Reduced density matrix sampling: Self-consistent embedding and multiscale electronic structure on current generation quantum computers

Jules Tilly [1,2,*] P. V. Sriluckshmy,[3] Akashkumar Patel [2,4] Enrico Fontana,[4,5] Ivan Rungger,[4] Edward Grant,[1,6]
Robert Anderson,[3] Jonathan Tennyson [2] and George H. Booth [3]

[1]*Rahko Limited, N4 3JP London, United Kingdom*
[2]*Department of Physics and Astronomy, University College London, WC1E 6BT London, United Kingdom*
[3]*Department of Physics, King's College London, Strand, WC2R 2LS London, United Kingdom*
[4]*National Physical Laboratory, TW11 0LW Teddington, United Kingdom*
[5]*Department of Computing, University of Strathclyde, G1 1XQ Glasgow, United Kingdom*
[6]*Department of Computer Science, University College London, WC1E 6BT London, United Kingdom*

We investigate fully self-consistent multiscale quantum-classical algorithms on current generation superconducting quantum computers, in a unified approach to tackle the correlated electronic structure of large systems in both quantum chemistry and condensed matter physics. In both of these contexts, a strongly correlated quantum region of the extended system is isolated and self-consistently coupled to its environment via the sampling of reduced density matrices. We analyze the viability of current generation quantum devices to provide the required fidelity of these objects for a robust and efficient optimization of this subspace. We show that with a simple error mitigation strategy these self-consistent algorithms are indeed highly robust, even in the presence of significant noises on quantum hardware. Furthermore, we demonstrate the use of these density matrices for the sampling of nonenergetic properties, including dipole moments and Fermi liquid parameters in condensed phase systems, achieving a reliable accuracy with sparse sampling. It appears that uncertainties derived from the iterative optimization of these subspaces is smaller than variances in the energy for a single subspace optimization with current quantum hardware. This boosts the prospect for routine self-consistency to improve the choice of correlated subspaces in hybrid quantum-classical approaches to electronic structure for large systems in this multiscale fashion.

I. INTRODUCTION

A solution to the quantum many-body problem is held up as a one of the most impactful and far-reaching applications of quantum computers [1–4]. Even in the current era of noisy intermediate-scale quantum (NISQ) devices [5], where the number of physical qubits is too small for error correction and subject to significant decoherence and quantum noise, significant progress in this area has been made, with developments also impacting the related field of quantum machine learning [6]. However, quantum resources are finite, and so treating entire systems of technological relevance is unlikely to be a realistic near-term proposition. Instead, efficient and practical hybrid quantum-classical and multiscale approaches are key to describing quantum systems with a desired accuracy within a limited quantum computational budget.

The general aim of these is to find a representation of the full system in which only a small number of "active" degrees of freedom are strongly entangled and require realization on a quantum device. The remaining degrees of freedom are then coupled to this quantum subsystem in a low-rank perturbative or even (potentially dynamical) mean-field representation of the quantum effects. These "external" weakly-interacting degrees of freedom can be efficiently described on a classical computer with polynomially-scaling resources with system size, while it is posited that an appropriate choice of strongly entangled "active" orbitals should not in general need to grow with system size. Self-consistency can then be used to update the initial choice of the active region, based on a multilevel description of the full system. One of the key technical considerations in devising these approaches is in the choice of how the description of the active space on the quantum processing unit (QPU) can be efficiently coupled to the external space. In this paper, we consider the sampling of reduced density matrices (RDMs) as the natural choice of low-rank quantum variable for NISQ computers, in order to describe the coupling of the classical and quantum realizations of the system and therefore straddle the length scales in the simulation of quantum matter. We present a unified description of this approach, with applications to both strongly correlated quantum chemistry and condensed matter problems on current generation quantum devices.

While multilevel descriptions of quantum systems are in themselves not a new proposition for quantum computers

*jules.tilly@rahko.ai

[7–15], the practical realization of fully self-consistent algorithms on quantum resources has proved to be a significant technical challenge. The ability of self-consistent coupled quantum-classical approaches to remain robust in the presence of the noise inherent in the sampling of the active-space quantum effects is a key practical consideration. We demonstrate that with a light-touch error mitigation strategy, noise resulting from gate infidelities does not preclude stable convergence of the algorithms presented in this paper. Furthermore, an efficient sampling and grouping scheme for the terms required in density matrices is presented, detailing a huge reduction in the number of these terms for a given active-space size, which tackles a key bottleneck for longer-term practical quantum-classical multiscale methods.

We demonstrate the feasibility of two of these self-consistent algorithms in a unified approach applicable to both quantum chemistry and condensed matter physics. The first is complete active-space self-consistent field (CASSCF) [16,17], a powerful approach for the simulation of molecular systems with strong quantum effects, such as those encountered routinely in inorganic chemistry, systems with competing spin states, excited states, and systems at bond-breaking geometries [18–20]. In these problems, the dominant strong quantum fluctuations can often be qualitatively captured within a small number of low-energy orbitals, where these orbitals are obtained from a prior mean-field calculation. However, since this starting point does not account for the correlated physics, these active orbitals are subsequently optimized in a variational fashion across the overall state as a product of the active-space wave function and mean-field electrons in the external space.

As systems get larger, and the spectrum of single-particle states transitions towards a continuous function of energy, the approach of choosing quantum regions using an energetic criteria becomes less well defined. In this case, it is more appropriate to consider the active quantum region in a real-space local picture, under the assumption that local quantum fluctuations dominate, such as those in the d orbitals of a transition metal oxide material. This is the approach taken in various quantum embedding methods such as dynamical mean-field theory [21], and in this paper we consider the related energy-weighted density matrix embedding theory (EwDMET) [22–24]. We perform the self-consistency at the level of energy weighted density matrices, denoting the moments of the local density of states, resulting in a systematic expansion of the zero-temperature DMFT physics [24]. For both of these approaches, we demonstrate the fidelity of the QPU sampling of the active-space RDMs required for a fully QPU-coupled self-consistent algorithm, and consider the scaling of sampling operations as the active space increases in size in future applications.

In Sec. II we review reduced density matrices and their sampling within QPUs as the self-consistent quantum variables in these multiscale methods. We demonstrate that judicious grouping of commuting terms allows even large active-space RDMs to be realistically sampled, with the proposed groupings opening the prospect for higher-rank RDMs and perturbative couplings to active spaces. In Sec. III we consider the CASSCF optimization of carbon monoxide on IBM quantum services (IBMQ) machines, where active-space

wave function are optimized on the QPU within the variational quantum eigensolver (VQE) [25]. We find that a naive implementation fails to achieve full self-consistency, however a light-touch error mitigation strategy aimed at compensating for gate errors is enough to address this issue on a 4-qubit system and to allow convergence of the active-space orbitals in the presence of the correlated physics. Our error mitigation strategy focuses on estimating the bias due to gate errors in the circuit execution via computing the deviation of the sampled electron number to the (known) number of electrons in the system (which can be found from the trace of the sampled one-body RDM). This approach is detailed in Appendix A. With this mitigation strategy, the sampling noise of the QPU does not prevent convergence, and CASSCF is found to be robust and reliable in its optimization. We stress that knowledge of the active-space RDMs also allows for the extraction of beyond-energetic first-order expectation values of the system, which are essential for a more complete description of the system properties, focusing on the effect of self-consistency and RDM fidelity on the dipole moment of a system.

Finally, in Sec. IV we focus on extended bulk systems, with the strongly correlated Bethe-Hubbard lattice considered. Specifically, we observe the QPU description of a local region to allow for the opening of Hubbard bands in the material within the QPU-coupled EwDMET approach. Similar error mitigation strategies on the sampled active-space RDMs allows for robust self-consistency in the method, resulting in excellent agreement for the local density of states and Matsubara self-energy of the system on current generation IBMQ machines.

II. SAMPLING REDUCED DENSITY MATRICES ON A QUANTUM COMPUTER

The reduced density matrices (RDM) used in this work are not defined by tracing out a subsystem, but rather tracing out the entire phase space of many electrons from the full N-particle density matrix of a pure state. For a m-body reduced density matrix, Γ_m, this integration over $N - m$ electronic variables can be written as

$$\Gamma_m(\mathbf{x}_1, \ldots, \mathbf{x}_m; \mathbf{x}'_1, \ldots, \mathbf{x}'_m)$$
$$= m! \binom{N}{m} \int \Gamma_N(\{\mathbf{x}\}_N; \{\mathbf{x}'\}_N) \Delta_{m+1}^N d\mathbf{x}_{m+1} \ldots$$
$$d\mathbf{x}_N d\mathbf{x}'_{m+1} d\mathbf{x}'_N, \tag{1}$$

where $\mathbf{x}_i$ represents the combined spatial and spin coordinate for electron i, and $\Delta_{m+1}^N = \prod_{i=m+1}^N \delta(\mathbf{x}_i - \mathbf{x}'_i)$. Fortunately, these reduced-body density matrices can be directly computed, rather than requiring tracing from higher-rank density matrices. By projecting the electronic coordinates into a basis set, we define the two-body RDM as

$$\Gamma_{ijkl} \equiv \langle \psi | \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_l \hat{a}_k | \psi \rangle, \tag{2}$$

with other rank RDMs defined equivalently, and where the indices $i, j, \ldots$ label spin-orbital degrees of freedom, $\hat{a}_i^{(\dagger)}$ are the fermionic annihilation (creation) operators, and we have omitted the explicit subscript denoting the rank of the RDM where it is obvious from the number of indices. In this example, the partial trace down to the one-body RDM can

then be written as

$$\gamma_{ik} = \frac{1}{N-1} \sum_j \Gamma_{ij,kj}. \tag{3}$$

Despite tracing out large numbers of degrees of freedom, these two-body RDMs still contain all the information about a quantum system required for physical observables of interest, which depend on (up to) pairwise operators, including the total energy. The rank of an operator defining a given observable determines the rank of the RDM required to compute its corresponding expectation value. For example, the electric dipole moment is a one-body quantity, requiring the one-body RDM, while the Hamiltonian defining the energy is a two-body expectation value, requiring the two-body RDM to evaluate. Nonobservable quantities of interest, such as entanglement entropies or mutual information, can also in general be computed from reduced-body density matrices [26].

Furthermore, using RDMs we can compute the probability of a given m-electron distribution, as the diagonal of the m-RDM. The sum over this distribution then gives the number of m tuples of particles in the system, which can be used as a normalization condition, e.g., the trace of the 2-electron distribution giving the number of pairs of electrons, as

$$\sum_{ij} \Gamma_{ij,ij} = \frac{N(N-1)}{2}. \tag{4}$$

Overall, these m-RDMs have all the information about the distribution and entanglement of m particles in a given state of an N particle system, which rationalizes their use as method-agnostic, low-rank quantities in order to couple quantum systems described at different levels of theory.

