

Computational supremacy in quantum simulation

Andrew D. King,^{1,*} Alberto Nocera,² Marek M. Rams,³ Jacek Dziarmaga,³ Roeland Wiersema,^{4,5} William Bernoudy,¹ Jack Raymond,¹ Nitin Kaushal,² Niclas Heinsdorf,^{2,6} Richard Harris,¹ Kelly Boothby,¹ Fabio Altomare,¹ Andrew J. Berkley,¹ Martin Boschnak,¹ Kevin Chern,¹ Holly Christiani,¹ Samantha Cibere,¹ Jake Connor,¹ Martin H. Dehn,¹ Rahul Deshpande,¹ Sara Ejtemaee,¹ Pau Farré,¹ Kelsey Hamer,¹ Emile Hoskinson,¹ Shuiyuan Huang,¹ Mark W. Johnson,¹ Samuel Kortas,¹ Eric Ladizinsky,¹ Tony Lai,¹ Trevor Lanting,¹ Ryan Li,¹ Allison J.R. MacDonald,¹ Gaelen Marsden,¹ Catherine C. McGeoch,¹ Reza Molavi,¹ Richard Neufeld,¹ Mana Norouzpour,¹ Travis Oh,¹ Joel Pasvolsky,¹ Patrick Poitras,¹ Gabriel Poulin-Lamarre,¹ Thomas Prescott,¹ Mauricio Reis,¹ Chris Rich,¹ Mohammad Samani,¹ Benjamin Sheldan,¹ Anatoly Smirnov,¹ Edward Sterpka,¹ Berta Trullas Clavera,¹ Nicholas Tsai,¹ Mark Volkmann,¹ Alexander Whitticar,¹ Jed D. Whittaker,¹ Warren Wilkinson,¹ Jason Yao,¹ T.J. Yi,¹ Anders W. Sandvik,⁷ Gonzalo Alvarez,⁸ Roger G. Melko,^{5,9} Juan Carrasquilla,^{4,5,10} Marcel Franz,² and Mohammad H. Amin^{1,11,†}

¹*D-Wave Quantum Inc., Burnaby, British Columbia, Canada*

²*Department of Physics and Astronomy and Quantum Matter Institute, University of British Columbia, Vancouver, British Columbia, Canada*

³*Jagiellonian University, Institute of Theoretical Physics, Łojasiewicza 11, PL-30348 Kraków, Poland*

⁴*Vector Institute, MaRS Centre, Toronto, Ontario, M5G 1M1, Canada*

⁵*Department of Physics and Astronomy, University of Waterloo, Ontario, N2L 3G1, Canada*

⁶*Max Planck Institute for Solid State Research, Stuttgart, 70569, Germany*

⁷*Department of Physics, Boston University, Boston, MA, USA.*

⁸*Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA*

⁹*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, N2L 2Y5, Canada*

¹⁰*Institute for Theoretical Physics, ETH Zürich, 8093, Switzerland*

¹¹*Department of Physics, Simon Fraser University, Burnaby, British Columbia, Canada*

(Dated: March 5, 2024)

Quantum computers hold the promise of solving certain problems that lie beyond the reach of conventional computers. Establishing this capability, especially for impactful and meaningful problems, remains a central challenge. One such problem is the simulation of nonequilibrium dynamics of a magnetic spin system quenched through a quantum phase transition. State-of-the-art classical simulations demand resources that grow exponentially with system size. Here we show that superconducting quantum annealing processors can rapidly generate samples in close agreement with solutions of the Schrödinger equation. We demonstrate area-law scaling of entanglement in the model quench in two-, three- and infinite-dimensional spin glasses, supporting the observed stretched-exponential scaling of effort for classical approaches. We assess approximate methods based on tensor networks and neural networks and conclude that no known approach can achieve the same accuracy as the quantum annealer within a reasonable timeframe. Thus quantum annealers can answer questions of practical importance that classical computers cannot.

INTRODUCTION

For decades, the theory of quantum computing has promised monumental speedups for select tasks. Only in the last few years has this come to fruition, with a growing variety of quantum processing units (QPUs) performing computations enormously faster than sophisticated classical approaches—these QPUs include photonic^{1,2}, neutral-atom^{3,4}, and superconducting^{5–8} systems. As these processors expand in size and precision, their growing capacity to tackle complex computations moves toward “quantum supremacy”⁹, when a QPU solves a problem that is impossible for any existing classical computer.

