Ground-state energy estimation of the water molecule on a trapped ion quantum computer

Yunseong Nam,^{1,*} Jwo-Sy Chen,¹ Neal C. Pisenti,¹ Kenneth Wright,¹ Conor Delaney,¹ Dmitri Maslov,²

Kenneth R. Brown,^{1,3} Stewart Allen,¹ Jason M. Amini,¹ Joel Apisdorf,¹ Kristin M. Beck,¹ Aleksey

Blinov,¹ Vandiver Chaplin,¹ Mika Chmielewski,^{1,4} Coleman Collins,¹ Shantanu Debnath,¹ Andrew M.

Ducore,¹ Kai M. Hudek,¹ Matthew Keesan,¹ Sarah M. Kreikemeier,¹ Jonathan Mizrahi,¹ Phil Solomon,¹

Mike Williams,¹ Jaime David Wong-Campos,¹ Christopher Monroe,^{1,4} and Jungsang Kim^{1,3,†}

¹IonQ, Inc., College Park, MD 20740, USA

²National Science Foundation, Alexandria, VA 22314, USA

³Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

⁴Joint Quantum Institute and Department of Physics,

University of Maryland, College Park, MD 20742, USA

Quantum computing leverages the quantum resources of superposition and entanglement to efficiently solve computational problems considered intractable for classical computers. Examples include calculating molecular and nuclear structure, simulating strongly-interacting electron systems, and modeling aspects of material function. While substantial theoretical advances have been made in mapping these problems to quantum algorithms, there remains a large gap between the resource requirements for solving such problems and the capabilities of currently available quantum hardware. Bridging this gap will require a co-design approach, where the expression of algorithms is developed in conjunction with the hardware itself to optimize execution. Here, we describe a scalable co-design framework for solving chemistry problems on a trapped ion quantum computer, and apply it to compute the ground-state energy of the water molecule. The robust operation of the trapped ion quantum computer yields energy estimates with errors approaching the chemical accuracy, which is the target threshold necessary for predicting the rates of chemical reaction dynamics.

Quantum computation has attracted much attention for its potential to solve certain computational problems that are difficult to tackle with classical computers. For example, integer factorization[1], unsorted database search[2], and the simulation of quantum systems[3] admit quantum algorithms that outperform the best-known classical algorithms given a sufficiently large problem size. However, these algorithms require substantial quantum resources to achieve a practical advantage over classical techniques, limiting their near-term utility on noisy intermediate-scale quantum (NISQ) devices[4] that are severely limited in the number of gates they can perform before errors dominate the output. Any useful quantum computation on a NISQ device will require further advances in hardware performance, as well as advances in algorithmic design.

Quantum chemistry is a promising application where quantum computing might overcome the limitations of known classical algorithms, hampered by an exponential scaling of computational resource requirements. One of the most challenging tasks in quantum chemistry is to determine molecular energies to within chemical accuracy, defined to be the target accuracy necessary to estimate chemical reaction rates at room temperature and generally taken to be $\approx 4 \text{ kJ/mol} = 1.6 \times 10^{-3}$ Hartree (Ha)[5]. Achieving chemical accuracy would allow computational methods to replace costly experimental procedures in chemical and materials engineering, augmenting these fields to accelerate the pace of discovery.

Early quantum computational techniques to simulate many-body Fermi systems[6] or calculate molecular energies [7] have dramatically improved over the past decade[8–10], but the resource requirements for useful chemical simulations still remain out of reach[11]. Hybrid approaches might relax these requirements, where a short quantum computation serves as a subroutine to calculate classically difficult quantities. The variational quantum eigensolver (VQE) method is one example, which estimates the ground state of a system by positing an ansatz state defined by a set of variational parameters and minimizing its energy. The quantum subroutine determines the energy for a particular set of ansatz parameters, and a classical optimization algorithm iteratively updates the ansatz to reduce the energy until it converges. Early demonstrations of the VQE method have been performed on different quantum architectures [12–15], but they rely on dramatic simplifications that work for small molecules but do not readily generalize to larger systems. Additionally, systematic errors in the experimental results far exceed chemical accuracy due to imperfections in the quantum computer (QC) hardware.

Here, we provide a generic VQE approach that scales to much larger molecular systems and use it to compute the ground-state energy of the water molecule (H_2O). We embrace co-design principles to fully optimize the quantum circuits for a trapped-ion QC, and experimentally compute the first three correction terms beyond the mean-field (Hatree-Fock) approximation. We achieve

^{*}Electronic address: nam@ionq.co

[†]Electronic address: kim@ionq.co

computational errors approaching chemical accuracy for the first time, without using any error mitigation techniques. These results establish a path for future computations on more complex systems as trapped-ion QCs continue to improve, eventually reaching beyond the capability of classical methods.

I. TRAPPED-ION QUANTUM COMPUTER

The trapped-ion system used in this study is a scalable, general-purpose programmable QC constructed at IonQ, Inc.[16] and illustrated schematically in Fig. 1a; see Methods for additional details. The computer consists of a linear chain of ¹⁷¹Yb⁺ ions on a surface trap operating at room temperature, where the qubit is implemented between the $|0\rangle \equiv |F = 0, m_F = 0\rangle$ and the $|1\rangle \equiv |F = 1, m_F = 0\rangle$ hyperfine levels of the ²S_{1/2} ground state of each ion, split by 12.6 GHz[17].

The qubit register is initialized to the $|0\rangle$ state using optical pumping, and measured at the end of the computation by state-dependent fluorescence on the dipoleallowed cycling transition between $|1\rangle$ and the $^2P_{1/2}$ excited state [18]. Scattered photons from the ions during detection are collected through a high numerical aperture lens (NA ≈ 0.6) and passed through a dichroic mirror to an array of photon detectors for simultaneous readout of the entire qubit register. State preparation and measurement (SPAM) errors are routinely characterized during computation, with typical data for a three-ion chain shown in Fig. 1b. Our system exhibits a small asymmetry in the SPAM error for $|0\rangle$ versus $|1\rangle$ (0.6% and 1.3%, respectively), which is well understood from an atomic model of the detection $\operatorname{process}[19]$. We observe no evidence of measurement crosstalk where the state of one qubit affects the readout of neighboring qubits, which allows SPAM correction to be performed with low overhead. SPAM errors can be readily improved by increasing the collection efficiency of the detection optics[18-20].

Quantum gates are implemented via two-photon Raman transitions driven by two laser beams from a modelocked pulsed laser at 355 nm, where the two laser beams generate a beatnote close to the qubit frequency [21]. One of the beams is a "global" beam, with a wide profile that uniformly illuminates all qubits in the chain. The other is an array of tightly focused beams, generated from a diffractive optical element and a multi-channel acoustooptic modulator (AOM), that address the ions individually. By controlling the phase, frequency, and amplitude of these beams, we can manipulate individual qubits to implement arbitrary quantum logic gates [22]. The AOM in our system has 32 independent channels, allowing us to scale the number of individually addressable and fullyconnected qubits to this number. Further scaling is possible with alternative optical setups or by sacrificing full connectivity and using ion-shuttling protocols[23].

