurement outcome $|b\rangle$. A correlation function α_{2l} has also been defined as

$$\alpha_{2l}(b, Q) = \text{tr}[|b\rangle\langle b| \Pi_{2l} (U_Q \rho_0 U_Q^\dagger)]. \quad (\text{E4})$$

We estimate $p(b|Q, m)$ from the measurement outcomes received from the quantum computer, whereas we compute $\alpha_{2l}(b, Q)$ analytically. Assuming the state preparation

and measurement (SPAM) are perfect, we have

$$\begin{aligned}
c_{2l}(m) &= \frac{1}{|B(2n)|^m} \sum_{\{Q\} \in B(2n)} \sum_{b \in \{0,1\}^n} \alpha_{2l}(b, Q_m \dots Q_1) \langle \langle b | \Lambda \mathcal{U}_{Q_m} \dots \Lambda \mathcal{U}_{Q_1} | \mathbf{0} \rangle \rangle \\
&= \frac{1}{|B(2n)|^m} \sum_{\{Q\} \in B(2n)} \sum_{b \in \{0,1\}^n} \langle \langle b \otimes b | (\Pi_{2l} \otimes \Lambda) \mathcal{U}_{Q_m}^{\otimes 2} (\mathcal{I} \otimes \Lambda) \mathcal{U}_{Q_{m-1}}^{\otimes 2} \dots (\mathcal{I} \otimes \Lambda) \mathcal{U}_{Q_1}^{\otimes 2} | \mathbf{0} \otimes \mathbf{0} \rangle \rangle \\
&= \sum_{b \in \{0,1\}^n} \langle \langle b^{\otimes 2} | (\Pi_{2l} \otimes \Lambda) \left[\sum_{k'=0}^{2n} |\Upsilon_{k'}^{(2)}\rangle \langle \Upsilon_{k'}^{(2)} | (\mathcal{I} \otimes \Lambda) \right]^{m-1} \sum_{k'=0}^{2n} |\Upsilon_{k'}^{(2)}\rangle \langle \Upsilon_{k'}^{(2)} | \mathbf{0}^{\otimes 2} \rangle \rangle \\
&= \sum_{b \in \{0,1\}^n} \langle \langle b^{\otimes 2} | (\mathcal{I} \otimes \Lambda) |\Upsilon_k^{(2)}\rangle \langle \Upsilon_k^{(2)} | (\mathcal{I} \otimes \Lambda) |\Upsilon_k^{(2)}\rangle \rangle^{m-1} \langle \langle \Upsilon_k^{(2)} | \mathbf{0}^{\otimes 2} \rangle \rangle. \tag{E5}
\end{aligned}$$

Reference [46] defined the term $\langle \langle \Upsilon_k^{(2)} | (\mathcal{I} \otimes \Lambda) |\Upsilon_k^{(2)} \rangle \rangle$ as the Majorana fidelities λ_{2l} . We also note that the first term in the last line is exactly what we have derived in Eq. (E2) if we assume SPAM error-free. Therefore we choose $m = 1$ and have

$$\begin{aligned}
c_{2l}(1) &= \sum_{b \in \{0,1\}^n} \langle \langle b^{\otimes 2} | (\mathcal{I} \otimes \Lambda) |\Upsilon_k^{(2)}\rangle \langle \Upsilon_k^{(2)} | \mathbf{0}^{\otimes 2} \rangle \rangle \\
&= \binom{2n}{2l}^{-1} \sum_{b \in \{0,1\}^n} \sum_{S, S' \in \binom{[2n]}{2l}} \langle \langle b | \gamma_S \rangle \rangle \langle \langle b | \Lambda | \gamma_S \rangle \rangle \langle \langle \gamma_{S'} | \mathbf{0} \rangle \rangle^2 \\
&= \frac{1}{2^n} \binom{2n}{2l}^{-1} \binom{n}{l} \sum_{T \in \binom{[n]}{l}} \langle \langle \gamma_{\text{pairs}(T)} | \Lambda | \gamma_{\text{pairs}(T)} \rangle \rangle \\
&= \frac{1}{2^n} \binom{n}{l} \tilde{f}_{2l}. \tag{E6}
\end{aligned}$$

We now have established the relationship between c_{2l} and \tilde{f}_{2l} . Therefore the noisy eigenvalues can be extracted either by fitting $c_{2l}(m)$ or simply measuring $c_{2l}(1)$. For m other than 1, we have

$$c_{2l}(m) = c_{2l}(1) \left[\binom{2n}{2l}^{-1} \sum_{S \in \binom{[2n]}{2l}} \langle \langle \gamma_S | \Lambda | \gamma_S \rangle \rangle \right]^{(m-1)}, \tag{E7}$$

which provides a microscopic origin for λ_{2l} .

