

Ground state energy and magnetization curve of a frustrated magnetic system from real-time evolution on a digital quantum processor

Aaron Szasz, Ed Younis, and Wibe Albert de Jong

Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

Models of interacting many-body quantum systems that may realize new exotic phases of matter, notably quantum spin liquids, are challenging to study using even state-of-the-art classical methods such as tensor network simulations. Quantum computing provides a promising route for overcoming these difficulties to find ground states, dynamics, and more. In this paper, we argue that recently developed hybrid quantum-classical algorithms based on real-time evolution are promising methods for solving a particularly important model in the search for spin liquids, the antiferromagnetic Heisenberg model on the two-dimensional kagome lattice. We show how to construct efficient quantum circuits to implement time evolution for the model and to evaluate key observables on the quantum computer, and we argue that the method has favorable scaling with increasing system size. We then restrict to a 12-spin star plaquette from the kagome lattice and a related 8-spin system, and we give an empirical demonstration on these small systems that the hybrid algorithms can efficiently find the ground state energy and the magnetization curve. For these demonstrations, we use four levels of approximation: exact state vectors, exact state vectors with statistical noise from sampling, noisy classical emulators, and (for the 8-spin system only) real quantum hardware, specifically the Quantinuum H1-1 processor; for the noisy simulations and hardware demonstration, we also employ error mitigation strategies based on the symmetries of the Hamiltonian. Our results strongly suggest that these hybrid algorithms present a promising direction for studying quantum spin liquids and more generally for resolving important unsolved problems in condensed matter theory and beyond.

Aaron Szasz: aszasz@google.com

1 Introduction:

Quantum spin liquids, insulating phases of matter with no magnetic order down to zero temperature and with long-range entanglement [1–3], have been the focus of intense computational effort over the past 50 years. The first theoretical proposal of a spin liquid came from Anderson in 1973, when he suggested that the true ground state of the antiferromagnetic Heisenberg model on a triangular lattice could be just such a state [4]. Although the true ground state of that model has since been shown to be magnetically ordered [5, 6], the current community consensus is that related models, including the triangular lattice Heisenberg model with both first- and second-neighbor interactions [7–11], the triangular lattice Hubbard model [12–20], and the kagome lattice Heisenberg model [21–29], do have spin liquid ground states; on the other hand, the precise nature of these spin liquid states remains in dispute. Other important problems in condensed matter physics, notably the mechanism for high-temperature superconductivity [30], have likewise not been conclusively resolved despite extensive study.

This remaining uncertainty is due to the limits of each of our classical computational methods for studying ground states: exact diagonalization is limited to small systems sizes, quantum Monte Carlo suffers from the sign problem for frustrated models, tensor networks are limited to low-to-medium entanglement and thus struggle with gapless ground states, and variational Monte Carlo methods are biased by the choice of ansatz. In principle, a ground state search on a quantum computer could avoid all of these pitfalls: a quantum algorithm could scale polynomially in the number of lattice sites, would not suffer from the sign problem, would have no limitation on allowable entanglement, and could be unbiased.

However, many proposed quantum algorithms such as quantum phase estimation [31–33] require very deep circuits or many qubits and are thus not suitable for near-term intermediate-scale quantum (NISQ) devices [34]. On the other hand, the most commonly discussed NISQ-

friendly method for ground state search, the variational quantum eigensolver (VQE) [35–37], may not realize a quantum advantage over classical algorithms. In the VQE approach, a parameterized quantum circuit prepares a state on the quantum device, and the parameters are optimized classically to find the parameters giving the lowest energy for the chosen circuit ansatz. Unfortunately, VQE suffers from the barren plateau problem [38–41]. Namely, for parameterized circuits that are sufficiently expressive to capture the true ground state of a complex quantum system, the classical parameter optimization has exponential cost due to exponentially small gradients of the variational energy with respect to the circuit ansatz parameters. Proposed strategies for getting around the barren plateau problem have thus far proven unsuccessful in general. For example, a circuit ansatz inspired by the Hamiltonian of interest might need many fewer parameters to approximately describe the ground state, and hence would be easier to optimize [42]; however, this approach only avoids barren plateaus for certain problems [43], and also an efficiently optimizable ansatz may be efficient to emulate classically [44]. Another hope would be to avoid barren plateaus by initializing the optimization with a state sufficiently close to the ground state; however, this “warm start” strategy requires increasingly good initial states with increasing system size [45]. Thus it remains unclear whether VQE can achieve quantum advantage for many problems of interest.

Our goal in this paper is to investigate a different class of hybrid quantum-classical algorithms, one based on real-time evolution and Krylov subspaces [46–52]. We demonstrate (a) that they have a potential advantage over existing classical algorithms and (b) that they could be practical on NISQ devices. In the algorithms we consider, we use a quantum computer to compute matrix elements $\langle \psi_0 | e^{-iHt} | \psi_0 \rangle$ for different values of the time delay t , where $|\psi_0\rangle$ is a state that ideally has a significant overlap with the true ground state. These overlap matrix elements are then classically post-processed to estimate the ground state energy. We specifically consider two such algorithms, unitary variational quantum phase estimation (UVQPE) [47, 50] and observable dynamic mode decomposition (ODMD) [51]. In each algorithm, the classical post-processing step is given by a small linear algebra problem: a generalized eigenvalue problem for UVQPE and a linear least squares problem for ODMD. Both algorithms can be made noise-resilient by a simple regularization of the linear algebra problem, making them suitable for noisy quantum devices before the advent of full fault-tolerance.

The particular model we use for our study is the aforementioned Heisenberg model on the kagome lattice shown in Fig. 1(a). The current consensus in the com-




Figure 1: (a) The kagome lattice, on which antiferromagnetic interactions between particles are believed to lead to exotic phases of matter such as quantum spin liquids. (b) A 12-spin star plaquette taken from the kagome lattice, which we use for a practical demonstration of hybrid algorithms. (c) An analogous 8-spin plaquette, also studied in the present work. Spins are numbered for later reference.

munity is that the ground state of the model is indeed a spin liquid, but there is debate as to the precise nature of the state. The leading candidate for the ground state is a Dirac spin liquid [24, 25, 27], which has no gap in the energy spectrum above the ground state, while other candidate ground states do have an energy gap [21–23]. Thus the nature of the ground state could be largely settled just by finding conclusively whether the low-energy spectrum is gapped or gapless. The model is also one of the hardest to study classically, as can be clearly seen in a recent survey paper [53]. Because the question of gaplessness only involves examining low-lying energy eigenvalues, an apparently simple task that is nevertheless very difficult classically, the kagome lattice Heisenberg model is an ideal problem to study via quantum computer. Indeed, VQE has also been demonstrated for studying this model [54–56] on clusters of up to 24 spins, finding accurate ground state energy with a classical emulator and moderately accurate energy with quantum hardware. However, due to the barren plateau problem discussed above, VQE will likely not scale well to system sizes large enough to answer the unresolved questions about the ground state of this classically challenging model.

We show how UVQPE and ODMD can be applied to find the ground state energy of the full 2D kagome lattice Heisenberg model with a lower computational cost compared with purely classical approaches. However, to reduce noise to a manageable level and to gain access to classical benchmarks, for a practical demonstration using quantum hardware or a classical emulator we need a smaller system. We thus opt for a proof-of-principle demonstration on a single 12-spin star plaquette of the

kagome lattice and an analogous smaller star plaquette with 8 spins, shown in Fig. 1(b) and (c), respectively. For both plaquettes, we run UVQPE and ODMD using an exact classical emulator both with and without statistical noise from measurement sampling. For the 8-spin plaquette, we also use a noisy classical emulator and the Quantinuum H1-1 processor. In all cases, both algorithms converge rapidly to the true ground state energy.

In addition to ground state energy, another interesting property of candidate spin liquid models is the magnetization curve. As an external magnetic field is applied and progressively strengthened, the zero-field spin liquid state can break down, giving rise to increasing magnetization and ultimately leading to a fully-polarized state. However, frustrated materials often demonstrate a series of “plateaus” in the spectrum, where a particular state with fixed magnetization is stable across a range of applied field strengths. These plateaus can also correspond to quantum spin liquids, or to a variety of interesting magnetic phases. Making use of the spin-symmetry of the Heisenberg model, we also use UVQPE and ODMD to find the magnetization curve for the 8- and 12-spin star plaquettes with a computational cost only polynomially higher than that of finding just the ground state energy.

We emphasize that the search for quantum spin liquids is far from the only condensed matter application of UVQPE, ODMD, and related algorithms. However, we select the kagome lattice Heisenberg model as our test case for several key reasons. First, the problem is one of the most challenging for classical computational methods, so a quantum algorithm that can make progress here would likely be useful for many other important problems. Second, the symmetries of the model, such as spin conservation, allow us to test algorithmic simplifications and error mitigation strategies that depend on those symmetries. Finally, the ground state of the model is itself an important unsolved problem, hence even a new algorithm specific to this problem would be important.

The organization of the rest of the paper is as follows. In Section 2 we introduce the hybrid quantum-classical algorithms in some detail, including a discussion of the “mirror circuit” approach to computing state overlaps. In Section 3 we discuss the application of the methods to the Heisenberg model on the full 2D kagome lattice. We focus on how to efficiently compute the expectation values $\langle\psi_0|e^{-iHt}|\psi_0\rangle$ on a quantum computer, including a model-specific discussion of an efficient symmetry-respecting implementation of the time evolution operator and a discussion of symmetry-based error mitigation. We also argue that the hybrid algorithms scale well with increasing system size and are thus a promis-

ing approach to solving this challenging model. In Section 4 we discuss the restriction of the model to the smaller star plaquettes, including how that restriction allows us to dramatically shorten the time-evolution circuits for use on quantum hardware. In Section 5, we show detailed results for the ground state energy from running UVQPE and ODMD on the star plaquettes, both with classical emulators and on hardware. In Section 6, we likewise present results for the magnetization curve. Finally, we conclude in Section 7.

2 Hybrid algorithms based on real-time evolution

We use two hybrid quantum-classical algorithms to obtain ground state energies from real-time evolution on a quantum computer, UVQPE [47] and ODMD [51]. For both algorithms, the quantum computer (or classical emulator) is used to find expectation values of the time evolution operator $U = e^{-iHt}$ in some initial state $|\psi_0\rangle$ for a range of times t . The expectation values are then used to construct either a generalized eigenvalue problem or a linear least squares problem, which is solved classically to obtain an approximate ground state energy and a representation of the corresponding eigenstate.

We provide practical introductions to UVQPE and ODMD in 2.1 and 2.2, respectively. We also provide, in 2.3, a thorough guide to using the mirror circuit approach [57] to measure expectation values of the time evolution operator, of the form $\langle\psi_0|e^{-iHt}|\psi_0\rangle$. In this section we focus on the general methods and postpone problem-specific details, such as how to find circuits to implement U or to prepare $|\psi_0\rangle$, to Sec. 3 below.

2.1 Unitary variational quantum phase estimation

The ground state of a Hamiltonian is far from generic. For example, in one dimension, the ground state of a gapped Hamiltonian has area law entanglement [58, 59], compared with volume law entanglement for typical states. Thus we might expect the ground state to have a high overlap with states that are relatively simple to prepare on a quantum computer. With this insight in mind, one promising approach to estimating ground state energy is to (a) find a set of states that can be feasibly generated on a quantum computer and which may have large overlap with the ground state, and (b) find the lowest energy of a state in the subspace spanned by those states by solving a generalized eigenvalue problem.

Variational quantum phase estimation (VQPE) [47] is a specific case of this approach, using a Krylov subspace in which the basis states are given by some initial state $|\psi_0\rangle$ that has significant overlap with the ground state

and time-evolved versions of the state:

$$|\psi_0\rangle, e^{-iH\Delta t}|\psi_0\rangle, e^{-2iH\Delta t}|\psi_0\rangle, e^{-3iH\Delta t}|\psi_0\rangle, \dots \quad (1)$$

for some fixed time interval Δt . The lowest energy within the subspace spanned by these states is the smallest eigenvalue of the generalized eigenvalue problem

$$\begin{pmatrix} \langle H \rangle & \langle e^{-iH\Delta t} H \rangle & \langle e^{-2iH\Delta t} H \rangle & \dots \\ \langle e^{iH\Delta t} H \rangle & \langle H \rangle & \langle e^{-iH\Delta t} H \rangle & \dots \\ \langle e^{2iH\Delta t} H \rangle & \langle e^{iH\Delta t} H \rangle & \langle H \rangle & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathbf{v} = \begin{pmatrix} 1 & \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \dots \\ \langle e^{iH\Delta t} \rangle & 1 & \langle e^{-iH\Delta t} \rangle & \dots \\ \langle e^{2iH\Delta t} \rangle & \langle e^{iH\Delta t} \rangle & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \lambda \mathbf{v} \quad (2)$$

where all expectation values are measured in the initial state $|\psi_0\rangle$ and we can truncate the basis (1) at any number of states that we choose. Note that the matrices have a Toeplitz structure, with all entries on each diagonal being equal. The energy estimate progressively improves as the number of basis states, i.e. the number of time evolution steps, increases.

A slight variant on this approach is to instead look for eigenvalues of the time evolution operator $U = e^{-iHT}$ for some fixed T . If T is small enough, so that the spectrum of HT is bounded by an interval of size less than 2π , then the eigenvalues of H can be unambiguously determined from the eigenvalues of U . In the special case that $T = \Delta t$, the same time step size used in generating the basis states, we get a significant improvement in the efficiency of the algorithm. The generalized eigenvalue problem becomes

$$\begin{pmatrix} \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \langle e^{-3iH\Delta t} \rangle & \dots \\ 1 & \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \dots \\ \langle e^{iH\Delta t} \rangle & 1 & \langle e^{-iH\Delta t} \rangle & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathbf{v} = \begin{pmatrix} 1 & \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \dots \\ \langle e^{iH\Delta t} \rangle & 1 & \langle e^{-iH\Delta t} \rangle & \dots \\ \langle e^{2iH\Delta t} \rangle & \langle e^{iH\Delta t} \rangle & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tilde{\lambda} \mathbf{v}. \quad (3)$$

We can then generate the matrices on both the left-hand side and the right-hand side just by measuring the expectation values

$$\langle \psi_0 | e^{-iH\Delta t} | \psi_0 \rangle, \langle \psi_0 | e^{-2iH\Delta t} | \psi_0 \rangle, \dots \quad (4)$$

whereas in the original formulation we also needed to compute Hamiltonian matrix elements such as $\langle \psi_0 | e^{-iH\Delta t} H | \psi_0 \rangle$. This reduces the number of measurements that need to be performed on the quantum computer by at least a factor of two, and most likely by more since H is in general non-unitary and thus requires multiple circuits to compute each expectation value. This modified version of the algorithm is called unitary variational quantum phase estimation, or UVQPE.

Note that, as remarked above, the eigenvalues of H can be unambiguously inferred from the eigenvalues of U only if the spectrum of $H \times \Delta t$ is limited to an interval of size 2π . As a result, in order to use UVQPE we must pick a time step size Δt that is inversely proportional to the spectral range of H . See Sec. IIIB of Ref. [51] for details.

Both variants of the method, VQPE and UVQPE, are susceptible to noise because the overlap matrix on

the right-hand side tends to have a very large condition number. This can be mitigated by performing a singular value decomposition of the overlap matrix and discarding all singular values below some threshold; in practice, we find good results when discarding singular values up to roughly 10x the noise level. After applying the singular value thresholding, UVQPE becomes relatively noise-resilient, as shown in Ref. [51], while VQPE remains highly susceptible to noise.

In addition to extracting the ground state energy, we also gain access to the corresponding eigenvector. The eigenvector for the lowest eigenvalue of the generalized eigenvalue problem describes a linear combination of the original basis states, known as a Ritz vector, that approximates the ground state of the original Hamiltonian. We will make use of this capability below to understand the convergence of the algorithm in the case of the Heisenberg model on the star plaquettes.