In this paper, we consider second-quantized Hamiltonians where spin symmetry is preserved, allowing further tracing out of spin degrees of freedom, defining the central spin-free two-body RDM of interest as

$$\overline{\Gamma}_{pqrs} \equiv \sum_{\sigma\tau} \Gamma_{p_\sigma q_\tau r_\sigma s_\tau}, \tag{5}$$

where $p, q, \dots$ denote spatial degrees of freedom and σ, τ denote spin labels. Further permutational symmetries can be used which reduce the number of independent quantities to evaluate, as

$$\Gamma_{p_\sigma q_\tau r_\sigma s_\tau} = -\Gamma_{p_\sigma q_\tau s_\tau r_\sigma} = \Gamma_{q_\tau p_\sigma s_\tau r_\sigma} = \Gamma_{r_\sigma s_\tau p_\sigma q_\tau}, \tag{6}$$

with time-reversal symmetry ensuring

$$\Gamma_{p_\sigma q_\tau r_\sigma s_\tau} = \Gamma_{p_{\bar{\tau}} q_{\bar{\sigma}} r_{\bar{\tau}} s_{\bar{\sigma}}}. \tag{7}$$

The resulting set of fermionic operators must be mapped to spin operators for sampling of the state on a QPU. For this, we use the Jordan-Wigner mapping [27], though other mappings (e.g., Bravyi-Kitaev mapping [28,29]) could be used, as long as the mapping is consistent. While each fermionic operator will in general be mapped to several spin operators, one can find several efficiencies to reduce this overall number of terms. We discuss scaling of number of terms for RDM sampling using efficient grouping methods in Appendix B, which will be essential for scaling to larger numbers of qubits or for the extraction of beyond two-body properties.
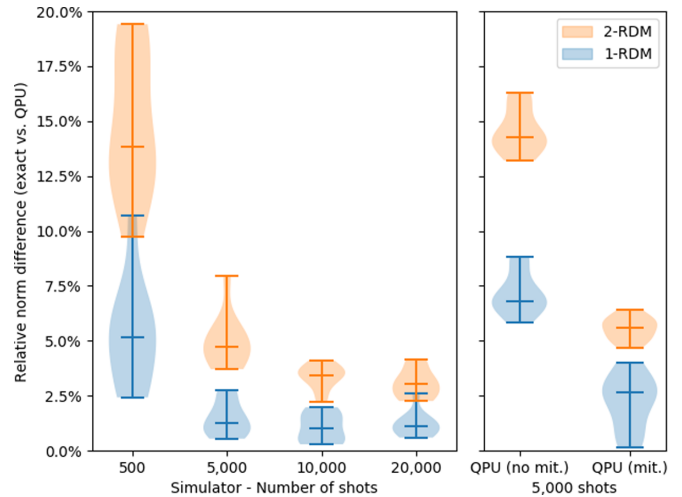

FIG. 1. Distribution of relative errors in the Frobenius norm for both QPU and quantum emulated sampling of the RDMs, compared to exact classical calculation. To obtain the distributions, we repeated the computation of the Frobenius norms differences over 20 realizations of the RDMs for each number of measurements considered. On the left, results are presented for a QPU simulator (assuming perfect qubits) and therefore displays the impact of finite sampling noise. On the right, the results computed on IBMQ Athens (4-qubits, depth 3 HEA), with and without error mitigation.

In order to test RDM sampling on a quantum computer, we computed the one- and two-body RDMs of magnesium porphyrin after optimization of a hardware efficient ansatz (HEA) wave function using the gradient free RotoSolve optimizer [30]. We use an active space restricted to 2 orbitals and 2 electrons (a 4-qubit Hamiltonian) for this test. It is worth noting that while the HEA is convenient for studying small systems due to its relatively small prefactor, it is expected to have difficulties scaling to larger active spaces [31], where other VQE ansatz are expected to be preferred such as the unitary coupled cluster [32]. The resulting RDMs are then compared to the RDMs obtained via exact methods for the same active space, and the distribution of relative errors in the elements are shown in Fig. 1. We conducted this test on a simulated QPU with different number of shots, as well as current quantum hardware (IBMQ Athens QPU, details about each QPU used in this paper can be found in Appendix C).

In our investigation of RDM observable sampling errors, we also considered approaches to reliably mitigate for these errors via physically justified extrapolation techniques. It is clear from the results of Fig. 1 that these can substantially ameliorate quantum noise and sampling errors. We present the simple extrapolation technique used in this paper in Appendix A; the technique relies on a binomial distribution of independent errors in the quantum circuit, and has the benefit of requiring no additional measurements (for a more involved approach to extrapolation, we recommend Ref. [33]).

This simple error mitigation technique significantly improves the overall accuracy of the QPU estimates, bringing it almost on par with the results of the simulator at an equivalent number of measurements. A key question that remains is whether the norm error presented above has a significant

impact on the ability to use these RDMs reliably within subsequent quantum chemical calculations, where manipulation of these noisy RDMs may prevent convergence or lead to unacceptable bias in desired quantities. In order to test this we apply this sampling to a QPU-solved complete active-space self-consistent field method (computing both energetics and molecular dipole moments following the optimization), as well as a QPU version of the energy-weighted density matrix embedding theory, as examples of multiscale approaches to allow quantum resources to be applied to realistic systems in electronic structure calculations.

III. QUANTUM CASSCF

The complete active-space self-consistent field (CASSCF) approach is generally the starting point in quantum chemistry for molecular systems exhibiting stronger correlation effects, and therefore a key step in the development of electronic structure methods suitable for quantum computation [16,20,34–48]. The central tenet of CASSCF is that the dominant strong quantum fluctuations required to qualitatively describe an electronic system are spanned by a small number of low-energy degrees of freedom about the chemical potential. The changes caused by explicitly considering interaction-driven virtual excitations in this space can change the occupation and induce entanglement of these orbitals, giving rise to correlated physics far from a mean-field description. The first step of CASSCF is therefore to partition the orbitals into three subspaces, denoted core, active, and virtual. Core orbitals are deep-lying orbitals, which are considered to be chemically-inert and fully occupied, while conversely, the virtual orbitals are considered high-energy states which remain unoccupied. Together, these denote the "external" space. The active space denotes the degrees of freedom, which are considered to span the dominant electron correlations corresponding to low-energy virtual excitations of the N_{act} electrons within it, with the full set of quantum fluctuations amongst this set to be considered. No entanglement or particle/spin fluctuations are considered between the external and active spaces. The overall CASSCF wave function at any point can therefore be written as

$$|\Psi_{CASSCF}\rangle = |\psi_{active}\rangle \otimes \det[\phi_c], \qquad (8)$$

where $|\psi_{active}\rangle$ denotes an N_{act}-electron wave function spanning the active degrees of freedom, while $\det[\phi_c]$ is a single product state over the core orbitals, accounting for the $N - N_{act}$ remaining electrons.

A key initial step for CASSCF is therefore to choose the orbitals in each set. These are selected from an initial mean-field calculation, where to a first approximation, the highest-energy occupied and lowest-energy virtual orbitals about the chemical potential are chosen as the active space. However, this choice is often augmented with other criteria for selection of the active space, including symmetry, locality and/or "chemical intuition", with approaches for automatic selection of this space, e.g., from quantum information arguments, a source of recent developments [1,49,50]. However, it is clear that selecting these orbitals from an initial mean-field calculation has an inherent flaw. The active space, designed to capture the strong correlations and dominant entanglement

between single-particle states, is chosen from a theory with no correlations or entanglement via simple mean-field orbital energetics, which can change substantially in the presence of electron correlation. To account for this, a self-consistency in the choice of the active-space orbitals is required for meaningful and qualitatively accurate results in the presence of strong correlation. This involves a variational optimization of the state given in Eq. (8), to account for an arbitrary mixing between all three classes of orbitals, defined by the exponential of an anti-Hermitian one-particle operator. This allows the character of the orbitals to change, by rotating core and virtual components into the active space in a variationally optimal way.

The CASSCF method from another perspective can be considered as an embedding of the correlated effects of the active space into a mean-field description of its "environment" (as given by the electrons in core orbitals), as presented in Ref. [14] for quantum emulation. However, this active "embedding region" is chosen largely on energetic criteria, with a strictly separable form and no entanglement with the core electrons. We contrast this with an alternative criteria based on locality in Sec. IV. The limitations of the approach come from the size of the active space, which for an exact treatment is often accepted to be 16 electrons in 16 orbitals [20], with some instances of computation up to 20 electrons in 20 orbitals [51]. This is due to the exponential scaling of classical resources with respect to this size in order to represent and optimize $|\psi_{active}\rangle$. Beyond this, approximate descriptions of the active-space wave function are increasingly being investigated, although all have their limitations [52–57]. This active size constraint stymies the application of CASSCF to systems with larger valence spaces, where a small active space is not sufficient and convergence of desired properties with respect to active size is not reached.

This limitation is a potential opportunity for NISQ computers to exhibit a quantum advantage in this keystone method in quantum chemistry, with the active-space paradigm often being touted as a near-term prospect for quantum computers [1,12] (see, e.g., Ref. [58] for a recent review of the limits of classical computers in this field and the requirements for quantum advantage). However, beyond simple analysis of gate depth and qubit number required is the question of the practical feasibility of a robust and convergent algorithm for the self-consistency of the full CASSCF method, which has to date been recently demonstrated for a single-orbital optimization step without full self-consistency in the paper of Takeshita *et al.* [12]. In the algorithm which we use, the coupling of the active-space correlations to the orbital rotations required for self-consistency is provided by the two-body RDM within the active space. Therefore, the faithful sampling of this two-body RDM with sufficient fidelity is critical for a well-behaved algorithm. This is especially important as the orbital rotation updates involve nonlinear functionals of the sampled two-body RDM elements, meaning that we expect noise from the QPU sampling to manifest as systematic error in the final results, even in the case that the sampling of the underlying RDM elements is unbiased. We investigate the two-body RDM active-space sampling for this purpose on QPUs as well as the importance of error mitigation, by using a parameterized gate circuit as the active-space wave function

optimized via the variational quantum eigensolver (VQE). However, the use of VQE in this paper could be replaced by quantum Krylov or imaginary-time solvers suitable for NISQ devices [59,60], as well as quantum phase estimation algorithms when suitable devices are available.