In actual hardware experiments, SPAM error is unavoidable. For measurement error, it satisfies the GTM assumption for noisy shadows and therefore it is not a concern. For state preparation error, its effect is absorbed into the $c_{2l}(1)$ prefactor [46] above, by replacing $|\mathbf{0}\rangle\rangle$ with $|\tilde{\mathbf{0}}\rangle\rangle$ in the first equality of Eq. (E5). We further note that the state preparation error in the shadow experiments, i.e., the noise effects in preparing ρ can also be taken into account in the above framework, as long as we could mimic as much as possible the state preparation circuit in determining \tilde{f}_{2l} . This sets the motivation for the revised robust Matchgate shadow protocol introduced in Sec. III A. We note that if the state preparation noise channel Λ' commutes with the twirl \mathcal{M} , i.e., $\Lambda' \circ \mathcal{M} = \mathcal{M} \circ \Lambda'$, our revised scheme would correct the state preparation error perfectly, assuming Λ' satisfies the GTM assumption. However, this condition may not be strictly satisfied in hardware.

3. Variance bound

The variance of $c_{2l}(m)$, defined in Eq. (E3), is important to determine the efficiency of the robust protocol. A detailed derivation can be found in Appendix of Ref. [31,46], assuming no noise.

We numerically test the sampling complexity to determine \tilde{f}_{2l} , as shown in Fig. 7, and find that it scales closely to $O(n^{2.95})$. The complexity comes from bounding the variance relative to f_n , the smallest coefficient among all the eigenvalues in the n -qubit Matchgate channel.

APPENDIX F: IMPLEMENTATION OF MATCHGATE SHADOWS

The Matchgate circuits are constructed using the open source code, see Ref. [77], which requires $O(n^2)$ gates implemented in $O(n)$ depth. Note that the current implementation assumes an even number of particles in the system studied. A generalization into odd-particle cases can be addressed with ancillary qubits, as discussed in Appendix A of Ref. [26].




FIG. 7. Sampling complexity for determining the Matchgate channel coefficients (with 1% error threshold) scales closely to $O(n^{2.95})$ in the system size limit.

TABLE II. IBM Q Hanoi hardware information, with data taken from IBM Quantum Platform [80].

	Oct.1st	Oct.2nd	Oct.14th
physical qubits	[1,2,3,5]	[1,2,3,5]	[6,7,10,12]
T_1 [μ s]	[132, 139, 156, 159]	[185, 114, 129, 130]	[137, 101, 112, 227]
T_2 [μ s]	[145, 266, 35, 199]	[175, 164, 33, 198]	[157, 252, 215, 262]
CNOT error rate (10^{-3})	[3.25, 11.6, 4.71]	[3.20, 8.26, 6.38]	[7.58, 5.58, 6.30]
X error rate (10^{-4})	[4.0, 1.3, 2.1, 1.8]	[3.4, 1.8, 2.0, 2.2]	[2.5, 2.2, 2.4, 1.8]
Readout error rate (10^{-2})	[2.1, 1.7, 1.0, 0.07]	[1.1, 1.2, 1.1, 0.08]	[2.0, 1.2, 0.09, 1.6]

APPENDIX G: QUANTUM DEFECT EMBEDDING THEORY

Quantum embedding theories [78] are frameworks to solve the time-independent Schrödinger equation for a system of electrons by separating the problem into the calculation of the energy levels or density of an active space and those of the remaining environment. Each part of the system is described at the quantum-mechanical level, with the active space being treated with a more accurate and computationally more expensive theoretical method than the environment.

Spin defects, e.g., NV center in diamond, are particular suitable for an embedding description as they can be naturally partitioned into the defect center and the host material. Recently, a Green's function-based quantum embedding theory was proposed for the calculation of defect properties, denoted as quantum defect embedding theory (QDET). We will provide a high-level description of this method in this section, and we direct interested readers to Refs. [51,79] for more details. QDET initiates from a mean-field electronic structure calculation, typically DFT, on a supercell (with hundreds/thousands of atoms) representing the point-defect of interest hosted in a pristine solid. Then an active space is defined by a sub-set of single particle orbitals localized around the defect center, whose excitations are described by an effective Hamiltonian H_{eff} . The effective potential entering H_{eff} is evaluated by computing the effect of the environment onto the active space with many-body perturbation theory techniques [51]. The Hamiltonian can be solved either using a full configuration interaction (FCI) approach on classical computers, or using quantum algorithms on quantum computers.