2.2 Observable dynamic mode decomposition

While UVQPE with singular value thresholding of the overlap matrix is reasonably noise-resilient, the robust-

ness of classically post-processing the expectation values in Eq. 4 to get the ground state energy can potentially be improved by replacing the generalized eigenvalue problem by a linear least squares problem. This approach is exemplified by a recently developed algorithm, observable dynamic mode decomposition (ODMD). [51]

ODMD is based on the classical dynamic mode decomposition (DMD) method, where linear least squares is used to find a linearization of the time-evolution operator for classical state vectors. In ODMD, we replace the classical state vectors by “observable vectors” of the form

$$\left(\langle e^{-miH\Delta t} \rangle \quad \langle e^{-(m+1)iH\Delta t} \rangle \quad \dots \quad \langle e^{-(m+d-1)iH\Delta t} \rangle \right)^T \quad (5)$$

consisting of expectation values of d successive powers of $e^{-iH\Delta t}$. We then look for a matrix A that approximately advances each observable vector to the next time step,

$$\left(\langle e^{-(m+1)iH\Delta t} \rangle \quad \langle e^{-(m+2)iH\Delta t} \rangle \quad \dots \quad \langle e^{-(m+d)iH\Delta t} \rangle \right)^T. \quad (6)$$

We can find the best linearization A across many time steps by solving the linear least square problem

$$\begin{pmatrix} \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \langle e^{-3iH\Delta t} \rangle & \dots \\ \langle e^{-2iH\Delta t} \rangle & \langle e^{-3iH\Delta t} \rangle & \langle e^{-4iH\Delta t} \rangle & \dots \\ \langle e^{-3iH\Delta t} \rangle & \langle e^{-4iH\Delta t} \rangle & \langle e^{-5iH\Delta t} \rangle & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} = A \begin{pmatrix} 1 & \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \dots \\ \langle e^{-iH\Delta t} \rangle & \langle e^{-2iH\Delta t} \rangle & \langle e^{-3iH\Delta t} \rangle & \dots \\ \langle e^{-2iH\Delta t} \rangle & \langle e^{-3iH\Delta t} \rangle & \langle e^{-4iH\Delta t} \rangle & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (7)$$

Note that in ODMD the matrices have a Hankel structure, with the entries on each anti-diagonal being equal. We then classically diagonalize A , which is a small and well-conditioned matrix, and the most negative eigenvalue phase is an approximation to the ground state energy. Like in VQPE or UVQPE, ODMD also gives access to the ground state as a linear combination of the basis states of Eq. (1), as shown in Ref. [51] Sec. 3C.

Here too we can make the solution more robust to noise by performing singular value thresholding on the matrix on the right-hand side. When using this thresholding approach, ODMD is slightly more resilient to statistical error/shot noise than is UVQPE. [51]

2.3 Mirror circuits for e^{-iHt} expectation values

Expectation values like $O \equiv \langle \psi_0 | e^{-iHt} | \psi_0 \rangle$ that are used as input for the classical post-processing portion of UVQPE and ODMD are commonly computed using the Hadamard test [60]. However, this approach can be prohibitively expensive because it involves controlled time evolution that replaces single-qubit gates by two-qubit ones and two-qubit gates by three-qubit ones.¹ Especially in the near-term, we want to reduce circuit depth as much as possible, so a better solution is needed, one that allows for lower gate depth.

Fortunately, there is precisely just such a method for computing matrix elements on a quantum com-

puter [57, 62]. The first key insight is that we can view the state $|\psi_0\rangle$ as $U_0|\mathbf{0}\rangle$ for some state preparation unitary U_0 , where $|\mathbf{0}\rangle = |0\rangle^{\otimes N}$ is the all-0 state. Note that finding a circuit implementing U_0 is problem-specific; we outline the construction for the example case of the Heisenberg model in Sec. 3.3 and Sec. 4.3 below. Then the expectation value becomes

$$O = \langle \psi_0 | e^{-iHt} | \psi_0 \rangle = \langle \mathbf{0} | U_0^\dagger e^{-iHt} U_0 | \mathbf{0} \rangle, \quad (8)$$

which we view as an inner product between the state on the left of the large vertical line, $|\mathbf{0}\rangle$, and the state on the right,

$$|0(t)\rangle \equiv U_0^\dagger e^{-iHt} U_0 |\mathbf{0}\rangle. \quad (9)$$

Then the matrix element $O = \langle \mathbf{0} | 0(t) \rangle$ is just the coefficient of the all-0 basis state when the state $|0(t)\rangle$ is written in the standard basis. Thus if we measure many copies of $|0(t)\rangle$ in the standard basis, we can approximate $|O|^2$ as the fraction of measured bit strings that are all zeros. Following Ref. [57], we refer to this sampled probability as F_1 .

So far we have only found $|O|$ and not the full expectation value O including phase, so we have not yet demonstrated a viable replacement for the Hadamard test. One solution is proposed in [57], making use of a reference state $|\text{Ref}\rangle$ that is easy to prepare, is orthogonal to $|\psi_0\rangle$, and is an eigenstate of H with known eigenvalue E_R . In other words, $|\text{Ref}\rangle$ must be an eigenstate of H whose energy can be derived classically. For the types

¹The multi-qubit gates can be avoided for certain nice models [61], but will be required in general.

of challenging problems where a quantum algorithm is needed to find the ground state, this type of classically solvable eigenstate can typically be found only by making use of symmetries of the particular Hamiltonian of interest; we give a concrete example for the Heisenberg model in Sec. 3 below. For now, we assume such a reference state exists. Then we can find a circuit that prepares

$$U_R|0\rangle = (|\psi_0\rangle + |\text{Ref}\rangle)/\sqrt{2}. \quad (10)$$

Using this circuit, we prepare the state

$$|0_R(t)\rangle = U_R^\dagger e^{-iHt} U_R|0\rangle \quad (11)$$

and measure in the standard basis. The observed probability that the measurement returns the all-0 bit string, which we call F_2 , approaches a theoretical value of

$$F_2 \rightarrow |(\langle\psi_0| + \langle\text{Ref}|)e^{-iHt}(|\psi_0\rangle + |\text{Ref}\rangle)|^2/4 \quad (12)$$

$$= (r^2 + 1 + 2r \cos(\theta + E_R t))/4$$

where r and θ are defined by $O \equiv r e^{i\theta}$.

Ref [57] stops here, with access to $|O|$ and the cosine of the phase $(\theta + E_R t)$. However, there is still an ambiguity of the phase, $\theta \leftrightarrow -(\theta + 2E_R t)$. To resolve this, we also want access to the sine of the phase, which we can find with

$$|(-i\langle\psi_0| + \langle\text{Ref}|)e^{-iHt}(|\psi_0\rangle + |\text{Ref}\rangle)|^2/4 \quad (13)$$

$$= (r^2 + 1 + 2r \sin(\theta + E_R t))/4.$$

This quantity is approximated by F_3 , the probability of measuring the all-0 bit string in the state

$$|0_{Ri}(t)\rangle = U_{Ri}^\dagger e^{-iHt} U_R|0\rangle \quad (14)$$

where U_{Ri} prepares the superposition

$$U_{Ri}|0\rangle = (i|\psi_0\rangle + |\text{Ref}\rangle)/\sqrt{2} \quad (15)$$

with an extra phase of $\pi/2$ between $|\psi_0\rangle$ and the reference state.

To summarize, F_1 , F_2 , and F_3 are given by the measured probability of the all-0 state in the three states $|0(t)\rangle$, $|0_R(t)\rangle$, and $|0_{Ri}(t)\rangle$ as given in Eqs. (9), (11), and (14), respectively. From these three measured quantities we can get the overlap according to

$$O = r e^{i\theta} = \left[2F_2 + 2iF_3 - (F_1 + 1) \cdot \left(\frac{i+1}{2} \right) \right] e^{-iE_R t}. \quad (16)$$

It turns out, however, that we can slightly improve this result. Here the magnitude of O , $r = |O|$, is computed from the combined results of three circuits, for F_1 , F_2 , and F_3 . In general the error from convolving the probability distributions for three independent quantities will

be larger than for any of the individual distributions, suggesting that we could do better by computing the magnitude from F_1 alone, $r = \sqrt{F_1}$, while keeping the angle from (16).

Indeed, using F_1 alone for the magnitude is already beneficial even if we consider just shot noise from sampling an otherwise noiseless simulation, as we show in Appendix A. A further benefit is that the hardware noise for the F_1 circuit will be slightly lower than for F_2 and F_3 , both because (a) the circuit is shorter since it does not involve the superposition with a reference state and (b) we can apply symmetry-based error mitigation that is much simpler than in the presence of a reference state (see Sec. 4.4).

Another question one could ask is, if we have a fixed measurement budget, i.e. a fixed total number of shots to be distributed between F_1 , F_2 , and F_3 , how should we distribute the shots among the three circuits? As we show in Appendix A, the optimal choice for a noiseless simulation is to use approximately the same number of shots for each of the three circuits, even when we only use F_1 to determine $|O|$. The reduced hardware noise for F_1 suggests that in practice we might want to allocate slightly more shots to that circuit. Hence in the experimental demonstration reported in Sec. 5 below, we divide our shots among the three circuits as 40% to F_1 and 30% each to F_2 and F_3 .

3 Application to the kagome lattice Heisenberg model

In this section, we first precisely define the Heisenberg model and discuss useful symmetries of the model. We then show how the expectation values $\langle\psi_0|e^{-iHt}|\psi_0\rangle$ can be efficiently computed for this model using the mirror circuit approach. Specifically, we provide an efficient quantum circuit implementation of the time evolution operator e^{-iHt} . We also show how to choose a good initial state $|\psi_0\rangle$ that should have reasonable overlap with the ground state while being easy to prepare on a quantum computer, and we furthermore show how to efficiently find a circuit to prepare the superposition of $|\psi_0\rangle$ with a reference state. We then derive bounds on the spectrum of the Hamiltonian in order to determine the maximum allowed value of the time step size Δt for use in UVQPE [Eq. (3)] and ODMD [Eq. (7)]. Finally, we argue that UVQPE and ODMD should scale well with increasing system size and thus are promising methods for finding the true ground state of the kagome lattice Heisenberg model.

3.1 The Heisenberg model

We consider the antiferromagnetic Heisenberg model,

$$H = J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j - h \sum_i S_i^z \quad (17)$$

where $\langle ij \rangle$ indicates nearest-neighbor pairs of sites on the kagome lattice, illustrated by the bonds between sites drawn in Fig. 1(a). $\mathbf{S}$ is the spin-1/2 operator $\boldsymbol{\sigma}/2 \equiv (\sigma^x, \sigma^y, \sigma^z)/2$, $J > 0$ is the antiferromagnetic coupling strength, and h is the strength of an external magnetic field. Note that with this definition of $\mathbf{S}$, we are effectively setting h to 1, and we will use those units throughout the paper.

The Heisenberg Hamiltonian conserves total S^z , $[H, \sum_i S_i^z] = 0$. This symmetry will be essential to our implementation of UVQPE and ODMD. Most importantly, each eigenstate of H lies within a subspace with fixed total spin (S), with several important consequences. First, if an initial state $|\psi_0\rangle$ has a well-defined total S^z , any time-evolved state $e^{-iHt}|\psi_0\rangle$ will also live in the same spin sector; consequently, since the ground state found by UVQPE and ODMD is in the span of the basis states (1), it will also be in the same fixed spin sector. We can thus use the hybrid algorithms to independently find the ground state energy within each spin sector. Since these subspaces are smaller than the full Hilbert space, it is easier to pick a good initial state with a high overlap with the ground state, and the algorithms should converge more quickly. Furthermore, in the important special case where there is no external field, $h = 0$, we need only consider a single spin sector, with total $S^z = 0$, since Lieb's theorem [63] guarantees that the ground state of an antiferromagnetic model will have total spin 0.

Second, the division of H eigenstates into orthogonal subspaces with fixed total spin enables the mirror circuit approach to measuring expectation values, state overlaps, and other observables on a quantum computer. As explained in Sec. 2.3 above, the mirror circuit method requires a reference state $|\text{Ref}\rangle$ that is guaranteed to be orthogonal to the state $|\psi_0\rangle$ in which expectation values are calculated, and that is furthermore an eigenstate of the Hamiltonian with known energy so that its time evolution is given by a simple known exponential. In the case of the Heisenberg model, we can use a fully polarized state, with all spins up or all spins down, as the reference. Each of these states is the only one in its symmetry sector, and hence must be an eigenstate of H and orthogonal to all states with different total S^z . Furthermore, the energy is easy to calculate analytically as simply the expectation value of H in a product state. On the kagome lattice, the energy of the fully polarized state with all spins pointing up is

$-(J/2)N_{\text{bonds}} - (h/2)N_{\text{spins}}$, where N_{bonds} is the number of nearest-neighbor pairs on the lattice and N_{spins} the total number of spins.

Finally, the conservation of total S^z during time evolution allows for simple approaches to error mitigation when using quantum hardware or a noisy emulator [64–66]. First, if an initial state $|\psi_0\rangle$ has a known total S^z , the states after time evolution should remain in the same symmetry sector, and thus measured bit strings not consistent with that symmetry sector can be thrown out. The use of state preparation circuits such as U_0 which takes $|\mathbf{0}\rangle$ to $|\psi_0\rangle$ complicates this logic, since U_0 in general does not conserve S^z . Fortunately, some bit strings can still be shown to be invalid even after applying U_0 , hence some errors can be detected and thrown out in post-processing. A less obvious consequence is that the magnitude of errors can be reduced by stochastically inserting single-qubit rotations that affect states in the wrong symmetry sector but not those in the correct symmetry sector, a method that was introduced in Ref. [66]. We provide further details on the error mitigation strategies after restricting to small spin plaquettes for our empirical demonstration; see Sec. 4 below.

In addition to the spin symmetry, the kagome lattice Heisenberg model also has spatial symmetries. The lattice, as shown in Fig. 1(a), is evidently symmetric under both horizontal and vertical reflections and under rotations by $\pi/3$, $2\pi/3$, and π (i.e. C_6 , C_3 , and C_2 rotations, respectively) depending on the center of the rotation. In principle these symmetries could also be used for error mitigation. Because exact time evolution of a spatially symmetric state would preserve the symmetry, the true probabilities of certain symmetry-related bit strings should be equal. Differences between the measured probabilities of such symmetry-related bit strings can be assumed to be due either to errors or to taking a finite number of samples, and either way the symmetry can be restored in classical post-processing at minimal cost. We do not use this error mitigation strategy in the practical demonstrations reported below because our mirror circuit method only requires the probability of the all-0 bit string, which transforms to itself under any spatial symmetry. However, we expect spatial symmetry-based error mitigation to play an important role in quantum studies of the kagome lattice that take other approaches.

3.2 Efficient circuits approximating e^{-iHt}

To exactly implement the time evolution operator $U = e^{-iHt}$ with a quantum circuit, we would first need to know the unitary U classically, which in general already requires solving for the spectrum of H , what we wanted to find in the first place. Even if we could somehow

find U without solving H , we would still have to solve the unitary synthesis problem to find a circuit that implements it, which also scales exponentially in system size.

In lieu of such an exact method, we take the standard approach of approximating the time evolution using a Suzuki-Trotter decomposition [67, 68]. We divide the terms of the Hamiltonian into groups, $H = \sum_i H_i$, where the time evolution of each H_i can be found exactly classically. In the lowest-order decomposition, the overall time evolution is approximated as

$$e^{-iHT} = \left[\prod_i e^{-iH_i T/m} \right]^m + \mathcal{O}(T^2/m). \quad (18)$$

Note that error per time step of size T/m scales as $(T/m)^2$, and the overall error T^2/m is the total, assuming $(T/m)^2$ is small. As the discretization in time becomes very fine, $m \rightarrow \infty$, the error term scaling as T^2/m goes to 0 and the Trotterized time evolution becomes exact.

For the Heisenberg model, a common choice is what we call a bond-by-bond Suzuki-Trotter decomposition. Each individual $\mathbf{S} \cdot \mathbf{S}$ term can be exponentiated classically, and in fact a simple quantum circuit with 3 CNOT operators can be found analytically for the operator $\exp(-i\theta \mathbf{S} \cdot \mathbf{S})$ for any value of θ . Any two of these exponentiated operators commute with each other if they act on disjoint pairs of sites, and thus if each H_i consists of $\mathbf{S} \cdot \mathbf{S}$ interactions on a collection of non-overlapping bonds, its exponential can easily be found classically and implemented as a quantum circuit. On the kagome lattice, the bonds can be organized into four disjoint groups, H_1 through H_4 , as shown in Fig. 1 of Ref. [69]. Another example is shown in Fig. 2. Note that because each vertex of the kagome lattice is connected to four edges, a bond-by-bond Trotter approach with three or fewer groups is not possible. If $\hbar$ is nonzero, a fifth step is needed, which includes all the on-site S^z terms.

This naive implementation of Trotter evolution for the kagome lattice Heisenberg model is valid and the corresponding quantum circuits are easy to generate, but it has some drawbacks. Notably, bond-by-bond Trotter evolution breaks all the spatial symmetries of the lattice, hence the symmetry or lack thereof in measured bit string counts can no longer be used to evaluate the reliability of a quantum computation or to perform symmetry-based error detection. Even more importantly, the total number of two-qubit gates required for each Trotter step can actually be reduced. For each time step, the bond-by-bond approach uses 3 CNOTs per bond, assuming the quantum hardware allows for gates to be applied between each pair of spins (i.e., qubits) connected by a bond. We will now show




Figure 2: Bond-by-bond Trotter decomposition for the kagome lattice Heisenberg model. The interactions are divided into four groups of non-overlapping bonds, indicated by thick solid lines, thick dashed lines, thin solid lines, and thin dashed lines. Note that the resulting Floquet operator for the Trotterized time evolution has a doubled unit cell indicated by the red parallelogram, containing 6 sites while the original model's unit cell has just 3. This decomposition also breaks mirror and rotation symmetries of the lattice.

an alternative approach that both preserves more of the spatial symmetries of the lattice and reduces the CNOT count to 8/3 CNOTs per bond, a more than 10% reduction.

