

What symmetry buys on hardware: compressing an $SU(2)$ -equivariant ansatz for the kagome Heisenberg antiferromagnet

Besart Lajci¹, Oleksa Hryniv², Maciej Koch-Janusz², and Richard D.P. East²

¹Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada

²Haiqu, San Francisco, California, USA

Abstract

The ground state of the kagome Heisenberg antiferromagnet, a candidate quantum spin liquid, is a demanding target for near-term hardware. The lattice is frustrated, and a generic variational circuit explores far more of Hilbert space than the physics requires. The $SU(2)$ -equivariant spin-network ansatz of East et al. [1] is built from $SU(2)$ -invariant gates, so it stays in one total-spin sector, and on the small plaquettes where every triangle can host a singlet dimer (8 and 12 spins) it prepares the ground state exactly in noiseless simulation. Here we ask what the symmetry buys on a real device, beyond the trainability arguments of geometric quantum machine learning [4,8]. We find it buys three things. First, the trained circuit can be compressed. An $SU(2)$ -aware compression [3] only has to reproduce each gate on the singlet sector, and it cuts the 8-spin circuit from 172 two-qubit gates to 46. In noise sweeps the compressed circuit wins above a two-qubit error an order of magnitude below today's rates, and a single mitigated run on `ibm_torino` gives an energy error of 0.47 (exact ground energy -12), in line with a device-faithful simulation of the chip (0.53). Second, the symmetry filters noise at readout, because the target state has zero magnetisation along every axis. Third, excited states come at no extra cost. A different initial singlet covering places the ansatz in a different magnetisation sector, and the same circuit then reaches the lowest state of each S_z sector of the 12-spin star. Three noiseless layers reach the frustrated 18-spin plaquette to within 2.5% of the exact energy, though under device noise neither circuit is usable at that size. At today's error rates the compression is already worth its noiseless cost, and at 18 spins the limit is the device rather than the ansatz.

1 The equivariant ansatz and what we ask of it

A circuit built entirely from $SU(2)$ -invariant gates commutes with global spin rotations, so it never leaves the total-spin sector it starts in [9]. The variational search [7] then runs over one sector instead of the full Hilbert space. The spin-network ansatz we study realises this for the kagome Heisenberg model (East et al. [1]). That paper sets out the construction and its noiseless behaviour, and ties the reduced search space to the trainability arguments of geometric quantum machine learning [4,8]. Our question is what happens on a device. How accurate is the circuit under realistic gate noise, and how many two-qubit gates does it cost to run? Beyond the ansatz itself, the symmetry also gives a compression of the trained circuit, excited states from re-initialisation, and a route to a fully frustrated cell.

The model is the nearest-neighbour Heisenberg antiferromagnet $H = \sum_{\langle i,j \rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j$, written here with Pauli matrices, on kagome plaquettes. The ground state is a spin singlet [5]. Each triangle (Fig. 1, left) is handled in the Schur basis [6], where three qubits split into one spin- $\frac{3}{2}$ and two spin- $\frac{1}{2}$ copies. The parametrised rotation mixes the two spin- $\frac{1}{2}$ copies and is then undone:

$$\text{Spin3}(\vec{\theta}) = S_3 U_2(\vec{\theta}) S_3^\dagger, \quad \text{Spin2}(\phi) = S_2 P_2(\phi) S_2^\dagger.$$

The starting state is a covering of the lattice by singlet dimers, which lies in the $S = 0$ sector. For readout we measure in the X , Y and Z bases and drop shots of nonzero total magnetisation. A singlet has zero magnetisation along any axis, so the filter removes only noise.




Figure 1: Left: kagome lattice of corner-sharing triangles. The shaded triangle is the three-spin unit handled in the Schur basis, where Spin3 acts (Spin2 acts on a pair). Right: the Spin2 vertex gate, the two-qubit $SU(2)$ -invariant primitive of the ansatz. An inverse Bell transformation (CX, then H) sends the singlet to a computational basis state. A controlled phase $P(\phi)$ marks that state, and the transformation is then undone. The result $U_{SU(2)}(\phi) = \text{Spin2}(\phi)$ commutes with the total spin and phases the singlet alone.

