

CHEMISTRY

Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer

Javier Robledo-Moreno^{1*}, Mario Motta^{1*}, Holger Haas¹, Ali Javadi-Abhari¹, Petar Jurcevic¹, William Kirby², Simon Martiel³, Kunal Sharma¹, Sandeep Sharma⁴, Tomonori Shirakawa^{5,6,7}, Iskandar Siddikov¹, Rong-Yang Sun^{5,6,7}, Kevin J. Sung¹, Maika Takita¹, Minh C. Tran², Seiji Yunoki^{5,6,7,8}, Antonio Mezzacapo^{1*}

A universal quantum computer can simulate diverse quantum systems, with electronic structure for chemistry offering challenging problems for practical use cases around the hundred-qubit mark. Although current quantum processors have reached this size, deep circuits and a large number of measurements lead to prohibitive runtimes for quantum computers in isolation. Here, we demonstrate the use of classical distributed computing to offload all but an intrinsically quantum component of a workflow for electronic structure simulations. Using a Heron superconducting processor and the supercomputer Fugaku, we simulate the ground-state dissociation of N₂ and the ground state properties of [2Fe-2S] and [4Fe-4S] clusters, with circuits up to 77 qubits and 10,570 gates. The proposed algorithm processes quantum samples to produce upper bounds for the ground-state energy and sparse approximations to the ground-state wave functions. Our results suggest that, for current error rates, a quantum-centric supercomputing architecture can tackle challenging chemistry problems beyond sizes amenable to exact diagonalization.

Copyright © 2025 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution License 4.0 (CC BY).

INTRODUCTION

The most common task in theoretical quantum chemistry is the computation of ground-state energies by solving the Schrödinger equation $H|\Psi\rangle = E|\Psi\rangle$ in the Born-Oppenheimer approximation. Exact numerical solutions in a finite basis set have a cost growing combinatorially in the number of electrons and orbitals. This limits exact diagonalization in the full configuration interaction (FCI) to system sizes close to 22 electrons in 22 orbitals (22e,22o) (1) and (26e,23o) (2). For system sizes beyond the reach of FCI, one must rely on approximate methods, e.g., diagrammatic techniques, wave function ansatzes, and Monte Carlo integration (3, 4).

Progress in quantum computing has triggered a flurry of theoretical proposals for computational chemistry over the past decade [e.g., (5–7)]. At the same time, attempts have been made at implementations on prefault-tolerant quantum processors (8–14), but these have so far been limited to small systems for two main reasons. First, despite numerous efforts to improve on the measurement problem [e.g., (15–17)], runtime for energy expectation value estimation on interesting systems remains out of any reasonable timescale. Second, the depths of chemically motivated quantum circuits for computations of chemistry are very high. For unitary coupled cluster (18) and a single step of time evolution, these quantities scale as M^4 (19) on a system with M spin-orbitals. Although this scaling

can be improved with various techniques (20), on prefault-tolerant devices, the signal emerging from circuits of such size is weakened by the accumulation of gate errors and qubit decoherence.

Here, we show that a quantum-centric supercomputing architecture and workflow—which we call sample-based quantum diagonalization (SQD)—allow us to tackle realistic electronic structure problems on system sizes beyond the reach of exact diagonalization on prefault-tolerant quantum processors. We conduct quantum experiments to study the ground-state properties of the N₂ molecule and the [2Fe-2S] and [4Fe-4S] clusters using 58, 45, and 77 qubits, respectively, and a maximum number of 3.5 K two-qubit gates.

The manuscript is structured as follows. In the Results section, we provide a brief description of the problem statement, the concerted quantum-classical workflow, and the configuration recovery technique, as well as the quantum circuits run in the experiments. This section ends with the presentation of the experiment results on the ground-state properties of the N₂ molecule in a correlation-consistent basis set and the active spaces of the [2Fe-2S] and [4Fe-4S] clusters. The Discussion section summarizes our findings and examines some conditions for the advantage with SQD or variations thereof. The Materials and Methods section provides detailed explanations on the subspace projection and diagonalization and approximate total spin symmetry restoration, the configuration recovery technique, and experimental details including the construction of the quantum circuits and the mapping into quantum processors.

RESULTS

We set up the discussion of our results by considering the quantum-centric supercomputing architecture (21) schematized in Fig. 1. The architecture enables scaling of computational capacity by leveraging quantum processors for their natural task: executing a limited number of large quantum circuits. We follow the workflow in Fig. 1 to summarize our methods.

¹IBM Quantum, IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA.

²IBM Quantum, IBM Research Cambridge, Cambridge, MA 02142, USA. ³IBM Quantum, IBM France Lab, Orsay, France.

⁴Department of Chemistry, University of Colorado, Boulder, CO 80302, USA. ⁵Computational Materials Science Research Team, RIKEN Center for Computational Science (R-CCS), Kobe, Hyogo 650-0047, Japan.

⁶Quantum Computational Science Research Team, RIKEN Center for Quantum Computing (RQC), Wako, Saitama 351-0198, Japan. ⁷RIKEN Interdisciplinary Theoretical and Mathematical Sciences Program (iTHEMS), Wako, Saitama 351-0198, Japan. ⁸RIKEN Center for Emergent Matter Science (CEMS), Wako, Saitama 351-0198, Japan.

*Corresponding author. Email: j.robledomoreno@ibm.com (J.R.-M.); mario.motta@ibm.com (M.M.); mezzacapo@ibm.com (A.M.)

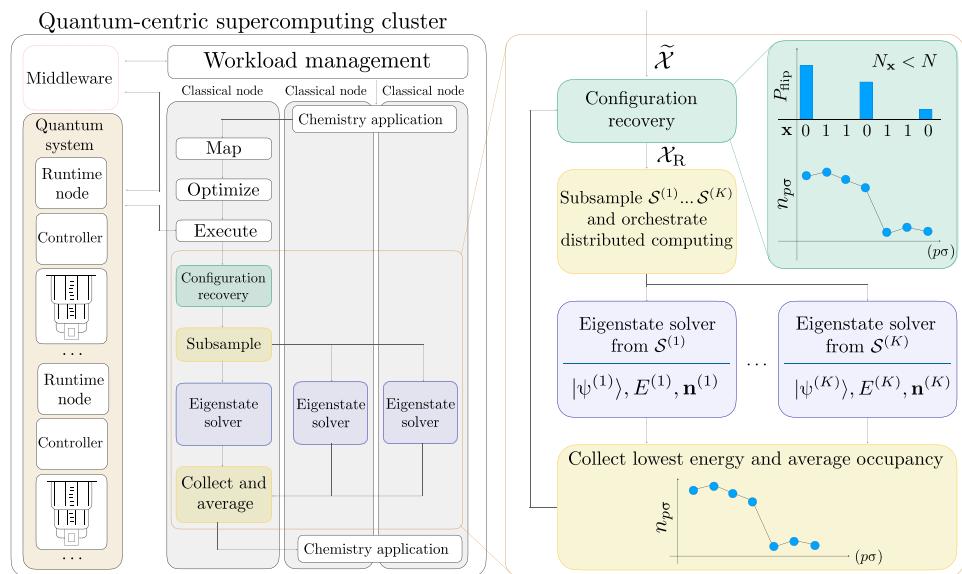


Fig. 1. Quantum-centric supercomputing architecture and SQD workflow diagram. (Left) We illustrate a simplified architecture used to execute our workflow. The architecture has a cluster with a quantum system alongside classical runtime nodes within an isolated environment. A workload management system controls hybrid quantum-classical jobs through middleware. Our workflow is distributed on a set of classical nodes. It includes standard quantum chemistry application routines such as computing electronic integrals, mapping to qubits, and preparing circuits to be executed. (Right) Details of the classical postprocessing step. The input is a set of noisy samples $\tilde{\mathcal{X}}$ from the quantum execution that are processed with our configuration recovery step, using information from a vector \mathbf{n} of reference orbital occupancies. The green inset shows an example where a configuration with $N_x < N$ is corrected. The set of recovered configurations \mathcal{X}_R is subsampled and distributed for projection and diagonalization on parallel classical nodes. A new average reference occupancy vector \mathbf{n} is computed from the results, and the configuration recovery loop is repeated self-consistently until convergence.

Our main goal is to find the ground state of chemistry Hamiltonians

$$\hat{H} = \sum_{pr} h_{pr} \hat{a}_{p\sigma}^\dagger \hat{a}_{r\sigma} + \sum_{prqs} \frac{(pr|qs)}{2} \hat{a}_{p\sigma}^\dagger \hat{a}_{q\tau}^\dagger \hat{a}_{s\tau} \hat{a}_{r\sigma} \quad (1)$$

expanded over a discrete basis set. Here, we have defined the fermionic creation/annihilation operator $\hat{a}_{p\sigma}^\dagger / \hat{a}_{p\sigma}$ associated to the p -th basis set element and the spin σ , whereas h_{pr} and $(pr|qs)$ are the one- and two-body electronic integrals, obtained from standard chemistry software (22). Throughout this manuscript, we use molecular orbitals as basis set elements. We map the degrees of freedom of Eq. 1 to qubits with a Jordan-Wigner (JW) transformation (23). We then construct a quantum circuit to be executed on quantum hardware, preparing a state $|\Psi\rangle$ on M qubits, which represents a molecular wave function on M molecular spin-orbitals. In the JW mapping, the single-qubit basis states $|0\rangle / |1\rangle$ represent empty/occupied spin-orbitals. These mapping and optimization steps are performed on classical nodes (see Fig. 1). We execute the circuit on a quantum computer and measure $|\Psi\rangle$ in the computational basis. Repeating this produces a set of measurement outcomes

$$\tilde{\mathcal{X}} = \{\mathbf{x} | \mathbf{x} \sim \tilde{P}_\Psi(\mathbf{x})\} \quad (2)$$

in the form of bitstrings $\mathbf{x} \in \{0, 1\}^M$ distributed according to some \tilde{P}_Ψ ; the bitstrings represent electronic configurations, also referred to as Slater determinants (SDs).

Configuration recovery

On a prefault-tolerant quantum computer, the action of noise alters the distribution from its ideal form $P_\Psi = |\langle \mathbf{x} | \Psi \rangle|^2$ to some other \tilde{P}_Ψ , which generates the noisy set of configurations $\tilde{\mathcal{X}}$, accessible to us via quantum measurement. Noise in the quantum system broadens the distribution P_Ψ over configurations that do not contribute to low-energy states, so-called deadwood (24). As a result, only a fraction of $\tilde{\mathcal{X}}$ contains a meaningful quantum signal. To improve this scenario, we introduce a self-consistent configuration recovery technique, which allows a probabilistic partial recovery of noiseless configuration samples from $\tilde{\mathcal{X}}$.

The configuration recovery scheme is inspired by the structure of chemistry problems. The Hamiltonian in Eq. 1 conserves the number of particles separately for each spin species. The recovery routine targets configurations \mathbf{x} that have the wrong particle number $N_x \neq N$ due to the accumulation of errors in the execution of the quantum circuit.

Repeated rounds of recovery can be carried out self-consistently. The first step of each recovery round is to iterate through the set $\tilde{\mathcal{X}}$ and find configurations \mathbf{x} with $N_x \neq N$ particles. If $N_x > N$ (or $N_x < N$), $|N_x - N|$ bits are sampled to be flipped from the set of occupied (or empty) spin-orbitals, according to a distribution proportional to a monotonically increasing function (see Materials and Methods section for further information) of $|x_{p\sigma} - n_{p\sigma}|$, the distance from the current value of the bit to the average occupancy of the spin-orbital $p\sigma$, obtained from the previous recovery round. This generates a new set of recovered configurations \mathcal{X}_R .

Following the next step of Fig. 1, we build K batches of d configurations $\mathcal{S}^{(1)} \dots, \mathcal{S}^{(K)}$ using samples from the set \mathcal{X}_R , according

to a distribution proportional to the empirical frequencies of each \mathbf{x} in \mathcal{X}_R . We project and diagonalize the Hamiltonian over each $S^{(k)}, k = 1, \dots, K$, as proposed recently in the quantum selected configuration interaction (SCI) method (25, 26), which draws inspiration from the classical SCI framework (27–33).

Each batch of sampled configurations spans a subspace $S^{(k)}$ in which the many-body Hamiltonian is projected

$$\hat{H}_{S^{(k)}} = \hat{P}_{S^{(k)}} \hat{H} \hat{P}_{S^{(k)}}, \text{ with } \hat{P}_{S^{(k)}} = \sum_{\mathbf{x} \in S^{(k)}} |\mathbf{x}\rangle\langle\mathbf{x}| \quad (3)$$

The ground states and energies of $\hat{H}_{S^{(k)}}$, which we label $|\psi^{(k)}\rangle$ and $E^{(k)}$, are then computed using the iterative Davidson method on multiple classical nodes. The computational cost—both quantum and classical—to produce $|\psi^{(k)}\rangle$ is polynomial in d , the dimension of the subspace.