In the alternative “triangle-by-triangle” Suzuki-Trotter decomposition, instead of exponentiating individual $\mathbf{S} \cdot \mathbf{S}$ interactions, we group together the three bonds around each triangle in the lattice to form a three-qubit operator H_Δ . We can still compute $U_\Delta = e^{-i\theta H_\Delta}$ analytically to get a three-qubit unitary, and then we can numerically solve the unitary synthesis problem to find a circuit that implements U_Δ . We can then divide the full Hamiltonian into two sets of non-overlapping triangles: H_1 containing all the up-pointing triangles and H_2 containing all the down-pointing ones. The two groups of terms are illustrated in Fig. 3.

Carrying out the numerical synthesis for U_Δ using the Berkeley Quantum Synthesis Toolkit (BQSKit) [70], we find that for any θ , it can be implemented using just 8 CNOTs, compared with 9 in total (3 per bond) when separately exponentiating the three spin-spin interactions.² The structure of the 8 CNOT circuit is shown in Fig. 4.

Evidently, the triangle-by-triangle Trotter evolution is both more efficient than the naive bond-by-bond approach, using 8/9 the number of CNOTs, and respects more of the spatial symmetries of the underlying lattice and Hamiltonian. For this reason, we will use the triangle-by-triangle approach in our experimental demonstration on the star plaquettes. More gener-

²While we do not have an analytical proof that an 8 CNOT circuit exactly implements the unitary for any θ , we numerically find that 8 CNOTs are sufficient to reproduce the unitary at any level of precision.




Figure 3: Triangle-by-triangle Trotter decomposition. The spin-spin interactions are organized into groups of three corresponding to triangles in the kagome lattice. The triangles are then grouped into up-pointing triangles (left panel) and down-pointing triangles (right). Efficient quantum circuits can be found classically for the exponential of H_1 and the exponential of H_2 . This Trotter decomposition preserves some symmetries of the underlying Hamiltonian, including one mirror symmetry and a C_3 rotation. As in the bond-by-bond decomposition, the unit cell is enlarged from 3 spins to 6.

ally, this approach to Trotterization, replacing bond-by-bond exponentiation with exponentiation of slightly larger groupings of qubits, should be beneficial for any lattice composed of small corner-sharing polygons, which are called bisimplex lattices [71].

We briefly note that the triangle-by-triangle approach remains efficient even when the connectivity among the qubits is restricted. For example, if the three qubits making up a triangle have only linear connectivity, the exact three-qubit time evolution can be still be implemented with just 12 CNOTs. In comparison, the bond-by-bond decomposition requires 15 CNOTs to implement all three separate spin-spin interactions: 9 for the interactions themselves, and 6 for SWAP gates. Thus the relative advantage of the triangle-by-triangle approach is maintained.

3.3 State preparation: how to choose and efficiently prepare a good initial state

Our next goal is first to choose a good initial state $|\psi_0\rangle$ that is likely to have high overlap with the ground state while being easy to prepare on quantum hardware; for concreteness, we focus on the case of $h = 0$. We will then show how to use the recently developed method of multi-state synthesis [72] to easily find circuits that implement the macroscopic superposition of $|\psi_0\rangle$ with the fully-polarized reference state.

3.3.1 Choosing and preparing the initial state

To find a good initial state that is likely to have high overlap with the ground state, we first consider once

again the symmetries of the Hamiltonian. As discussed in Sec. 3.1 above, H conserves total S^z . We will now specialize to the most-studied case, with $h = 0$, in which case H actually conserves not just S^z but also the total spin S . Then Lieb’s theorem guarantees that the initial state has the minimal possible total spin, which assuming an even number of sites in total will be 0 [63], so a good initial state should also be in this spin sector.

While restricting to $S = 0$ is a relatively strong constraint, the remaining Hilbert space still grows exponentially in the total number of spins; specifically, for a $2n$ -spin Heisenberg model, the dimension of the $S = 0$ eigenspace is $(2n)!/(n!(n+1)!)$ [73] $\sim 2^{2n}/n^{3/2}$.³ Furthermore, generic eigenstates of the total spin operator with small spin are highly entangled, and hence hard to prepare on a quantum computer [74].

Fortunately, there are some very simple, low-entangled states within the $S = 0$ space. Namely, a singlet dimer on two spins has total spin zero, and thus a tensor product of dimers does as well. So any covering of singlet dimers will be in the correct spin sector. A dimer covering is also easy to prepare on a quantum computer, since there is only entanglement between pairs of spins, and a singlet can be prepared with just one CNOT operator together with a few single-qubit rotations.

In fact, the dimer covering idea is even more promising once we consider the specifics of the Hamiltonian. Dimer coverings are known to be the ground states of some apparently frustrated models [75–77]. Furthermore, superpositions of dimer coverings, or resonating valence bond (RVB) states, are a type of spin liquid and have been proposed as the ground states of many frustrated models [4, 78, 79].

To understand why dimer coverings are promising on the triangular lattice, the starting point is the same triangle-by-triangle decomposition of the Hamiltonian that we used for our efficient implementation of Trotterized time evolution: we write $H = \sum_{\Delta} H_{\Delta}$, where for each triangle in the lattice H_{Δ} includes the three $\mathbf{S} \cdot \mathbf{S}$ interactions on the edges of the triangle. In the case we are considering, $h = 0$, H_{Δ} has just two energy eigenvalues, namely $\pm 3J/2$, each with multiplicity four. The ground state subspace corresponding to the eigenvalue $-3J/2$ is spanned by the states

$$\begin{array}{ccc}
 \begin{array}{c} \circ \\ \nearrow \quad \nwarrow \\ \circ \end{array} \uparrow, & \begin{array}{c} \uparrow \\ \nwarrow \quad \nearrow \\ \circ \end{array}, & \begin{array}{c} \uparrow \\ \nwarrow \quad \nearrow \\ \circ \end{array}, \\
 \begin{array}{c} \circ \\ \nwarrow \quad \nearrow \\ \circ \end{array} \downarrow, & \begin{array}{c} \downarrow \\ \nearrow \quad \nwarrow \\ \circ \end{array}, & \begin{array}{c} \downarrow \\ \nwarrow \quad \nearrow \\ \circ \end{array}
 \end{array} \quad (19)$$

³In comparison, the size of the $S^z = 0$ eigenspace is $(2n)!/2n!$, larger by a factor of $n + 1$.




Figure 4: Template for a quantum circuit implementing the exact time evolution for the Heisenberg Hamiltonian on one triangle of the kagome lattice. The empty boxes indicate single-qubit rotations; some are fully arbitrary while others are simpler, for example of the form $HR_z(\phi)$ where H is the Hadamard gate. The dashed vertical lines are included to clarify the structure of the circuit—the CNOT layouts of the first and second halves are the same. Note that the operator implemented by the circuit is symmetric with respect to permutations of the three qubits, so there are various equivalent templates. We use this particular qubit ordering in order to allow for a slightly increased degree of parallelism in applying gates when we consider just a single star plaquette, as discussed in Sec. 4.

where each ellipse denotes a pair of spins in the singlet state $|S\rangle = (|01\rangle - |10\rangle)/\sqrt{2}$. Note that this basis for the ground subspace is overcomplete: in each row, one of the states can be written as a linear combination of the other two.

Then for any state $|\psi\rangle$ on the full kagome lattice,

$$\langle\psi|H|\psi\rangle = \sum_{\Delta} \langle\psi|H_{\Delta}|\psi\rangle \geq \sum_{\Delta} -3J/2, \quad (20)$$

giving a lower bound on the ground state energy. If there was a singlet dimer covering of the kagome lattice such that each triangle had a dimer along one edge⁴, the state would saturate the lower bound, and thus would necessarily be an exact ground state of the model. Given the previously highlighted difficulty of solving the model, it should come as no surprise that no such dimer covering exists.

Nevertheless, we can find dimer coverings for which a large majority of the triangles do have a dimer on one edge, and these coverings are excellent candidates for the initial state $|\psi_0\rangle$. Two such coverings are shown in Fig. 5. In each, “defect” triangles with no dimer are shaded in red. A minimum density of defects is set by the fact that it is impossible to have two adjacent star plaquettes both with no defects. The dimer covering in the right panel of Fig. 5 is one of many that achieves this minimum.

Finally, we comment on how such a singlet dimer covering would be prepared on a quantum computer. Fortunately, this task is quite easy, since each dimer involves just two spins and thus all the dimers can be prepared from the state $|00\rangle$ in parallel, using a simple





Figure 5: Possible dimer coverings of the kagome lattice. (a) One simple covering has an enlarged 12-spin unit cell, as indicated by the dashed red lines. The triangles shaded blue have a dimer along one edge and hence the lowest possible local energy. “Defect” triangles with no dimer are located at the corners of the super-unit cell and are shaded in red. There are two defect triangles for six triangles with a dimer. (b) With a less regular pattern for the dimer covering, the density of defects (shaded in red and also marked with a red star) can be reduced, hence further reducing the energy of the state and presumably achieving a higher overlap with the true ground state. There must be at least one defect triangle for every two hexagon loops in the kagome lattice, and the pattern shown here achieves that minimum, so is a strong candidate for the initial state.

⁴More formally, the reduced density matrix of the state on each triangle should lie within the span of the states listed in Eq. 19.

circuit with a single CNOT gate:



Since this circuit is applied on each pair of qubits, the total CNOT count for the state preparation circuit for a system size with n total spins is $n/2$.


3.3.2 Preparing the superposition with the reference state

A remaining challenge is the preparation of the superposition of a dimer covering $|\psi_0\rangle$ with the reference state $|\text{Ref}\rangle$ for measuring the phase of observables in the mirror circuit approach. Recall from Sec. 3.1 that for the reference state we use the fully polarized state with all spins up, $|\text{Ref}\rangle = |\mathbf{0}\rangle$, since it is the unique state in its S^z symmetry sector and hence orthogonal to $|\psi_0\rangle$ and an eigenstate of H with known energy. Focusing on the circuit for the quantity F_2 , our goal is to prepare a state of the form (10), specifically

$$\frac{|\mathbf{0}\rangle + [(|\mathbf{01}\rangle - |\mathbf{10}\rangle)/\sqrt{2}]^{\otimes n/2}}{\sqrt{2}} \quad (22)$$

where n is the number of qubits, an equal weight superposition of the all-0 state with a tensor product of dimers.

We perform the state preparation in two easy steps. First, we prepare the Greenberger–Horne–Zeilinger (GHZ) state, $(|\mathbf{0}\rangle + |\mathbf{1}\rangle)/\sqrt{2}$, where $|\mathbf{1}\rangle = |\mathbf{1}\rangle^{\otimes n}$ is the all-1 state. With linear qubit connectivity, this state preparation can be achieved with $(n - 1)$ sequential CNOT gates following one Hadamard gate, as shown here with four qubits.




With a two-dimensional connectivity grid, the CNOT gates can be parallelized into $\mathcal{O}(\sqrt{n})$ layers; with all-to-all connectivity, parallelization can be further improved to give $\mathcal{O}(\log(n))$ layers. If we also make use of measurement and feed-forward, the GHZ state can be prepared in constant depth, though still requiring $\mathcal{O}(n)$ CNOT gates. [80]

In the second step, we apply a circuit that converts the GHZ state from Eq. 23 to the desired superposition (22). In general, finding a circuit that maps one state to another is exponentially hard in the number of qubits. However, in this case, even though both states

involve long-ranged entanglement, the transformation between the two can be accomplished entirely by local transformations on pairs of qubits. Specifically, we need to find a two-qubit circuit that solves the “multi-state mapping” problem

$$\begin{aligned} |00\rangle &\mapsto |00\rangle \\ |11\rangle &\mapsto (|01\rangle - |10\rangle)/\sqrt{2} \end{aligned} \quad (24)$$

with the images of $|01\rangle$ and $|10\rangle$ unconstrained. When the resulting circuit is applied, in the GHZ state, to each pair of qubits that should be in a singlet dimer in $|\psi_0\rangle$, the result is precisely the state (10). The multi-state mapping problem (24) can be easily solved using BQSKit [70, 72], giving the circuit



which implements the transformation

$$\begin{pmatrix} 1 & & \\ & 1/\sqrt{2} & 1/\sqrt{2} \\ & 1/\sqrt{2} & -1/\sqrt{2} \\ & & & 1 \end{pmatrix}. \quad (26)$$

Note that because we did not constrain the images of $|01\rangle$ and $|10\rangle$ under the transformation, there are other valid circuits for the desired state preparation that will map these two states differently.

Finally, to produce the slightly different superposition state (15), we can just add a single qubit rotation between the GHZ preparation and the application of the local transformation circuits (25). Specifically, acting with the gate $R_Z(-\pi/2)$ on any one qubit will transform the GHZ state as

$$\frac{|\mathbf{0}\rangle + |\mathbf{1}\rangle}{\sqrt{2}} \mapsto e^{-i\pi/4} \frac{|\mathbf{0}\rangle + i|\mathbf{1}\rangle}{\sqrt{2}}, \quad (27)$$

and then acting with (25) will give the desired superposition up to an irrelevant global phase.

In total, preparing the reference superposition states requires $(n - 1)$ CNOT operators to generate the GHZ state and n CNOTs to perform the transformation of $|\mathbf{1}\rangle$ to $|\psi_0\rangle$, for a relatively low total CNOT count of $2n - 1$.

3.4 Bounds on Hamiltonian spectrum

As noted in Sec. 2.1 above, in order to infer the eigenvalues of H from the eigenvalues of the time evolution operator $U_{\Delta t} = e^{-iH\Delta t}$, we need the spectrum of $H\Delta t$

to lie within an interval of size 2π . Thus to determine the maximum allowed size of Δt , we must derive bounds on the spectrum of the Hamiltonian.

With $h = 0$, we have already demonstrated that the ground state energy is bounded from below by $(-3J/2)N_\Delta$ where N_Δ is the number of triangles in the lattice. To bound the maximum energy, we can consider a single bond, with Hamiltonian $J\mathbf{S} \cdot \mathbf{S}$. This operator has eigenvalues $-3J/2$ and $J/2$ ($\times 3$), so the maximum energy for the whole system is bounded from above by $J/2$ times the number of bonds, which is $3N_\Delta$. This maximum energy is achieved by the fully polarized states with all spins up or all spins down. Thus we conclude that for $h = 0$, $\|H\|_2 = (3J/2)N_\Delta$.

If we turn on an external field (without loss of generality we take $h > 0$), the highest energy state will be the fully polarized state with spins pointing against the applied field, with total energy $(3J/2)N_\Delta + (h/2)n$ with n the total number of spins. A lower bound on the energy is $-(3J/2)N_\Delta - (h/2)n$. To prove this, we use the triangle inequality:

$$|E_0| = |\min_{|\psi\rangle} \langle \psi | H | \psi \rangle| \quad (28a)$$

$$\leq \max_{|\psi\rangle} |\langle \psi | H | \psi \rangle| \quad (28b)$$

$$\leq J \max \left| \left\langle \sum \mathbf{S} \cdot \mathbf{S} \right\rangle \right| + h \max \left| \left\langle \sum S^z \right\rangle \right| \quad (28c)$$

$$\leq (3J/2)N_\Delta + (h/2)n. \quad (28d)$$

In the final line, we used the bound on the $h = 0$ Hamiltonian from the previous paragraph and the fact that the maximal $|\langle S^z \rangle|$ on any given site is $1/2$, for the $|0\rangle$ and $|1\rangle$ states.

Finally, we conclude that $\|H\|_2 = (3J/2)N_\Delta + (h/2)n$, so we require $\Delta t < 2\pi/(3JN_\Delta + hn)$.