A. Fully self-consistent algorithm

We briefly summarize the key steps of the (two-step) CASSCF approach (sometimes also described as the related multiconfigurational self-consistent field method), with more details available in Ref. [17]. We start with the second quantized electronic Hamiltonian in a basis, as

$$\hat{H} = \sum_{ij} h_{ij} \hat{a}_i^\dagger \hat{a}_j + \frac{1}{2} \sum_{ijkl} g_{ijkl} \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_l \hat{a}_k + E_{\text{nuc}}, \qquad (9)$$

where h_{ij} and $g_{ijkl} = \langle ij|kl \rangle$ are the one and two-body integrals respectively, with E_{nuc} the scalar nuclear repulsion. We parametrize the orbitally-optimized CASSCF wave function of Eq. (8) as

$$|\Psi_{\text{CASSCF}}\rangle = |\mathbf{R}, \mathbf{c}\rangle = e^{-\sum_{ij} R_{ij} \hat{a}_j^\dagger \hat{a}_i} \sum_n c_n |n\rangle, \qquad (10)$$

where the one-body matrix operator $\mathbf{R}$ parametrizes the single-particle unitary rotation operator of the molecular orbital basis ($R_{ij} = -R_{ji}$), $|n\rangle$ the complete set of Slater determinants spanning the active space, and $\mathbf{c}$ defines the coefficients of the configurations indexed by n, spanning the selected active space. The full optimization problem can then be written as

$$E = \min_{\mathbf{R}, \mathbf{c}} \frac{\langle \mathbf{R}, \mathbf{c}| H |\mathbf{R}, \mathbf{c}\rangle}{\langle \mathbf{R}, \mathbf{c}|\mathbf{R}, \mathbf{c}\rangle}. \qquad (11)$$

Within the two-step algorithm, the optimization of $\mathbf{R}$ and $\mathbf{c}$ are treated separately and alternated, as the optimization of $\mathbf{R}$ can be efficiently performed for one-body unitary rotations on classical computers, given the knowledge of the active-space two-body RDM,

$$\Gamma_2 = \langle \psi_{\text{active}}| \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_l \hat{a}_k |\psi_{\text{active}}\rangle. \qquad (12)$$

The optimization of $\mathbf{R}$ then proceeds via construction of the gradient and Hessian of the energy with respect to these parameters, which can then be updated at modest computational expense via a quasi-second-order step, accelerated with iterative subspace methods as implemented in the PySCF package [17,61,62].

For a given rotation matrix parameterized by $\mathbf{R}$, the Hamiltonian can then be transformed into the new basis, and the Coulomb and exchange contribution from the static core electrons integrated out, resulting in an active-space Hamiltonian $H_{\text{act}}(\mathbf{R})$, which only spans the active-space degrees of freedom. The optimization of this active-space wave function is then amenable to implementation within a VQE minimization, as

$$E|_{\mathbf{R}} = \min_\theta \langle \psi_{\text{active}}(\theta)| H_{\text{act}}(\mathbf{R}) |\psi_{\text{active}}(\theta)\rangle, \qquad (13)$$

where θ denote the angles to optimize within the chosen quantum circuit parameters [25]. Once optimized, the 2RDM elements of the active space of Eq. (12) can be sampled, in

order to update $\mathbf{R}$ in the full space, until convergence. In practice, convergence can be triggered in a black-box fashion over a number of variables, such as the energy, orbital gradients, or density matrices themselves. However in this paper we run for a fixed number (ten) of orbital updates, which is sufficient to gauge convergence and to subsequently observe the fluctuations in the quantities of interest at this point.

B. Results

While the computational procedures for this coupled orbital optimization are well developed for exact or near-exact solvers in the domain of quantum chemistry, their utility in a fully self-consistent algorithm with a noisy quantum computer is far from clear (although there is some relevant recent work on noisy Monte Carlo solvers for active spaces [54,55,63,64]). We therefore consider the CASSCF algorithm with an active-space NISQ device solver, to determine the stability of the algorithm in the presence of sampling, gate noise, decoherence, and a parameterized quantum circuit for the state. This allows us to understand the feasibility of this multiscale approach, and develop practical strategies to ameliorate potential shortcomings from the noisy active-space sampling.

We apply the method to a carbon monoxide (CO) molecule in a cc-pVDZ basis set, and at a stretched bond length of 1.54 Å. This stretching of the multiple bond enhances the strong correlation in the electronic structure, as the atomic-like character of the constituent atoms is increased. An active space of two orbitals and two electrons, corresponding to the highest occupied and lowest unoccupied molecular orbitals, is selected to capture the dominant many-body entanglement in these lowest-energy quantum fluctuations. To ensure a significant level of orbital relaxation from the self-consistent procedure, and to test the stability of this noisy optimization in the case of a poor initial choice of orbitals, we select initial orbitals (and active space) from only a partially converged Hartree-Fock calculation. This was achieved by an early stopping of the mean-field self-consistent field procedure after only two updates of the Fock matrix prior to the CASSCF.

We first implemented the method on a quantum simulator with 500, 2000, 5000, and 10 000 shots to sample each mapped two-body RDM operator required, but in the absence of any additional noise model for the gates. For the quantum hardware experiments, we use IBMQ Bogota and IBMQ Santiago, which are both 5-qubit QPUs available through the the IBMQ platform, with equivalent levels of gate fidelity (details about each QPU used in this paper can be found in Appendix C). The initial calculations on IBMQ Bogota were performed without accounting for any error mitigation, before applying the light-touch error mitigation strategy presented in Appendix A on the IBMQ Santiago hardware to assess any improvements from this. For the QPU runs, we use a measurement ramp-up schedule whereby the number of measurements is increased if the output energy at a given iteration is higher than for the previous one (which should not be the case during the optimization). It is capped at 8000 shots, which is also the number of measurements used for RDM sampling after the state is optimized.

We used the same ansatz for all experiments, built on a four-qubit, three-layer version of the HEA [65,66], and the
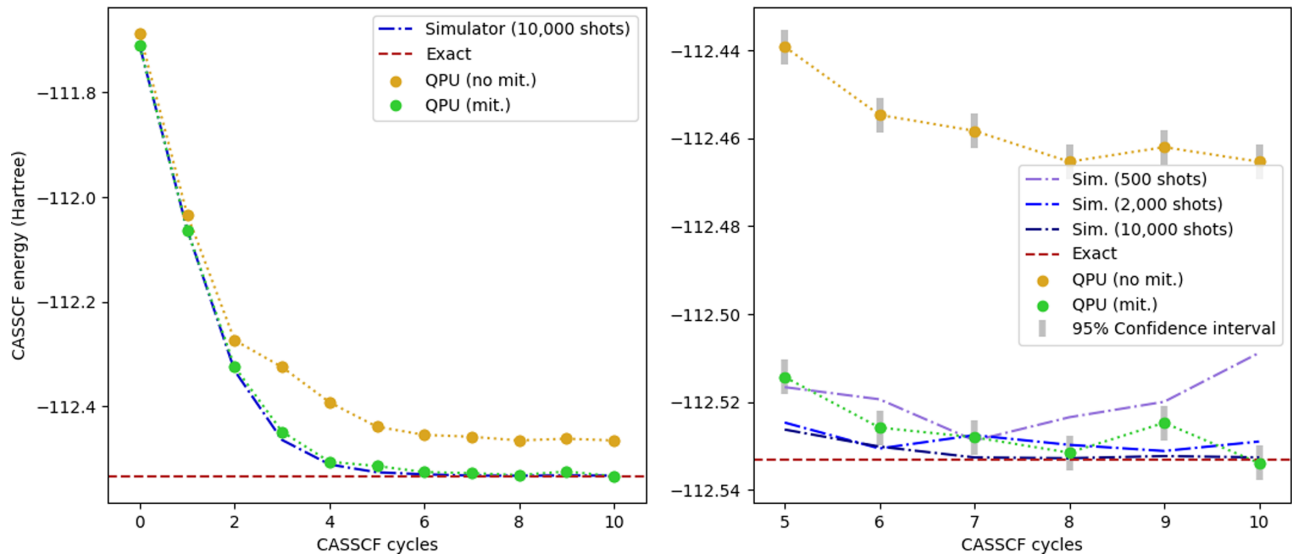
FIG. 2. Convergence of the energy of the CASSCF state for a 2-electron, 2-orbital active space of carbon monoxide for each orbital update step. Results are shown for a quantum simulator, on IBMQ Bogota (without error mitigation) and on IBMQ Santiago (with error mitigation). The plot on the right focuses on the final five orbital update cycles, showing the variation and bias in converged results, with additional simulated results for 500 and 2000 shots to illustrate the impact of finite sampling noise (results for 5000 shots are indistinguishable from those obtained with 10 000 shots and as such were not included). Error bars for QPU results on the right hand side plot represent 1.96 times the standard error spread of the measurement data, or the 95% confidence interval for the values estimated.

same optimizer: the gradient free RotoSolve methods [30]. This resulted in a total of 24 variational parameters in the model. We found it unnecessary to fully converge the ansatz each iteration, and therefore investigated varying the level of ansatz optimization each CASSCF step to improve efficiency. Five iterations of the VQE were in general sufficient on the first cycle, and we then used the parameters obtained to initialize the ansatz for the next cycle. A single iteration of the VQE for subsequent CASSCF steps after performing this warm start was sufficient to fully converge in a reasonable time, and to reach good accuracy.

The results of these CASSCF optimizations are presented in Fig. 2. Without error mitigation, the QPU results show significant systematic error at convergence of $\sim$60 mE_h, but nevertheless allow for a stable optimization. Including the error mitigation allows for significantly better results, with fluctuations of less that 10 mE_h from the exact CASSCF value from exact 2-step optimization of the same initial active space. As expected, the variance from the QPU experiment is significantly more than the corresponding quantum simulated results, even with error mitigation. This reflects the fact that the error mitigation effectively removes the bias in the sampled measurements, but does not materially improve on the variance resulting from quantum noises. In our quantum simulated results, we find strong convergence for any simulation without gate noise or decoherence above 5000 shots. Below that number, finite sampling noise prevents the algorithm reaching the sought after solution. At 500 shots, it fails to reach under 10 mE_h difference to the target state energy on average.

We can distinguish and isolate the effects of certain errors arising from the quantum solver on these results. Firstly, we have the systematic error in the VQE at each iteration, including the optimization, gate errors, and ansatz choice,

which lead to a nonexact energy and state for a given set of active orbitals. Secondly, we can consider the effect of stochastic noise in the RDM due to a finite number of samples. This second error will lead to incorrect orbital updates in the CASSCF macroiterations, and a loss of precision in the final CASSCF energy due to an inability to propagate to the optimal orbitals defining the active space and its Hamiltonian. Furthermore, since the orbital choices in CASSCF are not linear functionals of the sampled density matrix elements, even if the RDM elements are entirely unbiased and correct on average, this does not preclude a systematic error entering the orbital updates at any finite sampling.

