

Probing ground-state properties of the kagome antiferromagnetic Heisenberg model using the variational quantum eigensolver

Jan Lukas Bosse^{1,2,*} and Ashley Montanaro^{2,1,†}

¹*School of Mathematics, University of Bristol, Bristol, BS8 1QU, United Kingdom*

²*Phasecraft Ltd, Bristol, BS1 5DD, United Kingdom*



(Received 8 October 2021; accepted 5 January 2022; published 7 March 2022)

Finding and probing the ground states of spin lattices, such as the antiferromagnetic Heisenberg model on the kagome lattice (KAFH), is a very challenging problem on classical computers and only possible for relatively small systems. We propose using the variational quantum eigensolver (VQE) to find the ground state of the KAFH on a quantum computer. We find efficient ansatz circuits and show how physically interesting observables can be measured efficiently. To investigate the expressiveness and scaling of our ansatz circuits we used classical, exact simulations of VQE for the KAFH for different lattices ranging from 8 to 24 qubits. We find that the fidelity with the ground state approaches one exponentially in the circuit depth for all lattices considered, except for a 24-qubit lattice with an almost degenerate ground state. We conclude that VQE circuits that are able to represent the ground state of the KAFH on lattices inaccessible to exact diagonalization techniques may be achievable on near-term quantum hardware. However, for large systems, circuits with many variational parameters are needed to achieve high fidelity with the ground state.

DOI: [10.1103/PhysRevB.105.094409](https://doi.org/10.1103/PhysRevB.105.094409)

Simulating and investigating quantum mechanical systems will be one of the first applications of the noisy intermediate-scale quantum (NISQ) computing hardware available in the near term [1–3]. In particular, problems in quantum many-body physics and quantum chemistry are hard to intractable for today’s best supercomputers and of high scientific interest.

One algorithm that was devised with the low-circuit depths and relatively high noise rates of NISQ devices in mind is the variational quantum eigensolver (VQE) [4]. It is a hybrid quantum-classical algorithm to produce a ground state of a quantum Hamiltonian H . A classical optimizer is used to minimize the expectation value $\langle \psi(\theta) | H | \psi(\theta) \rangle$ over a family of states $|\psi(\theta)\rangle$. If this family of states includes a ground state of H , VQE aims to return this ground state.

In this work we focus on a quantum many-body system that is hard to study using classical supercomputers and that we hope is particularly suitable to be addressed using VQE on NISQ quantum computers: the antiferromagnetic Heisenberg model on the kagome lattice (KAFH). The phase of its ground state is still not entirely settled. Some analytical results indicate a valence bond crystal ground state [5,6] while other results indicate an algebraic spin-liquid ground state [7,8]. More recent numerical results predict the ground state to be a gapped [9–11] or gapless [12–14] spin-liquid state. Experimental results with herbertsmithite $\text{ZnCu}_3(\text{OH})_6\text{Cl}_2$ also indicate a gapless spin-liquid state [15,16]. The classical, numerical methods are limited in their ability to describe large systems [exact diagonalization (ED)] or to describe true two-dimensional (2D) systems instead of mimicking them using long strips or cylinders [density matrix renormalization

group (DMRG)]. In contrast, the KAFH maps well to NISQ quantum hardware because the qubits can directly represent the spin- $\frac{1}{2}$ degrees of freedom and many NISQ devices have a 2D connectivity close to that of the kagome lattice.

Here we develop efficient quantum circuits for VQE for KAFH targeted at nearest-neighbor quantum computing architectures, based on the popular and physically motivated Hamiltonian variational (HV) ansatz [17]. We then carry out exact, classical simulations of VQE for the KAFH with the HV ansatz on patches of the kagome lattice from 8 to 24 sites using variants of the HV ansatz with different levels of parametrization. These extensive experiments allow us to infer the possible scaling of VQE for larger system sizes beyond the capability of classical exact diagonalization. We also show that some physically interesting observables in the states produced by VQE match their values in the exact ground state.

We find that, similarly to previous work on VQE for the Hubbard model [18] and interacting, spinless fermions on a chain [19], the error as measured by infidelity decreases exponentially with the VQE circuit depth. Using one ansatz parametrization, our results are largely consistent with ground-state fidelity 0.999 being achievable using roughly half the number of qubits, $0.5n_{\text{qb}}$, VQE layers for strip-shaped lattices. Although we have insufficient data for lattices with equal length in both dimensions to determine scaling, consideration of entanglement spreading out linearly in the circuit depth suggests the scaling could even be as low as $\sqrt{n_{\text{qb}}}$. We show that each VQE layer can be implemented using a quantum circuit with two-qubit gate depth 7 on a quantum computer with square-lattice connectivity, or only 4 on an architecture with all-to-all connectivity.

However, this performance and the fastest exponential decay in infidelity is only achieved by a variant of the HV ansatz with one independent parameter per edge in the lattice and

*janlukas.bosse@bristol.ac.uk

†ashley@phasecraft.io

VQE layer ($\approx 2pn_{\text{qb}}$ parameters for a system of n_{qb} qubits and p VQE layers). The use of VQE with a large number of parameters may be expected to suffer from the “barren plateau” problem [20,21] of exponentially small gradients, and indeed we do find some evidence that these become small for large p . Also, in one case (a patch of the kagome lattice on 24 sites), VQE does not find the ground state in a reasonable number of layers. This is because the first excited state is very close in energy to the ground state; we find additionally that the overlap achieved with the subspace spanned by these two states is significantly better.

It is instructive to compare the complexity of VQE for KAFH with VQE for the Fermi-Hubbard model, which is another plausible early application for NISQ quantum computers, and was studied in detail in [18,22]. It was argued in [18] that a 5×5 Fermi-Hubbard instance (50 spins) might be solvable with high accuracy via a VQE circuit with overall two-qubit gate depth ≈ 550 in a square-lattice topology, which was substantially lower than other proposed near-term applications of quantum computers. If our numerical results for KAFH are representative of the performance of VQE for larger systems, a patch of the kagome lattice with 50 qubits could be solved with two-qubit gate depth only $\approx 0.5 \times 50 \times 7 = 175$ in the same topology, which is substantially lower. However, the variational ansatz used in [18] would use 125 parameters for the whole circuit, whereas the ansatz we use here would use 2500 parameters.

These points highlight the opportunities and challenges faced by VQE for condensed-matter systems. For small-scale versions of complex systems of significant physical interest, VQE ansätze can well represent the ground state and can be implemented efficiently. Yet, the VQE algorithm may struggle with finding the ground state where there is a small energy gap, and with handling optimization over a large number of variational parameters.