So far, nearly all attempts to showcase quantum supremacy have revolved around random-number generation, be it through boson sampling or random-circuit sampling^{7,10,11}. Classical simulation methods can faithfully replicate the quantum evolution of the circuits, but generally require time and memory resources that scale exponentially with problem size. However, the presence of noise and errors can substantially diminish the required resources¹². Moreover, great strides have been made in classical simulations of noisy quantum systems, curtailing some supremacy claims^{13–16}. While it is no longer controversial to say that near-term QPUs can solve some problems far faster than the best classical methods, this has not yet been firmly established for problems of practical interest. Ref. 8 marked the first attempt to demonstrate supremacy in a problem with potential applications across various domains—discretized dynamics of a kicked Ising model—although the claim was quickly

* aking@dwavesys.com

† amin@dwavesys.com




FIG. 1. **Sampling post-quench states in Ising spin glasses.** Random inputs are generated in topologies of varying dimension. For each input, we consider the task of sampling from the distribution of states following a quantum quench, i.e., rapid change of transverse field $\Gamma(t/t_a)$ and longitudinal field $\mathcal{J}(t/t_a)$ within time t_a . This task is performed using quantum annealing and classical methods based on tensor networks (MPS and PEPS) and neural networks (NQS).

brought into question^{16–20}.

Here we approach a more general and practical problem in the same category, simulating the continuous-time quantum dynamics of the transverse-field Ising model (TFIM), using superconducting quantum annealing (QA) processors^{21,22}. QA has already been employed to simulate quantum phase transitions (QPTs) in Ising-like systems, using both superconducting qubits^{22–24} and Rydberg arrays^{3,25,26}. Ref. 24 demonstrated quantum critical dynamics in cubic spin glasses on over 5000 qubits, with measured critical exponents closely matching estimates from the expected universality class—compelling evidence of Schrödinger evolution, which is not generally simulable with classical methods at this scale. Still, it remains important to establish that within the parametric range studied, despite the limited correlation length and finite experimental precision, approximate classical methods cannot match the solution quality of the QPU in a reasonable amount of time.

The aim of this paper is to evaluate the computational resources necessary to sample the state of a TFIM following a quench through a QPT, at a comparable accuracy to the QPU. Simulating quantum critical dynamics is an important problem with relevance to condensed matter physics^{27–33} and optimization^{3,24,34–36}. The ability of superconducting QA processors to solve this simula-

tion problem with high accuracy has been established for 1D chains²³ and small spin glasses²⁴. Here we simulate large programmable topologies of varying dimension as illustrated in Fig. 1.

To estimate QPU error for small problems, we compare simulations with ground truths computed using well-converged matrix product state (MPS) simulations run on the Summit and Frontier supercomputers at Oak Ridge National Laboratory. We then employ state-of-the-art approximate classical methods based on tensor networks and neural networks to estimate the computational cost of approximately simulating the quantum dynamics to match QPU accuracy. This direct, non-universal verification can only be achieved at small scales, where the ground truth can be computed. To extrapolate beyond the classically simulable regime, we verify universal quantum critical scaling in QPU results, and estimate classical resource requirements based on area-law scaling that is both expected and observed.

QUENCHING A QUANTUM SPIN GLASS

We consider a time-dependent Hamiltonian that interpolates between a driving Hamiltonian $\mathcal{H}_D$ and a classical Ising problem Hamiltonian $\mathcal{H}_P$:

$$\mathcal{H}(t) = \Gamma(t/t_a)\mathcal{H}_D + \mathcal{J}(t/t_a)\mathcal{H}_P, \quad (1)$$

$$\mathcal{H}_D = -\sum_i \sigma_i^x, \quad \mathcal{H}_P = \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z, \quad (2)$$

where $\sigma_i^{x,z}$ are Pauli matrices acting on qubit i , and Γ and $\mathcal{J}$ are the transverse-field and Ising energy scales³⁷, respectively. The evolution starts from time $t=0$ in a paramagnetic phase with $\Gamma(0) \gg \mathcal{J}(0)$ and ends at quench time $t=t_a$ with $\Gamma(1) \ll \mathcal{J}(1)$, deep in the spin-glass phase. These two phases are separated by a QPT whose universal behavior is dictated by the topology of the programmed Ising model. Nonzero couplings J_{ij} are chosen randomly within various topologies: square lattices, dimerized cubic lattices, diamond lattices, and dimerized biclique graphs³⁷. The first three represent lattice structures of real materials^{34,38} and the latter has applications in generative artificial intelligence (AI)³⁹.

We use two QPUs of distinct design: one Advantage system (ADV1) and one prototype Advantage2 system (ADV2). Although these processors have significantly different annealing schedules, we simulate the same model quench (Fig. 1) on both. Each QPU generates at least 1000 samples per second, with parallelization exploited for small problems³⁷.

CLASSICAL SIMULATION TECHNIQUES

We compare QPU results with classical algorithms designed to approximate the solution to the time-dependent Schrödinger equation—see Supplementary Material (SM,

Ref. 37) for details on these methods and other potential approaches. We use two tensor-network techniques—MPS and projected entangled-pair states (PEPS)^{40–45}—as well as neural quantum state (NQS)⁴⁶ simulations. MPS, PEPS, and NQS are currently recognized as the leading approximation frameworks in this context⁴⁷.