To separate these sources of error, we can consider the exact energy of each CASSCF iteration, but using the active-space orbitals obtained at each iteration from the noisy VQE update from the quantum solver. This eliminates errors due to the VQE optimization of a given active space, isolating the error due to convergence of the nonoptimal orbital rotations being found at each step, primarily due to the inherent sampling noise of the RDMs. These results are shown in Fig. 3, and show that the overwhelming majority of the error is arising from the bias in the VQE, while the convergence of the orbitals is highly robust to the errors in the active-space VQE description and sampling errors of the RDMs. Even without error mitigation of the RDMs, orbital optimization is accurate to within 10 mE_h, while the results of the simple error mitigation in the RDMs rendered an almost exact CASSCF energy. This demonstrates that the orbital optimization procedure is less susceptible to the errors in the VQE and RDM sampling than the inherent errors in the energy and wave function optimization for a given active space. This relative insensitivity to the fidelty of the RDMs bodes well for larger active-space calculations on QPUs and the practicality of orbital optimization through RDM sampling, as well as the
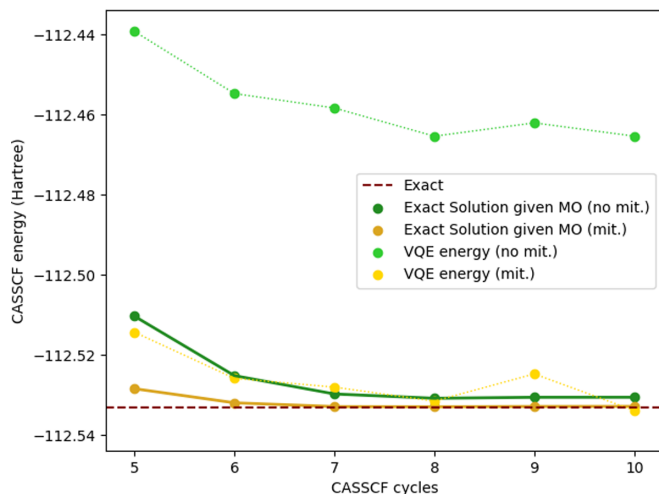
FIG. 3. Convergence of the final six orbital update cycles of the CASSCF energy, obtained with the IBMQ Bogota (with no error mitigation) and IBMQ Santiago (with error mitigation). Two series are presented for each QPU calculation: "VQE energy" results are equivalent to Fig. 2, while "exact solution" represent the exact energy from the current active space, that would have been obtained if the VQE was solved perfectly given the molecular orbitals found from the previous quantum VQE update step.
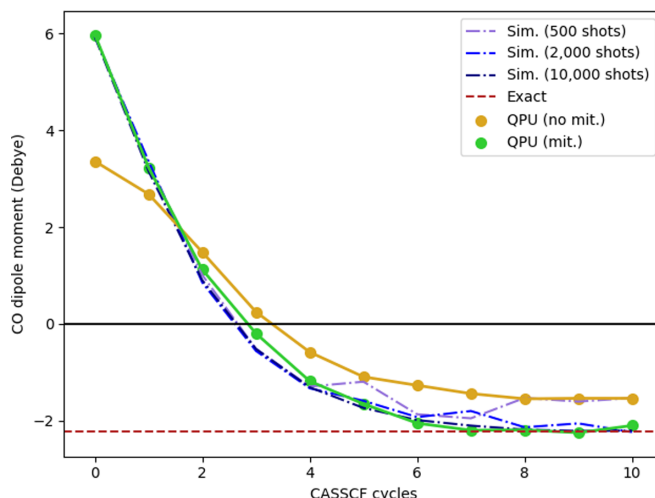

FIG. 4. Convergence of the dipole moment in Debye from QPU-CASSCF as the orbitals are optimized, both with (IBMQ Santiago) and without (IBMQ Bogota) error mitigation, as well as quantum simulated results from an RDM sampling of 500, 2000, and 10 000 shots. Positive dipole moments mean that the dipole moment points towards the oxygen (i.e., the oxygen atom has net negative charge), while the converged results flip the orientation of the dipole moment.

improvements which would transfer to this approach from improved active-space quantum algorithms [59,60].

A key question remains as to whether this robustness is a property just of orbital optimization, or whether this also extends to a broader set of expectation values which can be derived from the RDM sampling (other than the energy), as these also relax due to a more faithful description of the correlated wave function. We consider here the effect of orbital optimization and a correlated VQE wave function on the magnitude of the dipole moment of the carbon monoxide molecule in the same active space, which can be extracted from the sampled one-body reduced density matrix as a one-particle expectation value. This quantity characterizes the net charge distribution in the molecule, and from symmetry constraints can be described by a vectorial quantity which must be coincident with the carbon-oxygen bond. The magnitude of this vector is shown in Fig. 4, as the active-space orbitals are optimized in the presence of the correlated VQE state.

Since the dipole moments are linear functionals of the one-body RDM, we would expect an unbiased sampling of the RDM to give an unbiased estimate of the dipole moment from the optimized VQE-CASSCF state. We find that without error mitigation, there is still an error of $\sim 0.6D$, however the error mitigated results can effectively reduce the systematic error in the final dipole moment completely, with fluctuations in each cycle of a similar magnitude to the emulated values without quantum noise or decoherence with 5000 shots. It is also worth noting that the dipole moment of carbon monoxide is notoriously difficult to converge theoretically [67], hence the error with respect to experiment of 2 Debye is to be expected for CASSCF with a minimal active space. At this point, the fluctuations in the dipole moment agree with the magnitude of the fluctuations expected from the original density matrix sampling experiments of Fig. 1, and an unbiased estimate

of the exact CASSCF dipole moment is obtained. We also note that the correlation and orbital optimization reverses the direction of the charge imbalance in this system from the starting description.

The overall runtime of the full QPU-CASSCF calculations on IBMQ Bogota was ~ 14 QPU-hours, including 10 orbital updates, VQE optimization, and RDM sampling.

However, one must consider the potential for parallelization. In this case, all 8000 measurements conducted on the 49 operators could have been conducted in parallel, possibly reducing the overall runtime up to well under a second. While this is not a good indication for scaling and long term viability of the method (we encourage the reader to refer to Ref. [58] for an interesting assessment of the scalability of VQE), it does illustrate the potential for strong parallelization, and corresponding error mitigation techniques, for the viability of NISQ algorithms.

Finally, it is worth discussing the viability of extensions to CASSCF on quantum devices. In quantum chemistry, CASSCF is rarely the end of the story, as it neglects the contributions to expectation values arising from interactions between the active-space electrons and the external degrees of freedom. These can generally be treated at a perturbative level of theory [68], cumulant or energy-moment expansions [69] or via subspace expansions [12], and are required for quantitative accuracy for predictive calculations. These perturbative couplings between the spaces can be computed by considering the higher-rank RDMs in the active space. This approaches will significantly increase the number of terms which must be sampled. However, large reductions can be found with the appropriate groupings, and this is demonstrated in Appendix B, where 440 154 Pauli strings for the sampling of the four-body RDM within 6 qubits can be reduced to only

3182 commuting groups. We will explore the viability of this perturbative extension to CASSCF in future work.

IV. QPU-ENHANCED ENERGY-WEIGHTED DENSITY MATRIX EMBEDDING

The CASSCF method exploits the locality of correlation in the energy domain, choosing and optimizing a low-energy subspace for the correlated treatment. In this section, we demonstrate the utility of a faithful QPU sampling of RDMs in order to correlate and optimize a different subspace, which instead relies on spatial locality. This perspective is often more useful for strongly-correlated extended systems, where the atomic-level correlated degrees of freedom can be isolated, and where widely used methods such as density functional theory fail to provide accurate results [70]. These approaches fall under the umbrella of quantum embedding or quantum cluster methods, and are amongst the most promising for QPU-enhanced materials modeling [21,71,72]. We investigate the recently-developed "energy-weighted density matrix embedding theory" (EwDMET) as a promising candidate in this direction [22–24].

The EwDMET method connects the density matrix embedding theory (DMET) and dynamical mean-field theory (DMFT), two established approaches in quantum embedding [73–75]. Both of these "parent" approaches have recently been adapted for use with a quantum hardware solver, as well as related embedding techniques [7–11,13,76–78]. However, the EwDMET avoids a number of difficulties. Similar to DMET, it avoids any necessity to compute the single-particle Green's function of the resulting quantum cluster problem on the QPU, which is challenging for quantum hardware, although important progress is being made [79]. Instead, the method requires a desired number of one-particle spectral moments from the subspace problem, which can be obtained directly from the reduced density matrices of the ground state. The number of self-consistent spectral moments can then be systematically enlarged, to approach the complete dynamical character of DMFT as an orthogonal polynomial expansion. The method also removes all explicit numerical fitting steps, and constructs a rigorous self-consistency on these spectral moments, systematically extending the DMET formulation and connecting it to its fully dynamical limit. This rigorous and algebraic self-consistency enables nontrivial results to be obtained even at the lowest truncation of the spectral moment expansion. This requirement of only computing ground-state RDMs, while at the same time benefiting from a rigorous and algebraic self-consistency for nontrivial emergent physics makes it an ideal candidate for combination with QPU-derived RDMs in the NISQ era. We briefly review the salient features of EwDMET for this QPU formulation, with more details in Ref. [24].

As with all quantum cluster approaches, the algorithm begins with the choice of a local correlated space. This could be the d shell of a transition metal atom, or a cluster of sites for a discrete lattice model. The EwDMET method then allows for an improvable and self-consistent description of the one-particle quantum fluctuations between this fragment and its environment [22–24]. This information is contained within the self-consistently optimized (hole and particle) spectral moments of the fragment, defined as

$$T_{h,\alpha\beta}^{(n)} = \langle \Psi | \, \hat{a}_\alpha^\dagger (\hat{H} - E_0)^n \hat{a}_\beta \, | \Psi \rangle , \qquad (14)$$

$$T_{p,\alpha\beta}^{(n)} = \langle \Psi | \, \hat{a}_\alpha (\hat{H} - E_0)^n \hat{a}_\beta^\dagger \, | \Psi \rangle , \qquad (15)$$

where α, β index the degrees of freedom of this local fragment, $n \geqslant 0$ denotes the order of these moments, optimized up to a maximum desired value n_{mom}, and $| \Psi \rangle$ is the ground state of the constructed correlated subspace. As $n_{\text{mom}} \to \infty$, the method exactly reproduces the effective dynamics of DMFT, recast as a ground-state wave function theory, while systematic truncation to low n_{mom} will still faithfully describe the dominant low-energy fluctuations from the fragment into its environment. The EwDMET method rigorously maps the full system to a subspace consisting of the chosen fragment coupled to a "bath" space. The size of the bath is determined solely by the size of the fragment and the desired number of spectral moments to capture (and the correlated subspace is hence independent of the size of the full system). It is this correlated subspace problem which must then be solved on the QPU at each iteration, and the spectral moments of Eqs. (14) and (15) computed. With these computed spectral moments from the correlated fragment space, the one-particle description of the full system can be algebraically updated via the addition of noninteracting auxiliary states, to ensure that the fragment moments at the mean-field level over the full system exactly reproduce the correlated subspace ones. The procedure is iterated, updating the auxiliary space and bath space of the quantum cluster problem, until convergence.

A. Infinitely coordinated Bethe-Hubbard lattice

We apply this method to the paradigmatic Hubbard model of condensed matter physics, which describes a range of quantum phases and correlation-driven transitions. Specifically, the limit of an infinitely-coordinated extended Bethe-Hubbard lattice with local interactions defines our model of interest, which has the particular feature that correlation-driven changes to all one-particle properties are site local. This property was used to great effect to motivate the development of DMFT, by providing a nontrivial model for which it describes an exact limit [80,81]. The EwDMET has the same exact limit for this model as $n_{\text{mom}} \to \infty$.