Recently, we became aware of closely related work done by Kattemölle and van Wezel [23]. They also study the performance of VQE with the HV ansatz on the KAFH, including numerical simulations on 20 qubits. However, they consider a different embedding of the kagome lattice into the square lattice that produces one round of the HV ansatz with one layer of two-qubit gates less than ours at the cost of one extra qubit per unit cell of the kagome lattice. Consequently, their ansatz circuits are different from ours and, unlike ours, fully respect the translational symmetry of the kagome lattice. They report slightly better scaling of the fidelity with circuit depth for a lattice on 20 qubits, but more detailed comparisons on more lattices are needed to draw any definitive conclusions regarding scaling. Other differences are the inclusion in this work of the study of gradient scaling, to understand whether barren plateaus are a problem in our ansatz circuits, and comparison of local observables for VQE states and exact ground states, to understand how well VQE states represent ground-state properties other than the energy.

I. VARIATIONAL METHOD

The variational quantum eigensolver (VQE) [4,24] is a method for finding ground states (or possibly also excited states [25]) of quantum Hamiltonians by classically optimiz-




FIG. 1. The kagome lattice.

ing the parameters of a parametric unitary $U(\theta)$. It has been extensively studied using classical simulation [18,19,26] as well as on superconducting quantum computers [27–29] and with trapped ions [30–32].

At a high level, the goal is to minimize the objective function

$$f(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle = \langle \psi_i | U^\dagger(\theta) H U(\theta) | \psi_i \rangle, \quad (1)$$

where $U(\theta)$ is a family of unitaries parametrized by the classical parameters θ , $|\psi_i\rangle$ an easily prepared reference state, and H the Hamiltonian whose ground state $|\psi_0\rangle$ we wish to prepare. If and only if $U(\theta)$ is sufficiently expressive [21,33] [meaning there exists some θ^* such that $|\psi(\theta^*)\rangle = |\psi_0\rangle$], then $f(\theta)$ is minimal if and only if $|\psi(\theta)\rangle$ is the ground state of H .

To turn this high-level description into a concrete quantum algorithm, we need to specify its components: We need to encode the system of interest, and hence the Hamiltonian H , into a qubit system, we need to make some (informed) choice of a sufficiently expressive yet implementable ansatz circuit $U(\theta)$, and we need a way to efficiently estimate the expectation value $\langle \psi(\theta) | H | \psi(\theta) \rangle$ from measurement samples. Unless the goal of the experiment is to find the ground-state energy E_0 , we also need to specify what to do with the found ground state, e.g., measuring some observables of interest.

A. Encoding and lattice mapping

Our target Hamiltonian is the kagome antiferromagnetic Heisenberg model (KAFH)

$$H = \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j = \sum_{\langle i,j \rangle} X_i X_j + Y_i Y_j + Z_i Z_j, \quad (2)$$

where the sums run over neighboring sites $\langle i, j \rangle$ in the kagome lattice (Fig. 1), $\hbar = 1$, and the usual factor of $\frac{1}{2}$ in the definition of the spin operators \vec{S}_i is omitted for simplicity.

As a spin model, the Heisenberg model (2) has a canonical mapping to qubits: simply assign one spin to each qubit. Finding an efficient mapping from the kagome to the square lattice is more challenging. The HV ansatz requires time evolution generated by the terms in the Hamiltonian. Hence, we seek a mapping that can implement a Trotterized version of e^{-itH} , that is, to implement time evolution by all of the terms in H individually, with as few layers of parallel two-qubit gates as possible on a square lattice. The chosen mapping is shown in Fig. 2 and with it a Trotterized version of e^{-itH} can be implemented in seven layers of two-qubit gates on a




FIG. 2. Mapping the kagome lattice onto a square lattice. The colors of the lines show which interactions exist only in the kagome lattice, only in the square lattice and which exist in both. The line types show which pairs of qubits must be swapped to implement the interactions existing in the kagome lattice but missing in the square lattice while the numbers group the interactions that can be implemented in parallel.

nearest-neighbor architecture, as shown in Fig. 3. In a system with all-to-all connectivity, such as certain ion trap architectures, the two-qubit gate depth can even be reduced to four layers. This is because the edges of the kagome lattice can be colored with four colors such that no two edges with the same color share a vertex, as shown in Appendix C 3. Then, all terms in H corresponding to the same color can be implemented in parallel. Kattemölle and van Wezel [23] implemented this ansatz for periodic lattices on 12, 18, and 24 qubits and found results comparable to ours shown in Fig. 6.

B. ansatz circuits

The barren plateau problem [20,21] shows that for VQE to be practical even at larger system sizes or for deeper circuits it is imperative to make use of prior knowledge about the problem in the ansatz circuits $U(\theta)$. One such ansatz is the Hamiltonian variational (HV) ansatz [17]. It is based on the adiabatic theorem of quantum mechanics: If a system is prepared in the ground state of an (simple) initial Hamiltonian H_0 and then evolved under a time-dependent Hamiltonian $H(t)$, changing sufficiently slowly from H_0 into a target Hamiltonian H , the system will remain in the ground state and we end up with the ground state of H_1 . Setting $H(t) = (1 - \frac{t}{T})H_0 + \frac{t}{T}H$ the time evolution can be approximately implemented by Trotterization, i.e., alternately applying $e^{-iH_0\Delta t}$ and $e^{-iH\Delta t}$ for short times Δt and possibly also Trotterizing $e^{-iH_0\Delta t}$ and $e^{-iH\Delta t}$. To turn this intuition into a variational circuit the short Δt 's are replaced by longer evolution times θ_i which are then optimized. This makes the ansatz circuit with $2p$ layers:

$$U(\theta) = \prod_{i=1}^p e^{-i\theta_{i,0}H_0} e^{-i\theta_{i,1}H}. \quad (3)$$

We call this ansatz the ‘‘per hamiltonian’’ ansatz because it has one parameter per application of H_0 and H . In our case, e^{-iH} cannot be directly implemented because the individual terms in Eq. (2) do not commute. Instead, we split $H = \sum_{j=1}^5 H_j$ according to the five subgraphs labeled in Fig. 2. This yields