We drive high-fidelity single-qubit operations with a resonant Raman transition between $|0\rangle$ and $|1\rangle$ using

a composite pulse sequence [24, 25]. Two-qubit operations are mediated by the shared motional modes of the entire chain via an effective XX-Ising interaction using the Mølmer-Sørensen protocol[26, 27], and can be written in terms of Pauli X matrices on ions i and j as $XX(\theta) = \exp\left[-i\theta\sigma_x^i\sigma_x^j/2\right]$. Since the motional modes involve every ion in the chain, we can apply the XX gate between arbitrary pairs of ions with comparable speed and fidelity [22, 28–30]. This native all-to-all connectivity of two-qubit gates in the trapped ion QC provides complete flexibility to choose qubit mappings and gate configurations that maximize circuit performance on the hardware. Under typical operating conditions for this QC, the single-qubit gate fidelity can be maintained $\gtrsim 99.9\%,$ and the state fidelity of a maximally entangling $XX(\pi/2)$ gate is $\geq 96\%$. We estimate the fidelity of "small-angle" XX gates by concatenating $XX(\pi/2n)$ gates n times to approximate a full $XX(\pi/2)$ gate. The state fidelity \mathcal{F} is measured, and we estimate the per-gate error to be $\epsilon \lesssim (1-\mathcal{F})/n$. We show an example in Fig. 1c for n = 50, and calculate $\epsilon \lesssim 4 \times 10^{-3}$ for the XX($\pi/100$) gate. In general, small-angle XX gates have higher fidelities in a trapped-ion QC than maximally entangling XX gates, which can be used to improve the quality of a quantum computation.

II. MOLECULAR MODELING

We choose H_2O as a testbed for quantum co-design principles. The structure of H_2O is sufficiently complex to develop and test universal techniques for scalable quantum circuit synthesis, while simple enough to be accessible by currently available trapped-ion QCs. Simulations using classical hardware provide fully verified solutions to assess the performance of the quantum hardware, and build intuition about successful co-design strategies. What follows is a brief summary of the VQE co-design methodology, with further details supplied in the Methods.

We first write down a Hamiltonian under the Born-Oppenheimer approximation, where the atomic nuclei are fixed to their known equilibrium geometry. The Hamiltonian is represented in the second-quantized form

$$\hat{H} = \sum_{p,q} h_{pq} c_p^{\dagger} c_q + \sum_{p,q,r,s} h_{pqrs} c_p^{\dagger} c_q^{\dagger} c_r c_s , \qquad (1)$$

where $c_p^{\dagger}(c_p)$ are the creation (annihilation) operators for a molecular spin-orbital (SO) p. The SOs are spin-labeled molecular orbitals (MOs) obtained as a linear combination of atomic orbitals from the minimal STO-3G chemical basis[31] using the Hartree-Fock (HF) method[32]. The resulting 7 MOs (14 SOs) are shown schematically in Figure 2a, and the terms h_{pq} and h_{pqrs} from equation (1) are computed classically using a standard opensource tool based on *ab initio* methods[33]. The c_p and c_p^{\dagger} operators can be represented as Pauli operators acting on individual qubits using the Jordan-Wigner (JW)



000

100

Prepared State

. 111

111

FIG. 1: **Apparatus and performance.** a, Schematic representation of the trapped-ion QC. The qubit register is implemented in a linear chain of ¹⁷¹Yb⁺ ions residing inside an ultra-high vacuum chamber (not shown), and high-NA imaging optics enable individual addressing and readout of the ion qubits. The Raman beams (shown in red and purple) are generated from a pulsed laser at 355 nm and drive a two-photon transition between $|0\rangle$ and $|1\rangle$. Full control of the amplitude, frequency, and phase of the individual addressing beams enables implementations of arbitrary single– and two–qubit operators. **b**, SPAM characterization on a three-ion chain. From top to bottom, we show the SPAM error of each three-qubit state, the per-qubit SPAM error for $|0\rangle$ and $|1\rangle$, and a bar plot of the full SPAM matrix where the color is log-scaled for visibility. We see no indication of measurement crosstalk between qubits. **c**, Characterization of the small-angle XX(θ) gate performance. The state fidelity after 50 consecutive small-angle XX gates is $\approx 78\%$, and we estimate the per-gate error to be $\epsilon \lesssim 4 \times 10^{-3}$.

transformation [34], and we use the unitary coupledcluster (UCC) method to generate an ansatz state [35–37] with the first-order Trotter formula and one Trotter step. The expectation value of the Hamiltonian is computed by measuring projections of the prepared ansatz state in the combination of Pauli bases that correspond to each term in the JW-transformed Hamiltonian. To achieve meaningful accuracy, the circuit must be sufficiently sampled in each basis to reduce statistical errors [14], and systematic errors must be controlled.

а

For a small molecule like H₂O in the minimal basis set, it is possible to diagonalize the Hamiltonian in equation (1) to compute the full configuration-interaction (FCI) ground state energy (-75.0116 Ha). This energy is lower than the mean-field HF result (-74.9624 Ha)by ~ 49.2 mHa. From the FCI diagonalization, we generate a list of two-electron interaction terms $(c_p^{\dagger} c_q^{\dagger} c_r c_s)$ that contribute to modifications in the energy during the diagonalization process, with the degree of contribution characterized by the determinant. Some of these terms correspond to a pair of spin-up and spin-down electrons from the same filled MO being simultaneously excited to an empty MO (called the "bosonic" excitation terms hereafter), and the rest correspond to excitations of two electrons that are not paired in this way (see examples in Fig. 2a). Each term can be included in the preparation of the UCC ansatz in the form of $\exp \left[\theta_{pqrs}c_p^{\dagger}c_q^{\dagger}c_rc_s - c.c.\right]$ in the Trotter product formula, where θ_{pqrs} becomes the optimization parameter and h.c. denotes the Hermitian conjugate operator. We perform a numerical simulation

of the VQE process as more terms are added to the UCC ansatz, and estimate the lowest energy for each ansatz state as the parameters are optimized. This *in-silico* result serves as a reference to benchmark the computational outcome from the QC.