3.5 Favorable scaling with system size

Finally, we argue that ground state search for the kagome lattice Heisenberg model using UVQPE or ODMD has reasonably good scaling in system size. Good scaling is essential because very large systems are needed in order to definitively address the question of whether the low-lying energy spectrum is gapped or gapless. For a model which is gapless on the infinite two-dimensional plane, on a finite system with linear size L there will still be a gap of size scaling with $1/L$. Thus simulations on a finite size system will *always* observe a gap, and without considering the scaling of the gap size with L it will not be possible to distinguish a gap due to finite size from a gap that will persist as $L \rightarrow \infty$. Indeed, the largest exact classical simulation of this model considers clusters of up to 48 spins [28], or $L \approx 7$, which is still too small to determine whether the observed gap is fundamental or a finite size effect. For a

definitive answer, we will likely need to simulate much larger systems with hundreds or thousands of spins.⁵

We begin by estimating the total cost of measuring an expectation value $\langle \psi_0 | e^{-iHT} | \psi_0 \rangle$ using the mirror circuit approach. This involves three circuits, for measuring F_1 , F_2 , and F_3 , and each of those requires two copies of its respective state preparation circuit and an implementation of the time evolution e^{-iHT} . As established above, for an n -spin system, the state preparation for F_1 requires $n/2$ CNOT gates, while for F_2 and F_3 it requires $\approx 2n$ CNOTs. A single Trotter step using the triangle-by-triangle approach requires 8 CNOTs per triangle, and there are on average $3/2$ spins per triangle on the infinite 2D lattice, so in total the Trotter step requires $\approx (16/3)n$ CNOTs. Thus with m Trotter steps in total to implement e^{-iHT} , the total CNOT counts are approximately

$$F_1 : [1 + (16/3)m] \times n \quad (29a)$$

$$F_2, F_3 : [4 + (16/3)m] \times n \quad (29b)$$

The corresponding Trotter error scales as T^2/m , but also scales with the commutator of H_1 and H_2 , the two pieces of the Hamiltonian containing all the up-pointing triangles and all the down-pointing triangles, respectively. But the commutator of H_1 and H_2 can be written as a sum over all the individual commutators of one up-pointing triangle with one down-pointing triangle. Of these $\mathcal{O}(n^2)$ commutators, most are 0 because the triangles do not overlap; the number of nonzero commutators is $(3/2)$ times the number of triangles, which is proportional to n , and all are equal. This gives a total Trotter error scaling as $n \times T^2/m$. Then to achieve a fixed small Trotter error, the number of Trotter steps m should scale as nT^2 .⁶ In short, the total CNOT count for the circuits for measuring $\langle \psi_0 | e^{-iHT} | \psi_0 \rangle$ with fixed error should scale with n^2T^2 .

How small does the fixed error need to be? As shown for ODMD in App. A of Ref. [51], the singular value threshold used to filter out noise, which needs to be larger than the magnitude of the noise to be effective, must be smaller than the overlap between the chosen initial state $|\psi_0\rangle$ and the true ground state. In other

⁵Of course, this reasoning assumes that we have access only to energies and not to the corresponding eigenstates. Other signatures, including momentum of low-lying excitations, can help identify particular spin liquid phases even on somewhat smaller systems. [2, 11, 27] Nevertheless, at least hundreds of spins will still be required.

⁶The dependence on T^2 rather than T is somewhat unintuitive, since we typically fix a Trotter step size, so the number of Trotter steps scales linearly with the total evolution time. However, in that case the error accumulates over time, and here we require that the *total* Trotter error be approximately independent of the evolution time.

words, if $|\psi_0\rangle = \sqrt{p_0}|\text{GS}\rangle + \sqrt{1-p_0}|\phi\rangle$, then the Trotter error must be smaller than p_0 . This gives a required CNOT count of at least $n^2 T^2 p_0^{-1}$. We will return shortly to the question of estimating p_0 for reasonable choices of the initial state.

The next question is how many shots are required to get good estimates for F_1 , F_2 , and F_3 . As shown in Ref. [51], both UVQPE and ODMD are robust to a large degree of shot noise, which was modeled in that paper as Gaussian noise drawn from a distribution with fixed width independent of the size of the expectation value. However, we must still be able to resolve the scale of typical expectation values relative well. If we assume a worst case scenario that the overlap of $|\psi_0\rangle$ with the true ground state is small, $p_0 \ll 1$, then typical magnitudes of expectation values are approximately p_0 because e^{-iHT} will impart effectively random phases on the components of $|\phi\rangle$ that will cancel out on average. Then the magnitude of F_1 is approximately p_0^2 . Since shot noise for M samples scales as approximately $1/\sqrt{M}$, we need a number of shots that scales as p_0^{-4} .

The dependence on p_0 of both the required number of Trotter steps and the number of shots suggests that the algorithm will still fundamentally have an exponential scaling in system size. In particular, if the initial state $|\psi_0\rangle$ is made by tiling some translation-invariant unit cell as in Fig. 5(a), there is a fixed overlap per unit cell of $|\psi_0\rangle$ with the true ground state, hence the total overlap is exponentially small in the number of unit cells. If the overlap per unit cell is p and each unit cell has N_u spins, then the exponential cost p_0^{-1} will be proportional to $p^{-n/N_u} = (p^{-1/N_u})^n$. Fortunately, if the unit cell is large and if p is relatively close to 1, then $p^{-1/N_u} \ll 2$, so the hybrid algorithms will still have a significant exponential speedup compared with exact classical methods. Additionally, we find when we empirically consider small systems that the requirements on initial state overlap and number of shots are not too stringent. As we show in Sec. 5 below, for an 8-spin system with an initial overlap of 0.286, only around 1000 shots are required for UVQPE and ODMD to both rapidly converge to the correct ground state energy.

The last important question is how many time steps are needed for convergence. For UVQPE, the required number of time steps to achieve a specified level of error is given by Eq. (27) of Ref. [50]. We assume a very small size for the energy gap above the ground state, since otherwise the question of gaplessness of the model could be resolved without going to large system sizes. In that case, solving for the number of time steps j in the aforementioned equation gives

$$j \geq \log((E_{N-1} - E_0) \sin^2(\Xi) / \cos^2(\Xi) \varepsilon) / 2 \log(\tilde{\varepsilon}_{1,2}). \quad (30)$$

Here E_{N-1} is the maximal eigenvalue of the Hamilto-

nian, E_0 the minimal eigenvalue, $\cos^2(\Xi) = p_0$, ε is the desired level of error in the ground state energy, and $\tilde{\varepsilon}_{1,2} = 1 + 3(\Delta E \times \Delta t) / 2\pi$. ΔE is the energy gap above the ground state.

To find the scaling with system size, n , we examine each quantity in the expression. The full width of the energy spectrum, $E_{N-1} - E_0$, scales with n because energy is an extensive quantity. p_0 is exponentially small in n , though likely with a base close to 1, as explained above; we can therefore ignore $\sin^2(\Xi)$. Assuming we want a fixed level of error as we increase system size, the error ε is independent of n and can thus also be ignored in the thermodynamic limit. Furthermore, $\log(\tilde{\varepsilon}_{1,2}) \sim \Delta E \times \Delta t$. This gives

$$j \sim n \log(n) / (\Delta E \times \Delta t). \quad (31)$$

Finally, we consider the system-size dependence of ΔE and Δt . ΔE is the energy gap above the ground state, which as a worst-case scenario in the case of the gapless spectrum would go as the inverse of the linear size of the system, or $n^{-1/2}$. Δt is the time step size from Eq. (1), which as discussed in Sec. 2.1 needs to scale as the inverse of the spectral range of the Hamiltonian to avoid aliasing when extracting the eigenvalues of H from eigenvalues of $e^{-iH\Delta t}$; thus Δt is proportional to $1/n$. The result is that the required number of time steps scales as

$$j \sim n^{5/2} \log(n). \quad (32)$$

Note that if we required only that energy error per spin was constant when scaling up, rather than total error, so that $\varepsilon \sim n$, this would result in a slightly more favorable scaling of $n^{5/2}$ without the additional $\log(n)$.

Now recall that the cost, measured in number of CNOT gates, of finding one matrix element $\langle \psi_0 | e^{-iHT} | \psi_0 \rangle$ to sufficient accuracy using Trotter evolution scaled as $p_0^{-5} n^2 T^2$. To find the total cost across all time steps, we can sum T^2 up to the number of steps, which we estimate as $n^{5/2}$:

$$C = \sum_{s=1}^{n^{5/2}} (\Delta t \times s)^2 \sim \Delta t^2 \left(n^{5/2} \right)^3 \sim n^{11/2}. \quad (33)$$

Thus in total the cost of running UVQPE to find the ground state energy, in terms of CNOT gates, will scale as roughly $n^{15/2} p_0^{-5}$. The longest individual circuits will have on the order of $p_0^{-1} n^2 T^2 \sim p_0^{-1} n^2 (\Delta t)^2 (n^{5/2})^2 \sim p_0^{-1} n^5$ CNOT gates.

We can do the same type of analysis for ODMD. Starting from Eqs. (23) and (24) of Ref. [51] and again assuming that gap size $\Delta E = E_1 - E_0$ is small, we find that the required number of time steps scales as

$$d \sim (\Delta E \times \Delta t)^{-1} \sim n^{3/2}. \quad (34)$$

This gives a total CNOT cost of $n^{9/2}p_0^{-5}$ and longest individual circuits with $\mathcal{O}(p_0^{-1}n^3)$ CNOT gates.

We conclude the discussion of scaling with some important notes. (1) As mentioned in footnote 6, we have assumed that we need the same level of error in all expectation values. Possibly UVQPE and/or ODMD would still work even if the error increases in subsequent time steps, in which case the circuit depth would need to scale only with T rather than T^2 . This would significantly reduce the practical cost of the algorithms. (2) The potentially greatest contribution to the cost of the algorithms comes from the overlap between the initial state and the true ground state, p_0 . We argued that for a lattice model where both states are translation-invariant, p_0 is necessarily exponential in system size. We emphasize that other systems, such as molecules, where neither the ground state nor the initial state are translation invariant, may not be subject to the same exponential scaling. (3) Our analysis here was based on a first-order Trotter decomposition. A second-order Trotter decomposition [67] would result in a Trotter error scaling as nT^3/m^2 , which would result in a lower number of Trotter steps needed and hence a smaller number of CNOT gates. Other higher-order Trotter expansions might also reduce the resource requirements.

4 Restriction to star plaquettes

We now turn to a practical demonstration of the hybrid algorithms on small spin systems closely related to the full 2D kagome lattice model. We consider a single 12-spin star plaquette of the kagome lattice, shown in Fig. 1(b), as well as an analogous 8-spin system shown in Fig. 1(c), which can be viewed as one plaquette from the square kagome, or shuriken, lattice [81]. In each case, the Hamiltonian is the same as in Eq. (17) with spin-spin interactions of strength J between nearest neighbors and an applied field of strength h , just restricted to the smaller system.


In this section, we highlight important features of these smaller systems, including classically solvable exact ground states and a resulting technique for approximating the time evolution with shorter circuits. We also discuss some nuts-and-bolts details for an empirical demonstration with classical emulators and on quantum hardware, namely initial states and details of our symmetry-based error mitigation strategies. The actual results for ground state energy can be found in Sec. 5 and for the magnetization curve in Sec. 6.

4.1 Exact ground states when $h = 0$

Recall from Sec. 3.3, where we discussed the choice of initial state for the full two-dimensional lattice, that we

showed two key facts: (1) When $h = 0$, the Heisenberg Hamiltonian on a single 3-spin triangle has a 4-fold degenerate ground subspace spanned by the states shown in (19), which have two spins in a singlet dimer and the last spin free. (2) As shown in Eq. (20), a lower bound for the ground state energy is given by $-3J/2 \times N_\Delta$, where N_Δ is the number of triangles in the system.

We then concluded that a covering of the kagome lattice by nearest-neighbor singlet dimers would achieve the minimal energy $-3J/2$ on each triangle with a dimer along one edge, and hence would saturate the bound on the total energy and thus would necessarily be an exact ground state of the full system. While no such dimer covering exists for the full 2D lattice, it is possible on both star plaquettes. The resulting exact ground states of the star plaquette Heisenberg models (with $h = 0$) are the “pinwheel” states


(35)

12 spins
8 spins

and their mirror images. Note that the two mirror image pinwheel states for each plaquette are linearly independent but not orthogonal.⁷ This exact solvability of the $h = 0$ Heisenberg model on the star plaquettes is discussed further in Refs. [82, 83].


The exact solvability means that, although the kagome lattice has a very high degree of geometric frustration, the Heisenberg model on the star plaquettes is technically “frustration-free;” a Hamiltonian H is called frustration-free if it can be written as a sum of local terms $H = \sum H_j$, such that the ground state of H is simultaneously a lowest-energy eigenstate of each H_j . [84] The most obvious cases, such as Kitaev’s toric code model, clearly allow simultaneous diagonalization of all local terms because all the terms in the Hamiltonian commute. The present example is less obvious because the terms in H do not commute and furthermore the terms in the Hamiltonian as written in Eq. (17) do *not* all have a simultaneous ground state—rather, the frustration-free nature of the model only becomes apparent once the individual spin-spin interactions are grouped into triangles.

⁷Although we only explicitly discuss the 8- and 12-spin star plaquettes, in principle we could make larger loops of n corner-sharing triangles. Analogous pinwheel states are exact ground states for any n .

4.2 Trotterized time evolution


Just as for the full 2D kagome lattice, there are multiple choices for how to implement the time evolution operator e^{-iHt} for the star plaquettes. One option, since the systems are small enough to diagonalize exactly, would be to find the $2^n \times 2^n$ unitary for e^{-iHt} , then to perform unitary synthesis to find a circuit that implements the time evolution exactly up to some level of numerical precision. However, this approach has both practical and fundamental issues. On the practical side, “bottom-up” numerical synthesis is an exponentially expensive classical computation [85] while “top-down” matrix decomposition approaches produce extremely long circuits that are not practical on near-term devices [86, 87]. On the fundamental side, we would like to use a method such that our empirical approach at least in principle is scalable to the full 2D system.

We thus again consider possible Suzuki-Trotter decompositions. As in 2D, one option is a bond-by-bond decomposition, with bonds divided into four groups as shown here for the 8-spin plaquette: even-numbered bonds along the outside of the plaquette, odd-numbered outer bonds, even-numbered inner bonds, and odd-numbered inner bonds.



This decomposition is slightly better than the analogous one shown in Fig. 2 for the 2D lattice, since this decomposition still preserves one important spatial symmetry, namely the C_4 rotation for the 8-spin plaquette and C_6 rotation for the 12-spin plaquette.


The triangle-by-triangle decomposition is also a possibility. As shown below, as in 2D we have a decomposition $H = H_1 + H_2$, but here H_1 corresponds to even-numbered triangles around the plaquette [left column in (37)] and H_2 to odd-numbered triangles [right column].



The triangle-by-triangle approach maintains its most important advantage over the bond-by-bond approach,

namely the reduction in CNOT count from 9 to 8 per triangle. Furthermore, the triangle-by-triangle approach has two additional advantages specific to the star plaquettes, one minor and one which is important to the success of our demonstration on real quantum hardware.

The minor advantage comes from partial parallelization of Trotter steps to slightly reduce run time and thus noise. Consider three successive triangles around the edge of the plaquette. During one Trotter step, we first apply the circuit from Fig. 4 on the first and third triangles, then on the second:



For visual clarity we have not drawn the single-qubit rotations and have grouped the gates belonging to each copy of the circuit from Fig. 4. Evidently, the two pieces of the Trotter step can be made to overlap, resulting in a 1/8 reduction in the number of layers in the circuit.

More importantly, the triangle-by-triangle approach actually incurs exactly 0 error when time-evolving the pinwheel ground states, no matter the size of the Trotter time step. This remarkable property follows from the frustration-free nature of the Hamiltonian. To see this, consider one Trotter step of size t , given by the operator

$$F_t = e^{-iH_1 t} e^{-iH_2 t}, \quad (39)$$

where H_1 acts on the even triangles and H_2 on the odd triangles, as in (37). Then, since the exact ground state of H is also an eigenstate of the local Hamiltonians H_{Δ} acting on each triangle of the star plaquette, it is also an eigenstate of both H_1 and H_2 , each with eigenvalue $-3(J/2)N_{\Delta}/2$. The ground state is thus an eigenstate of F with eigenvalue $e^{-it \times -3(J/2)N_{\Delta}}$, exactly the same as its eigenvalue under the full time evolution $U_t = e^{-iHt}$. Intuitively, even though $e^{-iH_1 t} e^{-iH_2 t} \neq e^{-iHt}$, they are equal specifically when acting on the ground state because H_1 and H_2 each act on the state just as scalar multiplication. (Note, however, that the ground state of H is not an exact eigenstate of noisy approximations to H_1 and H_2 , hence on a real quantum device some Trotter error will remain.)

The lack of Trotter error for the ground state suggests a means of testing out UVQPE and ODMD on the frustration-free star plaquettes with shorter circuit depth than is required when trying to minimize Trotter error by taking many small Trotter steps, hence

with much lower noise when running on present-day and near-term quantum devices. Considering UVQPE in particular, since we know that for any time step size Δt , $U_{\Delta t}$ and $F_{\Delta t}$ have the same eigenvalue for the ground state we want to find, we can solve a generalized eigenvalue problem for the eigenvalues of the operator $F_{\Delta t}$ rather than $U_{\Delta t}$.

Then recall from Sec. 2.1 that we must choose Δt such that the spectrum of $H\Delta t$ lies in an interval of size 2π in order to unambiguously find eigenvalues of H from eigenvalues of $U_{\Delta t}$. Let us further assume that the eigenvalues of $H\Delta t$ specifically lie in the interval $(-\pi, \pi)$, and let the minimal and maximal values be $E_0\Delta t$ and $E_{N-1}\Delta t$. Then by the triangle inequality, the phases of the eigenvalues of $F_{\Delta t}$ will also lie in $[-E_{N-1}\Delta t, -E_0\Delta t]$. As a result, the ground state of H is guaranteed to still have the most positive phase for $F_{\Delta t}$ and thus will be found by the UVQPE variant where we look for eigenstates of $F_{\Delta t}$.