1. Let all $\text{---}1\text{---}$, $\leftarrow 1 \rightarrow$ and $\text{=}1\text{=}$ connected qubits interact
2. Let all $\text{---}2\text{---}$, $\leftarrow 2 \rightarrow$ and $\text{=}2\text{=}$ connected qubits interact
3. Let all 3 connected qubits interact
4. Swap all $\leftarrow 1 \rightarrow$ and $\leftarrow 2 \rightarrow$ connected pairs of qubits
5. Let 4 connected qubits interact by \parallel
6. Let $\text{---}5\text{---}$ connected qubits interact by $\text{=}1/2\text{=}$
7. Swap $\leftarrow 1 \rightarrow$ and $\leftarrow 2 \rightarrow$ back for initial configuration

FIG. 3. One full round of two-qubit interactions needed to implement all interactions of the kagome lattice in the square lattice. Here we usually assume that any two-qubit gate can be implemented as one elementary operation. With a \sqrt{i} SWAP native gate, as used in some architectures, the interaction gates take two two-qubit gates and the SWAPs take three two-qubit gates.

an alternative parametrization with $6p$ parameters:

$$U(\theta) = \prod_{i=1}^p \left[e^{-i\theta_{i,0}H_0} \prod_{j \in \{3,2,1,5,4\}} e^{-i\theta_{i,j}H_j} \right], \quad (4)$$

where $\prod_{j \in \{j_1, \dots, j_n\}} U_j = U_{j_1} \dots U_{j_n}$ denotes the product ordering. Because the decomposition $H = \sum_{j=1}^5 H_j$ corresponds to an edge coloring of the lattice we call this ansatz the ‘‘per edge color’’ ansatz. A variant of this ansatz is the ‘‘per edge color ii’’ ansatz, where we drop the $e^{-i\theta_{i,0}H_0}$ factors because we choose $H_0 = H_1$. The ‘‘per edge color’’ ansatz can be made even more expressive by allowing different evolution times for all interaction terms in Eq. (2) and using

$$U(\theta) = \prod_{i=1}^p \left[e^{-i\theta_{i,0}H_0} \prod_{(k,l)} e^{-i\theta_{i,kl}\tilde{S}_k \cdot \tilde{S}_l} \right]. \quad (5)$$

Here, the terms in the inner product are in the same order as in Eq. (4). Now each parameter corresponds to an edge in the lattice, hence, we dub this ansatz the ‘‘per edge’’ ansatz. Note that the circuits described by Eqs. (3)–(5) are all the same. The ansätze only differ in the number of independent parameters.

As an initial Hamiltonian H_0 we use the Heisenberg Hamiltonian on the dimer covering induced on the dimer covering induced by the terms acting in the first step of Fig. 3. Its ground state $|\psi_i\rangle$ is the product of singlets $|s=0\rangle = \frac{1}{\sqrt{2}}[|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle]$ on all connected pairs of spins. It can be prepared via

$$|s=0\rangle_{kl} = Z_k X_l \text{CNOT}_{kl} H_k |\downarrow\downarrow\rangle \quad (6)$$

and it is in the same $S^x = \sum_i X_i = 0$, $S^y = \sum_i Y_i = 0$, and $S^z = \sum_i Z_i = 0$ symmetry sector as the ground state for the lattices considered by us (shown in Figs. 4 and 13) with an even number of qubits. For lattices with an odd number of qubits the ground state was always twofold degenerate in the $S^\alpha = \pm \frac{1}{2}$ ($\alpha = x, y, z$) sectors and we prepared the initial state in the $S^\alpha = -\frac{1}{2}$ sectors. The exact ground states were found using the Lanczos algorithm.

All three ansätze (3)–(5) preserve all those three symmetries, meaning we stay in the correct symmetry sector

throughout the whole circuit. Additionally, this also gives a way to postselect samples: If a sample does not satisfy $S^\alpha = 0$, we know that an error must have occurred and we can discard that sample.

C. Measurements

Reducing the number of measurements needed to estimate expectation values in VQE has been extensively studied [34–36]. Most strategies are based on cleverly grouping the terms of the Hamiltonian such that all the terms in one group can be measured simultaneously, possibly after changing the measurement basis with some simple, local unitaries. In our case, the Hamiltonian can be decomposed as

$$H = \sum_{\langle i,j \rangle} X_i X_j + \sum_{\langle i,j \rangle} Y_i Y_j + \sum_{\langle i,j \rangle} Z_i Z_j \quad (7)$$

and each of the three terms can be estimated separately by measuring all qubits in the X , Y , or Z basis.

D. Probing observables

If VQE succeeds in preparing the ground state of the KAFH, that state can be further used to investigate the nature of this ground state. Many of the observables used in classical numerical or experimental studies are also readily measurable on quantum computers, such as the ground-state energy or correlation functions. Others, e.g., entanglement entropies or energy gaps, are more challenging to measure on a quantum computer.

1. Ground state and spin gap

The ground state energy is trivially measurable; $\langle H \rangle$ is exactly the cost function minimized by VQE. In contrast, the energy of the first excited state, and hence the energy gap, is not as easy to obtain. There exist several proposals for VQE algorithms to find excited energy states of Hamiltonians. The proposals most feasible on NISQ hardware are based on augmenting the Hamiltonian with projectors on the $k - 1$ lowest-lying states to find the k th eigenstate,

$$H_k = H + \sum_j^{k-1} \alpha_j |\psi_j\rangle \langle \psi_j| \quad \text{with } \alpha_j > 0 \quad (8)$$

such that the k th eigenstate is now the ground state of H_k and can be found by the usual VQE method. However, to estimate the expectation value of the second term in Eq. (8) one needs to either double the circuit depth [25] or double the number of qubits needed [37]. The subspace-search VQE algorithm [38] does away with the need of doubling the number of qubits or circuit depth, but potentially needs more complex ansatz circuits that are not as well motivated as the HV ansatz. Overall, none of the methods reviewed in [39] have as low requirements as the original VQE problem. Thus, if we are only just able to prepare the ground state with VQE we will be unable to prepare excited states for the same system.

But, because the Heisenberg model as well as our ansatz circuits conserve the total S^x , S^y , and S^z spin it is possible to obtain the ground states and energies within each symmetry sector separately by simply choosing $|\psi_i\rangle$ to be in that sym-

metry sector. This makes measuring the *spin gap*

$$\Delta E_S = \begin{cases} E_0(S^z = 1) - E_0(S^z = 0) & \text{even } N, \\ E_0(S^z = \frac{3}{2}) - E_0(S^z = \frac{1}{2}) & \text{odd } N \end{cases} \quad (9)$$

as easy (or hard) as the original VQE problem. For consistency with the literature we used the spin operators with eigenvalues $\pm \frac{1}{2}$ here, unlike in the rest of this paper.