The ground states are then used to obtain new occupancies

$$n_{p\sigma} = \frac{1}{K} \sum_{1 \leq k \leq K} \langle \psi^{(k)} | \hat{n}_{p\sigma} | \psi^{(k)} \rangle \quad (4)$$

for each spin-orbital tuple $(p\sigma)$, averaged on the K batches. These occupancies are sent back to the configuration recovery step, and this entire self-consistent iteration is repeated until convergence, realizing an SQD of the target Hamiltonian. The initial guess for \mathbf{n} used for the first round of recovery comes from running SQD using the raw quantum samples in the correct particle sector. The configuration recovery routine can be seen effectively as a problem-informed clustering of a noisy signal around the occupations \mathbf{n} . In general, the convergence of the configuration recovery procedure depends on the error rates and the physical properties of the system under consideration. With current error rates, and for the systems in this study, we have observed its convergence within three iteration steps in all systems. We always chose a maximum number of five recovery iterations. We foresee that the lowering of error rates in future quantum hardware will result in faster convergence. Additional details are provided in the Materials and Methods section.

To test the noise robustness, we perform numerical simulations that confirm the improvements of applying the configuration recovery routine to the dissociation of N_2 (6-31G basis set). In this test, we sample from the exact ground state $P_{\Psi_G}(\mathbf{x}) = |\langle \mathbf{x} | \Psi_G \rangle|^2$ and we set a subspace dimension of $d = 10^6$. We use a global depolarizing noise channel to model the effect of noise, $\tilde{P}_{\Psi_G}(\mathbf{x}) = \alpha P_{\Psi_G}(\mathbf{x}) + (1 - \alpha) \frac{1}{2^M}$, with $\alpha \in [0, 1]$ the parameter that controls the amount of noiseless quantum signal. Figure 2 shows the error in the ground-state energy relative to the noiseless case ($\alpha = 1$), as a function of the amount of signal α , for the estimator both with and without configuration recovery. On the N_2 model, errors below $10mE_h$ can be obtained from $\sim 20\%$ signal using the raw noisy samples. However, by using configuration recovery, we can tolerate an $\sim 2\%$ signal to reach the same error. This numerical experiment hints that the use of configuration recovery will be crucial for large-scale experiments.

Quantum circuits

Before presenting our experimental results, we discuss the circuits $|\Psi\rangle$ used to produce the candidate ground states. We use a truncated version of the local unitary cluster Jastrow (LUCJ) ansatz (34), shown in Fig. 3A

$$|\Psi\rangle = \prod_{\mu=1}^L e^{\hat{K}_\mu} e^{\hat{J}_\mu} e^{-\hat{K}_\mu} |\mathbf{x}_{\text{RHF}}\rangle \quad (5)$$

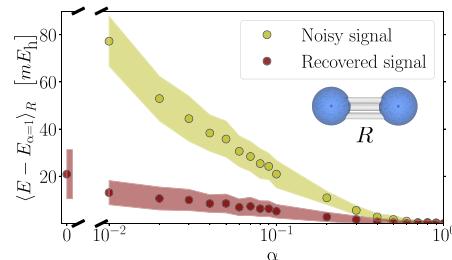


Fig. 2. Self-consistent configuration recovery. Energy error (mean and SD) in the dissociation of N_2 (6-31G), averaged over 24 bond lengths equally spaced between $R = 0.7$ Å and $R = 3.0$ Å, as a function of quantum signal to noise ratio, parameterized by α . The point at $\alpha = 0$ corresponds to sampling from the uniform distribution (no signal). The energies are obtained via projection and diagonalization of $d = 10^6$ raw noisy samples against a subspace of the same size obtained by the configuration recovery routine.

Here, $\hat{K}_\mu = \sum_{pr,\sigma} K_{pr}^\mu \hat{a}_{p\sigma}^\dagger \hat{a}_{r\sigma}$ are generic one-body operators, $\hat{J}_\mu = \sum_{pr,\sigma\tau} J_{pr,\sigma\tau}^\mu \hat{n}_{p\sigma} \hat{n}_{r\tau}$ are density-density operators restricted to spin-orbitals that are mapped onto adjacent qubits (34), and \mathbf{x}_{RHF} is the bitstring representing the restricted Hartree-Fock (RHF) state in the JW mapping. Through this local approximation, the LUCJ ansatz allows for moderate circuit depths. Its accuracy derives from the connection with unitary coupled cluster theory and adiabatic state preparation (34–36). The moderate depths of LUCJ are due to the use of exponentials of one-body operators, implementable in linear depth and a quadratic number of two-qubit gates, and density-density operators, implementable in constant depth and a linear number of ZZ rotations (due to the locality approximation) (34). The LUCJ circuit, compiled into one- and two-qubit gates, is shown in Fig. 3B.

In Fig. 3C, we show a numerical experiment comparing the potential energy curve of N_2 (6-31G basis) obtained by restricted coupled cluster with singles and doubles (CCSD) to one obtained from the LUCJ ansatz, numerically optimized using the subspace energy as the objective function (see Supplementary Materials for further information). The dimension of the diagonalization subspace for different bond lengths ranging from $d = 209764$ to $d = 1340964$, with a median of $d = 563250$. Because of the presence of strong static correlation, CCSD fails in the description of the dissociation curve, whereas the optimized LUCJ ansatz produces a qualitatively correct dissociation curve. We simulated the LUCJ ansatz using the ffsim library (37).

Throughout our experiments, we use the truncated LUCJ circuit $|\Psi\rangle = e^{-\hat{K}_2} e^{\hat{K}_1} e^{i\hat{J}_1} e^{-\hat{K}_1} |\mathbf{x}_{\text{RHF}}\rangle$, which is the result of considering the $L = 2$ circuit and removing the last orbital rotation and Jastrow operations. We parameterize the LUCJ circuits converting the CCSD wave function in Jastrow form and imposing a locality approximation to the resulting J^μ tensors, i.e., zeroing out the components $J_{p\sigma,rt}^\mu$ not corresponding to adjacent qubits (34). For systems where the $J_{p\sigma,rt}^\mu$ parameters obtained from CCSD have small amplitude, the $e^{-\hat{K}_1}$ and $e^{\hat{K}_1}$ terms in the ansatz approximately cancel: In such a situation, without $e^{\hat{K}_2}$, the resulting wave function is overconcentrated around the Hartree-Fock configuration (see the Materials and Methods section for additional information). Although we initialize the

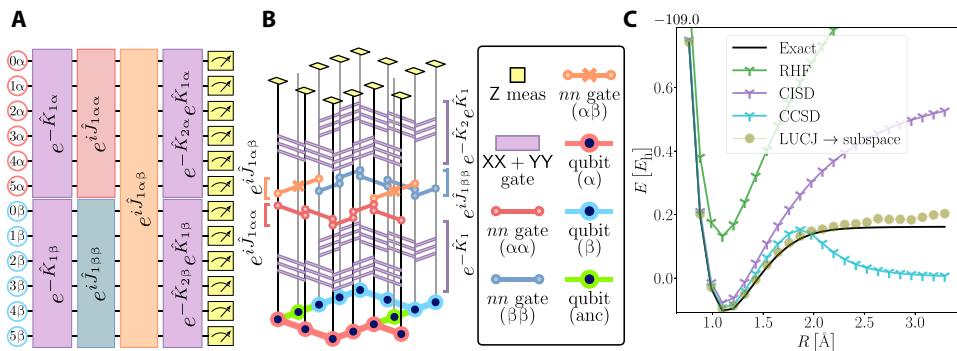


Fig. 3. Quantum circuits for chemistry: LUCJ. (A) Schematic representation of the truncated LUCJ circuit used to generate the set of samples. The circuit is composed of orbital rotation unitaries $\exp(-K_{\mu\alpha})$, same-spin cluster operators $\exp[i(J_{\mu\alpha\alpha} + J_{\mu\beta\beta})]$, and opposite-spin cluster operators $\exp[iJ_{\mu\alpha\beta}]$, with α/β denoting spin-up/down. (B) Single unit of the heavy-hex lattice with a compilation of the LUCJ circuit into single- and two-qubit gates, with the nn gate defined as $U_{nn}(\varphi) = \exp[i\varphi(1 - Z_0)(1 - Z_1)/4]$. (C) Comparison of the potential energy surface of N_2 (with a 6-31G basis) obtained from a noiseless subspace simulation of the LUCJ circuit in (A) with $L = 1$, and optimized parameters, against restricted CCSD, RHF, CISD, and exact energies are shown for reference.

LUCJ circuit with CCSD parameters, the nature of these two methods is very different. In particular, CCSD theory is nonvariational while SQD is. This adds an additional difficulty in the understanding of the relative performance of CCSD and SQD with LUCJ initialized from CCSD parameters. We have performed optimization-free experiments, exploiting the connection between LUCJ and classical coupled cluster theory, yet closing a quantum-classical optimization could further improve the quality of the solutions.

Implementation on a quantum-centric supercomputing platform

In the following, we present the experimental results obtained using the methods discussed so far, on Heron quantum processors and the Fugaku supercomputer. The largest experiment is run on a subset of 77 qubits of a 133-qubit Heron quantum processor. The median fidelities for this subset are 99.77% for two-qubits gates, 99.97% for single-qubit, and readout fidelity of 98.37%, with median coherence times $T_1 = 180\ \mu s$ and $T_2 = 150\ \mu s$. In SQD calculations, it is particularly important to have as many measurement outcomes with correct particle number as possible. However, qubits may be initialized imperfectly, i.e., not in the $|000 \dots 000\rangle$ state, resulting in more measurement outcomes with incorrect particle number. To mitigate this source of error, we use a reset-mitigation scheme by adding an additional measurement instruction before the circuit execution and postselecting outcomes based on this first measurement returning the initial state $|000 \dots 000\rangle$. This postselection results in $\sim 1/3$ retention rate of all the executions, i.e., qubits collapse initialized in the desired state upon the additional measurement with probability $\sim 1/3$.

The classical projection and diagonalizations are obtained with the Davidson method implemented in the library PySCF (22) on a single node, or DICE (30, 33) for distributed computing on multiple nodes. Convergence to the most accurate solution can be obtained in two ways: increasing the accuracy per diagonalization with the subspace size d , and increasing the number of batches K , which will reduce statistical errors in the analysis. For our largest experiment on the [4Fe-4S] cluster, we use up to $d = 100M$, distributing a single projection and diagonalization to 64 nodes of Fugaku, and $K = 100$ batches, for a total of 6400 nodes. We analyze runtime performance as a function of d and K versus the number of nodes used in the

Supplementary Materials. At 64 nodes per diagonalization on the largest experiments, classical runtimes are about 1.5 hours. The largest heat-bath configuration interaction (HCI) calculation that we performed on 16 nodes at $d = 2.3M$ took about 16 min.

We perform two classes of experiments: the breaking of the triple bond of N_2 (cc-pVDZ basis), in the top panel in Fig. 4A, and the ground states of [2Fe-2S] and [4Fe-4S] clusters (active spaces of the TZP-DKH basis), shown in the top panels in Fig. 4 (B and C). We study the ground-state properties of these molecular systems in the $S_z = 0$ and $S^2 = 0$ subspace, where S_z is the total \hat{z} component of the spin and $S^2 = S_x^2 + S_y^2 + S_z^2$. In this work, we used SDs to define the subspaces, which in general are not eigenfunctions of S^2 unlike configuration state functions (CSFs). In the closed-shell systems studied here, a source of spin contamination is the fact that sampled determinants are not closed under the spin inversion operation. Therefore, we achieved an approximate restoration of the S^2 symmetry by extending the set of sampled determinants to ensure closure under spin inversion, as detailed in the Materials and Methods section.