In essence, using QDET allows one to reduce the complexity of evaluating many-body states of a small guest region embedded in a large host system: the problem is reduced to diagonalizing a many-body Hamiltonian simply defined on an active space, where the number of degrees of freedom is much smaller than that required to describe the entire supercell of hundreds of atoms.

APPENDIX H: QUANTUM SIMULATION DETAILS

In this section, we provide detailed information on quantum simulations performed in this work, on both quantum hardware platforms and on simulators that can emulate the effects of noise.

1. IBM Q experiments for hydrogen

We study the hydrogen molecule with five different bond distances, ranging from 0.75 to 2.75 Å using the minimal

STO-3G basis set. Restricted HF orbitals are used as inputs for subsequent VQE calculation on a noisy quantum simulator (to obtain the quantum trial states) and QC-QMC calculation on quantum hardware. We used 4 qubits on IBM Quantum Hanoi superconducting qubit device, allocated based on the Jordan-Wigner mapping of spin-orbitals to qubits.

a. Quantum trial state

As we introduced in the main text, the quantum trial state adopted in the QC-QMC scheme has to satisfy two requirements: (i) $V_T|\Psi_T\rangle = |\Psi_T\rangle$ and (ii) $V_T|\mathbf{0}\rangle = |\mathbf{0}\rangle$. This implies that the trial circuit should conserve the particle number of the initial states at least in the $\{0, \zeta\}$ -particle subspace. We therefore chose the UCCSD ansatz. The issue with conventional UCCSD ansatz is that the circuit is too deep, with a $O(n^4)$ two-qubit gate count, and this is usually beyond the capabilities of NISQ hardware. We, therefore, designed an alternative circuit, shown in Fig. 8, that constructs effectively a double excitation. Such a circuit assumes only linear connectivity of the underlying hardware and does not require any swap of the physical qubits. It is worth re-emphasizing that this tailored UCCSD ansatz was chosen specifically for the hydrogen molecule, to be sufficient for enabling our QC-QMC experiments, and does not represent a universal solution. We leave the important task for constructing a general and efficient quantum trial state for future investigations.

b. Hardware information

The information of the physical qubits used on IBM Q Hanoi device is recorded in Table II.




FIG. 8. Quantum circuit for the performing Matchgate shadows of hydrogen. The first Hadamard and CNOT gate generates an equal superposition of $|0\rangle$ and $|\Psi_I\rangle$. They are followed by the effective UCCSD ansatz, and the Matchgate circuit U_Q for shadow tomography.

TABLE III. Data of the noisy eigenvalues of the four-qubit Matchgate channel, with experiments performed on the IBM Q Hanoi device on Oct 1st, 2nd and 14th. These noisy eigenvalues are used for the robust shadow protocol considering the state preparation error (SP). Specifically, shadows of hydrogen bond distance 0.75 Å are collected on Oct 14th; shadows for 2.25 Å are obtained on Oct 1st while the shadows for the rest three are obtained on Oct 2nd.

IBM Q Hanoi	Oct.1st	Oct.2nd	Oct.14th
physical qubits	[1,2,3,5]	[1,2,3,5]	[6,7,10,12]
\tilde{f}_0	1.0	1.0	1.0
\tilde{f}_2	0.1083	0.1146	0.1030
\tilde{f}_4	0.0629	0.0649	0.0565
\tilde{f}_6	0.0961	0.1042	0.0940
\tilde{f}_8	0.6461	0.7121	0.6417

c. Robust shadow experiments

The noise determination phase (to measure \tilde{f}_{2l}) and standard shadow experiments are performed employing 16 000

circuits each on the same set of physical qubits on the same day, to ensure the noise condition stays the same. Note that the original proposal of shadow tomography [18] assumes using a single shot for each shadow circuit. For better use of the quantum resources, we followed Ref. [48] and all the circuits are executed with 1024 shots. The circuit used for determining \tilde{f}_{2l} only differs in the absence of the first Hadamard gate from the Matchgate shadow circuits used to compute the overlap. This construction ensures that the noise in preparing ρ for QC-AFQMC algorithm is captured as much as possible in computing f_{2l} . The measured values of \tilde{f}_{2l} are summarized in Table III.