In principle, VQE could use coherent noise in the gates to break the spin conservation to get from a $S^\alpha = 1$ into a $S^\alpha = 0$ sector and thus lower the energy. However, postselection of the samples in the correct spin sectors ensures the energy cannot be lowered through this process.

2. Correlation functions and structure factors

The presence or absence of long-range order is indicated by the decay of the *spin-spin*

$$C_S(\vec{i}, \vec{j}) = \left| \langle S_i^z S_j^z \rangle \right| \quad (10)$$

and *dimer-dimer correlations*

$$C_D(\vec{i}\vec{j}, \vec{k}\vec{l}) = \langle \vec{S}_i \cdot \vec{S}_j \vec{S}_k \cdot \vec{S}_l \rangle - \langle \vec{S}_i \cdot \vec{S}_j \rangle \langle \vec{S}_k \cdot \vec{S}_l \rangle, \quad (11)$$

where \vec{i} and \vec{j} are on adjacent sites and similarly for $\vec{k}\vec{l}$. In this section, \vec{i} denotes the real-space location of the i th qubit on the lattice and difference vectors are written as $\vec{i}\vec{j}$. One sample of *all* spin-spin correlations can be obtained by measuring all qubits in the computational basis once. The dimer-dimer correlation between two edges that share no qubit can be measured by running the reverse circuit of Eq. (6) on each of the dimers and then measuring in the computational basis, the result $|\downarrow\downarrow\rangle_{ij}$ means $\vec{S}_i \cdot \vec{S}_j = -3$ while the three other outcomes mean $\vec{S}_i \cdot \vec{S}_j = +1$. Note that the first term in Eq. (11) cannot be reconstructed from the measurements in




FIG. 4. The six graphs used for classical simulations. The graphs are mapped to a square lattice as described in Figs. 2 and 3. Because we chose the dimer covering on all edges connected in step 1 of Fig. 3 connected qubits as the initial state for VQE and want to start in the right S^x , S^y , and S^z sectors we had to restrict ourselves to even numbers of columns. In the 2×6 , 2×10 , and 3×8 lattices we marked the qubits between which we computed the correlations in Sec. IID with bigger nodes.

the X , Y , and Z bases that were used to estimate $\langle H \rangle$ because it also contains terms of the form $X_i X_j Y_k Y_l$.

The spatial structure of the ground state is revealed by the *static spin structure factor*

$$S^z(\vec{q}) = \frac{1}{N} \sum_{i,j} e^{i\vec{q}\cdot(\vec{i}-\vec{j})} \langle S_i^z S_j^z \rangle, \quad (12)$$

which can also be obtained from measurements of all qubits in the computational basis. In fact, all data needed to compute this static spin structure factor and the spin-spin correlations are already available from the VQE optimization; to measure the $\sum_{(i,j)} Z_i Z_j$ part of the Hamiltonian we already needed to measure all qubits in the computational basis.

II. RESULTS

We carried out extensive numerical simulations to investigate the performance of our ansatz circuits. We were mainly interested in whether our ansatz circuits are expressive enough to represent the ground state and, if yes, at which depth. Since the end goal is to determine the nature of the ground state by measuring different observables, we also studied how closely the states found by VQE represent these observables. Because of this focus on expressibility we computed exact expectation values $f(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle$ from the full wave functions and also exact, analytical gradients $\partial_{\theta_\mu} f(\theta)$ using Yao.jl's [40] automatic differentiation algorithms. As a classical optimization algorithm, we used the L-BFGS algorithm from the NLOpt optimization suite. Note that this is only a good choice for exact, classical simulations. On real hardware, or if the expectation value is computed from measurement samples instead of from the full wave function, it does not work because it relies on the knowledge of exact gradients. As a reference, we also computed the exact ground state for each Hamiltonian using the Lanczos algorithm implemented in ARPACK from a sparse representation of the full Hamiltonian. The simulations were carried out on the Google Cloud Platform using 16×2.8 GHz Intel Xenon CPUs for the lattice sizes ≤ 16 qubits and on $2 \times$ NVIDIA Tesla T4 GPUs for the larger lattices.

To study the scaling of the attainable fidelity as a function of the number of ansatz layers p and lattice size we ran classical simulations of VQE on the 6 different lattices shown in Fig. 4, ranging in size from 8 to 24 qubits. For all simulations, we ran all three different ansatz circuits (3)–(5) with different, random initial parameters multiple times and also once with initial parameters corresponding to a discretized annealing schedule.

A. Different parametrizations

Figure 5 shows the difference between the three ansatz circuits (3)–(5) and a modified “per edge color ii” ansatz circuit where the $e^{-i\theta_{i,0}H_0}$ was omitted on the 2×8 lattice. The infidelity with the true ground state decays exponentially as a function of ansatz layers for the first three ansatz circuits. As expected, the more expressive “per edge” ansatz represents the ground state better at lower depths than the “per edge color” or “per hamiltonian” ansätze. The “per edge color ii” ansatz saves one layer of two-qubit gates per ansatz layer, but at the




FIG. 5. Infidelity with the ground state after parameter optimization as a function of ansatz layers p for the four different ansatz circuits on the 2×8 lattice. Results are shown for 20 runs per data point with the initial parameters chosen uniformly random within $[0, \frac{1}{p}]$. The error bars reflect the standard deviation between the 20 different runs.

cost of significantly worse fidelities, compared with the “per edge color” ansatz. This remained true for other lattice sizes and hence we omit the results for this ansatz circuit from now on.

Figure 6 compares the expressiveness of the ansatz circuits between the different lattice sizes for the “per edge color” ansatz in Fig. 6(a) and the “per edge” ansatz in Fig. 6(b). Generally, the ground state of smaller lattices is better represented by shallow circuits than for larger lattices. In fact, for the 3×8 lattice 16 ansatz layers were not enough to achieve substantial overlap with the ground state even for the most expressive ansatz circuits. This is due to the fact that for this particular lattice the gap between the ground state and the first excited state is very small compared to the energy. More details are found in Appendix A 1. Except for the 3×8 lattice, the exponential decay of the infidelity as a function of p holds still true.

B. Required circuit depths

In Fig. 6 we saw that larger lattices needed deeper ansatz circuits to accurately represent the ground state. This is made more quantitative in Fig. 7, where we show the depth required to reach a given fidelity as a function of qubits. For the $2 \times x$ lattices these data are consistent with the needed depth scaling linearly with the lattice diameter. Because VQE for the 3×8 lattice failed to converge to the ground state we do not have enough data to extrapolate the scaling of the required depth for lattices with equal length in both dimensions. But, if we assume that entanglement spreads linearly with the circuit depth, then a scaling as low as $\sqrt{n_{qb}}$ seems possible. Still, given that we obtained high fidelity for only five different lattices, Fig. 7 should be interpreted with caution.