 \times 50

Figure 2b shows the quantum resource requirements for each UCC ansatz circuit optimized for the trappedion QC. Relevant resource metrics include the number of qubits and the number of entangling gates. We also tabulate the ground state energy from our *in-silico* VQE simulation, as up to 21 terms are added to the ansatz beyond the HF calculations (see Extended Data Fig. 7). We see that the estimate of the ground state energy approaches the FCI value as more terms are added, reaching the FCI value within chemical accuracy once 17 or more terms are included in the ansatz. Inspecting the 21 most significant determinants in the FCI energy calculation, we observe that (1) the inner-most MO $1a_1$ (see Figure 2a) is always filled and therefore can be ignored, and (2) the $1b_2$ MO participates only once as a bosonic excitation. Ignoring $1a_1$ and $1b_2$ can reduce the qubit requirement without sacrificing much in absolute accuracy: the reduced Hamiltonian reaches within 2.1 mHa of the FCI ground state at HF+16 terms using 10 qubits and 140 entangling gates. Chemical accuracy for the full Hamiltonian is achieved at HF+17 terms with 11 qubits and 143 entangling gates, of which 89 are CNOT gates and 54 are small-angle $XX(\theta)$ gates that feature higher fidelity. These resource requirements are realistically within the near-term performance targets of a NISQ computer based



FIG. 2: Circuit design for quantum chemistry. a, Molecular and spin orbital diagram, and examples of bosonic and non-bosonic excitations from the Hartree-Fock (HF) ground state. b, Metrics for each circuit, labeled HF+N, as up to N of the most significant interaction terms are added to the ansatz state. The bosonic terms through HF+5 can be represented as pair excitations to reduce the qubit resource requirements, while the MO selection strategy prunes the two least significant molecular states (1a₁ and 1b₂) to reduce the qubit count slightly at the expense of \approx mHa accuracy. Energies should be compared to the (FCI) ground state energy, which is the exact result from diagonalizing the complete Hamiltonian in the minimal STO-3G chemical basis. c, Bosonic excitation template circuit. d, Non-bosonic excitation template circuit.

on trapped ions.

III. CIRCUIT OPTIMIZATION AND CO-DESIGN

We have implemented a number of circuit optimization techniques that take advantage of the unique features available in the IonQ trapped-ion QC, but are generic in the sense that they are applicable to any target molecule to be simulated. The strategies described here are executed by a full-stack, modularized software toolchain, which automatically produces optimized circuits[38] for generating the ansatz state of a molecular system.

Given a general UCC ansatz state, interaction terms take the form of a two-electron interaction $\theta_{pqrs}c_p^{\dagger}c_q^{\dagger}c_rc_s$. Since the indices p, q, r, s vary over the complete set of molecular states (which are represented by different qubits), implementing this interaction requires entangling gates between arbitrary pairs of qubits in the system. The all-to-all connectivity of trapped-ion QCs makes this a native operation, eliminating the overhead incurred by repeated SWAP gates to reorder qubits before an entangling operation can be applied between nearest-neighbor qubits. Given that the infidelity of these SWAP operations can dominate the quality of complicated computations, eliminating them from the optimized circuit dramatically increases the accuracy of the VQE result. This circuit optimization is a direct result of co-design for a particular hardware advantage.

Another general optimization strategy is to represent the bosonic excitations, where two electrons remain paired, as a single creation or annihilation operator. It is convenient to expand the SO label for the operator c_p to $c_{k\alpha}$, where k and α denote the MO and spin label, respectively. Then the bosonic operators are $d_k^{\dagger} = c_{k\alpha}^{\dagger} c_{k\beta}^{\dagger}$ and $d_j = c_{j\alpha} c_{j\beta} \ (\alpha \neq \beta)$, which can be directly translated to the Pauli raising/lowering operators σ_{\pm}^j on qubit j. The UCC operator corresponding to the bosonic excitation simplifies to $\exp[\theta_{jk} \sigma_{+}^j \sigma_{-}^k - h.c.]$, and a pair of arbitraryangle XX(θ) gates is sufficient to implement this interaction (see Fig. 2c). Thus ansatz states containing only bosonic excitations can be implemented very efficiently on our QC.

For all remaining terms, we must implement the twoelectron interaction via the JW transformation. Each of these terms looks like $\hat{V} = \exp[\theta_{pqrs}\sigma^p_+\sigma^q_+\sigma^r_-\sigma^s_-\bigotimes_k\sigma^k_z - h.c.]$, where the product $\bigotimes_k \sigma^k_z$ denotes the adequate JW string to reflect fermionic symmetries. A JW string with $m \sigma_z$ gates converts to m CNOT gates on either side of the subcircuit that would otherwise implement \hat{V} . Properly ordering these terms in the entire circuit can eliminate most of the CNOT gates so they represent a relatively low overhead as a function of terms in the UCC



FIG. 3: **Experimental Results. a**, Calibration curve for the small-angle XX(θ) gate. **b**, Bootstrap distribution of the HF+2 energy with mean and 1 σ uncertainty indicated by the orange diamond **b**, and experimentally determined energy surfaces for HF+1 **c** and HF+2 **d**. Each data point (orange diamond) represents an average of $\gtrsim 1000$ experimental runs, with the blue dash-dotted lines indicating *in-silico* results. **e**, Comparison of ground-state energy estimates as additional interactions are included in the UCC ansatz state (labeled HF+N, for N significant determinants). The orange diamonds indicate experimental results, with error bars indicating 1 σ uncertainty from the bootstrap distribution. The remaining points are from the *in-silico* VQE simulation as detailed in Figure 2, and show how the ansatz states converge to the full configuration-interaction ground state, indicated by the dot-dashed blue line. Chemical accuracy about the numeric results is indicated by the shaded green region in all figures.

ansatz[8]. The main portion of the quantum circuit is an implementation of a linear combination of eight terms, each containing a product of four σ_x and σ_y operators (with odd number of σ_x in each term). By optimizing the order of these eight operators and taking advantage of the all-to-all connectivity, we can implement this circuit with 13 CNOT gates (see Fig. 2d). When we concatenate several of these terms, some CNOT gates at the ends, including those that arise from a JW string, may cancel out.

Most ansatz states have both bosonic and non-bosonic excitation terms. For these situations, we start with the reduced representation where each qubit describes one MO, and run the quantum circuit that corresponds to all bosonic excitation terms first. Then, additional qubits (all prepared in the $|0\rangle$ state) are introduced, and each are entangled with a qubit representing a MO using a CNOT gate. Each entangled pair can now represent the two SOs corresponding to the MO (Extended Data Fig. 4e).

One last optimization takes advantage of the asymmetric SPAM error observed in our system. Normally, we encode a filled orbital (MO or SO) with $|1\rangle$ and an empty orbital with $|0\rangle$, but in a molecule with mostly closed molecular shells like H₂O, the filled orbitals in the HF-ground state remain mostly filled in the FCI ground state as well. Since our SPAM error is more than a factor of two smaller for $|0\rangle$ compared to $|1\rangle$, we encode the filled orbitals as $|0\rangle$ to reduce the systematic shift asso-

ciated with readout from the $|1\rangle$ state. This encoding has the ancillary benefit of requiring fewer single-qubit gates to initialize the circuit, but the advantage diminishes as measurement errors are suppressed or become more symmetric.