Having determined the values of \tilde{f}_{2l} , one can proceed to apply the robust Matchgate shadows procedure to overlaps, and overlap ratios. In the main text, we plotted the mean absolute error (MAE) of 120 overlap ratios with respect to the number of Matchgate shadow circuits used, and observed that the robust protocol was able to mitigate the effects of noise. Here, we plot the MAE of 36 separate overlap ratios, randomly chosen from the total of 120 tested in the main text,

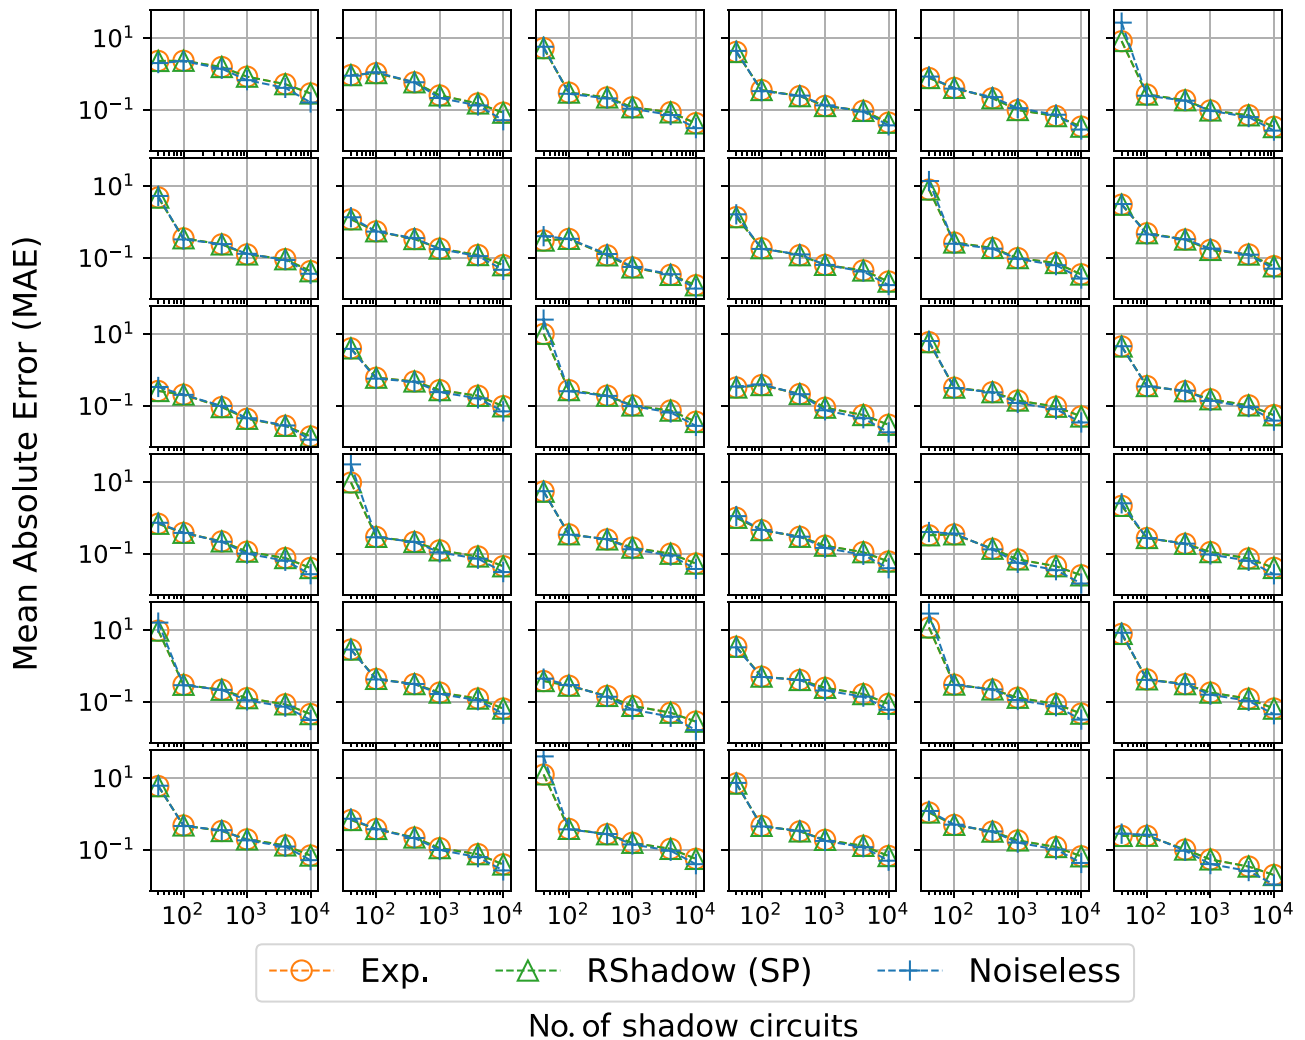


FIG. 9. The mean absolute error (MAE) of 36 overlap ratios randomly selected from the tested 120 overlap ratios' pool. For each overlap ratio, noisy results from IBM Hanoi, w/o the robust shadow correction (orange circle/green triangle) are all in good agreement with the noiseless reference value (blue cross). This verifies the noise resilience discussed in the main text.




FIG. 10. Quantum circuit for the performing Matchgate shadows of NV center. The first Hadamard and CNOT gate generates an equal superposition of $|0\rangle$ and $|\Psi_I\rangle$. They are followed by the effective UCCSD ansatz, and the Matchgate circuit U_Q for shadow tomography.

shown in Fig. 9. We can conclude that the noise resilience is present for each trial, and is not an artifact of averaging.

d. Classical postprocessing

As discussed in the main text, in the interest of conserving computational resources, only two of the five data points (solid points) in Fig. 4 were obtained using the scalable Matchgate shadows approach. The remaining three data points were obtained using an exponentially scaling approach used in Ref. [17] that is ultimately more efficient than the scalable Matchgate approach for small system sizes. We verified that the two schemes give the same results for overlap amplitudes. This exponentially scaling approach works by first computing and tabulating the local quantities, i.e., overlap amplitudes, force bias and local energies, for all of the canonical computational basis states (of which there are 4 for hydrogen). The local quantities of arbitrary walker states can then be computed by decomposing the Slater determinant into a linear combination of the four basis states of hydrogen, and then taking the corresponding linear combination of the tabulated data. For small molecules like hydrogen, this yields a very efficient simulation. However, for larger molecules an arbitrary Slater determinant would decompose into a number of canonical computational basis states scaling superpolynomially with n , rendering this method unscalable.

2. IonQ experiments for NV center