C. Barren plateaus?

The results in [41] indicate that the HV ansatz exhibits only mild barren plateaus in the case of the XXZ and transverse field Ising model (TFIM) on 1D chains. To investigate whether the same is true for our ansatz circuits for the KAFH, we evaluated the cost function gradients at the first five points during each optimization run, i.e., at essentially uni-

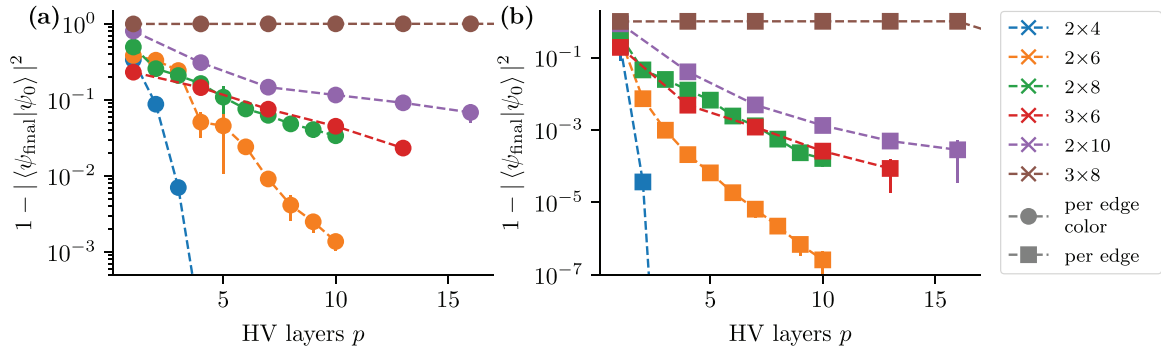


FIG. 6. Infidelity with the ground state after parameter optimization as a function of ansatz layers p for different lattice sizes. (a) The “per edge color” ansatz with one parameter per layer. (b) The more expressive “per edge” ansatz with one parameter for each gate. Results are shown for 10, 10, 4, 2, 2, 2, and 2 runs per data point for the 2×4 , 2×6 , 2×8 , 3×6 , 2×10 , and 3×8 lattice, respectively. As for Fig. 5 the initial parameters were chosen uniformly random within $[0, \frac{1}{p}]$ and the error bars reflect the standard deviation between the different runs. The 2×4 data are not fully shown because $p = 5$ was enough to produce the ground state with an infidelity of only 10^{-15} , the machine precision.

formly random points within $[0, 1/p]^{n_{\text{params}}}$ and long before the optimization had converged. The results for the “per edge” and the “per edge color” parametrization are shown in Fig. 8.

We see in the upper row of Fig. 8 that the variance of the first component of the gradient of the cost function does decay as a function of circuit depth. From the figures it is not clear if it is indeed an exponential decay, but the decay rate appears to lie between the fast decay reported by McClean *et al.* [20] for the random circuits and the much slower decay reported by Wiersema *et al.* for the TFIM. To understand the scaling of the magnitude of the gradient as a function of the number of qubits, we would need to consider much deeper circuits (cf. Fig. 4 of [20]) that were out of the scope of this work. It is also important to note that in the case of the “per edge” parametrization the number of parameters scales as $n_{\text{params}} \sim pn_{\text{qb}}$ compared with the other two parametrizations, where it scales as $n_{\text{params}} \sim p$. This means that each individual entry of the gradient in the lower right pane of Fig. 8 is much smaller than in the lower left pane.




FIG. 7. Number of ansatz layers p needed to represent the ground state with a threshold fidelity T as a function of qubits for the three different ansatz circuits and different thresholds. For the 3×6 and 2×10 lattices we ran the simulations only for $p = 1, 4, 7, \dots$. The missing intermediate points are obtained via linear interpolation from the data shown in Fig. 6. The 3×8 data are not shown here because significant fidelity with the ground state was only attainable with deep circuits and an informed choice of initial parameters. For more details see Appendix A 1.

D. Observables

The results in Figs. 6, 7, and 14 indicate that fairly deep circuits with many parameters are needed to obtain fidelities larger than 99.9% for larger lattices. However, the fidelity is a *global* figure of merit that gives error bounds on *all* observables, not only on *local* ones. And, ultimately, the goal of the experiments will be to measure mostly local observables, like those mentioned in Sec. 1D. To understand how well the states found by VQE reproduce local observables of the exact ground state, we compare the spin-spin correlation function and the static structure factor in the VQE states for different

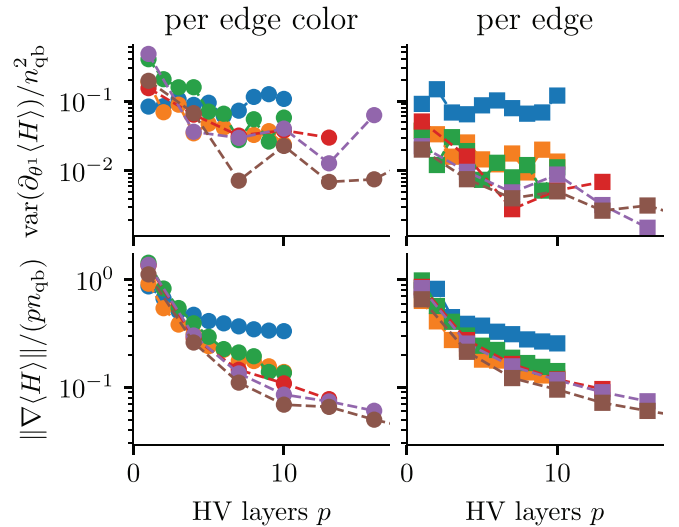


FIG. 8. The variance of the first component of the gradient (upper row) and norm of the whole gradient (lower row) as a function of the number of ansatz layers p for the “per edge color” (left column) and “per edge” (right column) parametrization for different lattices. In all cases, we rescaled with the number of qubits n_{qb} because the number of terms in the Hamiltonian scales with the number of qubits. For the norm of the whole gradient we also rescaled with $1/p$ because the number of parameters scales with p . The color and marker shape coding are the same as elsewhere in this paper (e.g., Fig. 6).