Combining these strategies, we achieve the quantum circuits for preparing the ansatz state with total entangling gate counts shown in Fig. 2b. These methods represent a fully general, scalable, and near-optimal framework that can be applied to simulating any physical systems using VQE with a UCC ansatz.

IV. RESULTS

Using our trapped-ion QC, we compute the first three bosonic excitation terms of the VQE ansatz for the H₂O molecule (Extended Data Fig. 4). In order to estimate the minimum energy for each circuit within chemical accuracy (fractional uncertainty of $\sim 10^{-5}$), all systematic errors in our QC must be carefully characterized and controlled. The intrinsic decoherence of a ¹⁷¹Yb⁺ trapped-ion qubit is negligible over the timescale of our computation[17], so the dominant errors arise in calibrating the angle of the XX(θ) gate and correcting for the systematic SPAM error of our ion chain. We accurately calibrate the angle θ using a circuit similar to that shown in Fig. 2c, where the parity varies as $\sin(2\theta)$. Fitting the parity to this functional form (Fig. 3a) compensates for non-linearities in the AOM and enables easy interpolation for arbitrary gate angles. Uncertainty in the SPAM correction can be made arbitrarily small given sufficient measurement statistics. Gate fidelity will begin to dominate as the computation length increases, but for the circuits experimentally demonstrated here we are not limited by this error and found no benefit to error mitigation techniques like Richardson extrapolation[39].

To compute the energy corresponding to a prepared ansatz state, we make a set of projective measurements in bases corresponding to the terms in the Hamiltonian, as previously described. Once a sufficient number of measurements are made, we use a statistical bootstrapping technique^[40] that accounts for SPAM error to estimate uncertainties from the resulting histograms (Fig. 3b). Figure 3c-d shows the experimentally determined energy surface for HF+1 and HF+2 as the ansatz parameters $\{\theta_i\}$ are scanned about their optimum values, and the data for HF+3 is shown in the Extended Data Figure 5. The experimentally determined groundstate energies for each of these three ansatz states is -74.977(1) Ha, -74.979(2) Ha, and -74.985(5) Ha, respectively, with parenthetical errors indicating 1σ uncertainty derived from the bootstrapped distribution. The dominant experimental uncertainty arises in the SPAM correction, which can be improved with upgrades to the hardware and new tomographic methods [18, 41]. A direct comparison to the *in-silico* VQE simulation can be found in Fig. 3e. The match to theory is very good both the absolute accuracy and precision are comparable to the chemical accuracy. This is critical to achieve quantum computational results that provide predictability for VQE-type optimization algorithms. Given the circuit requirements outlined in Fig. 2b and the gate fidelity achieved, we believe that executing a full simulation towards the FCI energy is within reach using current trapped-ion quantum computer technology.

V. SUMMARY AND OUTLOOK

Dramatic improvements must be made to both QC hardware and techniques to efficiently use the available quantum resources in order to perform meaningful quantum computations on a NISQ device. The work presented here is a framework for end-to-end optimization that maps useful problems in quantum chemistry to a

- Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.* 41, 303–332 (1999).
- [2] Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* **79**, 325 (1997).
- [3] Feynman, R. P. Simulating physics with computers. Int. J. Theor. Phys. 21, 467–488 (1982).
- [4] Preskill, J. Quantum computing in the NISQ era and

trapped-ion QC, fully leveraging the hardware-specific advantages. This framework yields near-optimal quantum circuits that can be run on existing NISQ hardware. To verify the performance of both the hardware and the optimization procedure, we compute the post-Hartree-Fock ground state energy of H_2O on a trapped-ion QC. Without any error mitigation, the experimental results for the first three correction terms are in excellent agreement with the theory at the level of chemical accuracy, both in the predicted values and their precision.

While these results are specific to a particular quantum chemistry problem and the trapped-ion QC hardware, the computational methodology we develop is completely general to simulating quantum systems. We anticipate that similar advances can be applied to other optimization problems that work on variational methods, such as the quantum approximate optimization algorithm[42] and various quantum machine learning applications[43, 44]. Increased attention to co-design principles like those demonstrated here will be necessary to push the boundary of possibility in near-term quantum computation.

Author Contributions

Experimental data collected and analyzed by N.C.P. and JS.C.; Y.N., C.D., D.M. K.B. and J.K. performed the circuit design and Y.N. performed *in-silico* VQE simulation; the apparatus was designed and built by K.W., J.M.A., K.M.B., JS.C., M.C., S.D., K.M.H., J.M., N.C.P., J.D.WC., S.M.K, S.A, J.A., P.S., M.W., A.M.D., A.B., V.C., M.K., C.C., C.M. and J.K.; Y.N., N.C.P, JS.C. and J.K. prepared the manuscript, with input from all authors.

Competing Interests

The authors declare no competing interests.

Correspondence

Correspondence and requests for materials should be addressed to Yunseong Nam (email: nam@ionq.co).

beyond. arXiv:1801.00862 (2018).

- [5] McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. & Yuan, X. Quantum computational chemistry. arXiv:1808.10402 (2018).
- [6] Abrams, D. S. & Lloyd, S. Simulation of many-body fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589 (1997).
- [7] Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-

Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707 (2005).