The model can be equivalently defined in this infinite-dimensional limit via its metallic noninteracting density of states [82], which is defined to have the form

$$A(\omega) = \frac{1}{2\pi} \sqrt{4 - \omega^2}, \qquad (16)$$

for a bandwidth of $|\omega| < 2$. This noninteracting spectrum was fit to a single central fragment site with 200 additional degrees of freedom, to approximate this full spectrum to a high-energy resolution [83,84]. The interacting Hamiltonian is then defined as resulting from the additional on-site Hubbard interaction term, $U \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}$, which is included on the fragment in the correlated subspace Hamiltonian at each iteration. The spectral moments of this central fragment site are then self consistently optimized, where we define the projection of the

noninteracting system Hamiltonian into this cluster subspace as h^{clust}.

In this paper, we truncate the spectral moment expansion at order $n_{\text{mom}} = 1$, defining the set of self-consistent fragment quantities. This simplifies their computation from the VQE solution for the ground state of the cluster Hamiltonian at each iteration, since these $n = 0$ and $n = 1$ moments over the fragment can be constructed from (parts of) the one- and two-body RDMs, for which we have efficient sampling as detailed previously. For instance, the $n = 1$ hole moment reduces to

$$T_{h,00}^{(1)} = \sum_{j \in \text{clust}} h_{0j}^{\text{clust}} \gamma_{0j} + U \Gamma_{0000}, \tag{17}$$

where 0 denotes the fragment site index. Physically, the restriction of $n_{\text{mom}} = 1$ means that the center of mass of the particle and hole spectral distributions can be self-consistently optimized on the lattice in the presence of the local correlation effects. This contrasts with DMET, where single site self-consistency cannot change the physics of the full system from the noninteracting picture in translationally symmetric systems, and so QPU emulations of this method with a single fragment site are restricted to single-shot computation without any self-consistency [9,10,78]. By using VQE as the solver on quantum hardware, we can identify the ground state of the cluster Hamiltonian, and subsequently sample the relevant RDMs to construct the required fragment spectral moments. We iterate this procedure until self-consistency, which we define to be when the sum of the squared update to the (four) parameters defining the auxiliary states varies by an energy of less than 10^{-4}. These self-consistent auxiliary states consist of individual poles in a self-energy, which modifies the spectral function of the system to match the correlated local moments from the VQE.

At this choice of spectral moment truncation, the cluster Hamiltonian consists of the single-fragment site and a single bath orbital, resulting in a four-qubit system to solve at each iteration of the EwDMET method. This cluster is solved via VQE on the QPU with the same three-layer HEA as applied in the CASSCF section, with a Jordan-Wigner mapping to the qubit representation. Additionally, the same error mitigation is used to control the noise inherent in the sampling of the RDMs required to construct the fragment spectral moments. Emulated QPU simulations without noise models were also performed for comparison to the QPU experiments of this algorithm.

B. Results

Figure 5 presents results for the single-particle spectrum for the model at self-consistency for VQE-EwDMET with $n_{\text{mom}} = 1$, at a strongly correlated limit with an on-site interaction of $U = 8$, which is twice the noninteracting bandwidth of the material. Self-consistency achieves the matching of the first two local spectral moments [Eqs. (14) and (15) for $n = 0$ and 1 over the fragment site] for both the mean-field state of the whole system, and the VQE results over the correlated cluster. At this point a spectrum can be obtained which is consistent through these local moments, via diagonalization of the resulting (dynamical) mean-field, with specifics found in Ref. [24]. The final converged spectra are presented for an
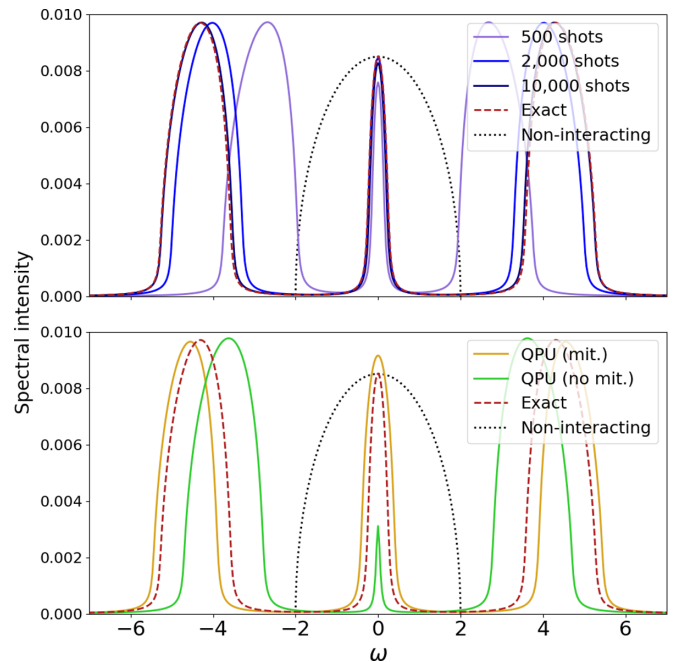

FIG. 5. The density of states for the Hubbard model on the Bethe lattice with infinite coordination for the EwDMET method with $n_{\text{mom}} = 1$. Upper results are performed on a classical QPU simulator with finite sampling noise and varying numbers of shots in the VQE solution to the cluster Hamiltonian at each iteration. Lower panel shows results from two QPU experiments, on IBMQ Santiago and IBMQ Bogota, with and without error mitigation, respectively, with 5000 shots for the sampling of the required RDMs. Grey dotted lines show the original noninteracting spectrum of the model, while the red dotted line shows the EwDMET($n_{\text{mom}} = 1$) results with an exact solution of the cluster Hamiltonian each iteration.

exact solver, an emulated quantum simulation without noise but with different numbers of shots for the sampling of the RDMs (500, 2000, and 10 000), and quantum hardware results on IBMQ machines with and without error mitigation, with 5000 shots for the sampling. All calculations converge within 8 iterations as the auxiliary and bath spaces adapt to the correlations described over the fragment site in each VQE cluster solution, demonstrating the robustness of the self-consistency in the presence of noise. At this level of theory, the spectrum shows significant qualitative changes from the noninteracting spectrum. Upper and lower Hubbard bands develop, splitting the original density of states, with a qualitatively correct charge gap between these bands shown. However, a small quasiparticle peak remains at the Fermi level, showing that the metallic character of the system is not entirely removed by the correlations, as is to be expected from numerically exact calculations on this system such as NRG+DMFT [85]. Consistency in higher orders of the spectral moments are required to get to a true Mott insulating state [24], which can be achieved at the expense of an increasing size of bath space and sampling higher spectral moments (which requires higher-body RDMs, described in Appendix B). Nevertheless, even at this low truncation, much of the true correlated spectral density is reproduced with significant physical correlation-driven redistribution of spectral weight observed.
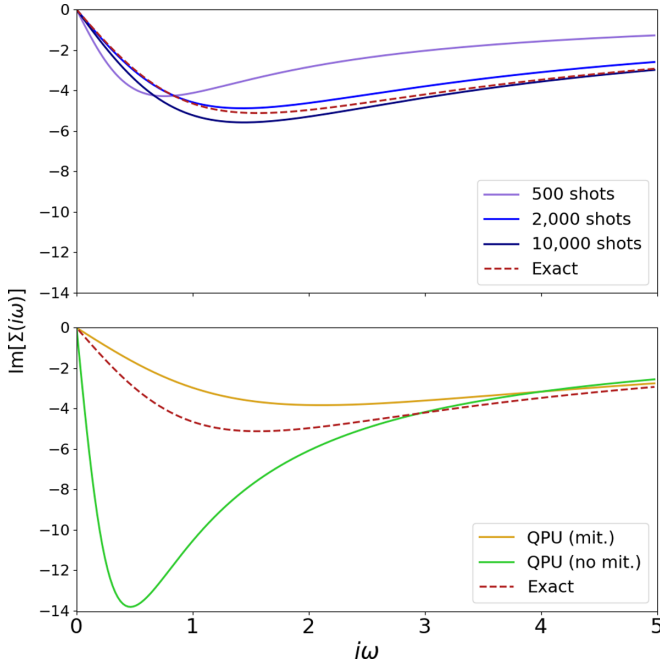
FIG. 6. Converged effective on-site Matsubara self-energy for the Bethe lattice Hubbard model with EwDMET at $n_{\mathrm{mom}} = 1$. Upper plot shows results from the quantum simulator with finite sampling of the RDM elements, compared to exact results (infinite sampling). Lower plot shows results from quantum hardware on IBMQ Santiago and IBMQ Bogota and 5000 shots in the RDM sampling, with and without error mitigation in the RDM sampling, respectively.

At 10 000 shots, the emulated results without further simulated quantum noises are are almost indistinguishable from the exact benchmark at all energies. At lower numbers of shots, the gap between the Hubbard bands is too small, as the variance in the sampled RDMs increases. As with the CASSCF method, the updated auxiliary space at each iteration is a nonlinear transformation of the spectral moments [which are themselves linear functionals of the RDM elements, as can be seen in Eqs. (17)]. The consequence of this is a systematic error in the resulting spectral functions at convergence due to the increasing RDM variance, rather than simply a manifestation in a noisy but unbiased spectrum. This behavior of an underestimated gap between the Hubbard bands is also present in the QPU results, where unmitigated results feature unrepresentative Hubbard bands. However, the performance is once again considerably improved with the error mitigation, with the Hubbard bands and low-energy peak resolved to higher accuracy, suggesting the method fits well with a QPU cluster solver, and removing the necessity for the full solution of the fragment Green's function at each point within a DMFT framework.

Properties of the system can also be observed from the effective self-energy of the system, which is obtained directly from the self-consistent auxiliary states, and allows access to quantities such as Fermi liquid parameters [24]. For the same system, Fig. 6 shows the imaginary part of the self energy on the Matsubara frequency axis. The finite sampling results are seen to approach the exact results with increasing shots, with 10,000 shots reaching a comparable performance

to the exact results, with the discrepancy far more visible in the self-energy than the original single-particle spectrum of Fig. 5. The quantum hardware unmitigated results correspondingly demonstrate significant overestimation of the resulting self-energy. Despite the fact that the Hubbard bands are closer, these unmitigated results show a larger effective mass and quasi-particle renormalization at the Fermi surface from the self-energy (a larger derivative at $i\omega \to 0$), which manifests in the smaller peak in Fig. 5 at that point. Error mitigated QPU self-energies are however more in line with exact results for the method, albeit now slightly underestimated at low-frequencies. Further improvement of the results can be obtained by increasing the moment order (n_{mom}) to which the dynamical quantities are all resolved. Similar to the perturbative corrections to CASSCF however, these will require sampling of higher-body RDMs, and is an avenue of continuing research.

V. CONCLUSIONS

We have presented a unified approach to self-consistent coupling of quantum and classical computational resources in quantum chemistry and condensed matter electronic structure problems. This coupling relies on the faithful and efficient sampling of reduced density matrices on quantum resources, where these objects span the correlated physics of an iteratively optimized subspace of the full system. We consider the required fidelity and sampling quality of these density matrices for robust optimization on current generation quantum hardware, developing a simple but effective approach to mitigate for gate errors and allow this full convergence. As well as converged energetics, we also analyze the viability of the sampled density matrices for nonenergetic quantities, including the dipole moment for *ab initio* simulation of chemical systems, and the self-energy and mass renormalization of strongly correlated extended models.