FIG. 9. Spin-spin correlation along the straight lines marked with bigger nodes in the 2×6 , 2×10 , and 3×8 lattices in Fig. 4 for the state found by VQE with different ansatz depths (colored) and for the exact ground state (black) as well as the first excited state (dashed), for the 3×8 lattice.

ansatz depths with the exact spin-spin correlation function and static structure factor of the exact ground state. The results are shown in Figs. 9,10.

As expected, the more expressive ansätze or deeper circuits reproduce the exact ground-state spin-spin correlations better. Reassuringly, the spin-spin correlations for $0 \leq j \leq 3$ of the VQE states for the 2×6 and the 2×10 lattices match the




FIG. 10. Comparison of the static structure factors $S^z(\vec{q}) = \frac{1}{N} \sum_{\vec{i}, \vec{j}} e^{i\vec{q} \cdot (\vec{i} - \vec{j})} \langle S_i^z S_j^z \rangle$ of the VQE states for different ansatz depths p and the exact ground state for the 2×10 and the 3×8 lattices with the “per edge color” parametrization. \vec{q} is measured in units of the inverse lattice spacing, hence, we show a little more than the first Brillouin zone here to make the periodicity more clear. As in Fig. 9 the 2×10 and the 3×8 lattices showed the biggest difference between the VQE state and the exact ground state.

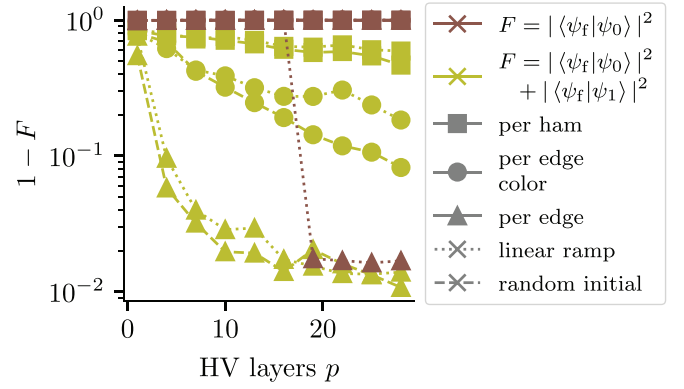


FIG. 11. Infidelity with the ground state and the ground state + first excited state as a function of ansatz layers p for the three different ansatz circuits on the 3×8 lattice. The “random” initial parameters are, again, chosen uniformly random within $[0, \frac{1}{p}]$ while the “linear ramp” initial parameters are chosen according to a discretized, Trotterized annealing schedule.

exact spin-spin correlations equally well, even though the global fidelity of the 2×6 VQE states with the exact ground state is much better than that of the 2×10 lattice (see Fig. 6). We found the same to be true for the 2×8 and the 3×6 lattices. This implies that even though it is hard to reach high fidelity with the exact ground state for large lattices, the VQE states may still represent local observables well. For the 3×8 lattice, on the other hand, there is a marked discrepancy between the spin-spin correlations of the VQE states and of the exact ground state, presumably because in this case the VQE state has a large fidelity with the first excited state and not with the ground state (cf. Fig. 11).

The VQE states reproduce the static structure factors of the ground state closely, as can be seen from Fig. 10. On larger lattices, we expect the peaks of the static structure factor to not lie on the corners of the first Brillouin zone, like they do here, but within the first Brillouin zone, if the periodicity of the ground state is larger than one unit cell, as it is for the ground state conjectured in [5]. It is important to note that the locations of the peaks in Fig. 10 are simply at the locations of the Fourier-transformed lattice. Instead, the features that indicate closeness between the VQE state and the exact ground state are the shapes of the peaks and the secondary peaks between the main peaks.

III. DISCUSSION AND OUTLOOK

In this work we investigated the performance of the VQE with the HV ansatz on the KAFH for different lattice sizes numerically. We also considered whether the fidelity of the states produced by VQE with the exact ground state is a too restrictive figure of merit by comparing different local, physically interesting observables in the VQE states and the exact ground state. Furthermore, we calculated the cost function gradients at random points in the parameter space to study how severe the barren plateau problem is with our ansatz.

In accordance with the results in [18,19], we find again that the relative energy error and infidelity with the ground state decay exponentially as a function of ansatz depth. We also

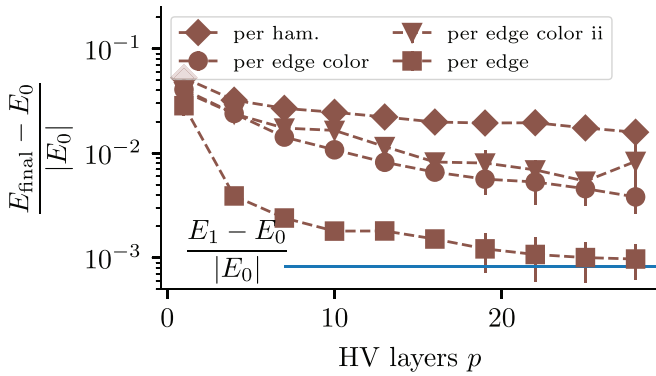


FIG. 12. Scaling of the relative energy error as a function of p for the 3×8 lattice with different ansatz circuits. Results are shown for three runs per data point and with the initial parameters chosen uniformly random within $[0, \frac{1}{p}]$. The error bars reflect the standard deviation between the different runs.

find, again in accordance with above results, that the number of layers needed to get to a fixed fidelity grows with the system size (see Fig. 7) and that in the presence of small energy gaps VQE may fail to find the ground state altogether (see Fig. 11). Furthermore, the results in Sec. IID show that even at lower fidelities the states found with VQE represent the values of local observables well. If this remains true for larger lattices and other local observables, this hints that requiring to find the ground state with a given fidelity is too strict a requirement, if one is mainly interested in the values of local observables.

Even though the barren plateaus of our ansatz are not as pronounced as for the hardware-efficient ansatz [20], the gradients still decay as a function of circuit depth (and presumably also as a function of system size). This is particularly true for the more expressive ansatz circuits with more free parameters. On real hardware, if the cost function is estimated from noisy samples, this opens up the question of the best classical optimization algorithm to use within VQE. To the best of our knowledge, there exist no classical optimization algorithms that are known to perform well for noisy, derivative-free problems with hundreds of parameters.