- [8] Hastings, M. B., Wecker, D., Bauer, B. & Troyer, M. Improving quantum algorithms for quantum chemistry. *Quant. Inf. Comput.* 15, 0001–0021 (2015).
- [9] Babbush, R. et al. Exponentially more precise quantum simulation of fermions in the configuration interaction representation. Quantum Sci. Technol. 3, 015006 (2017).
- [10] Berry, D. W. et al. Improved techniques for preparing eigenstates of fermionic hamiltonians. npj Quant. Inf. 4, 22 (2018).
- [11] Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *PNAS* **114**, 7555–7560 (2017).
- [12] Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. Nat. Commun. 5, 4213 (2014).
- [13] O'malley, P. et al. Scalable quantum simulation of molecular energies. Phys. Rev. X 6, 031007 (2016).
- [14] Kandala, A. *et al.* Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549, 242 (2017).
- [15] Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. Phys. Rev. X 8, 031022 (2018).
- [16] Ionq, inc. https://ionq.co/. Accessed: 2018-10-29.
- [17] Olmschenk, S. *et al.* Manipulation and detection of a trapped Yb⁺ hyperfine qubit. *Phys. Rev. A* 76, 052314 (2007).
- [18] Noek, R. et al. High speed, high fidelity detection of an atomic hyperfine qubit. Opt. Lett. 38, 4735–4738 (2013).
- [19] Acton, M. et al. Near-perfect simultaneous measurement of a qubit register. Quantum Inf. Comp. 6, 465 (2006).
- [20] Myerson, A. H. et al. High-fidelity readout of trapped-ion qubits. Phys. Rev. Lett. 100, 200502 (2008).
- [21] Hayes, D. et al. Entanglement of atomic qubits using an optical frequency comb. Phys. Rev. Lett. 104, 140501 (2010).
- [22] Debnath, S. et al. Demonstration of a small programmable quantum computer with atomic qubits. Nature 536, 63 (2016).
- [23] Kielpinski, D., Monroe, C. & Wineland, D. J. Architecture for a large-scale ion-trap quantum computer. *Nature* 417, 709–711 (2002).
- [24] Brown, K. R., Harrow, A. W. & Chuang, I. L. Arbitrarily accurate composite pulse sequences. *Phys. Rev. A* 70, 052318 (2004).
- [25] Mount, E. et al. Error compensation of single-qubit gates in a surface-electrode ion trap using composite pulses. *Phys. Rev. A* 92, 060301 (2015).
- [26] Sørensen, A. & Mølmer, K. Quantum computation with ions in thermal motion. *Phys. Rev. Lett.* 82, 1971–1974 (1999).
- [27] Choi, T. et al. Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement. Phys. Rev. Lett. 112, 190502 (2014).
- [28] Zhu, S.-L., Monroe, C. & Duan, L.-M. Arbitrary-speed quantum gates within large ion crystals through minimum control of laser beams. *Europhys. Lett.* 73, 485–491 (2006).
- [29] Zhu, S.-L., Monroe, C. & Duan, L.-M. Trapped ion quantum computation with transverse phonon modes. *Phys. Rev. Lett.* 97, 050505 (2006).
- [30] Linke, N. M. et al. Experimental comparison of two quan-

tum computing architectures. PNAS **114**, 3305–3310 (2017).

- [31] Taketa, H., Huzinaga, S. & O-ohata, K. Gaussianexpansion methods for molecular integrals. J. Phys. Soc. Jpn. 21, 2313–2324 (1966).
- [32] Szabo, A. & Ostlund, N. S. Modern quantum chemistry: introduction to advanced electronic structure theory (Courier Corp., 2012).
- [33] Turney, J. M. et al. Psi4: an open-source ab initio electronic structure program. Wiley Interdiscip. Rev.: Comput. Mol. Sci. 2, 556–565 (2012).
- [34] Jordan, P. & Wigner, E. Über das paulische äquivalenzverbot. Z. Physik 47, 631–651 (1928).
- [35] Čížek, J. On the correlation problem in atomic and molecular systems. calculation of wavefunction components in ursell-type expansion using quantum-field theoretical methods. J. Chem. Phys. 45, 4256–4266 (1966).
- [36] Hoffmann, M. R. & Simons, J. A unitary multiconfigurational coupledcluster method: Theory and applications. J. Chem. Phys. 88, 993–1002 (1988).
- [37] Bartlett, R. J., Kucharski, S. A. & Noga, J. Alternative coupled-cluster ansätze II. the unitary coupled-cluster method. *Chem. Phys. Lett.* **155**, 133–140 (1989).
- [38] Nam, Y., Ross, N. J., Su, Y., Childs, A. M. & Maslov, D. Automated optimization of large quantum circuits with continuous parameters. *npj Quant. Inf.* 4, 23 (2018).
- [39] Richardson, L. F. & Gaunt, J. A. VIII. the deferred approach to the limit. *Phil. Trans. R. Soc. Lond. A* 226, 299–361 (1927).
- [40] Efron, B. & Tibshirani, R. J. An Introduction to the Bootstrap (CRC press, 1994).
- [41] Keith, A. C., Baldwin, C. H., Glancy, S. & Knill, E. Joint quantum state and measurement tomography with incomplete measurements. *Phys. Rev. A* 98, 042318 (2018).
- [42] Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. arXiv:1411.4028 (2014).
- [43] Gao, X., Zhang, Z.-Y. & Duan, L.-M. A quantum machine learning algorithm based on generative models. *Sci*ence Advances 4 (2018).
- [44] Zhu, D. et al. Training of quantum circuits on a hybrid quantum computer. arXiv:1812.08862 (2018).
- [45] Knill, E. et al. Randomized benchmarking of quantum gates. Phys. Rev. A 77, 012307 (2008).
- [46] Shen, C. & Duan, L.-M. Correcting detection errors in quantum state engineering through data processing. *New J. Phys.* 14, 053053 (2012).
- [47] Ballance, C., Harty, T., Linke, N., Sepiol, M. & Lucas, D. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.* **117**, 060504 (2016).
- [48] Blume-Kohout, R. et al. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. Nat. Commun. 8, 14485 (2017).
- [49] Suzuki, M. General theory of fractal path integrals with applications to manybody theories and statistical physics. J. Math. Phys. 32, 400–407 (1991).
- [50] Childs, A. M., Maslov, D., Nam, Y., Ross, N. J. & Su, Y. Toward the first quantum simulation with quantum speedup. *PNAS* **115**, 9456–9461 (2018).

Methods

A. SPAM characterization and correction.

We monitor the SPAM error during the computation by interleaving experiments that prepare the all-bright $|11\cdots 11\rangle$ and all-dark $|00\cdots 00\rangle$ states of the ion chain. During detection, we relax the axial confinement such that the crosstalk error between adjacent ions is below the intrinsic dark count error of our photon detectors. This enables a fast measurement of the entire SPAM matrix, because it decomposes into a Kronecker product of SPAM matrices on the individual ions. The average single-ion SPAM errors, along with an example SPAM matrix of a three-ion chain, are summarized in Fig. 1b, which are independently verified via single-qubit randomized benchmarking[45]. The measured state vectors are then corrected for SPAM error via matrix inversion before computing expectation values of the circuit Hamiltonian [46, 47]. Note that for the QC performing computations here, the SPAM error is dominated by detection fidelity[19], which can be improved significantly by upgrading the photon collection efficiency of the hardware.

B. Single-qubit operation.

Single-qubit gates are implemented via Rabi oscillations between $|0\rangle$ and $|1\rangle$, using an SK1 composite pulse sequence[24]. The transition is driven by a two-photon Raman process from two separate beams generated by a pulsed 355 nm laser, where one beam is tightly focused onto the ion so it can be manipulated independently from the rest of the chain. We prepare arbitrary single-qubit states on the Bloch sphere by controlling the phase and pulse area delivered by the individual addressing beam. Conventional randomized benchmarking [45] and gate set tomography^[48] techniques are used to characterize our single-qubit gates, and we observe fidelities $\geq 99.9\%$ with good repeatability. A typical randomized benchmarking result is depicted in the Extended Data Fig. 6. During the computation, we interleave experiments to calibrate the laser power based on measured Rabi rates. Singlequbit gates are also used to measure SPAM error in the $|1\rangle$ state. Although we do not characterize the singlequbit gate fidelity during the computation, we are able to bound it $\geq 99\%$ by monitoring the SPAM error.