Overall, the picture is encouraging, with the self-consistent optimization found to be particularly robust to the presence of sampling noise on current generation quantum hardware. This self-consistency is found to be more reliable than the uncertainties resulting from the state optimization and energy obtained from the VQE at any single iteration. This points to a significant transfer from continuing improvements in both hardware and quantum algorithms for state preparation on quantum devices for self-consistent approaches. These conclusions, however, are restricted to relatively small correlated subspaces, and further work is required to understand the generality of these conclusions as we access QPU with larger qubit capacity. Furthermore, quantitative rather than qualitative accuracy in these application areas will require an efficient and compact description of higher-body density matrices, which will be the focus of future directions.

ACKNOWLEDGMENTS

APPENDIX A: LIGHT TOUCH ERROR MITIGATION

In order to improve the results and reduce the effects of the quantum noises of the IBMQ devices used, we employ a simple, but costless, extrapolation approach based on the assumption that the impact of the circuit errors is evenly distributed on the output state (for a more rigorous approach to extrapolation, see Ref. [33]). The aim is to recover an approximation of the true expectation value of a quantum circuit with respect to an operator, assuming a specific bias on the output results. It is worth noting that this method primarily focus on correcting gate and readout errors, and as such does not attempt to maintain or restore the purity of the quantum state produced. It is not expected to be an effective approach (when used in isolation) for mitigating errors on deep quantum circuits where the primary source of noise is decoherence.

We define the true (desired) expectation value of the circuit with respect to an operator, $\hat{O}$, as $\langle \hat{O} \rangle$, while the measured expectation value is denoted as $\langle \hat{O}_{\text{mes}} \rangle$. The outcome of a circuit can be associated with measurement eigenvalues ± 1, with the probability of measuring 1 equal to $P_1 = \text{Pr}(O_m = 1)$, and with O_m referring to a single measurement of operator $\hat{O}$ at the end of the quantum circuit. Similarly, we have $P_1 = 1 - P_{-1} = \text{Pr}(O_m = -1)$. We can associate the true expectation value with

$$\langle \hat{O} \rangle = P_1 - P_{-1} = 2P_1 - 1. \qquad (A1)$$

We now assume that there is a certain probability P_{err} that at least one gate error occurs during propagation and measurement of the quantum circuit. Any gate error changes the balance of probabilities between $\text{Pr}(O_m = 1)$ and $\text{Pr}(O_m = -1)$. One can assume that there exists a value between ± 1, representing expectation value of $\hat{O}$ given the error rate, which we denote $\langle \hat{O}_{\text{err}} \rangle$. This value, as well as the probability

distribution of the operator measurements O_m are unknown and cannot be recovered easily.

We can approximate the expectation value of the measured operator as $\langle \hat{O}_{\text{mes}} \rangle = (1 - P_{\text{err}})\langle \hat{O} \rangle + P_{\text{err}}\langle \hat{O}_{\text{err}} \rangle$, assuming a linear relationship between the gate errors and effect on the expectation value. We can rewrite this as

$$\langle \hat{O} \rangle = \frac{\langle \hat{O}_{\text{meas}} \rangle - P_{\text{err}}\langle \hat{O}_{\text{err}} \rangle}{(1 - P_{\text{err}})}. \qquad (A2)$$

We make the further assumption on the value of $\langle \hat{O}_{\text{err}} \rangle$ that, given a sufficiently large number of a random circuit errors, the probability of getting either eigenvalue when an error occurs is equal (or $\langle \hat{O}_{\text{err}} \rangle = 0$). This assumption is based on two observations: (1) there is no way to tell exactly what the impact of a gate error will be on measurement probability except that it will bias measurement averages towards 0 (since the dominant probability will on average be affected by more errors) and (2) given that the gate errors are random, these will not result in always measuring 1 or -1, ensuring that these are the least likely values for $\langle \hat{O}_{\text{err}} \rangle$.

This method is largely sufficient for the purpose of our experiments, and can act as a lower bound for the benefits error mitigation can achieve with no additional computing cost. With this approximation, and for eigenvalues ± 1, we can ignore this final term, and the expression simplifies to

$$\langle \hat{O} \rangle = \frac{\langle \hat{O}_{\text{meas}} \rangle}{(1 - P_{\text{err}})}, \qquad (A3)$$

while for binary eigenvalues of 1 and 0, we get

$$\langle \hat{O} \rangle = \frac{\langle \hat{O}_{\text{meas}} \rangle - 0.5 P_{\text{err}}}{(1 - P_{\text{err}})}. \qquad (A4)$$

In order to estimate the final bias on the true value of the operator, we have considered the electron number operator (trace of the one-body RDM). The deviation from the set number of electrons in the system gives us an estimate for the bias induced by quantum noise P_{err}. Given our assumptions, the bias factor can be recovered as follows:

$$\frac{1}{1 - P_{\text{err}}} = \frac{N_{\text{elec}}}{N_{\text{meas}}}, \qquad (A5)$$

with N_{elec} the target number of electrons, and N_{meas} the number of electron measured:

$$N_{\text{meas}} = \sum_i \gamma_{ii} = \sum_i \langle \psi | \hat{a}_i^\dagger \hat{a}_i | \psi \rangle. \qquad (A6)$$

This comes at no extra computational cost as the one-body RDM terms used are necessarily computed as part of the VQE process.

An alternative method would be to estimate P_{err} directly by computing it through the reported gate calibration data from the QPU provider (compounding the gate fidelities), but we found that in general this approach is less reliable. This is most likely due to the fact that using this latter method treats the bias resulting from circuit errors completely classically: It ignores any part of the bias that could be due to the reduced purity of the quantum state produced, which can otherwise be captured by the former method.
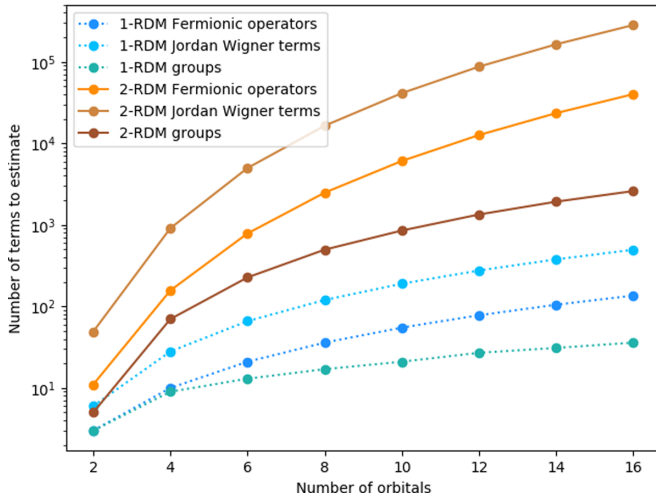
FIG. 7. Number of unique fermionic operators, corresponding set of unique Pauli strings (under Jordan-Wigner mapping), and commutative groups to be measured in order to compute all elements of the one-body (dotted line) and two-body RDM (solid line), for up to 16 orbitals (32 qubits). The groups were found via the largest-degree first coloring (LDFC) algorithm.

APPENDIX B: OPERATOR GROUPINGS FOR RDM ELEMENT SAMPLING

At most, sixteen Pauli strings result from each two-body fermionic operator. While this implies polynomial scaling ($\mathcal{O}(n^4)$) in the number of terms that need to be measured on the quantum computer, the total number of measurements can be further improved upon by using commutative features of the Pauli strings. In particular, Pauli strings that commute can be measured jointly. There are two main ways to approach commutativity: qubit-wise commutation (QWC) and general commutation (GC) of the operators. For QWC, we group two Pauli strings together if each operator in the first Pauli string commutes with the operator of corresponding index in the second Pauli string. GC is more general, and allows grouping of Pauli strings as long as they generally commute (for a review of Pauli strings commutation rules, we recommend [86–88]). In this Appendix, we study the scaling of measuring RDM using GC grouping for use in future research. These grouping strategies tend to be computationally expensive. However the final set of noncommuting Pauli strings required to measure a given rank of RDM will be identical and agnostic to the details of the Hamiltonian for a given number of orbitals.

TABLE II. Number of unique Pauli strings to be measured once grouped in order to sample all elements of the three- and four-body RDMs, as the number of (spatial) orbitals in the active space is enlarged. These numbers represent the terms in a direct Jordan-Wigner mapping (JW) and grouping of commuting terms from a largest-degree first coloring algorithm (JW-Groups).

	Active orbitals	4	6
	Fermionic operators	610	8400
Three-body RDM	JW	4928	71 742
	JW-Groups	189	2049
	Fermionic operators	939	40 065
Four-body RDM	JW	11 425	440 154
	JW-Groups	163	3182

As a result, once an optimal set of terms is established, it can be used for all systems, in a similar fashion to the energy measurement for a VQE problem of a given size (which is equivalent to the two-body RDM).

We present in Fig. 7 the number of commutative groups constructed for the one- and two-body RDMs as a function of the number of molecular orbitals in the active space (where each molecular orbital is mapped to two qubits). Our results are similar to those found previously in the literature (see, for instance Ref. [69]), showing a significant reduction in dimensionality and scaling of independant observable measurements as the active space increases in size. This results in a reduction of over two orders of magnitude for the number of terms to be sampled in a 16 orbital active space, with this factor increasing for larger active spaces. We present these numbers in Table I. In addition, we also present the groups for three- and four-body RDMs, up to 6 orbitals (Table II), as an investigation into the future feasibility of extended coupling schemes between the quantum region and environment such as multireference perturbative or subspace expansion approaches [12,55,68]. Finding these groups for the higher-body RDMs for larger numbers of orbitals became too computationally demanding for the current algorithm given our resources at time of writing. One point of note is that the symmetries used in Eqs. (5)–(7), combined with the Jordan-Wigner mapping, ensure that the one-body RDM only relies on half of the wave function (the same half of all Pauli strings required are identities, rendering half of the qubits used obsolete in the sampling). This feature can be used to easily

TABLE I. Number of unique Pauli strings to be measured once grouped in order to sample all elements of the one- and two-body RDMs, as the number of (spatial) orbitals in the active space is enlarged. These numbers represent the terms in a direct Jordan-Wigner mapping (JW) and grouping of commuting terms from a largest-degree first coloring algorithm (JW-Groups).