The effective Hamiltonian describing the strongly correlated electronic states of the NV center is generated by QDET.

The detailed information of classical calculations is documented in Ref. [52]. The NV center in diamond is simulated using four qubits on the IonQ Aria trapped ion quantum computer.

a. Quantum trial state

Since the ground state of NV center is an equal superposition of two Slater determinants [53], we can generalize the quantum trial state we used for the hydrogen molecule and apply it to the NV case. The circuit is shown in Fig. 10.

b. Hardware information

The IonQ Aria device has an average of 99.97% single-qubit gate fidelity and 98.6% two-qubit gate fidelity, with all-to-all qubit connectivity. The average readout fidelity is 99.55%. Detailed information about the device can be found on the IonQ website [81].

3. Noisy simulations of the water molecule

The water molecule is a more strongly correlated problem than hydrogen, with more than one correlated pair of electrons. We simulated it at equilibrium geometry with a (4e, 4o) active space based on the restricted Hartree Fock canonical orbitals, following Ref. [54], using eight qubits on the PennyLane quantum simulator [82] with various noise models. The eight qubits are associated with each orbital according to ascending order in energy with spin-up and down alternations.

a. Quantum trial state

Since the water molecule has a perfect pairing of electrons in the ground state, we consider a parameterized quantum circuit employing a four-qubit double-excitation gate $U_{DE}(\theta)$:

$$U_{DE}(\theta)|1100\rangle = \cos\left(\frac{\theta}{2}\right)|1100\rangle - \sin\left(\frac{\theta}{2}\right)|0011\rangle. \quad (\text{H1})$$

This gate is implemented between qubit indexed [0, 1, 4, 5], [2, 3, 6, 7], [1, 2, 7, 4], [0, 3, 6, 5]. These indices are chosen because they construct those Slater determinants with the largest coefficients from a CASCI calculation. This ansatz is then fed to a (noiseless) VQE calculation to produce a wave function with an energy error of approximately 10 mHa at the equilibrium geometry. This quantum trial state is only used for demonstrating the noise resilience in evaluating the overlap ratios using noise models, as the full QC-QMC calculations for the water molecule would be too costly to carry out.

- [1] X. Lin, X. Li, and X. Lin, A review on applications of computational methods in drug screening and design, *Molecules* **25**, 1375 (2020).
- [2] N. P. King, W. Sheffler, M. R. Sawaya, B. S. Vollmar, J. P. Sumida, I. André, T. Gonen, T. O. Yeates, and D. Baker, Computational design of self-assembling protein nanomaterials with atomic level accuracy, *Science* **336**, 1171 (2012).