Triple bond breaking in N_2

The breaking of the N_2 bond is a well-known test of the accuracy of electronic structure methods in the presence of static electronic correlation (38, 39). Restricted CCSD theory, a dominant paradigm for the accurate description of weakly correlated systems in quantum chemistry, fails in the description of N_2 dissociation due to static correlation effects: As correlations become stronger, RHF becomes unstable toward a symmetry-broken unrestricted Hartree-Fock (UHF) state. CCSD built from RHF predicts an artificial barrier to binding and overcorrelates at dissociation, whereas CCSD built from a UHF reference dissociates correctly at the cost of spin contamination, a manifestation of Löwdin's symmetry dilemma. We use a correlation-consistent cc-pVDZ basis set, to place emphasis on a theory's ability to treat both dynamic and static correlation in an accurate and balanced manner. We map the N_2 molecule onto a Heron processor as shown in the middle panel of Fig. 4A. We project and diagonalize a Hamiltonian using $d = 16 \times 10^6$ configurations. We consider $K = 10$ batches of configurations, and each point in the dissociation curve has $|\tilde{\mathcal{X}}| = 100 \times 10^3$ measurement outcomes and 10 iterations of recovery. The combined quantum runtime for all points in the dissociation curve is ~ 45 min. The experimental

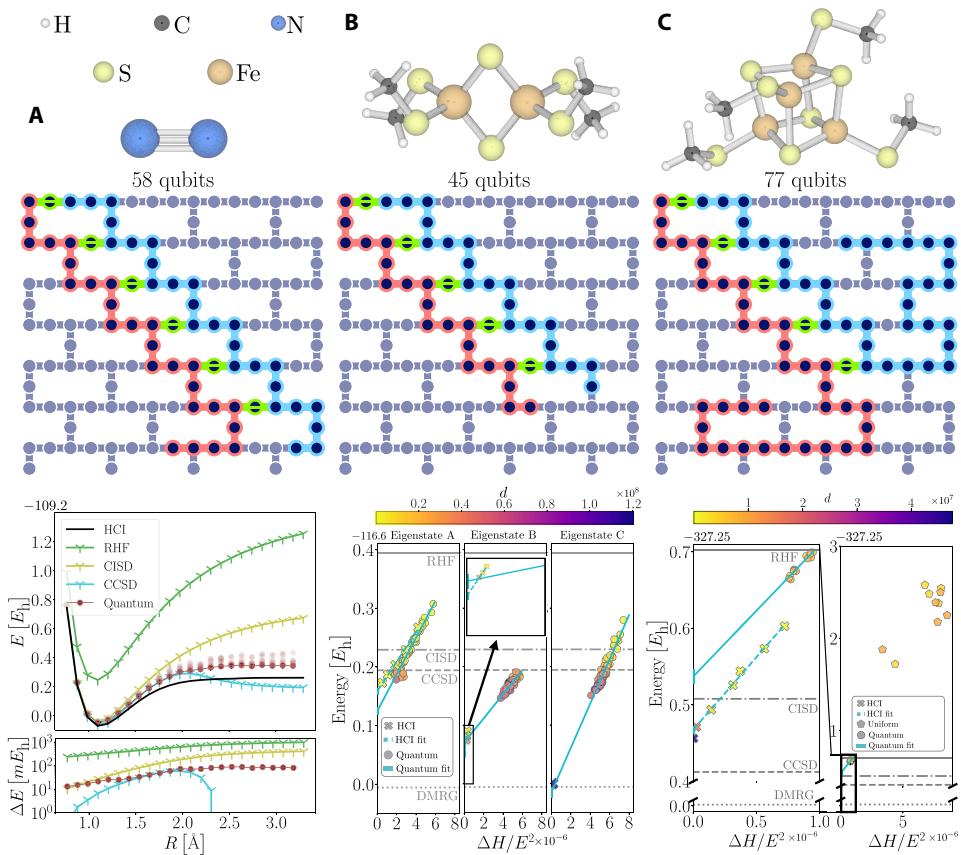


Fig. 4. Experiments: Chemistry on large basis sets. (A) 58 qubits are used to model the N_2 dissociation (cc-pVDZ basis set). (B) 45 qubits are used for the [2Fe-2S] cluster (TZP-DKH basis set) and (C) 77 qubits for the [4Fe-4S] cluster (TZP-DKH basis set). The top panels show a three-dimensional representation of the geometry of each molecule. The middle panels show the qubits selected on a Heron quantum processor layout, following the same color convention as (B) in Fig. 3. The bottom panel in (A) shows the potential energy surface comparison, as well as the energy difference ΔE between the HCl energy and the energies obtained from different methods. The brown scatterplot shows the value of $E^{(k)}$ for all batches of configurations, and the connected dots show $\min_k(E^{(k)})$. The bottom panel in (B) shows the energy-variance analysis for three different eigenstates that both HCl and our method find upon increasing the value of d , as labeled by the color bar. For each approximate eigenstate $|\psi^{(k)}\rangle$, the horizontal axis $\Delta H = \langle\psi^{(k)}|\hat{H}^2|\psi^{(k)}\rangle - \langle\psi^{(k)}|\hat{H}|\psi^{(k)}\rangle^2$. The bottom panel in (C) shows a comparison of the energy-variance analysis applied to quantum measurement outcomes and bitstrings (with the correct particle number) sampled from the uniform distribution. The DMRG energy in (B) and (C) is from (60).

data are reported in the bottom panel of Fig. 4A, showing the potential energy surface of N_2 compared to classical approximate methods. The data from our experiments are consistent with other classical methods except for CCSD, which fails in the description of the dissociation as seen for the smaller basis set considered in Fig. 4B. Among the classical SCI methods, HCl (30) obtains the best results for N_2 and will be our reference classical method in all the other experiments. The difference between our method and HCl energies is everywhere within tens of mE_h . We further analyze the accuracy of our experiments as a function of d and the effect of orbital optimizations in the accuracy of the predictions in the Supplementary Materials. This first test demonstrates that we are capable of addressing multireference ground states and builds confidence for the next set of experiments, which will focus on assessing the ability of the quantum-classical architecture to do precision many-body physics.

[2Fe-2S] cluster: Precision many-body physics

Iron-sulfur (FeS) clusters are molecular ensembles of sulfide-linked 1- to 8-iron centers in variable oxidation states. They are important

cofactors in biological processes ranging from nitrogen fixation to photosynthesis and respiration (40). Their electronic structure, with multiple low-lying states of differing electronic and magnetic character, is responsible for their rich chemistry. At the same time, they pose considerable challenges for experimental studies and numerical tools. For our experiments, we consider the synthetic $[\text{Fe}_2\text{S}_2(\text{SCH}_3)_4]^{2-}$ cluster (41), abbreviated [2Fe-2S] and used in numerical studies to mimic the oxidized dimers prominently found in ferredoxins (42).

The qubit mapping of the LUCJ circuit on the Heron processor for [2Fe-2S] is shown in the middle panel of Fig. 4B. We consider $K = 10$ batches of configurations and $|\tilde{\mathcal{X}}| = 2.4576 \times 10^6$ measurement outcomes. The quantum runtime for this system is ~ 45 min. For the [2Fe-2S] cluster, we perform an energy-variance analysis of the low-energy spectrum of the molecule. The energy-variance analysis is a tool routinely used in classical computational electronic structure to capture the convergence of the approximate eigenstate energy for different levels of accuracy of a computational method (43). Here, we use energy-variance analysis to assess the

convergence as a function of d , which is directly related to quantum and classical accuracy, runtimes, and costs. If one can statistically sample from a good approximation of an eigenstate, points at finite number of samples will be distributed linearly in the energy-variance plane (43). This gives us a tool to detect eigenstates for both quantum and classical methods.

The bottom panel in Fig. 4B shows an energy-variance comparison of HCI and SQD run for different subspace dimensions. As the subspace dimension is increased, three eigenstates can be identified, which we label A, B and, C. The extrapolation of the quantum data to the zero-variance limit is in good agreement with the HCI extrapolations for the same eigenstates. The extrapolated value of the energy for eigenstate C is in good agreement with the density matrix renormalization group (DMRG) calculation from (44).

[4Fe-4S] cluster: A stress test for methodology and quantum processors

The circuits considered for the N₂ experiments and [2Fe-2S] reached sizes of ~1 to 1.5k two-qubit gates. We now test the quality of the signal in noisy circuits that test the limits of Heron processors, using up to 6400 nodes of Fugaku for the classical processing. We consider the synthetic [Fe₄S₄(SCH₃)₄]⁻² cluster (41), abbreviated [4Fe-4S], a representative of nature's cubanes, whose ground-state deduction from experimental measurements was an early success of inorganic spectroscopy (45). The LUCJ circuit used for this molecular species contains ~3.5k two-qubit gates. The qubit mapping on the Heron processor for [4Fe-4S] is shown in the middle panel of Fig. 4C. As in the previous experiment, we consider $K = 10$ batches of configurations and $|\tilde{\mathcal{X}}| = 2.4576 \times 10^6$ measurement outcomes. The quantum runtime for this system is ~45 min. For $d > 250 \times 10^3$, configuration recovery is warm started with the \mathbf{n} obtained from the method at $d = 250 \times 10^3$, and only two iterations are then performed.

This last set of experiments sheds light on the quality of the quantum signal that is passed to the configuration recovery at these large circuit sizes. The bottom right panel in Fig. 4C shows a comparison of the energy-variance analysis from measurement outcomes obtained from the Heron processor and configurations sampled from the uniform distribution. We see that, even if the quantum solutions produced are worse than other classical methods, the energy and variance obtained from quantum data are notably lower than those obtained from uniformly distributed configurations (i.e., pure noise), on subspaces of the same size. This confirms that there is a valuable signal at circuit sizes of ~3.5k two-qubit gates.

DISCUSSION

Significance for quantum computing

Current quantum computers in isolation can perform calculations on systems sufficiently large that exact brute-force classical solutions are not available (46, 47). However, these studies have targeted spin systems, leading to circuits that match the connectivity and the measurement and coherence budgets of the quantum devices. In this work, we present electronic structure calculations on active spaces beyond the scale where FCIs are available. Key to achieve this result is the use of classical and quantum computers in concert to implement the SQD method.

SQD makes economical use of quantum computing resources by drawing samples from a single quantum circuit. Although, in principle, other estimators, such as the standard ones used in variational quantum eigensolvers, have bounded variance for any wave

function, the dire scaling of a number of measurements to estimate energies makes them impractical for the molecules targeted in this work (48). It is also more robust against quantum noise because reconstructing the exact ground-state probability distribution on a quantum computer is not required to get accurate energy approximations, as long as one is sampling relevant configurations (i.e., in the ground-state support).

We have used an LUCJ class of quantum circuits that can reproduce a low-rank decomposition and sparsification of the quantum unitary CCSD (qUCCSD), which allowed us to keep circuit depths manageable (34). Lower error rates on quantum operations will allow us to access deeper quantum circuits with higher connectivity, giving access to more general probability distributions. We have performed optimization-free experiments exploiting the connection between LUCJ and classical coupled cluster theory, for the purpose of assessing accuracy and scalability, yet closing a quantum-classical optimization in future work will further improve the quality of our samples.

Generalization

The SQD method can be applied to simulation tasks other than quantum chemistry, if the target ground-state wave function can be accurately approximated by a sparse vector. Developing quantum circuits with polynomially sized support in the computational basis will be an important element of the generalization of SQD as these circuits are the sources of samples processed by classical computers.

To counter wave function broadening on current quantum hardware, we have used a self-consistent configuration recovery method, exploiting a problem-inspired clustering that leverages the average occupation numbers of the molecular orbitals. We foresee generalizations of our configuration recovery technique to problems other than quantum chemistry that are not informed by the physics of the problem. Conversely, for specific applications, one could use even more information about the problem.

Implications in the search for quantum advantage

Computations can be ranked against three parameters: runtime, energy or cost, and accuracy. Although the first two are often easy to measure, ranking by accuracy is in general not straightforward. For methods that produce upper bounds to the ground-state energy, including SQD, the expectation value of the Hamiltonian defines an unconditional accuracy metric: A lower energy is ranked as a higher quality, all other conditions (e.g., total spin) being equal. Comparing SQD energies, for example, allowed us to benchmark our results against SCI and uniform configuration sampling on a 77-qubit experiment, without access to exact solutions. In addition, the approximate wave functions produced here can be stored in classical memory, which permits a classical prover to certify them, and allows their manipulation by further classical processing.

Using the expectation value of the Hamiltonian as an accuracy metric, one can easily and naturally rank SQD results along with those of variational classical methods, giving the search for this specific form of quantum advantage a quantitative meaning. Because every variational classical method has a specific domain of applicability (49), identifying areas where SQD may offer an accuracy advantage is a delicate problem. For example, variational quantum Monte Carlo methods (50, 51), of paramount importance in many-body physics, are sensitive to the structures of the probability distributions they are modeling, not just to their supports. Similarly,

methods based on tensor networks (52, 53) are successfully used to tackle strongly correlated problems in chemistry, granted the ability to converge their energies with bond dimensions, but convergence can be challenging in some cases (54, 55) because of its computational cost and its sensitivity to the nature and ordering of the basis-set orbitals. Developing an in-depth understanding of SQD through extensive numeric and methodological investigations is necessary to establish or rule out advantage in strictly variational ground-state simulations.

SQD shares with SCI the assumption that the ground state may be approximated by a sparse linear combination of determinants, i.e., with a number of determinants much smaller than the Hilbert space dimension. Note that this assumption does not necessarily underline variational Monte Carlo or tensor networks. Therefore, it is natural to look for conditions to improve over SCI. One such

condition is the existence of a quantum circuit that produces subspaces of better quality, and more efficiently, than classical heuristic selection methods. In the search for ground states, the quality of a subspace can be determined by a lower variational energy. We conduct numerical experiments that suggest that there exist LUCJ circuits whose samples produce subspaces of better quality than the HCI classical selection heuristic. The LUCJ circuit under consideration shares the same depth and connectivity as the circuits used in the experiments. In this study, we consider a particular flavor of SCI, HCI (29), and the [2Fe-2S] cluster (20 orbitals \rightarrow 40+ qubits).