C. Two-qubit operation.

Two-qubit entangling gates are implemented via the Mølmer-Sørenson interaction, where an amplitudemodulated laser pulse, composed of non-copropagating beams, achieves full spin-motion decoupling at the end of the gate[22, 27–30]. One of the laser beams uniformly illuminates the ions, while the other individually addresses the two particular ions involved in the gate. By varying the pulse area through the laser intensities, we control the geometric phase θ of the interaction, defined in terms of the Pauli-X matrices σ_x^i on the *i*-th ion as

$$XX_{ij}(\theta) \equiv e^{-i\theta\sigma_x^i \sigma_x^j/2}, \qquad (2)$$

where $\theta = \pi/2$ corresponds to a maximally entangling gate. We calibrate the small-angle XX gate ($\theta \ll \pi/2$) with the circuit presented in Extended Data Fig. 4a, projecting to the σ_x basis before measurement. An overall amplitude scale factor g, applied to the acoustooptic modulator (AOM) controlling the individual addressing beams, is scanned to vary the geometric phase $\theta \propto g^2 \equiv \Theta$, with $g^2 < 0$ realized by shifting the rf phase of the AOM thereby inverting $\Theta \rightarrow -\Theta$. The parity Π_{ij} of the output state, defined in terms of the observed two-qubit state probability P_{ij} as

$$\Pi_{ij} = P_{ij} (|00\rangle) + P_{ij} (|11\rangle) - P_{ij} (|10\rangle) - P_{ij} (|01\rangle), \quad (3)$$

is fit to $\Pi_{ij} = \sin(2k\Theta)$, and the factor k is applied to the AOM amplitude scale factor g to properly calibrate arbitrary-angle XX(θ) gates. An example calibration curve is shown in Fig. 3a. Small-angle two qubit gates have high fidelity, which we estimate to be $\gtrsim 99.6\%$ for the $\theta = \pi/100$ case (Fig. 1d), which is likely explained by the smaller absolute geometric phase errors and weaker light shifts. A method for rigorous characterization of an arbitrary-angle XX gate fidelity remains an open question for future study.

D. Water molecule.

We use the nuclear configuration of H₂O where the O–H bond length is fixed to $\sim 1.8a_0$ and the angle between the two O–H bonds to be $\sim 105^{\circ}$, where a_0 is the Bohr radius.

E. Molecular modeling with qubits.

We start with the second-quantized Hamiltonian of the form

$$\hat{H} = \sum_{p,q} t_{pq} a_p^{\dagger} a_q + \sum_{p,q,r,s} t_{pqrs} a_p^{\dagger} a_q^{\dagger} a_r a_s, \qquad (4)$$

where $a_p^{\dagger}(a_p)$ is the creation (annihilation) operator for an atomic orbital (AO) with a given spin. We evaluate $t_{pq} = \iint \chi_p^{*(1)} \hat{H}_1 \chi_q^{(2)} d\vec{r}^{(1)} d\vec{r}^{(2)}$ and $t_{pqrs} = \iint \chi_p^{*(1)} \chi_q^{*(2)} \hat{H}_2 \chi_r^{(2)} \chi_s^{(1)} d\vec{r}^{(1)} d\vec{r}^{(2)}$ using the minimal, STO-3G basis [31], where $\chi_p^{(i)}$ is the *p*-th AO for *i*-th electron and \hat{H}_j is the *j*-body electronic Hamiltonian. We then employ the HF method and solve the Roothan equation using the self-consistent approach[32] to obtain molecular orbitals (MOs) as a linear combination of AOs. Each MO consists of two spin orbitals (SOs) corresponding to each spin state of the electron (up and down). The resulting Hamiltonian in SOs is given in Eq. (1) of the main text.

We then apply Jordan-Wigner (JW) transform[34]

$$c_j^{\dagger} = 1^{\otimes j-1} \otimes \sigma_+^j \otimes \sigma_z^{\otimes N-j}, \quad c_j = 1^{\otimes j-1} \otimes \sigma_-^j \otimes \sigma_z^{\otimes N-j}$$
(5)

to transform our reduced, physical, SO-basis Hamiltonian to a qubit Hamiltonian. Additionally, we choose the unitary coupled-cluster (UCC) as our ansatz[36, 37] and apply the same transformation. Specifically, for the circuit-level implementation of the JW-transformed ansatz circuit $\exp(-iH_{c.c.})$, the effective Hamiltonian, coupled-cluster (CC) operator $H_{c.c.}$ is given by $H_{c.c.} = i(T - T^{\dagger})$, where T is the excitation operator. We use n-th order product-formula (PF) algorithms to further decompose the CC operator, according to

$$\exp\left(-i\sum_{j=1}^{L}\theta_{j}H_{j}\right) \approx [S_{n}(\lambda)]^{r}, \qquad (6)$$

where $H_{c.c.} = \sum_{j} \theta_{j} H_{j}$, $\lambda := 1/r$, and

$$S_{1}(\lambda) := \prod_{j=1}^{L} \exp(-i\theta_{j}H_{j}\lambda),$$

$$S_{2}(\lambda) := \prod_{j=1}^{L} \exp(-i\theta_{j}H_{j}\lambda/2) \prod_{j=L}^{1} \exp(-i\theta_{j}H_{j}\lambda/2),$$

$$S_{2k}(\lambda) := [S_{2k-2}(p_{k}\lambda)]^{2}S_{2k-2}((1-4p_{k})\lambda)[S_{2k-2}(p_{k}\lambda)]^{2},$$
(7)

with $p_k := 1/(4 - 4^{1/(2k-1)})$ for k > 1 [49]. We choose the PF algorithm among many available methods, as the PF resulted in the smallest algorithmic error for certain Hamiltonian simulations while requiring the least quantum resources[50]. We find it sufficient to use the first order Trotter formula (n = 1) with a single PF stage (r = 1), as the resulting ansatz leads to the FCI ground state energy within chemical accuracy.

F. Expectation values of the Hamiltonian.

Once the ansatz state is prepared, the expected energy of the state is evaluated by projecting the state onto a set of different multi-qubit Pauli bases, obtained from the JW transformation of the SO Hamiltonian in Eq. (1). Specifically, a set of basis-transformation operations, depicted in Extended Data Fig. 4d, are applied prior to measurement in the computational basis. The expectation of each Pauli string can be computed at the desired level of accuracy by sufficient statistical sampling. Due to the coefficients in the Hamiltonian, some terms contribute more than others to the final computed energy. We minimized the number of basis transformations required to evaluate the entire Hamiltonian by carefully partitioning the Pauli strings, grouping those that can be measured with a single, representative measurement circuit.