We can likewise generate the subspace used for the generalized eigenvalue problem by evolving the initial state $|\psi_0\rangle$ with F rather than U :

$$\begin{aligned} |\psi_m\rangle &= U_{m\Delta t}|\psi_0\rangle \\ \mapsto F_{m\Delta t}|\psi_0\rangle &= e^{-iH_1m\Delta t}e^{-iH_2m\Delta t}|\psi_0\rangle. \end{aligned} \quad (40)$$

Using this basis and solving for the ground state using F , the (j, k) element of the overlap matrix on the right-hand side of Eq. (3) would then become

$$\begin{aligned} \langle e^{-iH(k-j)\Delta t} \rangle &= \langle e^{iH_j\Delta t}e^{-iH_k\Delta t} \rangle \\ \mapsto \langle e^{iH_2j\Delta t}e^{iH_1(j-k)\Delta t}e^{-iH_2k\Delta t} \rangle. \end{aligned} \quad (41)$$

Likewise, the (j, k) element of the time-evolution oper-

ator matrix on the left-hand side becomes

$$\begin{aligned} \langle e^{-iH(1+k-j)\Delta t} \rangle &= \langle e^{iH_j\Delta t}e^{-iHt}e^{-iH_k\Delta t} \rangle \\ \mapsto \langle F_{j\Delta t}^\dagger F_{\Delta t} F_{k\Delta t} \rangle. \end{aligned} \quad (42)$$

Effectively, each matrix element requires at most three Trotter steps, substantially limiting the circuit depth while still producing a linear algebra problem that is, at least in theory, guaranteed to give the correct ground state energy.

We emphasize that this method, solving the generalized eigenvalue problem for $F_{\Delta t}$, is primarily useful in the special case that we already know the ground state. Note, however, that for small enough Δt , $F_{\Delta t}$ and $U_{\Delta t}$ only differ by an error of order Δt^2 , so even for real unsolved problems such as the full 2D kagome lattice Heisenberg model, it could potentially provide a good approximation to the eigenvalues.

Even in the specialized setting of a frustration-free Hamiltonian, this approach still has some drawbacks. First, we lose the Toeplitz structure of the matrices, since letting $j \rightarrow j+1$, $k \rightarrow k+1$ in Eq. (42) changes the value of the matrix element. Then unlike in UVQPE as presented in Sec. 2.1, we can no longer measure only the expectation values from the top row of the matrix, significantly increasing the number of different circuits that must be run. Second, even for the 8-spin plaquette, there are still 48 CNOT gates per Trotter step, so even three steps may lead to a high degree of noise on present-day quantum devices.

To make hardware demonstrations feasible we make a further approximation: we simply replace all occurrences of e^{-iHt} in Eq. (3) by a single Trotter step $F(t)$, giving the generalized eigenvalue problem

$$\begin{pmatrix} \langle F_{\Delta t} \rangle & \langle F_{2\Delta t} \rangle & \langle F_{3\Delta t} \rangle & \cdots \\ 1 & \langle F_{\Delta t} \rangle & \langle F_{2\Delta t} \rangle & \cdots \\ \langle F_{-\Delta t} \rangle & 1 & \langle F_{\Delta t} \rangle & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathbf{v} = \begin{pmatrix} 1 & \langle F_{\Delta t} \rangle & \langle F_{2\Delta t} \rangle & \cdots \\ \langle F_{-\Delta t} \rangle & 1 & \langle F_{\Delta t} \rangle & \cdots \\ \langle F_{-2\Delta t} \rangle & \langle F_{-\Delta t} \rangle & 1 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \lambda \mathbf{v} \quad (43)$$

with all expectation values again taken in an initial state $|\psi_0\rangle$. This restores the Toeplitz structure and hence the efficiency of the method in terms of the number of circuits to be run, and it limits every expectation value to a short circuit with a constant depth corresponding to a single Trotter step, thus bringing hardware noise down to a manageable level. On the other hand, this generalized eigenvalue problem no longer clearly represents energy minimization within a subspace, and we must reconsider why the method should work at all.

First, it is at least plausible that the approach *could*

give the correct ground state energy because, if the initial state $|\psi_0\rangle$ is chosen to be the exact ground state, then the matrix on the left-hand side is exact a scalar multiple $e^{-iE_0\Delta t}$ times the matrix on the right, giving the correct eigenvalue λ . For a more general initial state with an overlap with the true ground state, we should consider how UVQPE works; as shown in Ref. [50], with successively larger matrices the linear combinations of basis states allow phase cancellation of the coefficients of excited states, ultimately isolating the true ground state. This process still occurs for




Figure 6: Circuit to prepare initial states for generating the subspace bases for UVQE and ODMD. We show the circuit for the 8-spin plaquette; there is an analogous state for the 12-spin case. The state at the dashed line is the pinwheel of (35), and the CZ gates to the right add entanglement between dimers.

the generalized eigenvalue problem in Eq. (43). In fact, the effective increased phase randomness due to Trotter error actually makes the non-ground state components cancel out even faster, giving rapid convergence.


4.3 Initial states

For the full 2D kagome lattice, we argued in Sec. 3.3 that dimer coverings that come as close as possible to having one dimer along an edge of each triangle are good initial states, likely having significant overlap with the ground state. Applying the same logic on the star plaquettes gives the pinwheel states of (35), which as discussed above are in fact exact ground states.

Since our goal is to use the star plaquettes as test cases for the hybrid algorithms, we need to pick different initial states that are close to, but not equal to, these known ground states. One solution would be to use a different dimer covering, where dimers connect spins that are not adjacent. However, we would like to use an initial state that is cheap to prepare on any quantum hardware with connectivity similar to that of the model itself. Our solution is to take the known dimer pinwheel ground state, then apply controlled-Z gates clockwise on the outer bonds that do not currently feature a dimer, as shown in Fig. 6. The total overlap of this state with the ground state subspace is 0.286 for the 8-spin star and just 0.001 for the 12-spin star.

As noted in Sec. 3.5, the algorithms require the noise level to be below the initial state overlap, so the latter initial state would require on the order of 10^6 shots per expectation value even if the only noise comes from sampling. We therefore also use a state where every other CZ operator is omitted. The resulting state has an over-

lap of 0.016 with the ground state subspace. Schematically, the two initial states look like



where the thick lines indicate the locations of CZ operators applied to the pinwheel state. In Sec. 5.1, we use the two initial states with their dramatically different overlaps with the ground state to explore in more depth how the algorithms depend on the quality of the initial state.

Note that applying the layer of CZ gates takes us out of the spin-0 sector, but we do remain in the $S^z = 0$ sector. If we were aiming to find the ground state energy of an unsolved model, an initial state that is less than fully symmetric would be a bad choice. In this case it simply makes the computational task more challenging and hence actually makes for a better test of the algorithms.

We can also consider good initial states for other S^z sectors, which will be relevant for finding the overall ground state when we turn on the magnetic field, h , and hence for finding the magnetization curve. To find initial states that have reasonable overlap with the ground state in each S^z symmetry sector, we use a partial dimer covering, leaving enough free spins to reach the S^z value for the sector. We illustrate our initial states for the 8- and 12-spin stars and give their overlaps with the true ground states of their respective S^z sectors in Fig. 7.

4.4 Error mitigation details

Implementations of quantum algorithms on near-term devices benefit from error mitigation strategies to reduce noise. Standard methods include zero noise extrapolation [88–93], especially via Richardson extrapolation [94], dynamical decoupling [95, 96], and probabilistic error cancellation [89, 93]. In addition to these problem-agnostic methods, for Hamiltonian simulation we can use the symmetries of the Hamiltonian to reduce and detect errors [64–66]. As noted in Sec. 3.1, we use two error mitigation strategies based on the conservation of total S^z : detection of bit strings inconsistent with conservation of S^z during time evolution and twirling with single-qubit gates to average away leakage into incorrect symmetry sectors.

S^z error detection: First, we use symmetry-based error detection. When $h = 0$, our initial state $|\psi_0\rangle$ consisting of CZ operators applied to the pinwheel state




Figure 7: Initial states for the 8- and 12-spin plaquettes in different symmetry sectors with $S^z > 0$. For larger S^z , the states are chosen to be similar to the pinwheel states of Eq. (35) while being in the correct S^z sector and reasonably spatially symmetric. The circled dimers are in a spin singlet state, while the remaining spins are in the state $|\uparrow\rangle$. The number under each state gives the overlap with the true ground state in that symmetry sector, or if there are multiple degenerate ground states, the total of the overlap with the ground subspace.

(35) has total $S^z = 0$. The triangle-by-triangle Trotterized time evolution conserves S^z , so if measured bit strings after time evolution have a different total S^z , they can be safely discarded.

However, we have overlooked a significant complication. Recalling the mirror circuit method from Sec. 2.3, we sample bit strings not from the state $e^{-iHt}|\psi_0\rangle$ but rather from $|0(t)\rangle = U_0^\dagger e^{-iHt}U_0|0\rangle$ when measuring F_1 and the analogous states [Eqs. (11) and (14)] when measuring F_2 and F_3 . Focusing first on the state for F_1 , we note that the state-preparation circuit implementing U_0 explicitly does *not* conserve S^z , since it must convert from the all-0 state, which is in the $S^z = L$ sector, to a state in the $S^z = 0$ sector. Thus although the time evolved state $e^{-iHt}U_0|0\rangle$ should be in the $S^z = 0$ sector, allowing for error detection, the state $|0(t)\rangle$ that we actually measure will not be.

Fortunately, we can overcome this by careful scrutiny of the implementation of U_0 given in Eq. (21). On each two-qubit dimer, U_0 implements the transformation

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{pmatrix}. \quad (45)$$

Then when we apply $U_0^\dagger$ to a state that lies entirely in the $S^z = 0$ sector, we will map each pair of spins with labels $(j, j+1)$ by the Hermitian conjugate of (45). Looking at the rows of the matrix, we see that a pair of spins in the state $|00\rangle$ will map to a linear combination of $|01\rangle$ and $|11\rangle$, as will spins in the state $|11\rangle$. On the other hand, a state of the form $|01\rangle$ or $|10\rangle$ will map to a linear combination of $|00\rangle$ and $|10\rangle$.

But if we are correctly in the $S^z = 0$ sector, each basis state with weight in $e^{-iHt}U_0|0\rangle$ will have an equal number of 0s and 1s, and thus when the state is divided into pairs of spins, the number of pairs in the $|00\rangle$ state will be equal to the number of pairs in the $|11\rangle$ state. Thus the total number of pairs in either the state $|00\rangle$ or the state $|11\rangle$ will be *even*. Thus, after applying $U_0^\dagger$, any measured bit string in the state $|0(t)\rangle$ that has an *odd* total number of spin pairs in the states $|01\rangle$ and $|11\rangle$ corresponds to a symmetry sector error, and that measurement can be discarded. On any star plaquette with $4n$ spins, this condition eliminates exactly half of the possible bit strings as invalid measurement outcomes.

Other related approaches could further limit the number of valid bitstrings, thus allowing more errors to be detected and the corresponding shots discarded. For example⁸, U_0 could be divided into a layer of X gates acting on some qubits to transform $|0\rangle$ into a product state in the target S^z sector, followed by local entangling gates that transform $|01\rangle$, rather than $|00\rangle$, into a singlet dimer. The layer of X gates could be skipped when applying $U_0^\dagger$, with bit-flips done in classical post-processing of measured bitstrings. Then F_1 would be the probability to measure the product state from after the initial layer of X gates, rather than the all-0 state. In this case, only the n choose $n/2$ bitstrings with $S^z = 0$ would be valid, and any other measured bitstring would be discarded.

This approach to error detection scales well with system size and with error rate. For an n -qubit system, checking each measured bit string has cost linear in n , so the total cost is nM where M is the number of shots;

⁸As suggested by one of the anonymous referees.

there is negligible overhead compared with the existing cost of classically recording measurement outcomes. Unlike other error detection schemes, this method also does not rely on adding ancillas or extra CNOT gates. The method will also scale well with decreasing error rates as quantum hardware improves: the number of shots discarded will go to 0. Thus although this method is based on “post-selection,” in principle it incurs no overhead whatsoever in quantum resources.

We can perform a similar analysis for the circuits for F_2 and F_3 to find a condition to filter out invalid measurements on the states $|0_R(t)\rangle$ and $|0_{Ri}(t)\rangle$. However, the analysis of how basis states propagate through the GHZ preparation circuit is sufficiently complicated that we do not make use of filtering for F_2 and F_3 . Empirically, it remains true that around half of the possible bit strings are not valid measurement outcomes if the symmetry sector is conserved during time evolution.

S^z -preserving twirling: Our second symmetry-based mitigation method is a form of twirling, or randomized compiling, introduced in Ref. [66]. We use the fact that a well-chosen layer of single-qubit R_z rotation gates can leave desired symmetry sectors invariant while adding a phase to others, so that in averaging over various runs, the erroneous symmetry sector components will cancel out.

To be concrete, suppose we have a product state of $2n$ spins where half are in the state $|0\rangle$ and half are in the state $|1\rangle$. Then if we apply $R_z(\theta)^{\otimes 2n}$ to the state, where

$$R_z(\theta) = \begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix} \quad (46)$$

and importantly θ is the same for each qubit, the net phase acquired will be exactly zero. The same will be true for any superposition of such product states, i.e. any state, not just a product state, with $S^z = 0$. On the other hand, a state with total spin of $S^z = 2$, with $n + 1$ qubits in the state $|0\rangle$ and $n - 1$ in the state $|1\rangle$, would acquire an overall phase of 2θ .

Now let's suppose we have a superposition of states with $S^z = 0$ and $S^z = 2$, $a|0\rangle + b|2\rangle$. If we run half our shots with an R_z layer inserted with $\theta = \pi/2$, our final measurement outcomes will be effectively sampled from the density matrix

$$\begin{aligned} \rho &= \frac{1}{2}(a|0\rangle + b|2\rangle)(a^*\langle 0| + b^*\langle 2|) \\ &\quad + \frac{1}{2}(a|0\rangle - b|2\rangle)(a^*\langle 0| - b^*\langle 2|) \\ &= |a|^2|0\rangle\langle 0| + |b|^2|2\rangle\langle 2|. \end{aligned} \quad (47)$$

Thus if the state was supposed to be in the $S^z = 0$ sector, but there was a small leakage of order ϵ into

the $S^z = 2$ sector due to hardware noise, this twirling procedure, which costs only an inconsequential layer of 1-qubit gates, reduces the contribution from the erroneous symmetry sector to be of order ϵ^2 . Note that, although a direct measurement of ρ either with or without twirling would give symmetry sector $S^z = 0$ with probability $|a|^2$, when the measurements come after the further unitary operations such as inverted state preparation circuits, the linear-in- ϵ off-diagonal elements of the density matrix that are present without twirling would have an effect on the final measurements.

This twirling method works even when we take the superposition with a reference state for measuring F_2 and F_3 . In that case, the state during time evolution should be a superposition of $S^z = 0$ and $S^z = L$, so as long as we choose θ to be an integer multiple of $2\pi/L$, both symmetry sectors will be unaffected. On the other hand, the most likely leakage will be to $S^z = \pm 2$, $L - 2$, and we can choose $\theta = \pi/2$ so that state components in all of these symmetry sectors will acquire a relative negative sign. We can thus cancel out some errors by running each of our circuits half the time with a layer of $R_z(\pi/2)$ gates after the Trotter evolution and half without.

5 Ground state energy results

Finally, we present our results from running both UVQPE and ODMD on the star plaquettes to find the ground state energy, focusing only on the case $h = 0$. Following Lieb's theorem [63], we therefore consider only the $S^z = 0$ symmetry sector. We show results using exact classical state vector simulation both with and without shot noise due to taking a finite number of measurements to determine F_1 , F_2 , and F_3 for each expectation value. These simulations use the exact time evolution operator e^{-iHt} . For the 8-spin plaquette, we also show results using both a noisy classical emulator and real quantum hardware, namely the Quantinuum H1-1 processor; these simulations use the single Trotter step approximation discussed in Sec. 4.2 where U_t is replaced by F_t .