- [3] S. Hammes-Schiffer and G. Galli, Integration of theory and experiment in the modelling of heterogeneous electrocatalysis, *Nature Energy* **6**, 700 (2021).
- [4] W. Kohn, A. D. Becke, and R. G. Parr, Density functional theory of electronic structure, *J. Phys. Chem.* **100**, 12974 (1996).
- [5] T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, 2013).

- [6] W. Foulkes, L. Mitás, R. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
- [7] B. M. Austin, D. Y. Zubarev, and W. A. Lester, Jr., Quantum Monte Carlo and related approaches, *Chem. Rev.* **112**, 263 (2012).
- [8] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [9] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
- [10] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [11] M. Motta, C. Sun, A. T. Tan, M. J. O’Rourke, E. Ye, A. J. Minnich, F. G. Brandao, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [12] J. P. Misiewicz and F. A. Evangelista, Implementation of the projective quantum eigensolver on a quantum computer, *J. Phys. Chem. A* **128**, 2220 (2024).
- [13] U. Baek, D. Hait, J. Shee, O. Leimkuhler, W. J. Huggins, T. F. Stetina, M. Head-Gordon, and K. B. Whaley, Say no to optimization: A nonorthogonal quantum eigensolver, *PRX Quantum* **4**, 030307 (2023).
- [14] S. E. Smart and D. A. Mazziotti, Accelerated convergence of contracted quantum eigensolvers through a quasi-second-order, locally parameterized optimization, *J. Chem. Theory Comput.* **18**, 5286 (2022).
- [15] H. R. Grimsley, G. S. Barron, E. Barnes, S. E. Economou, and N. J. Mayhall, Adaptive, problem-tailored variational quantum eigensolver mitigates rough parameter landscapes and barren plateaus, *npj Quantum Inf.* **9**, 19 (2023).
- [16] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nat. Commun.* **12**, 6961 (2021).
- [17] W. J. Huggins, B. A. O’Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, Unbiasing fermionic quantum Monte Carlo with a quantum computer, *Nature (London)* **603**, 416 (2022).
- [18] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [19] X. Xu and Y. Li, Quantum-assisted Monte Carlo algorithms for fermions, *Quantum* **7**, 1072 (2023).
- [20] Y. Zhang, Y. Huang, J. Sun, D. Lv, and X. Yuan, Quantum computing quantum Monte Carlo, [arXiv:2206.10431](https://arxiv.org/abs/2206.10431).
- [21] A. Montanaro and S. Stanisic, Accelerating variational quantum Monte Carlo using the variational quantum eigensolver, [arXiv:2307.07719](https://arxiv.org/abs/2307.07719).
- [22] G. Mazzola, Quantum computing for chemistry and physics applications from a Monte Carlo perspective, *J. Chem. Phys.* **160**, 010901 (2024).
- [23] S. Kanno, H. Nakamura, T. Kobayashi, S. Gocho, M. Hatanaka, N. Yamamoto, and Q. Gao, Quantum computing quantum Monte Carlo with hybrid tensor network for electronic structure calculations, *npj Quantum Inf.* **10**, 56 (2024).
- [24] M. Amsler, P. Deglmann, M. Degroote, M. P. Kaicher, M. Kiser, M. Kühn, C. Kumar, A. Maier, G. Samsonidze, A. Schroeder *et al.*, Classical and quantum trial wave functions in auxiliary-field quantum Monte Carlo applied to oxygen allotropes and a CuBr₂ model system, *J. Chem. Phys.* **159**, 044119 (2023).
- [25] M. Kiser, A. Schroeder, G.-L. R. Anselmetti, C. Kumar, N. Moll, M. Streif, and D. Vodola, Classical and quantum cost of measurement strategies for quantum-enhanced auxiliary field quantum Monte Carlo, *New J. Phys.* **26**, 033022 (2024).
- [26] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Matchgate shadows for fermionic quantum simulation, *Commun. Math. Phys.* **404**, 629 (2023).
- [27] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [28] M. Scheurer, G.-L. R. Anselmetti, O. Oumarou, C. Gogolin, and N. C. Rubin, Tailored and externally corrected coupled cluster with quantum inputs, *J. Chem. Theory Comput.* **20**, 5068 (2024).
- [29] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, *PRX Quantum* **2**, 030348 (2021).
- [30] A. Zhao and A. Miyake, Group-theoretic error mitigation enabled by classical shadows and symmetries, *npj Quantum Inf.* **10**, 57 (2024).
- [31] B. Wu and D. E. Koh, Error-mitigated fermionic classical shadows on noisy quantum devices, *npj Quantum Inf.* **10**, 39 (2024).
- [32] D. E. Koh and S. Grewal, Classical shadows with noise, *Quantum* **6**, 776 (2022).
- [33] M. Motta and S. Zhang, Ab initio computations of molecular systems by the auxiliary-field quantum Monte Carlo method, *WIREs Comput. Mol. Sci.* **8**, e1364 (2018).
- [34] S. Zhang, Auxiliary-field quantum Monte Carlo for correlated electron systems, in *Emergent Phenomena in Correlated Matter*, Lecture Notes of the Autumn School Correlated Electrons (Forschungszentrum Jülich, Jülich, 2013), Vol. 3, Chap. 15.
- [35] A. Mahajan and S. Sharma, Taming the sign problem in auxiliary-field quantum Monte Carlo using accurate wave functions, *J. Chem. Theory Comput.* **17**, 4786 (2021).
- [36] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [37] A. K. Ekert, C. M. Alves, D. K. L. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek, Direct estimations of linear and nonlinear functionals of a quantum state, *Phys. Rev. Lett.* **88**, 217901 (2002).
- [38] G. Mazzola and G. Carleo, Exponential challenges in unbiasing quantum Monte Carlo algorithms with quantum computers, [arXiv:2205.09203](https://arxiv.org/abs/2205.09203).
- [39] J. Lee, D. R. Reichman, R. Babbush, N. C. Rubin, F. D. Malone, B. O’Gorman, and W. J. Huggins, Response to “Exponential challenges in unbiasing quantum Monte Carlo algorithms with quantum computers”, [arXiv:2207.13776](https://arxiv.org/abs/2207.13776).
- [40] J. Helsen and M. Walter, Thrifty shadow estimation: Reusing quantum circuits and bounding tails, *Phys. Rev. Lett.* **131**, 240602 (2023).
- [41] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [42] E. Tang, A quantum-inspired classical algorithm for recommendation systems, in *Proceedings of the 51st Annual ACM SIGACT*