G. Circuit for bosonic excitation terms.

For the cases where the pair of electrons from one MO are excited together to another MO, we can represent this process as a "bosonic" excitation. It is convenient to expand the SO index to reflect both MO (j, k, etc.) and spin (α or β) indices, e.g., $c_p \equiv c_{k\alpha}$. Then the bosonic creation and annihilation operators are $d_k^{\dagger} = c_{k\alpha}^{\dagger} c_{k\beta}^{\dagger}$ or $d_j = c_{j\alpha}c_{j\beta} \ (\alpha \neq \beta)$. In these bosonic excitations, we can use one creation (annihilation) operator to represent the occupation of both electrons in that orbital. Since these creation and annihilation operators now correspond to a pair of electrons, the anti-commutation relations that apply to fermionic particles no longer hold—thus the term "bosonic excitation"—and the σ_z^i strings in the JW transformation are not necessary. For these bosonic excitation terms, the operator d_k^{\dagger} (d_k) can be replaced by the qubit operator σ_{+}^{k} (σ_{-}^{k}), and the corresponding UCC operator reduces to $\exp[\theta_{kj}\sigma^k_+\sigma^j_- - h.c.]$. The implementation of this term in a quantum circuit requires two XX gates between the two qubits (k and j) with an angle of θ_{kj} , denoted as $XX_{kj}(\theta_{kj})$, as shown in Fig. 2c.

H. Circuit for non-bosonic excitation terms.

Several steps are taken to minimize the number of entangling gates in implementing the quantum circuit for generic two-electron interaction terms.

1. First, we need to allocate a qubit to represent each SO considered in the molecule. By putting frequentlyinteracting SOs closer to each other, we can reduce the product of σ_z operators introduced in the JW transformation (called JW strings). A simple greedy approach is used to identify the most frequent interaction between different SOs, and the identified SO pairs are mapped to nearest-possible qubits until no further mapping may be made.

2. An adequate ordering of the individual terms in the JW-transformed $H_{c.c}$ in the Trotterized ansatz circuit can lead to dramatic reduction in the JW strings between the terms [8]. We order the terms in the ansatz states such that the adjacent terms have maximal overlap in non-identity elements in the Pauli product, while making sure that at least one overlap is non- σ_z . This maximizes the cancellation of the JW strings between the adjacent terms, simplifying the resulting circuit.

3. A two-electron excitation term that consists of $\operatorname{Im}(c_p^p c_q^{\dagger} c_r c_s)$ translates to the product operator $\operatorname{Im}(\sigma_+^p \sigma_+^q \sigma_-^r \sigma_-^s \bigotimes_{\nu} \sigma_z^{\nu})$ by the JW transformation, where $\sigma_{\pm} = (\sigma_x \mp i \sigma_y)/2$ and the σ_z product is the JW string. The corresponding unitary operator takes the form $\exp[-i\theta(\bigotimes_{\nu} \sigma_i^{\nu})/2]$, where $\sigma_i^{\nu} \in \{1, \sigma_x, \sigma_y, \sigma_z\}$ is the Pauli operator for ν -th qubit. We implement the circuit as shown in (8)a below, instead of a nearest-neighborinspired, "staircase" construction (up to SWAP gates) such as in (8)b. Here, we assumed that all $\sigma_i^{\nu} = \sigma_z$ as a concrete example. The circuit implementation of the term $\exp[-i\theta \bigotimes_{\nu=0}^{N-1} \sigma_i^{\nu}/2]$ results in 2(m-1) CNOT gates, where *m* is the number of elements σ_i that are not identity, and the common target qubit for the CNOT gates can be chosen arbitrarily among the *m* qubits.



To optimize the resulting quantum circuit, we start by choosing the common target to be one of the four qubits to which either σ_x or σ_y is applied. Here, as a concrete example, we choose the last such qubit as the common target. Once the JW strings are expressed by the CNOT gates, the rest of the circuit consists of eight terms that can be ordered as $\sigma_x \sigma_x \sigma_x \sigma_y$, $\sigma_x \sigma_x \sigma_y \sigma_x$, $\sigma_y \sigma_x \sigma_y \sigma_y, \ \sigma_y \sigma_x \sigma_x \sigma_x, \ \sigma_y \sigma_y \sigma_x \sigma_y, \ \sigma_y \sigma_y \sigma_y \sigma_x, \ \sigma_x \sigma_y \sigma_y \sigma_y,$ and $\sigma_x \sigma_y \sigma_x \sigma_x$. This ordering, while not a unique example, allows us to better optimize the quantum circuit. First, we perform basis transformation $\sigma_i \mapsto \sigma_z$ by applying H or $S^{\dagger}H$ for i = x or y, respectively. Then, the resulting term always looks like a product of four σ_z gates, and the circuit shown in (8)a can be used to implement each term. The particular order has exactly two basis changes $(\sigma_x \leftrightarrow \sigma_y)$ between adjacent terms resulting in the cancellation of two pairs of CNOT gates between them, leaving only two CNOT gates between each term. This cancellation is not applicable for the staircase implementation shown in (8)b.

4. The two CNOT gates between the two terms discussed above can be further simplified, and reduced to a single CNOT gate using the following circuit identity:

$$\begin{array}{c} \bullet & \bullet \\ \bullet & \bullet \\$$

Noting that any z-rotations, including S and S[†], commute with the control of a CNOT gate, and HSH = $S^{\dagger}HS^{\dagger}$ and $HS^{\dagger}H = SHS$ up to a global phase, we can further simplify the single-qubit gates in the circuit by commuting the S, S[†], and H gates to either end of the circuit through the CNOT gates. The resulting circuit to simulate any two-electron term reduces to Fig. 2d, and consists of 13 CNOT gates, eight small-angle rotations about z-axis, and a few H, S, and S[†] gates to account for the relevant basis rotations.

I. Circuit efficiency.

For a two-electron interaction term over four qubits, we start with eight instances of the circuit shown in (8) containing 48 CNOT gates. Our final template circuit in Fig. 2d, which contains 13 entangling gates, is likely optimal. At least one multi-qubit gate needs to be expended to transform from one Pauli basis to another, and our design requires exactly one two-qubit gate for such transformation. If a quantum circuit is written with CNOT as the only available multi-qubit gate, it takes at least three CNOT gates to compute or uncompute a Pauli product of length four, such as $\sigma_x \sigma_x \sigma_x \sigma_y$. Given each two-electron excitation operator in the PF-based implementation includes eight such Pauli products, the 13 CNOT gates in Fig. 2d is likely a minimum. This circuit optimization is completely general for a two-electron interaction term applied to any molecule using any basis set and any order of PF (see (7)).