5.1 Exact classical simulation results

For both the 8-spin and 12-spin star plaquettes, we use an initial state $|\psi_0\rangle$ given by the pinwheel state (35) augmented by a layer of CZ gates as described in Sec. 4.3. For the 8-spin plaquette we use four CZ gates as illustrated in Fig. 6; the overlap of this state with the ground state subspace is 0.286. For the 12-spin plaquette we consider two different initial states of this type, shown in Eq. (44); their respective overlaps with the ground state subspace are 0.001 and 0.016. We




Figure 8: Convergence of UVQPE and ODMD estimate of the ground state energy with the number of time steps/Krylov basis states used for the 8-spin star plaquette. Time evolution and computation of expectation values are numerically exact, with no error or shot noise. The initial state is the pinwheel plus a layer of CZ operators, as shown in Fig. 6. Even though the simulations are exact, we use noise filtering with three levels of singular value threshold, because we need to invert nearly singular matrices so even noise at the level of machine precision can be problematic, and because we want to understand the effect of the filtering on convergence. We see that in the noiseless case, more aggressive noise filtering (larger δ) leads to slower convergence.




Figure 9: We perform the same numerical experiment as in Fig. 8, but for the 12-spin plaquette using six-CZ initial state from the left of Eq. (44), which has just a 0.1% overlap with the ground subspace. As a result of this low overlap, both algorithms initially converge towards a low-lying excited state instead of to the ground state. With small δ , both algorithms eventually escape the local minimum and find the ground state energy; with aggressive noise filtering using $\delta = 0.1$, the algorithms never converge to the ground state.




Figure 10: A closer examination of the energy convergence for the 12-spin plaquette with the six-CZ initial state, using UVQPE with $\delta = 10^{-6}$. The upper panel shows the energy estimate as a function of the number of time steps/basis states, similar to the UVQPE data from Fig. 9, rightmost panel; energy error is measured in units of $J/2$. In the lower panel, we show the convergence of $|\tilde{v}_0\rangle$, the UVQPE approximation to the ground state, by finding its overlap with each eigenstate of H . We plot the overlaps for the five eigenstates that contribute most to the approximation to the ground state across the simulation. On the right we show where the energies of those eigenstates fall on the spectrum of H ; the five states include the ground state and several low-lying excited states. Evidently, in the first 10 or so time steps, UVQPE converges towards the excited state marked with a green square. With this small δ , the algorithm escapes the local minimum and converges to the true ground state; with $\delta = 0.1$ as on the left of Fig. 9, the algorithm gets stuck at the excited state.

then compute the expectation values of Eq. (4) in two ways: (1) direct computation of the expectation value via matrix multiplication and (2) exactly generating the states $|0(t)\rangle$, $|0_R(t)\rangle$, and $|0_{Ri}(t)\rangle$ from Eqs. (9), (11), and (14), respectively, then sampling from the states and using Eq. (16) to find the expectation values.

We process the expectation values using both UVQPE and ODMD, truncating small singular values from the matrices on the right-hand side of Eq. (3) and Eq. (7), respectively. Specifically, we truncate all singular values below a threshold that is a factor δ times the largest singular value of the matrix; empirically, we need to truncate singular values up to around 10x the noise level on the individual matrix elements in order for the algorithms to converge correctly.

For the time step size in the simulations, Δt , we recall from Sec. 3.4 that for $h = 0$ the Hamiltonian spectrum is bounded by $\|H\|_2 = (3J/2)N_\Delta$, thus requiring $\Delta t < 2\pi/(3JN_\Delta)$. The 8- and 12-spin plaquettes have $N_\Delta = 4$ and 6, respectively. We choose a relatively conservative value of $\Delta t/(J/2) = 0.1$, satisfying the bound for both plaquettes. The resulting (unitless) time-step operator is $U_{\Delta t} = \exp(-i \times 0.1 \sum \sigma \cdot \sigma)$.

The results of the simulations are summarized in a series of figures, as follows:

- In Fig. 8, we show the convergence of the ground state energy for both UVQPE and ODMD for the 8-spin plaquette, with expectation values computed exactly. We specifically plot, on a log scale, the error in the energy estimate as a function of time step, or in other words of the number of Krylov states in the basis (1). Although the expectation values are accurate to machine precision, the linear algebra problems in Eqs. (3) and (7) are ill-conditioned, so some noise filtering is still needed; we use three singular value thresholds, $\delta = 10^{-1}$, 10^{-3} , and 10^{-5} . In this case, because the simulation is noiseless so not much singular value filtering is needed, both algorithms converge faster when δ is smaller.
- In Fig. 9, we show the same calculation but for the 12-spin plaquette, using the initial state with six CZ operators as shown on the left of Eq. (44), which has an overlap of only 10^{-3} with the true ground states. When the level of noise filtering δ is greater in magnitude than the initial state overlap, the algorithm converges to a low-lying excited state instead of to the ground state—essentially, the contribution of the ground state to the initial state gets filtered out. When δ is comparable to or smaller than the overlap, as in the middle panel, both algorithms plateau at the energy of the low-lying excited state before eventually converging.
- In Fig. 10, we focus on UVQPE (with $\delta = 10^{-6}$) to demonstrate that the plateau in energy does indeed correspond to finding a low-lying excited state. Recall from Sec. 2.1 that UVQPE gives an approximation to the ground state eigenvector, $|\tilde{v}_0\rangle$, in the form of a linear combination of the basis states (1). For the small star plaquettes, we can explicitly generate the basis states and thus the specified linear combination, and we can find the overlap of $|\tilde{v}_0\rangle$ with the true ground state and other eigenvectors of the Hamiltonian. In the figure, we show the convergence of energy in the upper panel and the decomposition of the corresponding ground state approximation into Hamiltonian eigenstates in the lower panel. After around 10 time steps, the UVQPE approximation has a very high overlap with the excited state marked with a green square. With a small value of δ , the algorithm eventually finds the true ground state, marked with a yellow circle, but with a larger δ it converges to the excited state instead.
- In Fig. 11, we show the same calculations as in




Figure 11: We perform the same numerical experiment on the 12-spin plaquette as in Fig. 9, but using the three-CZ initial state from the right of Eq. (44). This state has a much larger overlap with the ground subspace, 0.016 rather than 0.001. As a result, both algorithms converge to the ground state even when δ is large.




Figure 12: We perform the noisy version of the numerical experiment on the 8-spin plaquette from Fig. 8. Each expectation value from Eq. (4) is computed by sampling the states (9), (11), and (14) to find F_1 , F_2 , and F_3 ; We use 10^3 shots in total per expectation value, distributed among F_1 , F_2 , and F_3 . We perform the sampling 100 times at each time step, and we show here the mean energy estimate from UVQPE and ODM across the 100 shot noise realizations. Error bars show the standard deviation across noise realizations. A larger δ is needed to filter out the noise so that the algorithms can converge.




Figure 13: For each singular value threshold for the numerical experiment on the 8-spin plaquette (Fig. 12), we consider a single noise realization and show the convergence of energy and eigenvector in UVQPE in the style of Fig. 10. By looking at a single noise realization, we clearly see that when the singular value threshold is too low relative to the noise level, in this case $\delta = 10^{-2}$, the algorithm fails to converge at all. In contrast, convergence occurs quickly for a sufficiently large noise threshold, and occurs after a long energy plateau in the case of intermediate noise filtering.

Fig. 9, but using the initial state with three CZ operators as shown on the right of Eq. (44), which has a higher overlap with the ground state, around 0.016. In this case, the energy converges for all three levels of noise filtering, just as for the 8-spin plaquette. (We also show the convergence of the eigenvector for this initial state, analogous to Fig. 10, in the upper left panel of Fig. 21.)

- In Fig. 12, we show the convergence of the energy for both UVQPE and ODMD for the 8-spin plaquette with shot noise taken into account. We use 1000 shots in total for each expectation value, distributed with 40% to the circuit for F_1 and 30% each to the circuits for F_2 and F_3 . We run the sampling 100 times at each time step, and we view the n th experiment, across all time steps, as one noise realization. For each noise realization we compute all the expectation values and then the UVQPE and ODMD energy estimates as a function of time step. In the figure, the plotted value of the energy at each time step is the mean of the energies at that time step across all noise realizations, and the corresponding error bar is the standard deviation, indicating the range of energy estimates consistent with this level of shot noise. We use three singular value thresholds, $\delta = 10^{-1}$, $10^{-3/2}$, and 10^{-2} . The largest δ is above the scale of the shot noise, and both algorithms converge. The smallest δ is below the noise level, so the energy estimates do not converge. With an intermediate level of noise filtering, the energy starts to converge after a long plateau.

One question looking at the right panel of Fig. 12 is whether the algorithms fail to converge at all, or instead converge to some excited state rather than the ground state. The apparent convergence of the energy to a fixed value above that of the ground state is in fact a result of averaging over high-energy states and over 100 noise realizations. To show this, in Fig. 13 we plot the convergence in UVQPE of both energy and eigenvector, in the style of Fig. 10, for one representative shot noise realization for each of the three singular value thresholds. The convergence is clear for $\delta = 10^{-1}$. For the too-small value $\delta = 10^{-2}$, the estimated eigenvector continues to have large components of multiple high-energy eigenstates even at long times. In the case of intermediate δ , the estimated eigenvector fluctuates for a significant time before suddenly converging. (In some cases the convergence occurs through a series of such jumps after long plateaus in the energy estimate.) Thus the appearance in Fig. 12 of a long plateau followed by slow convergence apparently represents an average over many

trajectories in which the sudden convergence appears at different numbers of Krylov steps used.

- In Fig. 14, we show the same calculations but for the 12-spin plaquette. We use the higher-overlap initial state with three CZ operators; however, because the overlap is still smaller than for the 8-spin plaquette, we use 10^4 total shots per expectation value to reduce the noise level. The energy converges with $\delta = 0.1$. However, it converges to a low-lying excited state for $\delta = 0.3$, which is too large relative to the initial overlap, and it fails to converge at all for $\delta = 0.01$, which is too small relative to the shot noise.

To summarize, we find that both algorithms rapidly converge to the ground state energy even in the presence of (normally distributed) statistical noise, so long as we use a noise-filtering technique to regularize the linear algebra problems of Eqs. (3) and (7).

The appropriate level of noise filtering is a balancing act: if the threshold is below the true noise level (e.g. Fig. 12, right panel), the noise has too strong an effect and the algorithm doesn't converge; on the other hand, if the threshold is larger than the overlap between the initial state and the ground state (e.g. Fig. 9, left panel), the desired signal is also filtered out and the algorithm may converge to a low-lying excited state instead of to the ground state. However, as long as the initial state overlap is larger than the noise level, there will be a filtering threshold that allows both UVQPE and ODMD to converge.

5.2 Noisy emulator and hardware results

Finally, we turn to a demonstration of UVQPE and ODMD on a real quantum device, namely the Quantinuum H1-1 processor. We also use the corresponding noisy classical emulator provided by Quantinuum. By running on real hardware and on a noisy emulator, we show that the algorithm is robust not only to shot noise, as demonstrated above, but also to a variety of other coherent and incoherent noise sources such as state preparation and measurement errors, amplitude damping, and dephasing⁹.

⁹The quantity extracted from each circuit in our algorithms is just the probability of $|0\rangle$ in each of the states $|0(t)\rangle$, $|0_R(t)\rangle$, and $|0_{R_i}(t)\rangle$. These probabilities are affected by dephasing, which can be converted into bit flip errors both during time evolution and when applying the inverse state preparation unitaries $U_0^\dagger$, $U_R^\dagger$, and $U_{R_i}^\dagger$. Note that dephasing that occurs near the end of the circuit will not impact the final measurement, so in some cases dephasing will be less impactful than other noise sources. On the other hand, dephasing just before a layer of R_Z gates used for twirling will not be reduced, so such errors could be more problematic than expected despite error mitigation.




Figure 14: We perform the noisy version of the numerical experiment on the 12-spin plaquette from Fig. 11, using the same approach as in Fig. 12 with 10^4 shots per expectation value. We observe convergence for a range of δ around 0.1—large enough to filter out the shot noise, but small enough not to filter out the signal from the overlap between the initial state and the ground state.

As we saw in Fig. 14 above, for the 12-spin plaquette there is only a narrow range of noise-filtering thresholds δ that give robust convergence, even when we use 10^4 shots with no other source of noise. Given the high per-shot cost on the Quantinuum device, we therefore perform the hardware demonstration only on the 8-spin plaquette, where we can use 10^3 shots per expectation value. We use the initial state from Sec. 4.3 and a time step size of $0.1(J/2)$, the same as in the classical simulations reported in Fig. 12.

Unlike the noiseless simulations, where we computed the states (9), (11), and (14) exactly, here we are limited by (real or emulated) hardware noise to shorter circuits and thus make the $U_t \rightarrow F_t$ approximation discussed in Sec. 4.2, so that our longest circuits have at most 78 CNOT gates. For running on the Quantinuum hardware and emulator, we transpile our circuit to the correct gate set using BQSKit [70, 97].

Using this approximation, we compute the expectation values $\langle F_{k\Delta t} \rangle$ up to 20 time steps using the noisy classical emulator, which as we show in Fig. 15 is enough to converge the ground state energy to an accuracy of $10^{-2}(J/2)$. We run on hardware for five of the time steps, with spacing $2\Delta t = 0.2(J/2)$ through $10\Delta t = 1.0(J/2)$. Because the spacing $2\Delta t$ is still small enough to satisfy the bound $\Delta t < 2\pi/(3JN_\Delta)$, we can also estimate the ground state energy purely from the hardware results, with an accuracy of $10^{-1}(J/2)$ as shown in Fig. 16.

For both the noisy emulator and the hardware implementation, we use a total of 10^3 shots per expectation value: 400 for F_1 and 300 each for F_1 and F_2 . With the classical emulator, we also run half as many shots (200 for F_1 , 150 each for F_2 and F_3) with a layer of R_Z gates inserted as described in Sec. 4.4; replacing (a randomly selected) half of the original measurements with the R_Z -inserted, or “twirled” measurements slightly reduces the error in the expectation values as shown in Fig. 17. We




Figure 15: Convergence of energy estimate from UVQPE and ODMD for the 8-spin plaquette using expectation values measured with a noisy emulator that approximates the behavior of the Quantinuum quantum processor. Color corresponds to the algorithm used. Convex points (circle, square) show the results when no error mitigation is applied, while non-convex points ($\times$, $+$) show the results with two forms of symmetry-based error mitigation: symmetry sector post-selection and symmetry-based twirling. The effects of error mitigation are small, likely because shot noise is the dominant noise source (see Fig. 17).




Figure 16: Convergence of energy estimate from UVQPE and ODMD for the 8-spin plaquette using expectation values measured on the Quantinuum H1-1 processor, i.e. on quantum hardware. Even when measuring just five expectation values of the form $\langle F_{k\Delta t} \rangle$, we can obtain the correct ground state energy up to an error on the order of $10^{-1}(J/2)$. As shown in Fig. 17, with the number of shots used in these experiments, expectation values computed on hardware are statistically indistinguishable from those computed using the noisy emulator, so with more time steps the energy would further converge as in Fig. 15.

also show, for both noisy emulator and hardware, the effect of symmetry sector-based post selection.

Figs. 15 and 16 show the energy estimates both with and without the symmetry-based error mitigation. The effect on the convergence appears to be minimal, which we attribute to the fact that shot noise is actually the dominant noise source. This can be seen in Fig. 17 where, alongside the actual error in expectation values from noisy emulator and hardware, we also show an estimate of shot noise given 1000 total shots per time step. In particular, for each time step we run an exact classical simulation (including the $U \rightarrow F$ approximation) 100 times, and we take the standard deviation among the 100 experiments as the approximate shot noise. The large shot noise explains why, for some expectation values, the error mitigation strategies actually increase rather than decrease the error.

For the longer circuits and greater number of shots that would be needed to study a much larger two-dimensional system, we expect that the relative importance of noise channels like amplitude damping and dephasing would increase. In the noise filtering step of classical post-processing, in order to allow for a singular value threshold δ that is simultaneously larger than the error in the matrix elements and smaller than the overlap of the initial state with the ground state, an early fault-tolerant quantum computer might be needed to reduce such noise channels to manageable levels. A thor-

ough investigation of the relative importance of each type of error would require much larger systems and many more shots, and thus is beyond the scope of the current paper.

6 Magnetization curve results

As previously noted, for an antiferromagnetic model such as the Heisenberg model considered in this paper, when there is no magnetic field ($h = 0$) Lieb's theorem [63] guarantees that the ground state will be in the $S^z = 0$ sector. In other words, there is no net magnetic moment. On the other hand, if the field h is much larger than the interaction strength J , the ground state will be the fully polarized state with all spins pointing along the direction of the field, giving the maximum possible magnetization. At intermediate field strengths, there will be some tendency for spins to align with field, giving rise to the magnetization curve: the mean of $\langle S^z \rangle$ over all spins, measured in the ground state, as a function of the applied field h .

For the kagome lattice Heisenberg model, the magnetization curve is believed to have a series of plateaus, where the magnetization remains constant as the applied field increases [98–105]. The plateaus correspond to various phases of matter that are stable to small changes in the applied field. Most plateaus likely correspond to partial magnetic ordering, but some could be new exotic spin liquids different from the one expected at $h = 0$. Quantum computers may help to conclusively determine the nature of these plateau states, and finding the magnetization curve itself is the first step.