We first perform HCI calculations with different values of the selection cutoff ϵ_1 [see (56) for details on the definition of ϵ_1], as shown in Fig. 5. We observe that the larger and more restrictive values of ϵ_1 do not allow HCI to reach the DMRG energy reference. Instead, it converges to the first excited state (S_1) energy. Decreasing

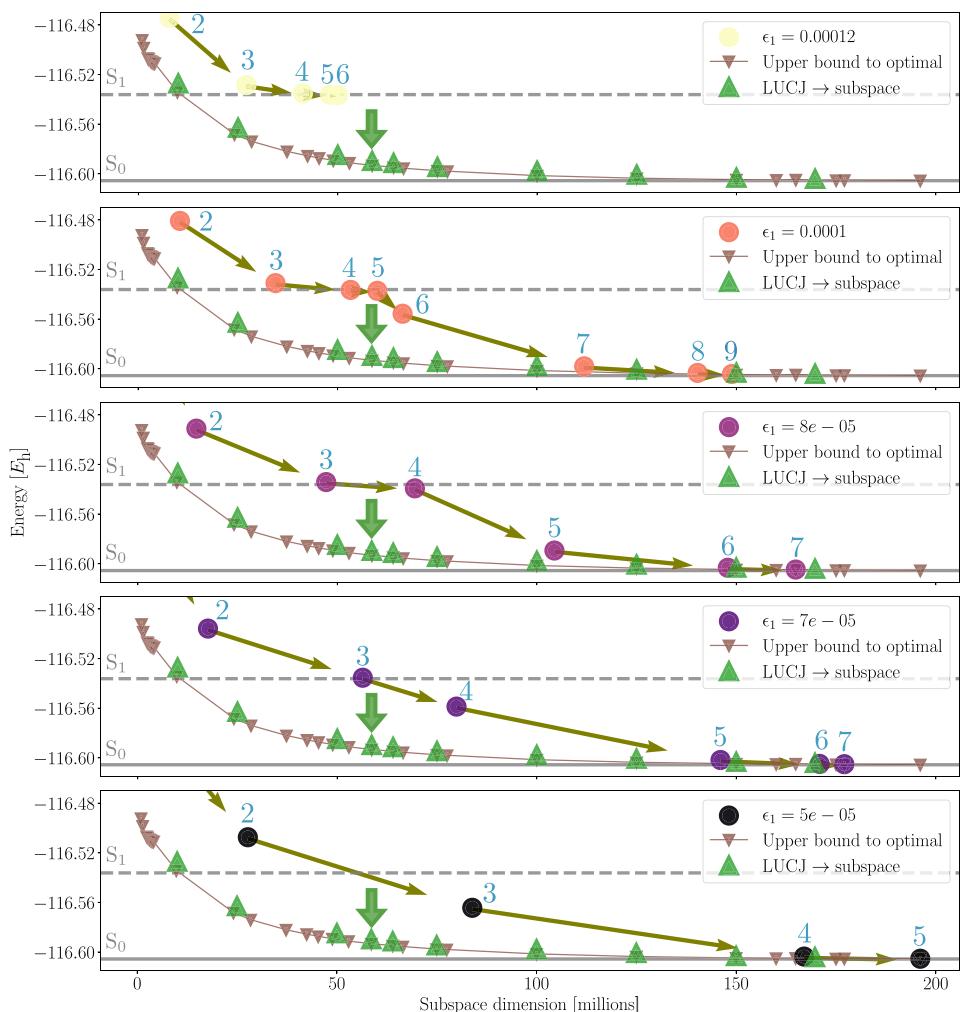


Fig. 5. Comparison of the quality of subspaces generated by HCI and an optimized LUCJ circuit in for the description of the ground state of [2Fe-2S]. Different panels show the energy as a function of subspace dimension of the corresponding diagonalization. The dots connected by arrows show the trajectory of HCI in the energy-subspace dimension plane. Each dot is labeled by the iteration it corresponds to. Different panels correspond to HCI calculations carried out with different values of the selection cutoff ϵ_1 as indicated in the legend. Brown triangles show energies and subspace dimensions obtained after truncating the most accurate HCI wave function by repeatedly removing the electronic configurations corresponding to the lowest wave function amplitudes. The green triangles show the energies for different subspace dimensions obtained from the optimized LUCJ circuit. The green arrows show the subspace dimension that the circuit was optimized for (see main text). The horizontal solid and dashed lines indicate the DMRG estimates for the S_0 and S_1 eigenstates.

the value of ϵ_1 allows the subspace dimension to become larger and improve the energy beyond that first excited state, at the cost of a notable increase in subspace dimension, which results in a higher computational cost. From the largest HCI wave function (last point for $\epsilon_1 = 5 \times 10^{-5}$), we repeatedly remove the electronic configurations whose wave function amplitude is below a given threshold and compute the energy in the resulting subspace, providing a collection of energy-subspace dimension pairs that we label as “upper bound to optimal” in Fig. 5. The removal of tens of millions of configurations (resulting in substantially smaller subspace dimensions) does not substantially deteriorate the quality of the ground-state approximation. These results indicate that HCI does not perform an optimal search of the relevant electronic configurations. Consequently, smaller subspaces exist that yield comparable energy values. In this particular molecule, HCI needs to explore and perform diagonalizations in subspaces larger than the optimal search.

We optimize an LUCJ circuit to produce samples in the identified optimal subspaces with high probability. The optimal parameters are found in a two-step optimization workflow. First, we optimize the Kullback-Leibler divergence between samples drawn from the LUCJ circuit and the amplitudes of 58 M-dimensional ground-state wave function estimation (green arrow in Fig. 5). Then, the circuit parameters are further fine-tuned to minimize the SQD energy using a differential-evolution strategy.

With the samples produced by the optimized LUCJ circuit, we proceed to evaluate the energy of resulting subspaces of varying dimensionality, including the 58 M-dimensional subspace that the circuit was optimized for. We observe good agreement between the LUCJ subspaces and those that are an upper bound to the optimal ones. It is worth noting that, despite being optimized for a 58 M-dimensional subspace, the resulting circuit produces subspaces of outstanding quality of dimensionalities both smaller and larger.

The active space considered for the [2Fe-2S] cluster is small enough for HCI to be able to find the relevant bitstrings by exploring and performing diagonalizations in subspaces notably larger than the optimal ones. This shows that, in some problem instances, the diagonalization itself is not the runtime bottleneck. Instead, the runtime bottleneck is a suboptimal proposal of electronic configurations. Larger and more strongly correlated systems will pose challenges for the classical heuristics. It is for these systems that one can look for a quantum advantage in sampling.

Characterizing the domain of applicability of different classical and quantum heuristics facilitates combining them in quantum-centric supercomputing environments, shifting the search for quantum advantage toward practical problems.

MATERIALS AND METHODS

Conventions and notation

Hamiltonian

Our starting point is the Born-Oppenheimer Hamiltonian, written in second quantization using a basis of N_{MO} orthonormal orbitals $\{\varphi_p\}_{p=1}^{N_{\text{MO}}}$, as shown in Eq. 1. We define the number operator $\hat{n}_{p\sigma} = \hat{a}_{p\sigma}^\dagger \hat{a}_{p\sigma}$, which describes the number of electrons with spin σ on orbital p . For nonrelativistic all-electron calculations, the quantities

$$\begin{aligned} E_0 &= \sum_{a < b} \frac{Z_a Z_b}{\|\mathbf{R}_a - \mathbf{R}_b\|} \\ h_{pr} &= \int d\mathbf{r} \varphi_p^*(\mathbf{r}) \left[-\frac{1}{2} \frac{\partial^2}{\partial \mathbf{r}^2} - \sum_{a=1}^{N_{\text{nuc}}} \frac{Z_a}{\|\mathbf{r} - \mathbf{R}_a\|} \right] \varphi_r(\mathbf{r}) \quad (6) \\ (pr \mid qs) &= \int d\mathbf{r}_1 \int d\mathbf{r}_2 \frac{\varphi_p^*(\mathbf{r}_1) \varphi_r(\mathbf{r}_1) \varphi_q^*(\mathbf{r}_2) \varphi_s(\mathbf{r}_2)}{\|\mathbf{r}_1 - \mathbf{r}_2\|} \end{aligned}$$

describe the internuclear electrostatic interaction energy and the one-electron and two-electron parts of the Hamiltonian, respectively [atomic units are used throughout, i.e., lengths and energies are measured in Bohr and Hartree units $a_B = \hbar^2 / (m_e e^2)$ and $E_h = e^2 / a_B$, respectively, where $-e$ and m_e are the electron charge and mass]. The symbols N_{nuc} , \mathbf{R}_a , and Z_a denote the total number of nuclei and their positions and atomic numbers, respectively.

For relativistic and/or active-space calculations, the indices p, r, q, s label active-space orbitals, and the quantities E_0 , h_{pr} , and $(pr \mid qs)$ are modified to account for relativistic effects and/or the potential generated by the inactive-electron density.

In this work, we use orbitals from a restricted closed-shell Hartree-Fock calculation (also called molecular orbitals and denoted MOs) as the basis functions φ_p . Furthermore, we denote N_σ the number of spin- σ electrons in the exact ground state. Furthermore $N_\uparrow + N_\downarrow = N$, the total number of electrons in the exact ground state.

Molecular species and active spaces

The molecules simulated in this work are listed in Table 1. For N_2 , we studied all non-core electrons and orbitals on quantum hardware. We computed the potential energy curve using (i) RHF at 6-31G, and cc-pVQZ level (57–59) using PySCF (22, 60) and enforcing $D_{\infty h}$ symmetry and, after projecting the nonrelativistic Born-Oppenheimer Hamiltonian in the space spanned by all non-core RHF orbitals with standard functionalities, with (ii) restricted and symmetry-preserving Moller-Plesset second-order perturbation theory (MP2), CCSD,

Table 1. Molecules studied in this work. For each molecule, we list the number of electrons and orbitals (N and N_{MO} , respectively) studied on quantum hardware, along with the underlying basis set and the dimension $D = \binom{N_{\text{MO}}}{N/2}$ of the Hilbert space of $N_\uparrow = N_\downarrow = N/2$ electrons in N_{MO} spatial orbitals (not considering molecular point-group symmetries).

Molecule	Basis	(N, N_{MO})	D
N_2	6-31G	(10e, 160)	1.91×10^7
N_2	cc-pVQZ	(10e, 260)	4.32×10^9
[2Fe-2S]	TZP-DKH	(30e, 200)	2.40×10^8
[4Fe-4S]	TZP-DKH	(54e, 360)	8.86×10^{15}

complete active-space configuration interaction (CASCI), for 6-31G and SCI, in its heat-bath flavor (HCI), for cc-pVDZ.

For [2Fe-2S] and [4Fe-4S], we used active spaces (41, 61), spanned by Fe[3d] and S[3p] orbitals, derived from a localized density functional theory calculation with BP86 functional (44, 62), TZP-DKH basis (63), and sf-X2C (spin-free exact two-component) Hamiltonian (64, 65) to include scalar relativistic effects. We computed approximations to the ground-state energy with restricted RHF, MP2, CCSD, and HCI. The RHF, MP2, CCSD, configuration interaction singles and doubles (CISD), CASCI, and HCI calculations were carried out with the PySCF library (22).

Electron configurations and qubit mapping

In this work, we represent many-electron states using qubit states with the JW transformation, which maps an electronic configuration, i.e., an SD of the form

$$|\mathbf{x}\rangle = \prod_{p\sigma} \left(\hat{a}_{p\sigma}^\dagger \right)^{x_{p\sigma}} |\emptyset\rangle \quad (7)$$

where $|\emptyset\rangle$ is the vacuum state (i.e., the state with zero electrons) and $x_{p\sigma} \in \{0, 1\}$, onto an element of the computational basis

$$|\mathbf{x}\rangle = \bigotimes_{p\sigma} |\mathbf{x}_{p\sigma}\rangle \quad (8)$$

labeled by a bitstring $\mathbf{x} = (x_{N_{\text{MO}}-1\downarrow} \dots x_{0\downarrow} x_{N_{\text{MO}}-1\uparrow} \dots x_{0\uparrow})$. The first half of the bitstring is denoted by $\mathbf{x}_\downarrow = (x_{N_{\text{MO}}-1\downarrow} \dots x_{0\downarrow})$ and the second half of the bitstring is denoted by $\mathbf{x}_\uparrow = (x_{N_{\text{MO}}-1\uparrow} \dots x_{0\uparrow})$. The JW mapping uses $M = 2N_{\text{MO}}$ qubits and allows computing the number of spin- σ electrons for a given configuration \mathbf{x} as $N_{\mathbf{x}\sigma} = \sum_p x_{p\sigma}$. The total number of electrons, $N_{\mathbf{x}} = \sum_\sigma N_{\mathbf{x}\sigma}$, is the Hamming weight of \mathbf{x} . At the single-configuration level, any \mathbf{x} in the right particle sector must satisfy: $N_{\mathbf{x}\sigma} = N_\sigma$. An important example is the RHF bitstring

$$|\mathbf{x}_{\text{RHF}}\rangle = \left| \underbrace{0 \dots 0}_{N_{\text{MO}}-N_1} \underbrace{1 \dots 1}_{N_1} \underbrace{0 \dots 0}_{N_{\text{MO}}-N_1} \underbrace{1 \dots 1}_{N_1} \right\rangle \quad (9)$$

which, by construction, has $N_{\mathbf{x}_{\text{RHF}}\sigma} = N_\sigma$.