2 Results

8-spin: symmetry-aware compression. The trained layer is built from three-qubit Spin3 gates and two-qubit Spin2 vertex gates (Fig. 1, right). In noiseless simulation it reaches the ground state to one part in 10^{13} , at a cost of 172 two-qubit gates. A generic transpiler [11] is blind to the symmetry and compiles each Spin3 as an arbitrary three-qubit unitary, spending about 40 two-qubit gates on each. The compression [3] exploits the fact that the ansatz only ever acts within the $SU(2)$ singlet sector. A trained Spin3 need only be reproduced *on that sector*, not as a full unitary, so we re-fit it as a short brickwork of the two-qubit vertex gate, with angles chosen to maximise the fidelity of the prepared state. Three of the four distinct Spin3 gates compress exactly. The fourth reaches fidelity 0.99 and leaves a fixed noiseless error of 0.041, the floor at which the compressed curve in Fig. 2 flattens. The layer shrinks to 46 two-qubit gates, about 70 after routing onto hardware connectivity. That is a 73% reduction.

8-spin: noise sweep. We study two simulated settings (Fig. 2). The first is an idealised all-to-all depolarising sweep, a stand-in for machines such as Quantinuum’s, with the one-qubit error fixed at 2×10^{-5} and readout error on every qubit. Here the compressed circuit overtakes the uncompressed one near a two-qubit error of 2×10^{-4} . At the 2.7×10^{-3} of today’s devices the errors are 0.12 compressed and 0.58 uncompressed. The second setting is a model of IBM Heron hardware. We use `ibm_marrakesh` (Heron r2) via Qiskit’s `FakeMarrakesh` model, keeping the chip’s connectivity and its median readout and one-qubit errors, and scaling only the two-qubit error. Routing inflates the uncompressed circuit from 172 to about 300 two-qubit gates. At the device’s current two-qubit error the compressed circuit gives 0.28 and the uncompressed 1.7. Even with perfect two-qubit gates the compressed circuit floors at 0.07 to 0.08, limited by readout.

8-spin: hardware run on `ibm_torino`. On this Heron r1 device we executed the compressed circuit once, with three measurement bases, 1000 shots each, the magnetisation filter, and the error mitigation of Haiqu’s execution pipeline. The uncompressed circuit was not run. The measured energy error is 0.47, several times the idealised prediction of about 0.12, but in line with an unmitigated, device-faithful `FakeTorino` simulation of the same circuit (0.53).

Comparison with real-time evolution. Ref. [2] reaches the same plaquettes with many short circuits rather than one trained circuit. Real-time evolution feeds a small classical linear-algebra problem over a Krylov basis of time-evolved states [10]. In noiseless statevector simulation both methods reach the exact energy to machine precision, but neither of those calculations runs on hardware. Theirs uses the exact evolution operator, and ours is the uncompressed circuit, about 300 two-qubit gates once routed. Looking to potential hardware suitable setups, their longest hardware circuit uses a single




Figure 2: 8-spin energy error against two-qubit gate error, compressed and uncompressed. Solid lines show an idealised all-to-all depolarising sweep (Quantinuum-like). The low-noise floor of the compressed curve is the 0.041 compression error. Dashed lines show a real Heron architecture (**FakeMarrakesh**, Heron r2, with its connectivity and median readout and one-qubit errors, and the two-qubit error scaled). Routing inflates the uncompressed circuit to ~ 300 gates, and the compressed one floors at 0.07 to 0.08, readout-limited. Dotted lines mark where Quantinuum H2 and **ibm_marrakesh** (Heron r2) sit today. The star is the measured **ibm_torino** run (Heron r1, compressed, mitigated). It need not lie on the **ibm_marrakesh** sweep.

Trotter step and totals at most 78 CNOTs on Quantinuum’s H1-1, and five measured time steps give a ground-state error of about 0.2 (adjusted to our units). Our compressed circuit routes to about 70 two-qubit gates, and its one run on **ibm_torino** gave 0.47, against the 0.041 compression floor. With one run for each method, on different machines, we read this as comparable circuit sizes and comparable accuracy, not a controlled benchmark.