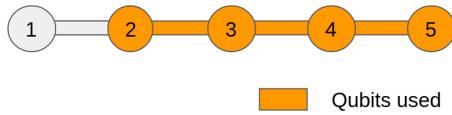|	Active orbitals	2	4	6	8	10	12	14	16
	Fermionic operators	3	10	21	36	55	78	105	136
One-body RDM	JW	6	28	66	120	190	276	378	496
	JW-Groups	3	9	13	17	21	27	31	36
	Fermionic operators	11	157	786	2 486	6 085	12 651	23 492	40 156
Two-body RDM	JW	49	910	4983	16 460	41 325	87 354	164 115	282 968
	JW-Groups	5	70	227	497	853	1342	1928	2601

FIG. 8. IBMQ Bogota lattice structure and qubits used for experiments. This QPU was used to compute CASSCF and EwDMET with error mitigation.

sample the energy of any tensor product state or one-body RDM functional [89].

It is worth noting that grouping of terms may entail additional costs. Firstly, the joint measurements of Pauli operators results in a covariance between terms, potentially increasing the overall variance of the observable expectation values. In exceptional cases, this can even increase the total number of samples required for a given fidelity [86]. In general, however, we should expect a reduction in number of measurements necessary for a given precision [87]. Secondly, joint measurements of Pauli strings groups require additional circuit depth to rotate the measurement basis appropriately. This additional circuit scales $O(N^2)$, with N the number of qubits [87], and therefore should be considered small on larger gate depth circuits (by comparison, the "generalized unitary coupled cluster ansatz" scales $O(N^3)$ in depth; see for instance Ref. [32]). In the case of current generation QPUs however, this additional circuit length (largely composed of entangling gates) results in quantum noise that would arguably out-weigh the benefits obtained from reduction of finite sampling noise from operator joint measurements. For this reason, while this reduction in terms to sample is a promising feature for longer-term viability of density matrix sampling, we are unlikely to benefit from this for small qubit arrays on current generation QPUs, and therefore leave the use of operator grouping in actual experiments for future work.

We present below further details on the approach we have used to group operators, as presented in Fig. 7. In order to find these groupings, we require a graph of commutative relationships between all the Pauli strings required to measure the elements of the RDMs. To find a low number of groups of fully connected sub-graphs, we employed the "largest-degree first coloring" (LDFC) algorithm (similar to what is proposed in Ref. [88]), a graph coloring heuristic. As an example for an alternative to the LDFC algorithm, one can start by grouping Pauli strings according to the frequency of identity operators in the string (as done, for instance in Ref. [69]).

The steps required to complete grouping of Pauli terms (using LDFC), joint measurements and measurement results aggregation are outlined at a high level below. For a more detailed description, we recommend Ref. [87].
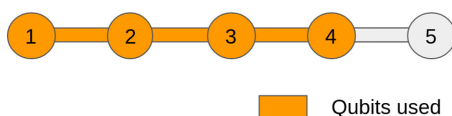

FIG. 9. IBMQ Santiago lattice structure and qubits used for experiments. This QPU was used to compute CASSCF and EwDMET without error mitigation.
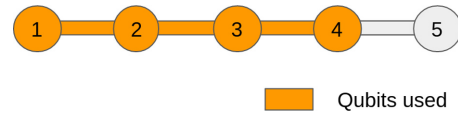

FIG. 10. IBMQ Athens lattice structure and qubits used for experiments. This QPU was used to compute the RDM sampling studies presented in Fig. 1, with and without error mitigation

(i) Initialization: From the list of Pauli terms that require grouping, define a graph $G(V, e)$, with V the vertices, corresponding to each Pauli operator, and e the edges representing anti-commuting relationship between Pauli operators.

(ii) LDFC step 1: Rank the elements of V according to their degree, i.e. number of edges they are connected to. Colors are represented by integers, the color of each vertex is initialized to 0 (unallocated).

(iii) LDFC step 2: First allocate the color 1 to the element of V with the highest degree. Continue by allocating to the next element of V the lowest color that is not already attributed to one of its neighbours. Iterate likewise until all vertices have a color.

(iv) Joint-measurement basis identification: The groups have now been defined. From each group, identify a basis (multiplicative), from which all the other elements of the groups can be computed. The basis size should be N, with N the number of qubits.

(v) Joint-measurement circuit construction: Once a basis is identified, construct the circuit required for joint measurements of the operators (following for instance the instructions set in Ref. [87], aiming to map each of the operators in the basis to a single qubit Z-operator measurement.

(vi) Reconciliation: From the results of the measurement, reconstruct the expectation value of each element in each group that can then be used to compute the one- and two-body RDMs.

APPENDIX C: IBM QPU LATTICE STRUCTURES AND ADDITIONAL INFORMATION

In this Appendix, we present additional information about the IBM QPU used during the experiment. The information provided below is sourced from IBMQ Experience reported calibration of the machines at time of running the experiment and may change slightly over time.

TABLE III. Selected calibration metrics from IBMQ experience. These values are averaged for all qubits/connections and taken at a point in time near the experiment was run. They may change over time.

QPU	Bogota	Santiago	Athens
Single-qubit Pauli-X error	2.95e–4	4.55e–4	4.16e–4
Qubit frequency (GHz)	4.89	4.78	5.094
Two-qubit gate error	1.40e–2	1.22e–2	1.043e–2
Two-qubit Gate time (ns)	536.89	408.89	346.67
Read-out length (ns)	5048.89	4017.78	3022.22
Read-out error	3.77e–2	1.82e–2	1.82e–2

1. Lattice structures

IBMQ Bogota, Santiago, and Athens are all 5-qubit QPUs, following IBM's *Canary* r3 processor type, with reported quantum volume of 32 [90]. The lattice structure, as well as the qubits used are presented in Figs. 8–10, respectively.

2. Calibration information

The information presented in Table III summarizes the calibration data of the QPU used. It is directly taken from the IBMQ Experience Portal and may change over time as IBM recalibrates the processors.

[1] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, Proc. Natl. Acad. Sci. USA **114**, 7555 (2017).

[2] M. Kühn, S. Zanker, P. Deglmann, M. Marthaler, and H. Weiss, Accuracy and resource estimations for quantum chemistry on a near-term quantum computer, J. Chem. Theory Comput. **15**, 4764 (2019).

[3] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, and R. C. Pooser, Quantum chemistry as a benchmark for near-term quantum computers, npj Quantum Inf. **5**, 99 (2019).

[4] B. Bauer, S. Bravyi, M. Motta, and G. Kin-Lic Chan, Quantum algorithms for quantum chemistry and quantum materials science, Chem. Rev. **120**, 12685 (2020).

[5] J. Preskill, Quantum computing in the nisq era and beyond, Quantum **2**, 79 (2018).

[6] R. Xia and S. Kais, Quantum machine learning for electronic structure calculations, Nat. Commun. **9**, 4195 (2018).

[7] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, Hybrid Quantum-Classical Approach to Correlated Materials, Phys. Rev. X **6**, 031045 (2016).

[8] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, npj Comput. Mater. **6**, 85 (2020).

[9] T. Yamazaki, S. Matsuura, A. Narimani, A. Saidmuradov, and A. Zaribafiyan, Towards the practical application of near-term quantum computers in quantum chemistry simulations: A problem decomposition approach, arXiv:1806.01305.

[10] N. C. Rubin, A hybrid classical/quantum approach for large-scale studies of quantum systems with density matrix embedding theory, arXiv:1610.06910.

[11] D. Dhawan, M. Metcalf, and D. Zgid, Dynamical Self-energy Mapping (DSEM) for quantum computing, arXiv:2010.05441v2.

[12] T. Takeshita, N. C. Rubin, Z. Jiang, E. Lee, R. Babbush, and J. R. McClean, Increasing the Representation Accuracy of Quantum Simulations of Chemistry without Extra Quantum Resources, Phys. Rev. X **10**, 011004 (2020).

[13] Y. Yao, F. Zhang, C.-Z. Wang, K.-M. Ho, and P. P. Orth, Gutzwiller hybrid quantum-classical computing approach for correlated materials, Phys. Rev. Res. **3**, 013184 (2021).

[14] M. Rossmannek, P. K. Barkoutsos, P. J. Ollitrault, and I. Tavernelli, Quantum HF/DFT-embedding algorithms for electronic structure calculations: Scaling up to complex molecular systems, J. Chem. Phys. **154**, 114105 (2020).

[15] S. Yalouz, B. Senjean, J. Günther, F. Buda, T. E. O'Brien, and L. Visscher, A state-averaged orbital-optimized hybrid quantum–classical algorithm for a democratic description of ground and excited states, Quantum Sci. Technol. **6**, 024004 (2021).

[16] B. O. Roos, P. R. Taylor, and P. E. M. Sigbahn, A complete active space scf method (CASSCF) using a density matrix formulated super-ci approach, Chem. Phys. **48**, 157 (1980).

[17] Q. Sun, J. Yang, and G. Kin-Lic Chan, A general second order complete active space self-consistent-field solver for large-scale systems, Chem. Phys. Lett. **683**, 291 (2017).

[18] D. Reta Mañeru, A. K. Pal, I. de P. R. Moreira, S. N. Datta, and F. Illas, The triplet-singlet gap in the *m*-Xylylene radical: A not so simple one, J. Chem. Theory Comput. **10**, 335 (2014).

[19] Y. Li, L. Wang, X. Guo, and J. Zhang, A CASSCF/CASPT2 insight into excited-state intramolecular proton transfer of four imidazole derivatives, J. Comput. Chem. **36**, 2374 (2015).

[20] J. Olsen, The CASSCF method: A perspective and commentary, Int. J. Quantum Chem. **111**, 3267 (2011).

[21] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, Quantum cluster theories, Rev. Mod. Phys. **77**, 1027 (2005).

[22] E. Fertitta and G. H. Booth, Rigorous wave function embedding with dynamical fluctuations, Phys. Rev. B **98**, 235132 (2018).

[23] E. Fertitta and G. H. Booth, Energy-weighted density matrix embedding of open correlated chemical fragments, J. Chem. Phys. **151**, 014115 (2019).

[24] P. V. Sriluckshmy, M. Nusspickel, E. Fertitta, and G. H. Booth, Fully algebraic and self-consistent effective dynamics in a static quantum embedding, Phys. Rev. B **103**, 085131 (2021).

[25] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2013).

[26] K. Boguslawski, P. Tecmer, G. Barcza, Ö. Legeza, and M. Reiher, Orbital entanglement in bond-formation processes, J. Chem. Theory Comput. **9**, 2959 (2013).

[27] P. Jordan and E. Wigner, Uber das paulische aquivalenzverbot, Z. Phys. **47**, 631 (1928).

[28] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, Ann. Phys. (NY) **298**, 210 (2002).

[29] P. J. Love J. T. Seeley, and M. J. Martin, The Bravyi-Kitaev transformation for quantum computation of electronic structure, J. Chem. Phys. **137**, 224109 (2012).

[30] M. Ostaszewski, E. Grant, and M. Benedetti, Structure optimization for parameterized quantum circuits, Quantum **5**, 391 (2021).

[31] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).

[32] J. Lee, W. J. Huggins, M. Head-Gordon, and K. B. Whaley, Generalized unitary coupled cluster wave functions for quantum computation, J. Chem. Theory Comput. **15**, 311 (2019).

[33] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, Phys. Rev. X **8**, 031027 (2018).