The Fock space is the vector space containing all possible electronic configurations for N_{MO} orbitals with all possible filling factors for each spin sector. The terms determinant and electronic configuration are used interchangeably in the manuscript.

Sample-based quantum diagonalization

This section provides a detailed description of the SQD procedure. We describe the subspace projection and diagonalization and the self-consistent configuration recovery scheme. We also motivate the energy-variance analysis presented in Fig. 4.

In what follows, we use \mathcal{X} and $\tilde{\mathcal{X}}$ to denote a set of configurations sampled from probability distributions $P_\Psi(\mathbf{x}) = |\langle \mathbf{x} | \Psi \rangle|^2$ and $\tilde{P}_\Psi(\mathbf{x}) = \langle \mathbf{x} | \tilde{\rho} | \mathbf{x} \rangle$, where $\tilde{\rho}$ is a density operator corresponding to a noisy counterpart of $|\Psi\rangle \langle \Psi|$. We denote $\mathcal{X}_N \subset \tilde{\mathcal{X}}$ the subset of configurations with the right particle number. The set of configurations recovered by the configuration recovery procedure is denoted by $\mathcal{X}_{\rightarrow N}$ and $\mathcal{X}_R = \mathcal{X}_N \cup \mathcal{X}_{\rightarrow N}$ the set of configurations output by self-consistent configuration recovery, and $\mathcal{S}^{(k)} \subset \mathcal{X}_R$ with $k = 1 \dots K$ a set of approximately d configurations sampled from \mathcal{X}_R . The wave function $|\Psi^{(k)}\rangle$ is obtained by projection and diagonalization of \hat{H} in the subspace spanned by the configurations in $\mathcal{S}^{(k)}$.

Eigenstate solver

Given a set of d electronic configurations, the many-electron Hamiltonian is projected and diagonalized in the subspace spanned by the single-particle states defined by the electronic configurations, as proposed recently in (25). We begin by considering a set of configurations $\mathcal{X} = \{\mathbf{x}_{(i)}\}$, all with the right-particle number for each spin sector. Subindices between parenthesis label configurations in a set and not configuration components. We use the generic label \mathcal{X} for a set of configurations with the right particle number. In practice, the configuration recovery procedure will execute the eigenstate solver on the sets \mathcal{X}_N and \mathcal{X}_R .

Conservation of spin. In this study, we perform an approximate restoration of the total spin symmetry labeled by the \hat{S}^2 quantum number. In particular, we are interested on wave functions $|\Psi^{(k)}\rangle$ that are as close as possible to singlet states, i.e., eigenfunctions of total spin with eigenvalue 0, $\hat{S}^2 |\Psi^{(k)}\rangle = 0$. Conservation of symmetries is notoriously important because molecular eigenstates are joint eigenfunctions of \hat{H} and its symmetries, including \hat{S}_z , \hat{N} , molecular point-group symmetries (if any), and \hat{S}^2 . Conservation of \hat{S}_z and \hat{N} (and molecular point-group symmetries that are isomorphic to $\mathbb{Z}_2^{N^n}$) can be achieved in a relatively easy way because the eigenfunctions of these operators are SDs $|\mathbf{x}\rangle$, e.g.

$$\hat{N} |\mathbf{x}\rangle = (N_{\alpha x} + N_{\beta x}) |\mathbf{x}\rangle, \quad \hat{S}_z |\mathbf{x}\rangle = (N_{\alpha x} - N_{\beta x}) |\mathbf{x}\rangle \quad (10)$$

Therefore, in this work, we ensure that the configurations used to span the subspace all have the desired eigenvalue of \hat{S}_z and \hat{N} as part of the configuration recovery procedure in the presence of noise. Conservation of \hat{S}^2 is more difficult to achieve in CI methods because the eigenfunctions of \hat{S}^2 are not SDs. The proper way of ensuring conservation of \hat{S}^2 is expanding $|\Psi^{(k)}\rangle$ on a set of CSFs, i.e., spin symmetry-adapted linear combinations of SDs. The use of CSFs was common in early CI codes because, in addition to the obvious benefits of reducing the memory footprint of the CI vector, automatic conservation of total spin enhanced the stability of the CI iterations. Modern CI codes tend to use SDs as opposed to CSFs because the formation of the σ vector (the most memory- and rate-limiting step of CI algorithms) is considerably more efficient and easily parallelized in SD-based algorithms and in part because system memory and disk are more plentiful than in previous machines. However, in SD-based CI codes, the conservation of total spin is no longer guaranteed. In the case of SQD, sampling from a quantum computer may return sets of SDs that do not allow constructing eigenfunctions of total spin. For example, in a (2e,2o) system, one may sample the configuration $|1001\rangle$ (having a single spin-down excitation over the RHF state $|0101\rangle$), which is a linear combination of the open-shell singlet and triplet states, respectively $(|1001\rangle \pm |0110\rangle)/\sqrt{2}$. If the configuration $|0110\rangle$ is not sampled, one can construct neither eigenfunction of total spin, leading to spin contamination or redundancy (i.e., the configuration $|1001\rangle$ is involved in a CI calculation but has coefficient 0 in the CI vector). In this work, to facilitate conservation of total spin, we relied on the following procedure: Instead of collecting directly d independent and identically distributed (i.i.d.) samples from \mathcal{X} to make the batch $\mathcal{S}^{(k)}$, we collect $\sqrt{d}/2$ samples and identify all unique configurations \mathbf{x}^u (u for unique) of length $M/2$ obtained from $\mathbf{x}_{\uparrow(i)}$ and $\mathbf{x}_{\downarrow(i)}$ for $1 \leq i \leq \sqrt{d}/2$, forming

the set $\mathcal{U}^{(k)} = \{\mathbf{x}^u\}$. The size of the set is upper bounded by $|\mathcal{U}^{(k)}| \leq \sqrt{d}$. From $\mathcal{U}^{(k)}$, we obtain the batch set $\mathcal{S}^{(k)}$ as

$$\mathcal{S}^{(k)} = \left\{ \mathbf{x} \mid \mathbf{x} = \mathbf{x}_{(i)}^u \oplus \mathbf{x}_{(j)}^u \text{ for all } \mathbf{x}_{(i)}^u, \mathbf{x}_{(j)}^u \in \mathcal{U}^{(k)} \right\} \quad (11)$$

The size of the set above is upper bounded by $|\mathcal{S}^{(k)}| \leq d$. This procedure facilitates total spin conservation (for example, from the configuration $|1001\rangle$, one can build the set $\{|1001\rangle, |1010\rangle, |0101\rangle, |0110\rangle\}$, which contains two closed-shell configurations and allows constructing an open-shell singlet state) but does not enforce it. Therefore, in combination with the sampling strategy mentioned above, we achieve the conservation of total spin by a soft constraint in the eigenstate solver, i.e., by adding a penalty term to mitigate spin contamination

$$\left\{ H + \lambda [S^2 - s(s+1)]^2 \right\} |\psi\rangle = E |\psi\rangle \quad (12)$$

where λ can be understood as a Lagrange multiplier that penalizes contributions from $S^2 \neq s(s+1)$. In this work, we used a soft constraint with $\lambda = 0.2$.

Projection and diagonalization. For each subsampled set $\mathcal{S}^{(k)}$, the Hamiltonian is projected into the corresponding subspace spanned by the configurations in $\mathcal{S}^{(k)}$ in Eq. 3.

We then diagonalize this projected Hamiltonian (solving $\hat{H}_{\mathcal{S}^{(k)}} |\psi^{(k)}\rangle = E^{(k)} |\psi^{(k)}\rangle$), and its ground state forms an approximation to the ground state of \hat{H} . The approximate ground state $|\psi^{(k)}\rangle$ is defined by its amplitudes in the subspace

$$|\psi^{(k)}\rangle = \sum_{\mathbf{x} \in \mathcal{S}^{(k)}} c_{\mathbf{x}}^{(k)} |\mathbf{x}\rangle \quad (13)$$

The k th estimate of the ground-state energy is given by

$$E^{(k)} = \langle \psi^{(k)} | \hat{H}_{\mathcal{S}^{(k)}} | \psi^{(k)} \rangle \quad (14)$$

Self-consistent configuration recovery

After a quantum state $\hat{\rho}$, corresponding to a noiseless state $|\Psi\rangle$, is prepared in our prefault-tolerant quantum processor, we measure it in the computational basis, obtaining the set of measurements

$$\tilde{\mathcal{X}} = \{\mathbf{x} | \mathbf{x} \sim \tilde{\rho}_{\Psi}\} \quad (15)$$

The class of noiseless states $|\Psi\rangle$ considered in this work are eigenstates to the total particle number operator and the total number operator for each spin species

$$\sum_{p=1}^{N_{\text{MO}}} \sum_{\sigma} \hat{n}_{p\sigma} |\Psi\rangle = N |\Psi\rangle, \quad \sum_{p=1}^{N_{\text{MO}}} \hat{n}_{p\uparrow} |\Psi\rangle = N_{\uparrow} |\Psi\rangle, \quad \sum_{p=1}^{N_{\text{MO}}} \hat{n}_{p\downarrow} |\Psi\rangle = N_{\downarrow} |\Psi\rangle \quad (16)$$

From the measurements on the quantum processors, we observe that there are a number of configurations in $\tilde{\mathcal{X}}$ whose $N_{x\uparrow}$ and $N_{x\downarrow}$ do not match the N_{\uparrow} and N_{\downarrow} of the ground state. In table S1, we report typical values of the fraction of sampled configurations with the wrong particle number. Because the circuits we use to produce $|\Psi\rangle$ are particle-number preserving, we are certain that configurations with wrong particle numbers have been corrupted by noise. It is this subset of configurations that the configuration recovery scheme is applied to. The configurations in $\tilde{\mathcal{X}}$ whose $N_{x\uparrow} = N_{\uparrow}$ and $N_{x\downarrow} = N_{\downarrow}$ are not subject to the configuration recovery subroutine.

We then probabilistically flip bits and restore the correct particle number using information obtained from observables of the system.

We use the spin-orbital occupancy averaged over all collected batches of subsamples \mathbf{n} , whose components are defined in Eq. 4.

Consider a configuration \mathbf{x} with $N_{x\uparrow}$ spin-up electrons, where $N_{x\uparrow} > N_{\uparrow}$. From the set of occupied spin-orbitals in the first half of \mathbf{x} , $|N_{x\uparrow} - N_{\uparrow}|$ bits are sampled to be flipped. If instead $N_{x\uparrow} < N_{\uparrow}$, the bits to be flipped are instead sampled from the unoccupied spin-orbitals. The same procedure applies to the spin-down orbitals.

The probability of flipping bit $x_{p\sigma}$ depends on the distance between the value of the bit and the reference orbital occupancy $n_{p\sigma}$. The simplest approach would be to define the distribution proportional to $|x_{p\sigma} - n_{p\sigma}|$. However, this introduces an undesirable effect: If for some spin-orbital $p\sigma$, $n_{p\sigma} \approx 0.5$, then $|x_{p\sigma} - n_{p\sigma}| \approx 0.5$ as well, i.e., we assign roughly 50% probability weight to flip the bit, regardless of its initial value. On the other hand, the initial value $x_{p\sigma}$ will, in general, retain some correlation with the other values in the bit-string, even in the presence of noise. Hence, a better approach is to deweight the probability of flipping when $|x_{p\sigma} - n_{p\sigma}|$ is small by using a modified rectified linear unit (ReLU) function $w(|x_{p\sigma} - n_{p\sigma}|)$, defined as

$$w(y) = \begin{cases} \delta \frac{y}{h} & \text{if } y \leq h \\ \delta + (1-\delta) \frac{y-h}{1-h} & \text{if } y > h \end{cases} \quad (17)$$

The parameter $h \in (0, 1)$ defines the location of the “corner” of the ReLU function, whereas the parameter $\delta \in [0, c]$ defines the value of the ReLU function at the corner. w becomes a true ReLU function when $\delta = 0$, and for values of $\delta > 0$, the ReLU is modified so that it is not identically zero except at $y = 0$. In the specific cases in this work, we chose the values $\delta = 0.01$ and $h = N / M$ (the filling factor) in all experiments.