12-spin: excited states by re-initialisation. The initial dimer covering fixes the magnetisation, and removing one singlet raises S_z by one. The same ansatz therefore reaches the lowest state of each S_z sector by re-optimising from a different covering. The circuit is grown layer by layer with warm starts. Accuracy is not monotonic in S_z (Table 1). It runs from 10^{-2} at $S_z = 1$ to 10^{-11} at $S_z = 5$. Ref. [2] treats the same 12-spin star by real-time evolution, with one exact-statevector Krylov calculation per sector. Reading their per-sector plots (see Fig. 21 of Ref. [2]), their errors in our units run from about 2×10^{-2} in the slowest, small-overlap sectors to about 2×10^{-9} at $S_z = 5$. Ours match theirs at $S_z = 1$ and are one to two orders lower elsewhere, though their curves are still falling where the plots end.

S_z	1	2	3	4	5
energy error	9.9×10^{-3}	5.4×10^{-4}	5.4×10^{-6}	6.4×10^{-5}	1.6×10^{-11}

Table 1: 12-spin star: noiseless energy error per magnetisation sector, against the exact sector minimum.

18-spin: A true frustrated cell. The 18-spin periodic plaquette is the smallest kagome cell that cannot be fully dimerised. Nine singlets cover its eighteen sites, but its twelve triangles cannot all host one, so no covering minimises every triangle at once. Here the frustration is complete and exact preparation is out of reach. Three noiseless layers get to -31.47 against an exact -32.26 , a 2.5% error. Ref. [1] finds the best converged energy on this cell keeps falling with depth, though the optimisation grows harder. Under noise the 8-spin picture repeats. In the idealised all-to-all model at two-qubit error 8.6×10^{-4} the compressed circuit reaches about 6% energy error against 12% uncompressed, trading a higher noiseless floor (4.8%) for $4.7\times$ fewer two-qubit gates (294 against 1374). Simulated on a realistic IBM Heron device instead, with the circuits routed onto its heavy-hex connectivity and its real error rates, neither circuit gives a usable energy. The errors are about 33% compressed and 77% uncompressed. Compression still halves the error but no longer decides success. For this circuit on today’s heavy-hex hardware the limit is the device, not the gate count.

References

- [1] R. D. P. East, G. Alonso-Linaje, and C.-Y. Park, “All you need is spin: SU(2) equivariant variational quantum circuits based on spin networks,” *Quantum Sci. Technol.* **11**, 025025 (2026), arXiv:2309.07250.
- [2] A. Szasz, E. Younis, and W. A. de Jong, “Ground state energy and magnetization curve of a frustrated magnetic system from real-time evolution on a digital quantum processor,” *Quantum* **9**, 1704 (2025), arXiv:2401.03015.
- [3] O. Hryniv, “Utilization of SU(2) Symmetry for Efficient Simulation of Quantum Systems,” arXiv:2506.19879 (2025).
- [4] M. Larocca et al., *PRX Quantum* **3**, 030341 (2022), arXiv:2205.02261; J. J. Meyer et al., *PRX Quantum* **4**, 010328 (2023), arXiv:2205.06217; Q. T. Nguyen et al., *PRX Quantum* **5**, 020328 (2024), arXiv:2210.08566.
- [5] M. R. Norman, “Colloquium: Herbertsmithite and the search for the quantum spin liquid,” *Rev. Mod. Phys.* **88**, 041002 (2016), arXiv:1604.03048.
- [6] D. Bacon, I. L. Chuang, and A. W. Harrow, “Efficient quantum circuits for Schur and Clebsch-Gordan transforms,” *Phys. Rev. Lett.* **97**, 170502 (2006), arXiv:quant-ph/0407082.
- [7] A. Peruzzo et al., “A variational eigenvalue solver on a photonic quantum processor,” *Nat. Commun.* **5**, 4213 (2014), arXiv:1304.3061; 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), arXiv:1509.04279.
- [8] 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), arXiv:1803.11173.
- [9] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, “Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm,” *npj Quantum Inf.* **6**, 10 (2020), arXiv:1904.10910.
- [10] N. H. Stair, R. Huang, and F. A. Evangelista, *J. Chem. Theory Comput.* **16**, 2236 (2020), arXiv:1911.05163; K. Klymko et al., *PRX Quantum* **3**, 020323 (2022), arXiv:2103.08563; M. Motta et al., *Nat. Phys.* **16**, 205 (2020), arXiv:1901.07653.
- [11] A. Javadi-Abhari et al., “Quantum computing with Qiskit,” arXiv:2405.08810 (2024).