With the depth of 228 two-qubit gates reported in [42] as a circuit depth budget and ~ 50 qubits as a system size that is not feasible to simulate using exact, classical methods [43,44]; useful quantum advantage seems in reach with VQE for the KAFH. It should be noted, nevertheless, that approximate methods like DMRG [9,10,14] have been used to investigate the ground-state properties of the KAFH for systems with hundreds of sites, while other tensor-network-based approaches claim to handle thousands [11] of sites or even give infinite-size results [13]. However, most tensor-network-based approaches are biased towards low-entanglement solutions and they all have, by design, limited expressibility. Therefore, we hope that VQE can serve as an intermediate tool between exact diagonalization, which is limited to small system sizes, but gives exact results for these and can express arbitrary states, and tensor-network methods, which work for much larger systems, but have limited expressibility and known biases towards certain solutions.




FIG. 13. Three lattices with 1, 2, and 3 completely enclosed triangles on 15, 19, and 23 sites, respectively.

To get close to the ground state with relatively shallow circuits, one will almost certainly need parametrizations with one parameter per gate which yields hundreds of parameters to be optimized. Previous VQE experiments have not considered cost functions with that many free parameters and, hence, the best classical optimization algorithm for such problems is yet to be found.

Computing properties of the ground state accurately will require the use of error-mitigation techniques or small-scale quantum error correction. As discussed above, the symmetries of the KAFH allow a simple notion of error detection, by checking the total spin in each direction. Many other techniques targeted at NISQ-era quantum computers are now known [45].

On the 3×8 lattice we found that VQE consistently finds instead of the ground state the first excited state, which happens to be closer to the dimer covering state that we chose as an initial state. This hints that choosing an initial state that is already close to the ground state (if such a state is known and easily preparable) can help the performance of VQE. Kattemölle and van Wezel [46] observed similar behavior for a periodic lattice with six spins. However, minimizing the infidelity instead of energy, the recovered exponential decay of the infidelity with p is observed on all other lattices. Although measuring the fidelity with the true ground state is not possible in the real experiment, this still shows that the circuits are expressive enough to represent the true ground state with high fidelity, although it may be hard to find if the gap to the first excited state is only small.

Independently, Kattemölle and van Wezel [23] also studied VQE with the HV ansatz on the KAFH, albeit with a different ansatz circuit and initial state. Unlike ours, their ansatz circuit does not break the translational symmetry of an infinite lattice. Still, their initial state necessarily does because there exists no dimer covering of the kagome lattice that does not break the translational symmetry. They report slightly better scaling of the infidelity as a function of p for a 20-site lattice. This hints at the possibility that the performance of VQE with the HV ansatz also depends on the exact ansatz circuits and one might want to experiment with different mappings of the kagome lattice onto the square lattice of the hardware.




FIG. 14. Infidelity with the ground state as a function of ansatz layers p for the three different ansatz circuits on the lattices shown in Fig. 13. As before, the initial parameters are chosen uniformly random within $[0, \frac{1}{p}]$.

ACKNOWLEDGMENTS

The authors would like to thank F. Gambetta, R. Santos, and other members of the Phasecraft team for helpful discussions and feedback on the early drafts of this paper. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 817581) and from the EPSRC Grant No. EP/S516090/1. Google Cloud credits were provided by Google via the EPSRC Prosperity Partnership in Quantum Software for Modeling and Simulation (Grant No. EP/S005021/1). All data are available at the University of Bristol data repository [47].

APPENDICES

1. Details about the 3×8 lattice

VQE was only able to find the ground state of the 3×8 lattice with the most expressive “per edge” ansatz with $p \geq 19$ ansatz layers and using initial parameters corresponding to a Trotterized annealing schedule. However, changing the y axis in Fig. 6 from infidelity to relative energy error $\frac{E_{\text{final}} - E_0}{|E_0|}$ we observed the same exponential decay for the 3×8 lattice as for the other lattices, until the energy of the first excited state E_1 is reached (see Fig. 12). The reason is the very small relative energy gap $\frac{E_1 - E_0}{|E_0|} \approx 0.0008$ between the ground state and the first excited state. Among all other lattices considered

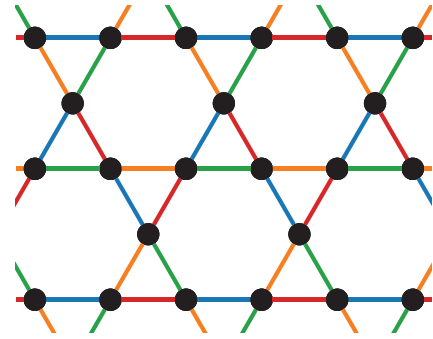


FIG. 15. An edge 4-coloring of the kagome lattice.

by us, “lattice 3” from Fig. 13 had the smallest energy gap of $\frac{E_1 - E_0}{|E_0|} \approx 0.005$. And, in fact, when considering not only the projection onto the ground state, but also onto the first excited state, we get much better fidelities, as is shown in Fig. 11. Moreover, we found that the first excited state is close to a dimer covering that shares many dimers with the dimer covering that we use as an initial state.

2. Results for other large lattices

Outside of the six lattices shown in Fig. 3 that fit into a rectangular part of a square lattice, we also considered three other large lattices shown in Fig. 13. For these lattices the terms acting in the first step of Fig. 3 do not form a dimer covering of the lattice and hence we added edges to H_0 such that it acts on $\frac{n_{\text{qb}} - 1}{2}$ edges and has its ground states in the $S^\alpha = \pm \frac{1}{2}$ sectors, the same symmetry sectors as the ground state of \hat{H} . We ran the same simulations as for Fig. 6 for these three graphs as well and show the results in Fig. 14.

For lattices 2 and 3 we see the same exponential decay in infidelity as before for all three parametrizations. But for lattice 1 and the “per edge” parametrization the infidelity decays much slower for $p \geq 7$, similar to what we see in Fig. 11 for $p \geq 10$.

3. An edge 4-coloring of the kagome lattice

Figure 15 shows an edge 4-coloring of the kagome lattice which can be used to Trotterize time evolution with H in only 4 layers on hardware with all-to-all connectivity.

-
- [1] J. I. Cirac and P. Zoller, Goals and opportunities in quantum simulation, *Nat. Phys.* **8**, 264 (2012).
- [2] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [3] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [4] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [5] P. Nikolic and T. Senthil, Physics of low-energy singlet states of the kagome lattice quantum Heisenberg antiferromagnet, *Phys. Rev. B* **68**, 214415 (2003).
- [6] R. R. P. Singh and D. A. Huse, Ground state of the spin-1/2 kagome-lattice Heisenberg antiferromagnet, *Phys. Rev. B* **76**, 180407(R) (2007).
- [7] Y. Ran, M. Hermele, P. A. Lee, and X.-G. Wen, Projected-Wave-Function Study of the Spin-1/2 Heisenberg Model on the Kagomé Lattice, *Phys. Rev. Lett.* **98**, 117205 (2007).
- [8] M. Hermele, Y. Ran, P. A. Lee, and X.-G. Wen, Properties of an algebraic spin liquid on the kagome lattice, *Phys. Rev. B* **77**, 224413 (2008).