We do not assume that we know \mathbf{n} a priori, and instead, we compute it and improve it self-consistently, following the procedure:

1) Setup phase: (i) Find the subset of configurations of $\tilde{\mathcal{X}}$ that live in the correct particle sector for both spin species, which we denote by \mathcal{X}_N : $\mathcal{X}_N \subset \tilde{\mathcal{X}}$. (ii) Obtain batches of samples $(\mathcal{S}^{(1)}, \dots, \mathcal{S}^{(K)})$ from \mathcal{X}_N as described in the previous section. (iii) Run the eigenstate solver on the batches and obtain approximate eigenstates $|\psi^{(1)}\rangle, \dots, |\psi^{(K)}\rangle$ (Eqs. 3 and 13). (iv) From the approximate eigenstates construct the first guess for \mathbf{n} , according to Eq. 4.

2) Self-consistent iterations (repeat until stopping criterion is met): (i) \mathbf{n} is used to correct the configurations with the wrong particle number in $\tilde{\mathcal{X}}$ (we give this subset the label $\mathcal{X}_{\mathbf{n}}$). The resulting set of recovered configurations is labeled $\mathcal{X}_{\mathbf{n}}$. (ii) From $\mathcal{X}_{\mathbf{n}} = \mathcal{X}_N \cup \mathcal{X}_{\mathbf{n}}$, batches of samples $(\mathcal{S}^{(1)}, \dots, \mathcal{S}^{(K)})$ are obtained as described in Discussion. (iii) Run the eigenstate solver on the batches and obtain approximate eigenstates $|\psi^{(1)}\rangle, \dots, |\psi^{(K)}\rangle$ (Eqs. 3 and 13). (iv) From the approximate eigenstates construct refined guess for \mathbf{n} , according to Eq. 4. (v) If the stopping criterion is not met, go back to step 2a.

We direct the reader to Fig. 2 in the main text for a numerical emulation of the effect of the configuration recovery on the accuracy of SQD in the presence of noise. This method is implemented in the package (66).

Energy-variance extrapolation

It is guaranteed that the accuracy of SQD (with or without configuration recovery) increases as the number of configurations used for the subspace expansion d is increased. However, the convergence of

the ground-state properties with d is not expected to follow any specific functional relation. Therefore, attempting to analyze the convergence as a function of d is not well motivated. Instead, the different eigenstate approximations obtained for different values of d and different batches of samples $S^{(k)}$ are used for an energy-variance extrapolation (43, 67–71). Consider the approximate eigenstate $|\psi\rangle$, whose energy is given by $E = \langle\psi|\hat{H}|\psi\rangle$, and consider the exact eigenstate energy E_T . The difference between the approximate energy and E_T

$$\delta E = \langle\hat{H}\rangle - E_T \quad (18)$$

vanishes linearly with the Hamiltonian variance divided by the square of the variational energy (43)

$$\frac{\Delta H}{E^2} = \frac{\langle\psi|\hat{H}^2|\psi\rangle - \langle\psi|\hat{H}|\psi\rangle^2}{E^2} \quad (19)$$

This linear relation is satisfied as long as $|\psi\rangle$ is sufficiently close to an eigenstate of the Hamiltonian, as measured by the state fidelity. The least-squares fit of a collection of energy-variance points yields an estimate of E_T , as the intersect of the fit with the ordinates. This point is the extrapolation of the estimate of the energy to the limit where $|\psi\rangle$ coincides with the exact eigenstate. Besides the energy extrapolation, the energy-variance analysis may also reveal the existence of multiple eigenstates close in energy to the ground state.

In the main text and in the Supplementary Materials, we apply the energy-variance analysis to two different sets of energy-variance pairs. The first one is the energy-variance pairs obtained by the HCI procedure where different energies and variances are obtained by changing the cutoff parameter that indirectly controls the number of determinants in the subspace projection and diagonalization, resulting in different levels of accuracy. The second set of energy-variance pairs are those obtained from SQD for different numbers of configurations d as well as for different batches of sampled configurations $S^{(k)}$.

Experimental details

In this section, we describe the quantum circuits that are used to produce the configurations to which we apply the eigenstate solver. In addition, we provide details on the setting of circuit parameters from an efficient classical CCSD calculation and the mapping of the circuits to quantum processors with heavy-hex connectivity.

Quantum circuits

In this work, we used the LUCJ ansatz (34) to sample randomly distributed electronic configurations. LUCJ derives from the UCJ ansatz, which has the form (35) of a product of L layers, as defined in Eq. 5. We recall that

$$\hat{K}_\mu = \sum_{pq\sigma} K_{pq}^\mu \hat{a}_{p\sigma}^\dagger \hat{a}_{q\sigma} \quad , \quad \hat{J}_\mu = \sum_{pr\sigma\tau} J_{pr,q\tau}^\mu \hat{n}_{p\sigma} \hat{n}_{q\tau} \quad (20)$$

In Eq. 21, $p, q = 0 \dots N_{\text{MO}} - 1$ label molecular spatial orbitals and σ, τ label spin polarizations (α, β for spin-up and spin-down electrons, respectively). $K_{pq}^\mu / J_{pq,\sigma\tau}^\mu$ has complex/real matrix elements and is anti-Hermitian/symmetric. The UCJ ansatz can be derived from a twice-factorized low-rank decomposition of the QUCCD ansatz (35, 72), and the L -product form is such that the exact FCI wave function can be obtained via Eq. 5 (35, 36).

The local UCJ or LUCJ (34) introduces a “local” approximation of the UCJ ansatz, which makes the following modifications for opposite-spin and same-spin number-number terms

$$\begin{aligned} \sum_{pq} J_{pq,\alpha\beta} \hat{n}_{p\alpha} \hat{n}_{q\beta} &\rightarrow \sum_{p \in S} J_{p\alpha,p\beta} \hat{n}_{p\alpha} \hat{n}_{p\beta} \\ \sum_{pq} J_{pq,\alpha\sigma} \hat{n}_{p\alpha} \hat{n}_{q\sigma} &\rightarrow \sum_{pq \in S'} J_{p\alpha,q\sigma} \hat{n}_{p\alpha} \hat{n}_{q\sigma} \end{aligned} \quad (21)$$

where $\sigma = \alpha, \beta$ and the sets S, S' are such that the quantum circuit implementing $e^{\hat{J}_\mu}$ has depth $O(1)$ and only comprises $O(|S| + |S'|)$ number-number “nn gates,” i.e., two-qubit unitaries of the form $U_{nn}(\varphi) = e^{-i\frac{\varphi}{4}(Z_p + Z_q - Z_p Z_q)}$, acting on adjacent qubits p, q in the topology of a certain processor. For example, on a heavy-hex processor, $S = \{4k, k = 0 \dots (N_{\text{MO}} - 1)/4\}$ and $S' = \{(p, p+1), p = 0 \dots N - 2\}$ (34). The circuits $e^{\pm\hat{K}_\mu}$, on the other hand, can be implemented by a Bogolyubov circuit acting on N_{MO} qubits and comprising $O[N_\alpha(N_{\text{MO}} - N_\alpha)]$ gates and depth $O(N_{\text{MO}})$ (73–75). Through the local approximation, the LUCJ Ansatz balances hardware friendliness and accuracy, the latter ultimately deriving from its connection to coupled cluster theory and adiabatic state preparation (34).

Unless otherwise specified, we use the truncated LUCJ circuit $|\Psi\rangle = e^{\hat{K}_2} e^{-\hat{K}_1} e^{\hat{J}_1} e^{\hat{K}_1} |\mathbf{x}_{\text{RHF}}\rangle$, which is the result of considering the two-layer LUCJ circuit and removing the last orbital rotation and last Jastrow operations. The resulting state is implemented by a circuit whose depth is identical to the single-layer LUCJ circuit. The addition of the $e^{\hat{K}_2}$ operation to the circuit can have a large impact on the configurations generated by the circuit when the parameters are set from the t_2 tensor from a classical restricted closed-shell CCSD calculation (as described in the next section). For dynamically correlated species, the $J_{pq,\sigma\tau}^\mu$ parameters obtained from t_2 can have a small amplitude, resulting in the approximate cancellation of the $\exp(-\hat{K}_1)$ and $\exp(\hat{K}_1)$ terms in the ansatz. Without $\exp(\hat{K}_2)$, the resulting wave function can be overconcentrated around the Hartree-Fock configuration. Therefore, for dynamically correlated species, the action of $\exp(\hat{K}_2)$ is to remove some of the excessive concentration of the $|\Psi\rangle$ wave function, when parameters are set from a restricted closed-shell CCSD calculation.

Initialization of LUCJ parameters

In our experiments, we parameterize the LUCJ circuits using the following procedure:

1) First, we carry out a classical restricted closed-shell CCSD calculation, yielding amplitudes $t_{1,ai}$ and $t_{2,abij}$, where ij/ab labels occupied/unoccupied orbitals in the RHF state.

2) We reshape the t_2 tensor into the matrix $(t_2)_{ai,bj}$ and diagonalize it, $(t_2)_{ai,bj} = \sum_y \tau_y U_{ai,y} U_{bj,y}$, where the eigenvectors τ_y are sorted in decreasing order of absolute value.

3) We extend the unitaries to the following matrices

$$\tilde{U}_{y,pr} = \delta_{pa} \delta_{ri} U_{ai,y} \quad (22)$$

i.e., matrices where only the occupied/unoccupied block is nonzero.

4) We define the Hermitian operators

$$X_{\pm,y} = \frac{1 \mp i}{2} \left(\tilde{U}_y \pm i \tilde{U}_y^T \right) \quad (23)$$

and their eigenpairs

$$X_{\pm,y} V_{\pm,y} = g_{\pm,y} V_{\pm,y} \quad (24)$$

5) We define the operators

$$\begin{aligned} (J_{2y})_{pr}^{\sigma\tau} &= \tau_y (g_{+,y})_p (g_{+,y})_r \\ (K_{2y})_{pr}^{\sigma} &= \log(V_{+,y})_{pr} \\ (J_{2y+1})_{pr}^{\sigma\tau} &= -\tau_y (g_{-,y})_p (g_{-,y})_r \\ (K_{2y+1})_{pr}^{\sigma} &= \log(V_{-,y})_{pr} \end{aligned} \quad (25)$$

6) We retain the first L matrices J, K . This allows us to refine a nonlocal UCJ wave function (35, 72).

7) We zero out the entries of the J matrices leading to quantum gates acting on nonadjacent qubits. For a heavy-hex lattice, we retain the following elements

$$\begin{aligned} (J_{\mu})_{p,p+1}^{\alpha\alpha}, \quad p &= 0 \dots N_{\text{MO}} - 2 \\ (J_{\mu})_{p,p+1}^{\beta\beta}, \quad p &= 0 \dots N_{\text{MO}} - 2 \\ (J_{\mu})_{p,p}^{\alpha\beta}, \quad p &= 0 \dots N_{\text{MO}} - 1, \quad p \% 4 = 0 \end{aligned} \quad (26)$$

The final sparsification allows us to construct an LUCJ wave function. In a conventional LUCJ calculation, these parameters are the starting point of a variational optimization. For the hardware experiments reported in this study, we used these parameters, without further optimization, to define an LUCJ circuit, which we used to sample randomly distributed configurations.