- Symposium on Theory of Computing* (ACM, New York, 2019), pp. 217–228.
- [43] G. H. Low, Classical shadows of fermions with particle number symmetry, [arXiv:2208.08964](https://arxiv.org/abs/2208.08964).
- [44] B. O’Gorman, Fermionic tomography and learning, [arXiv:2207.14787](https://arxiv.org/abs/2207.14787).
- [45] P. Jordan and E. Wigner, Über das paulische äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [46] J. Helsen, S. Nezami, M. Reagor, and M. Walter, Matchgate benchmarking: Scalable benchmarking of a continuous family of many-qubit gates, *Quantum* **6**, 657 (2022).
- [47] T. E. O’Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, Error mitigation via verified phase estimation, *PRX Quantum* **2**, 020317 (2021).
- [48] Y. Zhou and Q. Liu, Performance analysis of multi-shot shadow estimation, *Quantum* **7**, 1044 (2023).
- [49] S. Hsieh, C. Zu, T. Mittiga, T. Smart, F. Machado, B. Kobrin, T. Höhn, N. Rui *et al.*, Imaging stress and magnetism at high pressures using a nanoscale quantum sensor, *Science* **366**, 1349 (2019).
- [50] P. Bhattacharyya, W. Chen, X. Huang, S. Chatterjee, B. Huang, B. Kobrin, Y. Lyu, T. Smart, M. Block, E. Wang *et al.*, Imaging the meissner effect in hydride superconductors using quantum sensors, *Nature (London)* **627**, 73 (2024).
- [51] N. Sheng, C. Vorwerk, M. Govoni, and G. Galli, Green’s function formulation of quantum defect embedding theory, *J. Chem. Theory Comput.* **18**, 3512 (2022).
- [52] B. Huang, N. Sheng, M. Govoni, and G. Galli, Quantum simulations of fermionic hamiltonians with efficient encoding and ansatz schemes, *J. Chem. Theory Comput.* **19**, 1487 (2023).
- [53] B. Huang, M. Govoni, and G. Galli, Simulating the electronic structure of spin defects on quantum computers, *PRX Quantum* **3**, 010339 (2022).
- [54] 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).
- [55] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush *et al.*, Quantum chemistry calculations on a trapped-ion quantum simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [56] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu *et al.*, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, *Nat. Commun.* **14**, 1952 (2023).
- [57] H. Q. Pham, R. Ouyang, and D. Lv, Scalable quantum Monte Carlo with direct-product trial wave functions, *J. Chem. Theory Comput.* **20**, 3524 (2024).
- [58] Y. Chen, L. Zhang, W. E, and R. Car, Hybrid auxiliary field quantum Monte Carlo for molecular systems, *J. Chem. Theory Comput.* **19**, 4484 (2023).
- [59] J. L. Weber, H. Vuong, R. A. Friesner, and D. R. Reichman, Expanding the design space of constraints in auxiliary-field quantum Monte Carlo, *J. Chem. Theory Comput.* **19**, 7567 (2023).
- [60] J. Shee, J. L. Weber, D. R. Reichman, R. A. Friesner, and S. Zhang, On the potentially transformative role of auxiliary-field quantum Monte Carlo in quantum chemistry: A highly accurate method for transition metals and beyond, *J. Chem. Phys.* **158**, 140901 (2023).