[34] B. H. Lengsfield, General second order MCSCF theory: A density matrix directed algorithm, J. Chem. Phys. **73**, 382 (1980).

[35] H.-J. Werner and W. Meyer, A quadratically convergent multiconfiguration-self-consistent field method with simultaneous optimization of orbitals and ci coefficients, J. Chem. Phys. **73**, 2342 (1980).

[36] H.-J. Werner and P. J. Knowles, A second order multiconfiguration scf procedure with optimum convergence, J. Chem. Phys. **82**, 5053 (1985).

[37] H.-J. Werner, Matrix-formulatoin direct multiconfiguration self-consistent field and multiconfiguration reference configuration-interaction methods, Adv. Chem. Phys. **69**, 1 (1987).

[38] J. Olsen, D. L. Yeager, and P. Jørgensen, Optimization and characterization of a multiconfigurational self-consistent field (MCSCF) state, Adv. Chem. Phys. **54**, 1 (1983).

[39] H.-J. A. Jensen and P. Jørgensen, A direct approach to second-order MCSCF calculations using a norm extended optimization scheme, J. Chem. Phys. **80**, 1204 (1984).

[40] R. Shepard, The multiconfiguration self-consistent field method, Adv. Chem. Phys. **69**, 63 (1987).

[41] S. Ten-no and S. Iwata, Multiconfiguration self-consistent field procedure employing linear combination of atomic-electron distributions, J. Chem. Phys. **105**, 3604 (1996).

[42] M. W. Schmidt and M. S. Gordon, The construction and interpretation of MCSCF wavefunctions, Annu. Rev. Phys. Chem., **49**, 233 (1998).

[43] H. Nakano and K. Hirao, A quasi-complete active space self-consistent field method, Chem. Phys. Lett. **317**, 90 (2000).

[44] F. Aquilante, T. B. Pedersen, R. Lindh, B. O. Roos, A. Sanchez de Meras, and H. Koch, Accurate ab initio density fitting for multiconfigurational self-consistent field methods, J. Chem. Phys. **129**, 024113 (2008).

[45] P. G. Szalay, T. Müller, G. Gidofalvi, H. Lischka, and R. Shepard, Multiconfiguration self-consistent field and multireference configuration interaction methods and applications, Chem. Rev. **112**, 108 (2012).

[46] W. Győrffy, T. Shiozaki, G. Knizia, and H.-J. Werner, Analytical energy gradients for second-order multireference perturbation theory using density fitting, J. Chem. Phys. **138**, 104104 (2013).

[47] I. Kim, S. M. Parker, and T. Shiozaki, Orbital optimization in the active space decomposition model, J. Chem. Theory Comput. **11**, 3636 (2015).

[48] E. G. Hohenstein, N. Luehr, I. S. Ufimtsev, and T. J. Martinez, An atomic orbital-based formulation of the complete active space self-consistent field method on graphical processing units, J. Chem. Phys. **142**, 224103 (2015).

[49] WS. Jeong, S. J. Stoneburner, D. King, R. Li, A. Walker, R. Lindh, and L. Gagliardi, Automation of active space selection for multireference methods via machine learning on chemical bond dissociation, J. Chem. Theory Comput. **16**, 2389 (2020).

[50] E. R. Sayfutyarova, Q. Sun, G. K.-L. Chan, and G. Knizia, Automated construction of molecular active spaces from atomic valence orbitals, J. Chem. Theory Comput. **13**, 4063 (2017).

[51] K. D. Vogiatzis, D. Ma, J. Olsen, L. Gagliardi, and W. A. de Jong, Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations, J Chem Phys. **147**, 184111 (2017).

[52] D. S. Levine, D. Hait, N. M. Tubman, S. Lehtola, K. B. Whaley, and M. Head-Gordon, CASSCF with extremely large active spaces using the adaptive sampling configuration interaction method, J. Chem. Theory Comput. **16**, 2340 (2020).

[53] J. E. T. Smith, B. Mussard, A. A. Holmes, and S. Sharma, Cheap and near exact casscf with large active spaces, J. Chem. Theory Comput. **13**, 5468 (2017).

[54] R. E. Thomas, Q. Sun, A. Alavi, and G. H. Booth, Stochastic multiconfigurational self-consistent field theory, J. Chem. Theory Comput. **11**, 5316 (2015).

[55] R. J. Anderson, T. Shiozaki, and G. H. Booth, Efficient and stochastic multireference perturbation theory for large active spaces within a full configuration interaction quantum Monte Carlo framework, J. Chem. Phys. **152**, 054101 (2020).

[56] D. Ghosh, J. Hachmann, T. Yanai, and G. K.-L. Chan, Orbital optimization in the density matrix renormalization group, with applications to polyenes and β-carotene, J. Chem. Phys. **128**, 144117 (2008).

[57] D. Zgid and M. Nooijen, The density matrix renormalization group self-consistent field method: Orbital optimization with the density matrix renormalization group method in the active space, J. Chem. Phys. **128**, 144116 (2008).

[58] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. D. Bochevarov, How will quantum computers provide an industrially relevant computational advantage in quantum chemistry? arXiv:2009.12472v1.

[59] D. Wang, O. Higgott, and S. Brierley, Accelerated Variational Quantum Eigensolver, Phys. Rev. Lett. **122**, 140504 (2019).

[60] K. Yeter-Aydeniz, B. T. Gard, J. Jakowski, S. Majumder, G. S. Barron, G. Siopsis, T. Humble, and R. C. Pooser, Benchmarking quantum chemistry computations with variational, imaginary time evolution, and Krylov space solver algorithms, arXiv:2102.05511.

[61] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan, PySCF: the Python-based simulations of chemistry framework, WIREs Comput. Mol. Sci. **8**, 1340 (2018).

[62] Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui *et al.*, Recent developments in the pySCF program package, J. Chem. Phys. **153**, 024109 (2020).

[63] G. Li Manni, D. Smart, S, and A. Alavi, Combining the complete active space self-consistent field method and the full configuration interaction quantum Monte Carlo within a super-CI framework, with application to challenging metal-porphyrins, J. Chem. Theory Comput. **12**, 1245 (2016).

[64] J. J. Halson, R. J. Anderson, and G. H. Booth, Improved stochastic multireference perturbation theory for correlated systems with large active spaces, Mol. Phys. **118**, e1802072 (2020).

[65] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, Nature (London) **549**, 242 (2017).

[66] K. Mitarai, T. Yan, and K. Fujii, Generalization of the Output of a Variational Quantum Eigensolver by Parameter Interpolation with a Low-Depth Ansatz, Phys. Rev. Appl. **11**, 044087 (2019).

[67] G. E. Scuseria, M. D. Miller, F. Jensen, and J. Geertsen, The dipole moment of carbon monoxide, J. Chem. Phys. **94**, 6660 (1991).

[68] C. Angeli, M. Pastore, and R. Cimiraglia, New perspectives in multireference perturbation theory: the n-electron valence state approach, Theor. Chem. Acc. **117**, 743 (2007).

[69] H. J. Vallury, M. A. Jones, C. D. Hill, and L. C. L. Hollenberg, Quantum computed moments correction to variational estimates, Quantum **4**, 373 (2020).

[70] J. Cohen, P. Mori-Sánchez, and W. Yang, Challenges for density functional theory, Chem. Rev. **112**, 289 (2012).

[71] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Electronic structure calculations with dynamical mean-field theory, Rev. Mod. Phys. **78**, 865 (2006).

[72] Q. Sun and G. K.-L. Chan, Quantum embedding theories, Acc. Chem. Res. **49**, 2705 (2016).

[73] G. Knizia and G. K.-L. Chan, Density Matrix Embedding: A Simple Alternative to Dynamical Mean-Field Theory, Phys. Rev. Lett. **109**, 186404 (2012).

[74] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, Rev. Mod. Phys. **68**, 13 (1996).

[75] B.-X. Zheng, J. S. Kretchmer, H. Shi, S. Zhang, and G. K.-L. Chan, Cluster size convergence of the density matrix embedding theory and its dynamical cluster formulation: A study with an auxiliary-field quantum monte carlo solver, Phys. Rev. B **95**, 045103 (2017).

[76] T. Keen, T. Maier, S. Johnston, and P. Lougovski, Quantum-classical simulation of two-site dynamical mean-field theory on noisy quantum hardware, Quantum Sci. Technol. **5**, 035001 (2019).

[77] I. Rungger, N. Fitzpatrick, H. Chen, C. H. Alderete, H. Apel, A. Cowtan, A. Patterson, D. Munoz Ramo, Y. Zhu, N. H. Nguyen *et al.*, Dynamical mean field theory algorithm and experiment on quantum computers, arXiv:1910.04735.

[78] Y. Kawashima, M. P. Coons, Y. Nam, E. Lloyd, S. Matsuura, A. J. Garza, S. Johri, L. Huntington, V. Senicourt, A. O. Maksymov *et al.*, Efficient and accurate electronic structure simulation demonstrated on a trapped-ion quantum computer, arXiv:2102.07045.

[79] S.-N. Sun, M. Motta, R. N. Tazhigulov, A. T. K. Tan, G. K.-L. Chan, and A. J. Minnich, Quantum computation of Finite-Temperature Static and Dynamical Properties of Spin Systems Using Quantum Imaginary Time Evolution, PRX Quantum **2**, 010317 (2021).

[80] W. Metzner and D. Vollhardt, Correlated Lattice Fermions in $d = \infty$ Dimensions, Phys. Rev. Lett. **62**, 324 (1989).

[81] P. D. Gujrati, Bethe or Bethe-Like Lattice Calculations Are More Reliable Than Conventional Mean-Field Calculations, Phys. Rev. Lett. **74**, 809 (1995).

[82] M. Eckstein, M. Kollar, K. Byczuk, and D. Vollhardt, Hopping on the Bethe lattice: Exact results for densities of states and dynamical mean-field theory, Phys. Rev. B **71**, 235119 (2005).

[83] M. Nusspickel and G. H. Booth, Frequency-dependent and algebraic bath states for a dynamical mean-field theory with compact support, Phys. Rev. B **101**, 045126 (2020).

[84] M. Kollar, Construction of a dispersion relation from an arbitrary density of states, Int. J. Mod. Phys. B **16**, 3491 (2002).

[85] R. Bulla, Zero Temperature Metal-Insulator Transition in the Infinite-Dimensional Hubbard Model, Phys. Rev. Lett. **83**, 136 (1999).

[86] J. R McClean, J. Romero, R. Babbush, and Alán Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, New J. Phys. **18**, 023023 (2016).

[87] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families, arXiv:1907.13623.

[88] I. Hamamura and T. Imamichi, Efficient evaluation of quantum observables using entangled measurements, npj Quantum Inf. **6**, 56 (2020).

[89] Google AI Quantum and Collaborators, Hartree-Fock on a superconducting qubit quantum computer, Science **369**, 1084 (2020).

[90] J. Gambetta and D. McClure, Hitting a quantum volume chord: IBM quantum adds six new systems with quantum volume 32, (2020), IBM Research Blog.