Mapping to heavy-hex processors

The choice of retaining the elements $(J_{\mu})_{p,p}^{\alpha\beta}$ with $p = 0 \dots N_{\text{MO}} - 1, p \% 4 = 0$ when implementing LUCJ on a heavy-hex processor has an important technical motivation: On such devices, assuming a number of qubits greatly exceeding N_{MO} , spin-up and spin-down orbitals can be mapped on two segments of adjacent qubits forming a “zigzag” pattern and connected through an auxiliary qubit for $p = 0, 4, 8, \dots$ as shown in the three rightmost panels of Fig. 4. Such a qubit layout allows implementing LUCJ with a minimal overhead of SWAP gates (two per auxiliary qubit and layer of LUCJ). However, on current processors with up to 133 qubits, for $N_{\text{MO}} > 21$ one cannot couple $N_{\text{MO}}/4$ spin-orbitals with opposite spins through auxiliary qubits without incurring a substantial overhead of SWAP gates, for the simple reason that a chain of 22 or more qubits is longer than the “diagonal” of the processor. In such a situation, as shown in the rightmost panel of Fig. 4, the segments on which spin-up and spin-down qubits are mapped form two “tails” that are not connected by auxiliary qubits. This fact has two implications: (i) no more than six spin-orbitals with opposite spins can be coupled through auxiliary qubits, and (ii) if one retains the elements $(J_{\mu})_{p,p}^{\alpha\beta}$ with $p = 0 \dots N_{\text{MO}} - 1, p \% 4 = 0$ and $p \leq 16$, the largest elements of J_{μ} may be discarded, yielding a lower-accuracy wave function. The first problem is a fundamental one, which can only be resolved with a substantially different mapping of fermionic degrees of freedom onto qubits and/or through the availability of larger processors. The second problem, on the other hand, has a simple solution, which we now describe. First, for any pair of spatial orbitals $p, r = 0 \dots N_{\text{MO}} - 1$, the orbital rotation

$$\hat{S}_{pr} = e^{-i\frac{\pi}{2}\sum_{\sigma}(\hat{a}_{pr}^{\dagger}\hat{a}_{r\sigma} + \hat{a}_{r\sigma}^{\dagger}\hat{a}_{pr})} \quad (27)$$

implements the permutation $S_{pr} \in S_{N_{\text{MO}}}$ exchanging orbitals p and r , in the sense that

$$\hat{S}_{pr}^{\dagger} \left(\sum_{qs\tau} M_{qs} \hat{a}_{q\tau}^{\dagger} \hat{a}_{s\tau} \right) \hat{S}_{pr} = \sum_{qs\tau} M'_{qs} \hat{a}_{q\tau}^{\dagger} \hat{a}_{s\tau}, \quad M'_{qs} = M_{S_{pr}(q), S_{pr}(s)} \quad (28)$$

Three immediate implications of this fact are as follows:

$$1) \hat{S}_{pr}^{\dagger} \hat{n}_{q\tau} \hat{S}_{pr} = \hat{n}_{S_{pr}(q)\tau}$$

2) for any density-density operator $\hat{J}_{\mu} = \sum_{qs\sigma\tau} (J_{\mu})_{qs}^{\sigma\tau} \hat{n}_{q\sigma} \hat{n}_{s\tau}$, one has

$$\hat{S}_{pr}^{\dagger} \hat{J}_{\mu} \hat{S}_{pr} = \sum_{qs\tau} (J'_{\mu})_{qs}^{\sigma\tau} \hat{n}_{q\sigma} \hat{n}_{s\tau} = \hat{J}'_{\mu}, \quad (J'_{\mu})_{qs}^{\sigma\tau} = (J_{\mu})_{S_{pr}(q), S_{pr}(s)}^{\sigma\tau} \quad (29)$$

3) that for any permutation $S \in S_{N_{\text{MO}}}$ there exists an orbital rotation $e^{\hat{K}_S}$ implementing the permutation S (this is true because permutations can be written as products of exchange permutations, an exchange permutation can be implemented by an orbital rotation, and orbital rotations are closed under multiplication).

Consider now the tensor $(J_{\mu})_{qs}^{\sigma\tau}$ resulting from the low-rank decomposition of the CCSD operator described in the previous subsection. Let $p_0 \dots p_{\ell}$ be the ℓ elements of $(J_{\mu})_{qs}^{\alpha\beta}$ with the largest absolute values, and let $S \in S_{N_{\text{MO}}}$ be the permutation such that $S(p_0) = S(p_0) = 0 \dots S(p_{\ell}) = 4\ell$. Then

$$e^{-\hat{K}_S} \hat{J}_{\mu} e^{\hat{K}_S} = \hat{J}'_{\mu} \quad (30)$$

where the elements of $(J'_{\mu})_{qs}^{\alpha\beta}$ with the largest absolute values are at positions $0, \dots, 4\ell$. Before sparsifying the tensor \hat{J}_{μ} , one can use the identity

$$e^{\hat{K}_{\mu}} e^{\hat{J}_{\mu}} e^{-\hat{K}_{\mu}} = e^{\hat{K}_{\mu}} e^{\hat{K}_S} e^{\hat{J}'_{\mu}} e^{-\hat{K}_S} e^{-\hat{K}_{\mu}} = e^{\hat{K}'_{\mu}} e^{\hat{J}'_{\mu}} e^{-\hat{K}'_{\mu}} \quad (31)$$

to obtain a UCJ operator with transformed orbital rotations $e^{-\hat{K}_{\mu}}$ and an opposite-spin density-density interaction whose largest elements in absolute value act on spatial orbitals $p = 0, 4, 8, \dots, 4\ell$. Sparsification of $(J'_{\mu})_{qs}^{\alpha\beta}$ then leads to retaining the ℓ dominant opposite-spin density-density interaction terms (as many as allowed by the size and topology of the available heavy-hex processor) with a minimal overhead of SWAP gates.

In this work, we retained the elements $(J_{\mu})_{p,p}^{\alpha\beta}$ with $p = 0 \dots N_{\text{MO}} - 1, p \% 4 = 0$ and $p \leq 16$. In future work, the procedure described here could be used to modify the LUCJ wave function and the resulting probability distribution for electronic configurations.

Supplementary Materials

This PDF file includes:

Supplementary Text

Figs. S1 to S20

Table S1

REFERENCES AND NOTES

1. K. D. Vogiatzis, D. Ma, J. Olsen, L. Gagliardi, W. A. de Jong, Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations. *J. Chem. Phys.* **147**, 184111 (2017).
2. H. Gao, S. Imamura, A. Kasagi, E. Yoshida, Distributed implementation of full configuration interaction for one trillion determinants. *J. Chem. Theory Comput.* **20**, 1185–1192 (2024).

3. J. P. F. Le Blanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, E. Gull, Solutions of the two-dimensional Hubbard model: Benchmarks and results from a wide range of numerical algorithms. *Phys. Rev. X* **5**, 041011 (2015).
4. M. Motta, D. M. Ceperley, G. K.-L. Chan, J. A. Gomez, E. Gull, S. Guo, C. Jimenez-Hoyos, T. N. Lan, J. Li, F. Ma, A. J. Millis, N. V. Prokof'ev, U. Ray, G. E. Scuseria, S. Sorella, E. M. Stoudenmire, Q. Sun, I. S. Tupitsyn, S. R. White, D. Zgid, S. Zhang, Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods. *Phys. Rev. X* **7**, 031059 (2017).
5. Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing. *Chem. Rev.* **119**, 10856–10915 (2019).
6. S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, X. Yuan, Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020).
7. B. Bauer, S. Bravyi, M. Motta, G. K.-L. Chan, Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* **120**, 12685–12717 (2020).
8. A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
9. Google AI Quantum and Collaborators, Hartree-Fock on a superconducting qubit quantum computer. *Science* **369**, 1084–1089 (2020).
10. W. J. Huggins, B. A. O'Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, J. Lee, Unbiasing fermionic quantum Monte Carlo with a quantum computer. *Nature* **603**, 416–420 (2022).
11. M. Motta, G. O. Jones, J. E. Rice, T. P. Gujarati, R. Sakuma, I. Liepuoniute, J. M. Garcia, Y. Ohnishi, Quantum chemistry simulation of ground- and excited-state properties of the sulfonium cation on a superconducting quantum processor. *Chem. Sci.* **14**, 2915–2927 (2023).
12. L. Zhao, J. Goings, K. Shin, W. Kyoung, J. I. Fuks, J.-K. Kevin Rhee, Y. M. Rhee, K. Wright, J. Nguyen, J. Kim, S. Johri, Orbital-optimized pair-correlated electron simulations on trapped-ion quantum computers. *NPJ Quantum Inf.* **9**, 60 (2023).
13. T. E. O'Brien, G. Anselmetti, F. Gkrtsis, V. E. Elfving, S. Polla, W. J. Huggins, O. Oumarou, K. Kechedzhi, D. Abanin, R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, K. Anderson, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, J. C. Bardin, A. Bengtsson, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, J. Campero, Z. Chen, B. Chiaro, D. Chik, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, S. Demura, I. Drozdov, A. Dunsworth, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, R. Gosula, A. G. Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, P. Heu, M. R. Hoffmann, S. Hong, T. Huang, A. Huff, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, T. Khattar, M. Khezri, M. Kieferová, S. Kim, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. T. Lill, W. Liu, W. P. Livingston, A. Locharla, F. D. Malone, S. Mandrà, O. Martin, S. Martin, J. R. McClean, T. McCourt, M. McEwen, X. Mi, A. Mieszala, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, R. Movassagh, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, S. Omonije, A. Opremcak, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, C. Rocque, P. Roushan, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, N. Shutt, V. Shvarts, J. Skrzyni, W. C. Smith, R. D. Somma, G. Sterling, D. Strain, M. Szalay, D. Thor, A. Torres, G. Vidal, B. Villalonga, C. V. Heidweiller, T. White, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, N. Zobrist, D. Bacon, S. Boixo, Y. Chen, J. Hilton, J. Kelly, E. Lucero, A. Megrant, H. Neven, V. Smelyanskiy, C. Gogolin, R. Babbush, N. C. Rubin, Purification-based quantum error mitigation of pair-correlated electron simulations. *Nat. Phys.* **19**, 1787–1792 (2023).
14. T. Weaving, A. Ralli, P. J. Love, S. Succi, P. V. Coveney, Contextual subspace variational quantum eigensolver calculation of the dissociation curve of molecular nitrogen on a superconducting quantum computer. arXiv:2312.04392 [quant-ph] (2023).
15. V. Verteletskyi, T.-C. Yen, A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *J. Chem. Phys.* **152**, 124114 (2020).
16. W. J. Huggins, J. R. Mc Clean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers. *npj Quantum Inf.* **7**, 23 (2021).
17. A. Dutt, W. Kirby, R. Raymond, C. Hadfield, S. Sheldon, I. L. Chuang, A. Mezzacapo, Practical Benchmarking of Randomized Measurement Methods for Quantum Chemistry Hamiltonians. arXiv:2312.07497 [quant-ph] (2023).
18. A. Anand, P. Schleich, S. Alperin-Lea, P. W. K. Jensen, S. Sim, M. Díaz-Tinoco, J. S. Kottmann, M. Degroote, A. F. Izmaylov, A. Aspuru-Guzik, A quantum computing view on unitary coupled cluster theory. *Chem. Soc. Rev.* **51**, 1659–1684 (2022).
19. M. B. Hastings, D. Wecker, B. Bauer, M. Troyer, Improving quantum algorithms for quantum chemistry. *Quantum Inf. Comput.* **15**, 1–21 (2015).
20. J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. Mc Clean, N. Wiebe, R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction. *PRX quantum* **2**, 030305 (2021).
21. Y. Alexeev, M. Amsler, P. Baity, M. A. Barroca, S. Bassini, T. Battelle, D. Camps, D. Casanova, Y. J. Choi, F. T. Chong, C. Chung, C. Codella, A. D. Corcoles, J. Cruise, A. D. Meglio, J. Dubois, I. Duran, T. Eckl, S. Economou, S. Eidenbenz, B. Elmegreen, C. Fare, I. Faro, C. S. Fernández, R. N. B. Ferreira, K. Fuji, B. Fuller, L. Gagliardi, G. Galli, J. R. Glick, I. Gobbi, P. Gokhale, S. de la Puent Gonzalez, J. Greiner, B. Gropp, M. Grossi, E. Gull, B. Healy, B. Huang, T. S. Humble, N. Ito, A. F. Izmaylov, A. Javadi-Abhari, D. Jennewein, S. Jha, L. Jiang, B. Jones, W. A. de Jong, P. Jurcevic, W. Kirby, S. Kister, M. Kitagawa, J. Klassen, K. Klymko, K. Koh, M. Kondo, D. M. Kurkcuoglu, K. Kurowski, T. Laino, R. Landfield, M. Leininger, V. Leyton-Ortega, A. Li, M. Lin, J. Liu, N. Lorente, A. Luckow, S. Martiel, F. Martin-Fernandez, M. Martonosi, C. Marvinney, A. C. Medina, D. Merten, A. Mezzacapo, K. Michielsen, A. Mitra, T. Mittal, K. Moon, J. Moore, M. Motta, Y.-H. Na, Y. Nam, P. Narang, Y.-y. Ohnishi, D. Ottaviani, M. Otten, S. Pakin, V. R. Pascuzzi, E. Penault, T. Piontek, J. Pitera, P. Rall, G. S. Ravi, N. Robertson, M. Rossi, P. Rydlchowski, H. Ryu, G. Samsonidze, M. Sato, N. Saurabh, V. Sharma, K. Sharma, S. Shin, G. Slessman, M. Steiner, I. Sitiidikov, I.-S. Suh, E. Switzer, W. Tang, J. Thompson, S. Todo, M. Tran, D. Trenev, C. Trott, H.-H. Tseng, E. Tureci, D. G. Valinas, S. Vallecorsa, C. Wever, K. Wojciechowski, X. Wu, S. Yoo, N. Yoshioka, V. W.-z. Yu, S. Yunoki, S. Zhuk, D. Zubarev, Quantum-centric Supercomputing for Materials Science: A Perspective on Challenges and Future Directions. arXiv:2312.09733 [quant-ph] (2023).
22. Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. Mc Clain, E. R. Sayfutyarova, S. Sharma, S. Wouters, G. K.-L. Chan, PySCF: The Python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.* **8**, e1340 (2018).
23. P. Jordan, E. Wigner, Über das Paulische Äquivalenzverbot. *Z. Phys.* **47**, 631–651 (1928).
24. J. Ivanic, K. Ruedenberg, Identification of deadwood in configuration spaces through general direct configuration interaction. *Theor. Chem. Acc.* **106**, 339–351 (2001).
25. K. Kanno, M. Kohda, R. Imai, S. Koh, K. Mitarai, W. Mizukami, Y. O. Nakagawa, Quantum-Selected Configuration Interaction: Classical diagonalization of Hamiltonians in subspaces selected by quantum computers. arXiv:2302.11320 [quant-ph] (2023).
26. Y. O. Nakagawa, M. Kamoshita, W. Mizukami, S. Sudo, Y.-y. Ohnishi, ADAPT-QSCI: Adaptive Construction of Input State for Quantum-Selected Configuration Interaction. arXiv:2311.01105 [quant-ph] (2023).
27. F. A. Evangelista, Adaptive multiconfigurational wave functions. *J. Chem. Phys.* **140**, 124114 (2014).
28. J. B. Schriber, F. A. Evangelista, Communication: An adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *J. Chem. Phys.* **144**, 161106 (2016).
29. A. A. Holmes, H. J. Changlani, C. Umrigar, Efficient heat-bath sampling in Fock space. *J. Chem. Theory Comput.* **12**, 1561–1571 (2016).
30. A. A. Holmes, N. M. Tubman, C. J. Umrigar, Heat-bath configuration interaction: An efficient selected configuration interaction algorithm inspired by heat-bath sampling. *J. Chem. Theory Comput.* **12**, 3674–3680 (2016).
31. N. M. Tubman, J. Lee, T. Y. Takeshita, M. Head-Gordon, K. B. Whaley, A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **145**, 044112 (2016).
32. J. B. Schriber, F. A. Evangelista, Adaptive configuration interaction for computing challenging electronic excited states with tunable accuracy. *J. Chem. Theory Comput.* **13**, 5354–5366 (2017).
33. S. Sharma, A. A. Holmes, G. Jeanmairet, A. Alavi, C. J. Umrigar, Semistochastic heat-bath configuration interaction method: Selected configuration interaction with semistochastic perturbation theory. *J. Chem. Theory Comput.* **13**, 1595–1604 (2017).
34. M. Motta, K. J. Sung, K. B. Whaley, M. Head-Gordon, J. Shee, Bridging physical intuition and hardware efficiency for correlated electronic states: The local unitary cluster Jastrow ansatz for electronic structure. *Chem. Sci.* **14**, 11213–11227 (2023).
35. Y. Matsuzawa, Y. Kurashige, Jastrow-type decomposition in quantum chemistry for low-depth quantum circuits. *J. Chem. Theory Comput.* **16**, 944–952 (2020).
36. F. A. Evangelista, G. K.-L. Chan, G. E. Scuseria, Exact parameterization of fermionic wave functions via unitary coupled cluster theory. *J. Chem. Phys.* **151**, 244112 (2019).
37. The ffsim developers, ffsim: Faster simulations of fermionic quantum circuits (2024); <https://github.com/qiskit-community/ffsim>.
38. P.-D. Fan, P. Piecuch, The usefulness of exponential wave function expansions employing one- and two-body cluster operators in electronic structure theory: The extended and generalized coupled-cluster methods. *Adv. Quantum Chem.* **51**, 1–57 (2006).
39. I. W. Bulik, T. M. Henderson, G. E. Scuseria, Can single-reference coupled cluster theory describe static correlation? *J. Chem. Theory Comput.* **11**, 3171–3179 (2015).
40. H. Beinert, R. H. Holm, E. Münck, Iron-sulfur clusters: Nature's modular, multipurpose structures. *Science* **277**, 653–659 (1997).
41. S. Sharma, K. Sivalingam, F. Neese, G. K.-L. Chan, Low-energy spectrum of iron-sulfur clusters directly from many-particle quantum mechanics. *Nat. Chem.* **6**, 927–933 (2014).
42. P. Venkateswara Rao, R. Holm, Synthetic analogues of the active sites of iron-sulfur proteins. *Chem. Rev.* **104**, 527–560 (2004).