- [9] S. Depenbrock, I. P. McCulloch, and U. Schollwöck, Nature of the Spin-Liquid Ground State of the $s = 1/2$ Heisenberg Model on the Kagome Lattice, *Phys. Rev. Lett.* **109**, 067201 (2012).
- [10] S. Yan, D. A. Huse, and S. R. White, Spin-liquid ground state of the $s = 1/2$ kagome Heisenberg antiferromagnet, *Science* **332**, 1173 (2011).
- [11] J.-W. Mei, J.-Y. Chen, H. He, and X.-G. Wen, Gapped spin liquid with z_2 topological order for the kagome Heisenberg model, *Phys. Rev. B* **95**, 235107 (2017).
- [12] Y. Iqbal, F. Becca, S. Sorella, and D. Poilblanc, Gapless spin-liquid phase in the kagome spin- $\frac{1}{2}$ Heisenberg antiferromagnet, *Phys. Rev. B* **87**, 060405(R) (2013).
- [13] H. J. Liao, Z. Y. Xie, J. Chen, Z. Y. Liu, H. D. Xie, R. Z. Huang, B. Normand, and T. Xiang, Gapless Spin-Liquid Ground State in the $s = 1/2$ Kagome Antiferromagnet, *Phys. Rev. Lett.* **118**, 137202 (2017).
- [14] H. C. Jiang, Z. Y. Weng, and D. N. Sheng, Density Matrix Renormalization Group Numerical Study of the Kagome Antiferromagnet, *Phys. Rev. Lett.* **101**, 117203 (2008).
- [15] M. Fu, T. Imai, T.-H. Han, and Y. S. Lee, Evidence for a gapped spin-liquid ground state in a kagome Heisenberg antiferromagnet, *Science* **350**, 655 (2015).
- [16] P. Khuntia, M. Velazquez, Q. Barthélemy, F. Bert, E. Kermarrec, A. Legros, B. Bernu, L. Messio, A. Zorko, and P. Mendels, Gapless ground state in the archetypal quantum kagome antiferromagnet $\text{ZnCu}_3(\text{OH})_6\text{Cl}_2$, *Nat. Phys.* **16**, 469 (2020).
- [17] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
- [18] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, Strategies for solving the Fermi-Hubbard model on near-term quantum computers, *Phys. Rev. B* **102**, 235122 (2020).
- [19] A. Uvarov, J. D. Biamonte, and D. Yudin, Variational quantum eigensolver for frustrated quantum systems, *Phys. Rev. B* **102**, 075104 (2020).
- [20] 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).
- [21] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and Barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [22] Z. Cai, Resource Estimation for Quantum Variational Simulations of the Hubbard Model, *Phys. Rev. Appl.* **14**, 014059 (2020).
- [23] J. Kattemölle and J. van Wezel, Variational quantum eigensolver for the Heisenberg antiferromagnet on the kagome lattice, [arXiv:2108.02175](https://arxiv.org/abs/2108.02175)
- [24] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [25] O. Higgott, D. Wang, and S. Brierley, Variational quantum computation of excited states, *Quantum* **3**, 156 (2019).
- [26] S. I. Mohtashim, T. Chatterjee, A. Das, R. Gupta, and S. S. Kale, Disordered phase in ising and metastability in cellular potts models hint at glassy dynamics, [arXiv:2106.11298](https://arxiv.org/abs/2106.11298).
- [27] 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).
- [28] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, E. Farhi *et al.*, Hartree-fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [29] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp, Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer, *Phys. Rev. Appl.* **11**, 044092 (2019).
- [30] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [31] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure, *Phys. Rev. A* **95**, 020501(R) (2017).
- [32] Y. Nam, J.-S. Chen, N. C. Pisenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, A. M. Ducore, K. M. Hudek, M. Keesan, S. M. Kreikemeier *et al.*, Ground-state energy estimation of the water molecule on a trapped ion quantum computer, *npj Quantum Infor.* **6**, 1 (2020).
- [33] T. Hubregtsen, J. Pichlmeier, P. Stecher, and K. Bertels, Evaluation of parameterized quantum circuits: On the relation between classification accuracy, expressibility, and entangling capability, *Quantum Mach. Intell.* **3**, 9 (2021).
- [34] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).
- [35] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, *J. Chem. Theory Comput.* **16**, 2400 (2020).
- [36] O. Crawford, B. van Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [37] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, Variational quantum algorithms for discovering hamiltonian spectra, *Phys. Rev. A* **99**, 062304 (2019).
- [38] K. M. Nakanishi, K. Mitarai, and K. Fujii, Subspace-search variational quantum eigensolver for excited states, *Phys. Rev. Research* **1**, 033062 (2019).
- [39] K. Kuroiwa and Y. O. Nakagawa, Penalty methods for a variational quantum eigensolver, *Phys. Rev. Research* **3**, 013197 (2021).
- [40] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, Yao.jl: Extensible, efficient framework for quantum algorithm design, *Quantum* **4**, 341 (2020).
- [41] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [42] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. B. Buckley,

- D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Y.-A. Chen, B. Chiaro, R. Collins, S. J. Cotton, W. Courtney *et al.*, and Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, [arXiv:2010.07965](https://arxiv.org/abs/2010.07965).
- [43] P. Sindzingre and C. Lhuillier, Low-energy excitations of the kagomé antiferromagnet and the spin-gap issue, *Europhys. Lett.* **88**, 27009 (2009).
- [44] A. M. Läuchli, J. Sudan, and R. Moessner, $s = \frac{1}{2}$ kagome Heisenberg antiferromagnet revisited, *Phys. Rev. B* **100**, 155142 (2019).
- [45] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Hybrid quantum-classical algorithms and quantum error mitigation, *J. Phys. Soc. Jpn.* **90**, 032001 (2021).
- [46] J. Kattmölle (private communication), the data are available at https://github.com/barbireau/HVQE/tree/main/data/kagome/periodic/own_graph/opg.
- [47] J. L. Bosse and A. Montanaro, Data from “Probing ground state properties of the kagome antiferromagnetic Heisenberg model using the variational quantum eigensolver,” <https://data.bris.ac.uk/data/dataset/25sbud0vk99lz2ohjch1cm1vwg>.