Fortunately, for a Hamiltonian that conserves S^z , such as the Heisenberg model, we can find the magnetization curve simply by computing the $h = 0$ ground state energy within each S^z symmetry sector. The reason is that the field term in H is exactly the conserved quantity, total S^z . Therefore every eigenstate at $h = 0$ is also an eigenstate at finite h , with its change in energy just given by $-h \sum S^z$. In particular, the energies of all states in a given symmetry sector change in the same way as h is increased, so for each sector the $h = 0$ ground state remains the lowest in energy for any h . Furthermore, the energy eigenvalues themselves are easily computed from the $h = 0$ eigenvalues.

We show this computation of the magnetization curve using the exact eigenvalue spectra of the 8- and 12-spin plaquettes in Figs. 18 and 19, respectively. In each figure, the lines in the upper panel show how the energy of the ground state within each spin sector evolves with increasing h . At each crossover point, where a new spin sector becomes the ground state, the magnetization jumps, giving the magnetization curve shown in




Figure 17: Error in computing the expectation values $O = \langle F_{k\Delta t} \rangle$ at time steps $k = 1$ through 20, using the Quantinuum H1-1 processor and the corresponding noisy classical emulator for the 8-spin plaquette, with 1000 shots per expectation value. Letting $O \equiv r e^{i\theta}$, the left panel shows the error in the magnitude and the right panel the error in the angle. Emulator results are shown by convex symbols (circle, triangle, diamond), while hardware results at the five time steps indicated by vertical lines are denoted by non-convex markers ($\times$ and $+$). Color shows the number of symmetry-based error mitigation techniques applied: none for blue, one for green, and two for yellow. For comparison, a classical estimate of the shot noise is shown as a dashed gray line. Shot noise appears in this case to be the dominant noise source.

the lower panel.

Here we used the exact eigenvalues to illustrate the approach to finding the magnetization curve. However, our goal is to show that this calculation is also feasible using UVQPE and/or ODMD. Once we have the ground state energy in each S^z symmetry sector at $h = 0$, the computation of the magnetization curve is identical to the exact calculation above. We therefore simply demonstrate finding the ground state energy in all symmetry sectors using UVQPE.

In Fig. 20, in the upper half of each panel, we show the convergence of ground state energy for the four distinct non-trivial S^z symmetry sector of the 8-spin plaquette. We use the initial states discussed in Sec. 4.3; each panel is labeled by the overlap between the initial state $|\psi_0\rangle$ and the ground state(s) $|v_0\rangle$. We run UVQPE using exact classical simulation, with no Trotter error or shot noise, and we use a singular value threshold of $\delta = 10^{-6}$. All symmetry sectors converge to the true ground state energy with high precision within 20 steps.

Also recall from Sec. 2.1 that UVQPE gives an approximation to the ground state eigenvector, $|\tilde{v}_0\rangle$, in the form of a linear combination of the basis states (1). For the small 8- and 12-spin plaquettes, we can explicitly generate the basis states and thus the specified linear combination, and we can find the overlap of $|\tilde{v}_0\rangle$ with the true ground state and other eigenvectors of the Hamiltonian. To help understand the convergence of the ground state energy from UVQPE, we also show in Fig. 20, in the lower half of each panel, the convergence of the approximate ground state eigenvector. Specifically, we plot the overlap of five eigenvectors of H with $|\tilde{v}_0\rangle$, observing how the overlap with the true ground state $|v_0\rangle$ grows as the number of time steps/Krylov ba-




Figure 18: Demonstration of how to derive the magnetization curve from $h = 0$ ground state energies, illustrated using the 8-spin plaquette. Top: Each color corresponds to one S^z symmetry sector: the solid line indicates the ground state within that sector, and the shaded region shows the full range of eigenvalues for the sector. The energies simply vary linearly in h starting with $h = 0$ on the left. The ground state switches from one magnetization sector to the next when the ground state lines cross at the locations indicated by dashed vertical lines. Bottom: We find the magnetization curve by plotting the magnetization of the symmetry sector with the lowest-energy ground state for each h . The value jumps at the crossing points because there are only a few possible values of magnetization, but the curve could be continuous for an infinite 2D system.




Figure 19: Same as Fig. 18 but for the 12-spin plaquette.

sis vectors increases, while the overlaps with the other eigenvectors decrease.

To gain additional insight into the relationship between the convergence of the approximations to the energy and to the eigenvector, it is helpful to consider worse initial states that have a much lower overlap with the true ground state. As we show for the $S^z = 1$ and $S^z = 2$ symmetry sectors in App. B, the algorithm will often converge towards a low-lying excited state first, leading to a plateau in the energy convergence, before ultimately reaching a good approximation to the ground state.

We also test out UVQPE for finding the ground state in each S^z sector, and hence the magnetization curve, for the 12-spin plaquette, with results shown in Fig. 21. As in the 8-spin case, the UVQPE simulations are carried out using exact time evolution and with no shot noise, using initial states from Sec. 4.3. Note that for $S^z = 0$ we use the initial state with three CZ operators, from the left of Eq. (44). Here the convergence behavior, which we can understand by looking at the eigenvector convergence in the lower half of each panel, shows some interesting features. In the $S^z = 0$ sector, the eigenvector initially has a significant component of an excited state, though less prominently than in the case of the six-CZ initial state as shown in Fig. 10 above. Although it looks like weight remains in two different eigenstates in the long-time limit, those are actually two orthogonal states both within the ground state subspace. Interestingly, convergence is slower for $S^z = 1$ even though the initial state overlap is much larger, which occurs because the initial state $|\psi_0\rangle$ has similarly high overlap with several low-lying excited states.

The most important takeaway message, however, is that for both plaquettes and in all symmetry sectors, UVQPE eventually converges to the correct ground state energy. As a result, the algorithm (likewise, ODMD) can be used to compute magnetization curves effectively following the prescription from Figs. 18 and 19.

7 Discussion

In this paper, we have proposed that hybrid quantum-classical algorithms based on real-time evolution provide a promising route to resolving open problems in frustrated magnetism and quantum spin liquid physics, including on the challenging kagome lattice Heisenberg model. We focused particularly on two recent algorithms, UVQPE [47] and ODMD [51], that use a series of expectation values of the time evolution operator, $\langle\psi_0|e^{-iHt}|\psi_0\rangle$ for $t = \Delta t, 2\Delta t, \dots$, to find an approximation to the ground state energy and give access to the ground state itself as a superposition of the states $e^{-ik\Delta t}|\psi_0\rangle$. The expectation values are measured on a quantum computer and are post-processed classically via a small generalized eigenvalue or linear least squares problem to find the ground state.

We have made three main contributions. First, we provided a compact and practical summary of UVQPE and ODMD in Sec. 2. We also give a detailed description of the mirror circuit method [57] for computing expectation values and wave function overlaps on a quantum computer. We show how to resolve a phase ambiguity in previous versions of the mirror circuit method, and we perform a careful numerical analysis of how to distribute a measurement budget among the various required quantum circuits.

Second, we outlined a detailed approach to running UVQPE and ODMD on the 2D kagome lattice Heisenberg model. We specifically highlighted the uses of S^z spin symmetry, namely to enable the mirror circuit method, to reduce the effective Hilbert space size, and to perform error mitigation. We then described an efficient “triangle-by-triangle” implementation of Trotterized time evolution, and we showed how to choose and prepare a good initial state $|\psi_0\rangle$ in which to calculate the expectation values $\langle\psi_0|e^{-iHt}|\psi_0\rangle$. Finally, we argued that the cost of the algorithms does not scale too quickly with system size: the scaling is exponential due to translation invariance of the ground state, but the exponential has a very small base.

Third, we provided an empirical demonstration of UVQPE and ODMD applied to a single 12-spin star plaquette of the kagome lattice and an analogous 8-spin plaquette. Remarkably, these apparently frustrated small systems are actually “frustration-free” with an exact classically solvable ground state. We make use of this exact solvability to design good initial states and reduce the required circuit depth for Trotter evolution. For the 8-spin plaquette this allows a compact and efficient simulation on both quantum hardware, namely the Quantinuum H1-1 processor, and a corresponding noisy classical emulator. We find that even in the presence of both shot noise and (real or emulated) hardware noise,




Figure 20: Convergence of ground state energy and corresponding approximation to the ground state eigenvector in the different spin sectors of the 8-spin star plaquette. Each panel corresponds to one S^z symmetry sector, and is also labeled with the overlap between the initial state $|\psi_0\rangle$ and the true ground state in that sector, $|v_0\rangle$. We measure the convergence of the UVQPE approximation to the ground state by finding its overlap with each eigenstate of H , and we plot the overlaps for the five eigenstates that contribute most to the approximation to the ground state across the simulation. The energies of those eigenstates are shown at the right of each panel along with the full spectrum in the corresponding S^z symmetry sector. Energy error is measured in units of $J/2$.




Figure 21: Here we show the same analysis as in Fig. 20, but for the 12-spin plaquette. Note that convergence is slower both when the initial overlap is smaller and when the initial state has significant overlap with several eigenstates that all have relatively low energy.

both algorithms converge rapidly to the true ground state energy.

Although the latter two points are specific to one model, UVQPE and ODMD are effective in a wide variety of systems. Even in models with fewer symmetries to exploit, the basic algorithms of UVQPE and ODMD are still effective, for example by using the Hadamard test to find matrix elements, rather than using the mirror circuit approach.

Conclusively determining the true ground state of the kagome lattice Heisenberg model and other models of frustrated magnetism will require more qubits than are available in present-day quantum processors, as well as lower noise levels. However, we have shown, with theoretical analysis of 2D systems and practical demonstrations on small systems, that hybrid algorithms based on real-time evolution such as UVQPE and ODMD are strong candidates for solving these important problems in condensed matter and materials physics. Even before the advent of fully fault-tolerant error-corrected quantum computers, these algorithms may allow for new physical insights that have not been possible via classical computation.

Acknowledgments

We thank Katie Klymko, Yizhi Shen, Roel Van Beumen, Daan Camps, and Siva Darbha for their insights into real-time evolution-based hybrid algorithms. The use of the dimer ground states was inspired partially by a conversation with Ehud Altman. We thank Yan Wang for discussions of the exact ground state of the star plaquette model and Mike Kolodrubetz for a discussion about the resulting Floquet evolution. Circuit diagrams were prepared using Quantikz [106]. This work was supported by the “Embedding QC into Many-body Frameworks for Strongly Correlated Molecular and Materials Systems” project, which is funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences (BES), the Division of Chemical Sciences, Geosciences, and Biosciences, and by the Office of Science, Office of Advanced Scientific Computing Research Accelerated Research for Quantum Computing Program of the U.S. Department of Energy. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. Our hardware demonstration used the Quantinuum H1-1 processor.

A Shot noise distribution for overlap from mirror circuits

As noted in the main text, Sec. 2.3, when computing expectation values using the mirror circuit method, there are two important choices we need to make. First, with a fixed measurement budget, we must decide how to allocate shots between the three circuits for measuring F_1 , F_2 , and F_3 . Second, we can choose whether to compute the magnitude of the expectation value using Eq. 16 or instead to use $\sqrt{F_1}$.

We empirically investigate both of these questions using the expectation values $\langle \psi_0 | e^{-iHt} | \psi_0 \rangle$ for the 8-spin plaquette with the $S^z = 0$ initial state from Sec. 4.3, with 10 different times t . For each time, we compute the states from which F_1 , F_2 , and F_3 are sampled. We then run 100 experiments for each t , where in each one we sample from these states m times to measure the three quantities and hence compute the overlap, O_m . We estimate the typical error due to shot noise as the standard deviation of the error in the overlap, $|O - O_m|$, over all experiments and all t for a given number of shots m . (Note that O and O_m are complex, and this error takes into account the errors in both magnitude and phase.) We use $m = 10^2$, 10^3 , and 10^4 .

As we show in Fig. 22, the most accurate results for the expectation values are found when the number of shots devoted to F_1 , F_2 , and F_3 are roughly equal. This is true regardless of which method is used to compute $|O|$. Furthermore, we find that the typical error is substantially lower when computing $|O|$ as $\sqrt{F_1}$ rather than with Eq. (16), which is reasonable since a convolution of probability distributions has a larger variance than the constituent distributions.

B UVQPE ground state eigenvector convergence with bad initial states

In Sec. 6, we saw that the convergence of ground state energy from UVQPE can be understood via the convergence of the approximation to the corresponding eigenvector. For the 8-spin plaquette, where the initial state in each S^z symmetry sector had a high overlap with the true ground state, both the energy and corresponding approximate eigenvector converged quickly. In contrast, on the 12-spin plaquette we observed more complex behavior, especially when the initial state had larger overlap with low-lying excited states than with the ground state. Here we show similar behavior on the smaller 8-spin plaquette, in the $S^z = 1$ and $S^z = 2$ sectors, by using initial states whose overlaps with the true ground states of those symmetry sectors are just 3% and 0.3%, respectively. We observe that the energy first decreases




Figure 22: Investigation of the optimal allocation of shots between the circuits for F_1 , F_2 , and F_3 used in the mirror circuit with reference state approach to computing matrix elements. For each data point, we fix a total number of shots to be divided among the three circuits, indicated by color in the figure. Error is measured as $|O - O_m|$, where O is the exact overlap and O_m the estimated one from the sampled circuit output. For each data point, we run the experiment for 100 different time steps and for each one perform the sampling with the specified number of shots 100 times; the “typical error” is the standard deviation of the error over all 10^4 shots. Error bars indicate the standard deviation of the typical error over batches of such experiments. In (a), we vary the fraction of shots assigned to the circuit for F_1 , with the F_2 and F_3 circuits evenly dividing the remaining measurement budget. Squares indicate that we have used only F_1 to compute the magnitude, while circles compute O_m using Eq. (16). Evidently the optimal choice is to recompute the magnitude using just F_1 , and furthermore to use approximately the same number of shots for each of the three circuits, indicated by the dashed vertical line. In (b), we fix the fraction of shots assigned to F_1 at 0.3, then vary the fraction assigned to F_2 with F_3 taking the remainder. The optimal choice is clearly to take an equal number of shots for F_2 and for F_3 .




Figure 23: Here we show the same analysis as in Fig. 20, but with initial states for the 8-spin plaquette that have much lower overlap with the ground states of their respective symmetry sectors. Convergence takes longer because UVQPE first populates multiple low-lying states, after which the energy estimate plateaus while weight shifts from low excited states to the ground state.

as UVQPE populates multiple low-lying states, plateaus for some time, then ultimately decreases again as the weight shifts towards only the overall ground state. If a larger threshold δ is used, the convergence takes much longer; for $S^z = 2$ with $\delta = 0.1$, UVQPE converges to the first excited state rather than to the ground state. Since a cutoff δ of this magnitude is needed in a noisy simulation, a very small overlap of less than a percent may be insufficient for UVQPE to be used in practice.