43. T. Kashima, M. Imada, Path-integral renormalization group method for numerical study on ground states of strongly correlated electronic systems. *J. Phys. Soc. Jpn.* **70**, 2287–2299 (2001).

44. A. D. Becke, Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **38**, 3098–3100 (1988).

45. V. Papaefthymiou, M. M. Millar, E. Muenck, Moessbauer and EPR studies of a synthetic analog for the iron-sulfur Fe_4S_4 core of oxidized and reduced high-potential iron proteins. *Inorg. Chem.* **25**, 3010–3014 (1986).

46. Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, A. Kandala, Evidence for the utility of quantum computing before fault tolerance. *Nature* **618**, 500–505 (2023).

47. K. Shinjo, K. Seki, T. Shirakawa, R.-Y. Sun, S. Yunoki, Unveiling clean two-dimensional discrete time quasicrystals on a digital quantum computer. arXiv:2403.16718 [quant-ph] (2024).

48. D. Wecker, M. B. Hastings, M. Troyer, Progress towards practical quantum variational algorithms. *Phys. Rev. A* **92**, 042303 (2015).

49. S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell, E. F. Valeev, L. Lin, G. K.-L. Chan, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nat. Commun.* **14**, 1952 (2023).

50. F. Becca, S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge Univ. Press, 2017).

51. G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).

52. S. R. White, R. L. Martin, Ab initio quantum chemistry using the density matrix renormalization group. *J. Comp. Phys.* **110**, 4127–4130 (1999).

53. G. K.-L. Chan, S. Sharma, The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.* **62**, 465–481 (2011).

54. C. Hubig, J. Haegeman, U. Schollwöck, Error estimates for extrapolations with matrix-product states. *Phys. Rev. B* **97**, 045125 (2018).

55. Z. Li, J. Li, N. S. Dattani, C. Umrigar, G. K. Chan, The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations. *J. Chem. Phys.* **150**, 024302 (2019).

56. J. E. Smith, B. Mussard, A. A. Holmes, S. Sharma, Cheap and near exact CASSCF with large active spaces. *J. Chem. Theory Comput.* **13**, 5468–5478 (2017).

57. W. J. Hehre, R. F. Stewart, J. A. Pople, Self-consistent molecular-orbital methods. I. Use of Gaussian expansions of slater-type atomic orbitals. *J. Comp. Phys.* **51**, 2657–2664 (1969).

58. W. J. Hehre, R. Ditchfield, J. A. Pople, Self-consistent molecular orbital methods. XII. Further extensions of Gaussian-type basis sets for use in molecular orbital studies of organic molecules. *J. Comp. Phys.* **56**, 2257–2261 (1972).

59. T. H. Dunning, Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* **90**, 1007–1023 (1989).

60. Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui, J. J. Eriksen, Y. Gao, S. Guo, J. Hermann, M. R. Hermes, K. Koh, P. Koval, S. Lehtola, Z. Li, J. Liu, N. Mardirossian, J. D. McClain, M. Motta, B. Mussard, H. Q. Pham, A. Pulkin, W. Purwanto, P. J. Robinson, E. Ronca, E. R. Sayfutyarova, M. Scheurer, H. F. Schurkus, J. E. T. Smith, C. Sun, S.-N. Sun, S. Upadhyay, L. K. Wagner, X. Wang, A. White, J. D. Whitfield, M. J. Williamson, S. Wouters, J. Yang, J. M. Yu, T. Zhu, T. C. Berkelbach, S. Sharma, A. Y. Sokolov, G. K.-L. Chan, Recent developments in the PySCF program package. *J. Chem. Phys.* **153**, 024109 (2020).

61. Z. Li, G. K.-L. Chan, Spin-projected matrix product states: Versatile tool for strongly correlated systems. *J. Chem. Theory Comput.* **13**, 2681–2695 (2017).

62. J. P. Perdew, Density-functional approximation for the correlation energy of the inhomogeneous electron gas. *Phys. Rev. B* **33**, 8822–8824 (1986).

63. F. Jorge, A. Canal Neto, G. Camilleti, S. Machado, Contracted Gaussian basis sets for Douglas-Kroll-Hess calculations: Estimating scalar relativistic effects of some atomic and molecular properties. *J. Chem. Phys.* **130**, 064108 (2009).

64. W. Liu, Ideas of relativistic quantum chemistry. *Mol. Phys.* **108**, 1679–1706 (2010).

65. Z. Li, Y. Xiao, W. Liu, On the spin separation of algebraic two-component relativistic Hamiltonians. *J. Chem. Phys.* **137**, 154114 (2012).

66. A. A. Saki, S. Barison, B. Fuller, J. R. Garrison, J. R. Glick, C. Johnson, A. Mezzacapo, J. Robledo-Moreno, M. Rossmannek, P. Schweigert, I. Stidikov, K. J. Sung, Qiskit add-on: Sample-based quantum diagonalization (2024); <https://github.com/Qiskit/qiskit-addon-sqd>.

67. Y. Kwon, D. M. Ceperley, R. M. Martin, Effects of backflow correlation in the three-dimensional electron gas: Quantum Monte Carlo study. *Phys. Rev. B* **58**, 6800–6806 (1998).

68. M. Imada, T. Kashima, Path-integral renormalization group method for numerical study of strongly correlated electron systems. *J. Phys. Soc. Jpn.* **69**, 2723–2726 (2000).

69. Y. Nomura, A. S. Darmawan, Y. Yamaji, M. Imada, Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Phys. Rev. B* **96**, 205152 (2017).

70. S. Sorella, Generalized Lanczos algorithm for variational quantum Monte Carlo. *Phys. Rev. B* **64**, 024512 (2001).

71. T. Mizusaki, M. Imada, Extrapolation method for shell model calculations. *Phys. Rev. C* **65**, 064319 (2002).

72. M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, G. K.-L. Chan, Low rank representations for quantum simulation of electronic structure. *npj Quantum Inf.* **7**, 83 (2021).

73. M. Reck, A. Zeilinger, H. J. Bernstein, P. Bertani, Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.* **73**, 58–61 (1994).

74. W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, I. A. Walmsley, Optimal design for universal multiport interferometers. *Optica* **3**, 1460–1465 (2016).

75. Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, S. Boixo, Quantum algorithms to simulate many-body physics of correlated fermions. *Phys. Rev. Appl.* **9**, 044036 (2018).

Acknowledgments: We thank the IBM Quantum Service and Data team for help with the workflow execution. We acknowledge feedback and insightful conversations with A. Baiardi, S. Bravyi, G. Carleo, G. K.-L. Chan, A. Córcoles, O. Dial, J. Gambetta, A. Izmaylov, C. Johnson, Y. Kawashima, D. Kremer, I. Lauer, P. Love, G. Mazzola, T. Nakajima, P. Nation, H. Paik, E. Pritchett, M. Rossmannek, P. Schweigert, J. E. T. Smith, M. Stoudenmire, N. Yoshioka, A. Wack, and C. Zoufal. **Funding:** A part of this work is supported by project JPNP20017, funded by the New Energy and Industrial Technology Development Organization (NEDO), Japan, and by the RIKEN TRIP initiative (RIKEN Quantum). **Author contributions:** Design of the workflow and experiments: J.R.-M., M.M., W.K., K.S., M.C.T., and A.M. Implementation of the workflow: J.R.-M., M.M., A.J.-A., S.M., S.S., T.S., I.S., R.-Y.S., and K.J.S. Tuning and calibration of the Heron processor: H.H. Execution of the quantum part of the workflow on Heron: M.M., P.J., and M.T. Execution of the classical part of the workflow on Fugaku: T.S., R.-Y.S., and S.Y. Numerical benchmarks: J.R.-M., M.M., T.S., and K.J.S. Analytical derivations: J.R.-M., M.M., W.K., K.S., M.C.T., and A.M. All authors contributed to the manuscript writing and data analysis. **Competing interests:** J.R.-M., M.M., W.K., K.S., M.C.T., W.K., K.J.S., and A.M. are coauthors on a filed patent for the self-consistent configuration recovery technique, filed by International Business Machines Cor. (no. 99226518, filed 25 January 2024). The authors declare that they have no other competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials and in https://github.com/jrm874/sqd_data_repository or <https://doi.org/10.5281/zenodo.15324153>. The code to reproduce all results reported is available as the open-source package qiskit-addon-sqd (66) <https://github.com/Qiskit/qiskit-addon-sqd>.

Submitted 30 November 2024

Accepted 15 May 2025

Published 18 June 2025

10.1126/sciadv.adu9991