- [61] A. Mahajan and S. Sharma, Efficient local energy evaluation for multi-slater wave functions in orbital space quantum Monte Carlo, *J. Chem. Phys.* **153**, 194108 (2020).
- [62] T. Jiang, B. O’Gorman, A. Mahajan, and J. Lee, Unbiasing fermionic auxiliary-field quantum Monte Carlo with matrix product state trial wavefunctions, [arXiv:2405.05440](https://arxiv.org/abs/2405.05440).
- [63] B. Posey, C. Gropp, B. Wilson, B. McGeachie, S. Padhi, A. Herzog, and A. Apon, Addressing the challenges of executing a massive computational cluster in the cloud, in *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)* (IEEE, Piscataway, NJ, 2018), pp. 253–262.
- [64] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Duskocil, C. M. Abuan, and J. Romero, Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis, *Phys. Rev. Res.* **4**, 033154 (2022).
- [65] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [66] J. Goings, L. Zhao, J. Jakowski, T. Morris, and R. Pooser, Molecular symmetry in vqe: A dual approach for trapped-ion simulations of benzene, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2023), Vol. 2, pp. 76–82.
- [67] R. J. Anderson, T. Shiozaki, and G. H. Booth, Efficient and stochastic multireference perturbation theory for large active spaces within a full configuration interaction quantum Monte Carlo framework, *J. Chem. Phys.* **152**, 054101 (2020).
- [68] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, *Proc. Natl. Acad. Sci. USA* **114**, 7555 (2017).
- [69] J. Hubbard, Calculation of partition functions, *Phys. Rev. Lett.* **3**, 77 (1959).
- [70] R. Stratonovich, On a method of calculating quantum distribution functions, *Soviet Physics Doklady* **2**, 416 (1957).
- [71] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [72] M. Suzuki, Relationship between d-dimensional quantal spin systems and (d+1)-dimensional ising systems: Equivalence, critical exponents and systematic approximants of the partition function and spin correlations, *Prog. Theor. Phys.* **56**, 1454 (1976).
- [73] D. J. Thouless, Stability conditions and nuclear rotations in the hartree-fock theory, *Nucl. Phys.* **21**, 225 (1960).
- [74] D. Thouless, Vibrational states of nuclei in the random phase approximation, *Nucl. Phys.* **22**, 78 (1961).
- [75] H. Zhu, Multiqubit clifford groups are unitary 3-designs, *Phys. Rev. A* **96**, 062336 (2017).
- [76] A. Hashim, R. K. Naik, A. Morvan, J.-L. Ville, B. Mitchell, J. M. Kreikebaum, M. Davis, E. Smith, C. Iancu, K. P. O’Brien, I. Hincks, J. J. Wallman, J. Emerson, and I. Siddiqi, Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor, *Phys. Rev. X* **11**, 041039 (2021).
- [77] A. Zhao, Fermionic classical shadows, <https://github.com/zhao-andrew/symmetry-adjusted-classical-shadows>.

- [78] Q. Sun and G. K.-L. Chan, Quantum embedding theories, *Acc. Chem. Res.* **49**, 2705 (2016).
- [79] H. Ma, N. Sheng, M. Govoni, and G. Galli, Quantum embedding theory for strongly correlated states in materials, *J. Chem. Theory Comput.* **17**, 2116 (2021).
- [80] <https://quantum.ibm.com>.
- [81] <https://ionq.com>.
- [82] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968).