J. Higher fidelity two-qubit gates.

The number of entangling gates is not the sole factor that determines the quality of a quantum computation: it also depends on the type of entangling gate used. Specifically, a small-angle XX gate performs better in our trapped-ion QC than a CNOT gate, which requires the maximally entangling $XX(\pi/2)$ gate. Therefore, it is advantageous to convert CNOT gates to small-angle XX gates wherever possible.

Shown below is a circuit replacement rule that may be used to convert CNOT gates to $XX(\theta)$ gates. The rule can be applied to the target circuit in Fig. 2d to replace four CNOT gates with four XX gates. Assuming the infidelity associated with a small-angle XX gate is much lower than that of a CNOT gate, applying this procedure replaces four out of 13 CNOT gates with small-angle XX gates, bringing about a ~30% reduction in the twoelectron interaction infidelity.



K. Combining bosonic and non-bosonic terms.

A typical UCC ansatz state contains both bosonic and non-bosonic excitation terms. In this general case, we first assign N qubits to represent the occupation of N MOs with a pair of spin-up and spin-down electrons, and execute the circuit corresponding to all bosonic excitation terms first. Then, we introduce an additional N ancilla qubits all prepared in $|0\rangle$ state. We apply N CNOT gates, each employing a qubit used in simulating the bosonic excitation terms as the control qubit and a fresh ancilla qubit as the target. Each pair can now be used to represent the two SOs corresponding to the MO in executing the remaining non-bosonic excitation terms in the circuit. An example circuit for N = 4 is shown in Extended Data Figure 4e.

The savings obtained from employing this hybrid approach in executing the bosonic and non-bosonic excitation terms diminish as the problem size increases, as there are more non-pair excitations than pair excitations in general. Nevertheless, for NISQ devices where every saving matters, we find the reduction to be non-negligible.

L. Selection of spin-orbitals.

In the simulation for small molecules such as H_2O where the solution can be simulated classically, we have full knowledge over which SOs, each mapping to a qubit, participate in our simulations. In such cases where the information regarding which SOs are significant in finding the ground state of the molecule, we may drop those SOs that are least significant from our simulation. Such information is likely accessible even in the superclassical regime since, for instance, we may reasonably expect that core orbitals are less likely to participate in the excitation than valence orbitals.

Restricting the simulation to the most significant SOs results in qubit count savings, at the small cost of leaving out those less significant interactions that should in principle be accounted for in obtaining the ground state energy. In terms of implementation, one may allocate the freed-up quantum resources to simulating interactions between the more significant SOs. This is thus useful whenever the number of qubits is a limiting factor.

M. Bootstrap error analysis.

In deriving error estimates for the computed energies, we employ the empirical Bootstrap technique[40]. All single-shot samples for a particular implementation of the circuit are binned, and a random bootstrap sample S^* of the same size as the original dataset is drawn with replacement from the data. Similarly, we bin all SPAM characterization data acquired during a particular run, and draw a bootstrap sample from that data as well. This is repeated five hundred times, and with each bootstrap sample we compute the circuit energy expectation $\langle H_0 \rangle$ to build a histogram of possible measurements consistent with the empirical data. The mean of this distribution is reported as the measured value, and the standard deviation provides a 1σ error estimate.



Extended Data Figure 4: VQE ansatz circuits for **a**, HF+1, **b**, HF+2, and **c**, HF+3. Each circuit ends in a measurement M_i , which is one of three measurement circuits shown in **d**, to generate expectation values of each Pauli operator. For HF+1, shown here in the bosonic representation (see text), the electron-pair excitation from the $1b_1$ MO to the $2b_1$ MO is achieved by applying two XX(θ) gates between qubits denoting the occupation of electrons in these MOs. The second bosonic excitation added in HF+2 (corresponding to electron pairs from the $3a_1$ MO being excited to the $4a_1$ MO) involves a completely independent set of MOs, so it amounts to parallel execution of the circuit in **a** with different values of the ansatz parameters θ_0 and θ_1 . For HF+3, the circuit involves six XX(θ) gates with three independent parameters. All interaction terms are listed in Extended Data Figure 7. The NOT gates can be commuted through the entire circuit and applied classically after the measurement which, due to asymmetric SPAM error in the $|1\rangle$ state, improves the fidelity of the computation. **e**, A sample circuit mixing bosonic and non-bosonic excitation terms. The bosonic excitations are simulated first on a subset of the qubits representing the electron pairs in each MO, then CNOT gates between each MO qubit and an ancilla serve to re-encode the MOs as SOs and the circuit proceeds to execute the remaining non-bosonic interactions. For all XX gates, the two-qubit interaction is applied between the top and bottom qubit lines enclosed by the box.



Extended Data Figure 5: Bootstrap distribution for the experimentally computed energy in HF+3, compared to the *in-silico* result (blue dash-dotted line) and bounds of chemical accuracy (green shaded region). We determine $\langle H_0 \rangle = -74.985(5)$ Ha, indicated by the orange diamond. Error bars signify 1σ uncertainty.



Extended Data Figure 6: Typical randomized benchmarkting data for single-qubit gates, using an SK1 composite pulse sequence. The extracted fidelity is 99.964(5)%.

Orbital							
$2b_1$	↑ ↓	— ♦	★ —			↑ ↓	- +
$4a_1$		▲ —	- +	+ +	+ +		▲ —
$1\mathrm{b}_{2}$	↑ ↓	↑ ↓	↑ ↓	+ +	↑ ↓	↑ ↓	↑ ↓
$3a_1$	★ ↓	— ♦	— ♦		↑ ↓		★ —
$1b_1$		♠ —	★ —	↑ ↓		↑ ↓	- +
$2a_1$	↑ ↓	+ +					
$1a_1$	↑ ↓	+ +					
	▲ —				- +	▲ —	
	-+	★ ↓	+ +	★ ↓	★ —	-+	↑ ↓
	↑ ↓						
	- +	— ♦	★ —	↑ ↓	↑ ↓	↑ ↓	↑ ↓
	▲ —	↑ ↓	↑ ↓	↑ ↓	♠ —	— ♦	★ ↓
	★ ↓	★ —	— ♦		— ♦	♠ —	★ ↓
	+ +	↑ ↓	↑ ↓	↑ ↓	↑ ↓	↑ ↓	+ +
	- +	▲ —	+ +			- +	▲ —
	▲ —	- +		- +	▲ —	- +	▲ —
	↑ ↓	↑ ↓	↑ ↓	↑ ↓	▲ ♦	↑ ↓	↑ ↓
	↑ ↓	↑ ↓	↑ ↓	♠ —	— ♦	★ —	- +
	- +	★ —	↑ ↓	↑ ↓	↑ ↓	★ —	- +
	★ —	— ♦		↑ ↓	↑ ↓	↑ ↓	★ ↓
	↑ ↓	↑ ↓	↑ ↓	↑ ↓	+ +	↑ ↓	★ ↓

Extended Data Figure 7: Most significant interaction terms beyond Hartree-Fock.