References

- [1] L. Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199 (2010).
- [2] L. Savary and L. Balents, Quantum spin liquids: a review, *Reports on Progress in Physics* **80**, 016502 (2016).
- [3] Y. Zhou, K. Kanoda, and T.-K. Ng, Quantum spin liquid states, *Rev. Mod. Phys.* **89**, 025003 (2017).
- [4] P. Anderson, Resonating valence bonds: A new kind of insulator?, *Materials Research Bulletin* **8**, 153 (1973).
- [5] D. A. Huse and V. Elser, Simple Variational Wave Functions for Two-Dimensional Heisenberg Spin- $\frac{1}{2}$ Antiferromagnets, *Phys. Rev. Lett.* **60**, 2531 (1988).
- [6] S. R. White and A. L. Chernyshev, Néel Order in Square and Triangular Lattice Heisenberg Models, *Phys. Rev. Lett.* **99**, 127004 (2007).
- [7] R. Kaneko, S. Morita, and M. Imada, Gapless Spin-Liquid Phase in an Extended Spin 1/2 Triangular Heisenberg Model, *Journal of the Physical Society of Japan* **83**, 093707 (2014).
- [8] W.-J. Hu, S.-S. Gong, W. Zhu, and D. N. Sheng, Competing spin-liquid states in the spin- $\frac{1}{2}$ Heisenberg model on the triangular lattice, *Phys. Rev. B* **92**, 140403 (2015).
- [9] Z. Zhu and S. R. White, Spin liquid phase of the $S = \frac{1}{2}$ $J_1 - J_2$ Heisenberg model on the triangular lattice, *Phys. Rev. B* **92**, 041105 (2015).
- [10] Y. Iqbal, W.-J. Hu, R. Thomale, D. Poilblanc, and F. Becca, Spin liquid nature in the Heisenberg $J_1 - J_2$ triangular antiferromagnet, *Phys. Rev. B* **93**, 144411 (2016).
- [11] S. Hu, W. Zhu, S. Eggert, and Y.-C. He, Dirac Spin Liquid on the Spin-1/2 Triangular Heisenberg Antiferromagnet, *Phys. Rev. Lett.* **123**, 207203 (2019).
- [12] O. I. Motrunich, Variational study of triangular lattice spin-1/2 model with ring exchanges and spin liquid state in κ -(ET)₂Cu₂(CN)₃, *Phys. Rev. B* **72**, 045105 (2005).
- [13] H.-Y. Yang, A. M. Läuchli, F. Mila, and K. P. Schmidt, Effective Spin Model for the Spin-Liquid Phase of the Hubbard Model on the Triangular Lattice, *Phys. Rev. Lett.* **105**, 267204 (2010).
- [14] R. V. Mishmash, J. R. Garrison, S. Bieri, and C. Xu, Theory of a Competitive Spin Liquid State for Weak Mott Insulators on the Triangular Lattice, *Phys. Rev. Lett.* **111**, 157203 (2013).
- [15] T. Shirakawa, T. Tohyama, J. Kokalj, S. Sota, and S. Yunoki, Ground-state phase diagram of the triangular lattice Hubbard model by the density-matrix renormalization group method, *Phys. Rev. B* **96**, 205130 (2017).
- [16] A. Szasz, J. Motruk, M. P. Zaletel, and J. E. Moore, Chiral Spin Liquid Phase of the Triangular Lattice Hubbard Model: A Density Matrix Renormalization Group Study, *Phys. Rev. X* **10**, 021042 (2020).
- [17] L. F. Tocchio, A. Montorsi, and F. Becca, Magnetic and spin-liquid phases in the frustrated $t-t'$ Hubbard model on the triangular lattice, *Phys. Rev. B* **102**, 115150 (2020).
- [18] A. Szasz and J. Motruk, Phase diagram of the anisotropic triangular lattice Hubbard model, *Phys. Rev. B* **103**, 235132 (2021).
- [19] L. F. Tocchio, A. Montorsi, and F. Becca, Hubbard model on triangular N -leg cylinders: Chiral and nonchiral spin liquids, *Phys. Rev. Res.* **3**, 043082 (2021).
- [20] B.-B. Chen, Z. Chen, S.-S. Gong, D. N. Sheng, W. Li, and A. Weichselbaum, Quantum spin liquid with emergent chiral order in the triangular-lattice Hubbard model, *Phys. Rev. B* **106**, 094420 (2022).
- [21] 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).
- [22] 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).
- [23] H.-C. Jiang, Z. Wang, and L. Balents, Identifying topological order by entanglement entropy, *Nature Physics* **8**, 902 (2012).
- [24] 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 (2013).
- [25] Y. Iqbal, D. Poilblanc, and F. Becca, Vanishing spin gap in a competing spin-liquid phase in the kagome Heisenberg antiferromagnet, *Phys. Rev. B* **89**, 020407 (2014).

- [26] 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).
- [27] Y.-C. He, M. P. Zaletel, M. Oshikawa, and F. Pollmann, Signatures of Dirac Cones in a DMRG Study of the Kagome Heisenberg Model, *Phys. Rev. X* **7**, 031020 (2017).
- [28] A. M. Läuchli, J. Sudan, and R. Moessner, $S = \frac{1}{2}$ kagome Heisenberg antiferromagnet revisited, *Phys. Rev. B* **100**, 155142 (2019).
- [29] J. Motruk, D. Rossi, D. A. Abanin, and L. Rademaker, Kagome chiral spin liquid in transition metal dichalcogenide moiré bilayers, *Phys. Rev. Res.* **5**, L022049 (2023).
- [30] X. Zhou, W.-S. Lee, M. Imada, N. Trivedi, P. Phillips, H.-Y. Kee, P. Torma, and M. Eremets, High-temperature superconductivity, *Nature Reviews Physics* **3**, 462 (2021).
- [31] A. Y. Kitaev, Quantum measurements and the Abelian Stabilizer Problem, arXiv (1995), [arXiv:quant-ph/9511026 \[quant-ph\]](#).
- [32] D. S. Abrams and S. Lloyd, Simulation of Many-Body Fermi Systems on a Universal Quantum Computer, *Phys. Rev. Lett.* **79**, 2586 (1997).
- [33] D. S. Abrams and S. Lloyd, Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors, *Phys. Rev. Lett.* **83**, 5162 (1999).
- [34] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [35] 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).
- [36] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martínez, Quantum Computation of Electronic Transitions Using a Variational Quantum Eigensolver, *Phys. Rev. Lett.* **122**, 230401 (2019).
- [37] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [38] 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).
- [39] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nat. Commun.* **12**, 6961 (2021).
- [40] Z. Holmes, A. Arrasmith, B. Yan, P. J. Coles, A. Albrecht, and A. T. Sornborger, Barren Plateaus Preclude Learning Scramblers, *Phys. Rev. Lett.* **126**, 190501 (2021).
- [41] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, A Review of Barren Plateaus in Variational Quantum Computing, arXiv (2024), [arXiv:2405.00781 \[quant-ph\]](#).
- [42] 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).
- [43] R. Wiersema, E. K okc u, A. F. Kemper, and B. N. Bakalov, Classification of dynamical Lie algebras of 2-local spin systems on linear, circular and fully connected topologies, *npj Quantum Information* **10**, 110 (2024).
- [44] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, E. R. Anschuetz, and Z. Holmes, Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing, arXiv (2024), [arXiv:2312.09121 \[quant-ph\]](#).
- [45] H. Mhiri, R. Puig, S. Lerch, M. S. Rudolph, T. Chotibut, S. Thanasilp, and Z. Holmes, A unifying account of warm start guarantees for patches of quantum landscapes, arXiv (2025), [arXiv:2502.07889 \[quant-ph\]](#).
- [46] N. H. Stair, R. Huang, and F. A. Evangelista, A multireference quantum Krylov algorithm for strongly correlated electrons, *J. Chem. Theory Comput.* **16**, 2236 (2020).
- [47] K. Klymko, C. Mejuto-Zaera, S. J. Cotton, F. Wudarski, M. Urbanek, D. Hait, M. Head-Gordon, K. B. Whaley, J. Moussa, N. Wiebe, W. A. de Jong, and N. M. Tubman, Real-Time Evolution for Ultracompact Hamiltonian Eigenstates on Quantum Hardware, *PRX Quantum* **3**, 020323 (2022).
- [48] Z. Ding and L. Lin, Even Shorter Quantum Circuit for Phase Estimation on Early Fault-Tolerant Quantum Computers with Applications to Ground-State Energy Estimation, *PRX Quantum* **4**, 020331 (2023).
- [49] Z. Ding and L. Lin, Simultaneous estimation of multiple eigenvalues with short-depth quantum circuit on early fault-tolerant quantum computers, *Quantum* **7**, 1136 (2023).
- [50] Y. Shen, K. Klymko, J. Sud, D. B. Williams-

- Young, W. A. d. Jong, and N. M. Tubman, Real-Time Krylov Theory for Quantum Computing Algorithms, *Quantum* **7**, 1066 (2023).
- [51] Y. Shen, D. Camps, A. Szasz, S. Darbha, K. Klymko, D. B. Williams-Young, N. M. Tubman, and R. V. Beeumen, Estimating Eigenenergies from Quantum Dynamics: A Unified Noise-Resilient Measurement-Driven Approach, arXiv (2023), [arXiv:2306.01858 \[quant-ph\]](#).
- [52] M. Motta, W. Kirby, I. Liepuoniute, K. J. Sung, J. Cohn, A. Mezzacapo, K. Klymko, N. Nguyen, N. Yoshioka, and J. E. Rice, Subspace methods for electronic structure simulations on quantum computers, arXiv (2023), [arXiv:2312.00178 \[quant-ph\]](#).
- [53] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, X. Cao, J. Carrasquilla, F. Ferrari, A. Georges, M. Hibat-Allah, M. Imada, A. M. Läuchli, G. Mazzola, A. Mezzacapo, A. Millis, J. R. Moreno, T. Neupert, Y. Nomura, J. Nys, O. Parcollet, R. Pohle, I. Romero, M. Schmid, J. M. Silvester, S. Sorella, L. F. Tocchio, L. Wang, S. R. White, A. Wietek, Q. Yang, Y. Yang, S. Zhang, and G. Carleo, Variational benchmarks for quantum many-body problems, *Science* **386**, 296 (2024).
- [54] J. L. Bosse and A. Montanaro, Probing ground-state properties of the kagome antiferromagnetic Heisenberg model using the variational quantum eigensolver, *Phys. Rev. B* **105**, 094409 (2022).
- [55] J. Kattemölle and J. van Wezel, Variational quantum eigensolver for the Heisenberg antiferromagnet on the kagome lattice, *Phys. Rev. B* **106**, 214429 (2022).
- [56] Y. Wang, Prepare Ground States of Highly Frustrated Magnetic Clusters on Quantum Computers, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 02 (2023) pp. 397–398.
- [57] C. L. Cortes and S. K. Gray, Quantum Krylov subspace algorithms for ground- and excited-state energy estimation, *Phys. Rev. A* **105**, 022417 (2022).
- [58] M. B. Hastings, An area law for one-dimensional quantum systems, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P08024 (2007).
- [59] I. Arad, A. Kitaev, Z. Landau, and U. Vazirani, An area law and sub-exponential algorithm for 1D systems, arXiv (2013), [arXiv:1301.1162 \[quant-ph\]](#).
- [60] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proc. Roy. Soc. Lond. A* **454**, 339 (1998).
- [61] E. Kökcü, D. Camps, L. B. Otfelie, W. A. de Jong, R. V. Beeumen, and A. F. Kemper, Algebraic Compression of Free Fermionic Quantum Circuits: Particle Creation, Arbitrary Lattices and Controlled Evolution, arXiv (2023), [arXiv:2303.09538 \[quant-ph\]](#).
- [62] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [63] E. H. Lieb, Two theorems on the Hubbard model, *Phys. Rev. Lett.* **62**, 1201 (1989).
- [64] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O’Brien, Low-cost error mitigation by symmetry verification, *Phys. Rev. A* **98**, 062339 (2018).
- [65] Z. Cai, Quantum Error Mitigation using Symmetry Expansion, *Quantum* **5**, 548 (2021).
- [66] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, Faster Digital Quantum Simulation by Symmetry Protection, *PRX Quantum* **2**, 010323 (2021).
- [67] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, *Journal of Mathematical Physics* **32**, 400 (1991).
- [68] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter Error with Commutator Scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [69] J. Kattemölle and S. Hariharan, Line-graph qubit routing: from kagome to heavy-hex and more, arXiv (2023), [arXiv:2306.05385 \[quant-ph\]](#).
- [70] E. Younis, C. C. Iancu, W. Lavrijsen, M. Davis, and E. Smith, Berkeley Quantum Synthesis Toolkit (BQSKit) v1 (2021), <https://www.osti.gov/biblio/1785933>.
- [71] C. L. Henley, Effective Hamiltonians and dilution effects in Kagome and related anti-ferromagnets, *Canadian Journal of Physics* **79**, 1307 (2001).
- [72] A. Szasz, E. Younis, and W. De Jong, Numerical Circuit Synthesis and Compilation for Multi-State Preparation, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 01 (2023) pp. 768–778.
- [73] OEIS Foundation Inc., The On-Line Encyclopedia of Integer Sequences (2023), published electronically at <http://oeis.org/A000108>.
- [74] A. Carbone, D. E. Galli, M. Motta, and B. Jones, Quantum Circuits for the Preparation of Spin Eigenfunctions on Quantum Computers, *Symmetry* **14**, 624 (2022).
- [75] B. Sriram Shastry and B. Sutherland, Exact ground state of a quantum mechanical antiferromagnet, *Physica B+C* **108**, 1069 (1981).
- [76] D. S. Rokhsar and S. A. Kivelson, Superconduct-

- tivity and the Quantum Hard-Core Dimer Gas, *Phys. Rev. Lett.* **61**, 2376 (1988).
- [77] R. Moessner and K. S. Raman, Quantum Dimer Models, in *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, edited by C. Lacroix, P. Mendels, and F. Mila (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 437–479.
- [78] L. Capriotti, F. Becca, A. Parola, and S. Sorella, Resonating Valence Bond Wave Functions for Strongly Frustrated Spin Systems, *Phys. Rev. Lett.* **87**, 097201 (2001).
- [79] G. Baskaran, Resonating Valence Bond States in 2 and 3D: Brief History and Recent Examples, arXiv (2006), [arXiv:cond-mat/0611553](https://arxiv.org/abs/cond-mat/0611553) [cond-mat.str-el] .
- [80] H. J. Briegel and R. Raussendorf, Persistent Entanglement in Arrays of Interacting Particles, *Phys. Rev. Lett.* **86**, 910 (2001).
- [81] R. Siddharthan and A. Georges, Square kagome quantum antiferromagnet and the eight-vertex model, *Phys. Rev. B* **65**, 014417 (2001).
- [82] F. Monti and A. Sütő, Spin-12 Heisenberg model on Δ trees, *Physics Letters A* **156**, 197 (1991).
- [83] S. Miyahara, Exact Results in Frustrated Quantum Magnetism, in *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, edited by C. Lacroix, P. Mendels, and F. Mila (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 513–536.
- [84] S. Bravyi and B. Terhal, Complexity of Stoquastic Frustration-Free Hamiltonians, *SIAM Journal on Computing* **39**, 1462 (2010).
- [85] M. G. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, and C. Iancu, Towards Optimal Topology Aware Quantum Circuit Synthesis, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2020) pp. 223–234.
- [86] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Quantum Circuits for General Multiqubit Gates, *Phys. Rev. Lett.* **93**, 130502 (2004).
- [87] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of Quantum Logic Circuits, in *Proceedings of the 2005 Asia and South Pacific Design Automation Conference*, ASP-DAC '05 (Association for Computing Machinery, New York, NY, USA, 2005) p. 272–275.
- [88] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [89] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [90] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature* **567**, 491 (2019).
- [91] A. He, B. Nachman, W. A. de Jong, and C. W. Bauer, Zero-noise extrapolation for quantum-gate error mitigation with identity insertions, *Phys. Rev. A* **102**, 012426 (2020).
- [92] M. Krebsbach, B. Trauzettel, and A. Calzona, Optimization of Richardson extrapolation for quantum error mitigation, *Phys. Rev. A* **106**, 062436 (2022).
- [93] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
- [94] L. F. Richardson and J. A. Gaunt, VIII. The deferred approach to the limit, *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **226**, 299 (1927).
- [95] L. Viola, E. Knill, and S. Lloyd, Dynamical Decoupling of Open Quantum Systems, *Phys. Rev. Lett.* **82**, 2417 (1999).
- [96] N. Ezzell, B. Pokharel, L. Tewala, G. Quiroz, and D. A. Lidar, Dynamical decoupling for superconducting qubits: A performance survey, *Phys. Rev. Appl.* **20**, 064027 (2023).
- [97] E. Younis and C. Iancu, Quantum Circuit Optimization and Transpilation via Parameterized Circuit Instantiation, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2022) pp. 465–475.
- [98] K. Hida, Magnetization Process of the $S=1$ and $1/2$ Uniform and Distorted Kagomé Heisenberg Antiferromagnets, *Journal of the Physical Society of Japan* **70**, 3673 (2001).
- [99] J. Schulenburg, A. Honecker, J. Schnack, J. Richter, and H.-J. Schmidt, Macroscopic Magnetization Jumps due to Independent Magnons in Frustrated Quantum Spin Lattices, *Phys. Rev. Lett.* **88**, 167207 (2002).
- [100] A. Honecker, J. Schulenburg, and J. Richter, Magnetization plateaus in frustrated antiferromagnetic quantum spin models, *Journal of Physics: Condensed Matter* **16**, S749 (2004).
- [101] T. Shimokawa and H. Nakano, Magnetization Curve of the Kagome-Strip-Lattice Antiferromagnet, *Journal of Low Temperature Physics* **170**, 328 (2013).
- [102] S. Nishimoto, N. Shibata, and C. Hotta, Controlling frustrated liquids and solids with an applied

- field in a kagome Heisenberg antiferromagnet, [Nature Communications](#) **4**, 2287 (2013).
- [103] T. Nakamura, Machine learning as an improved estimator for magnetization curve and spin gap, [Scientific Reports](#) **10**, 14201 (2020).
- [104] H. Schlüter, J. Richter, and J. Schnack, Melting of Magnetization Plateaus for Kagomé and Square-Kagomé Lattice Antiferromagnets, [Journal of the Physical Society of Japan](#) **91**, 094711 (2022).
- [105] X.-Y. Liu, Y. Gao, H. Li, W. Jin, J. Xiang, H. Jin, Z. Chen, W. Li, and G. Su, Quantum spin liquid candidate as superior refrigerant in cascade demagnetization cooling, [Communications Physics](#) **5**, 233 (2022).
- [106] A. Kay, Tutorial on the Quantikz Package, arXiv (2023), [arXiv:1809.03842 \[quant-ph\]](#) .