Tensor-network ansätze are designed to leverage the area law of entanglement entropy^{48,49} to reduce the complexity of state representation. In particular, for a system of N spins in d spatial dimensions, MPS (or, more generally, tree tensor network) methods reduce the complexity from exponential in N to exponential in $N^{(d-1)/d}$. The benefits diminish as d becomes larger, e.g., for bi-clique problems, where d diverges. The tree-like nature of such network geometries allows for the existence of canonical representations—and easy transformations between them—that underlie the time-dependent variational principle (TDVP)⁵⁰, a general algorithm to simulate time evolution. The main trade-off between approximation quality and computational cost is controlled by the maximum MPS bond dimension χ . Required memory scales as $O(N\chi^2)$ and time of TDVP simulation scales as $O(N\chi^3)$ (other geometries of tree tensor networks beyond MPS have higher powers of χ appearing in the corresponding scalings). The access to canonical forms enables efficient and monotonic control over Hilbert-space truncation, making MPS a suitable technique for ground-truth estimation.

While PEPS ansätze are better aligned with the geometry of the problem and so can more efficiently encode states with area law beyond one dimension, in general they do not benefit from the existence of canonical forms. Inference from PEPS representation, or its exact contraction, is known to be a $\#P$ -complete problem in general^{51,52}. All this makes it susceptible to truncation error during the time evolution³⁷. In principle, for short-ranged quantum correlations, PEPS can employ local or approximate heuristic contractions to reduce computational cost^{42–44}. This can make PEPS efficient for sufficiently fast quenches with short-ranged correlations, such as for dynamics of a kicked Ising model on a “heavy-hex” lattice^{16–18}. However, we demonstrate that the cost of expanding the neighborhood for the longest experimental annealing times becomes prohibitive compared to MPS. Moreover, as the size of the tensors is exponential in d , this methodology quickly becomes impractical, with extreme restrictions on feasible simulations even in three dimensions⁵³.

NQS can also serve as ansätze for approximate wave functions. A variety of results have been demonstrated for calculation of ground states, but with few demonstrations relating to time evolution of disordered systems^{47,54–56}. We examine various time-integration techniques and architectures, including complex restricted Boltzmann machines (cRBMs) and autoregressive models such as transformers and recurrent neural networks (RNNs). We conclude that these approaches cannot produce results with acceptable accuracy beyond the short-

est time scales and smallest models presented in this paper³⁷.

TWO-DIMENSIONAL SYSTEMS

To make the case for computational supremacy, we must first establish accurate and well-understood behavior on nontrivial problems that can still be simulated classically. We start with 2D spin glasses on cylindrical $L \times L$ square lattices, for which we calculate ground truths for up to $L=8$. Square lattices, relative to higher-dimensional topologies, are amenable to classical simulations: MPS and PEPS have been applied to related 2D lattice simulation problems^{26,47,57}.

The computational task at hand is to fairly sample from the quenched state projected in the computational basis. We start with a single exemplary 6×6 spin glass and compare QPU results with an MPS-based ground truth using two statistics: a spin-glass order parameter

$$\langle q^2 \rangle = \frac{2}{N(N-1)} \sum_{i < j} c_{ij}^2, \quad (3)$$

and residual energy $E_{\text{res}} = \langle \mathcal{H}_P \rangle - E_0$. Here, $c_{ij} = \langle \sigma_i^z \sigma_j^z \rangle$ is the two-point correlation function and E_0 is the ground-state energy of $\mathcal{H}_P$. As shown in Fig. 2A, the close agreement between the QPU data (symbols) and the ground truth (solid lines) extends across nearly three orders of magnitude of t_a .

For evaluating the estimated samples in comparison to ground-truth MPS results, our main figure of merit is based on spin-spin correlations. Fig. 2B shows all c_{ij} in the same problem. We compare the ground truth (converged at $\tilde{\chi}=256$) with QPU data (left) and approximate MPS output using a smaller $\chi=64$ (right). Agreement can be summarized as a correlation error

$$\epsilon_c = \left(\frac{\sum_{i,j} (c_{ij} - \tilde{c}_{ij})^2}{\sum_{i,j} \tilde{c}_{ij}^2} \right)^{1/2}, \quad (4)$$

where “ $\sim$ ” denotes a ground-truth value.

Various other distributional-divergence measures exist, such as Kullback–Leibler divergence and classical fidelity (Bhattacharyya distance), but they have the disadvantage that when evaluated by sampling, the number of samples must scale exponentially with the distribution entropy, which grows with N in our application. Nevertheless, we use classical fidelity $\mathcal{F} = \sum_z \sqrt{p_z \tilde{p}_z}$, where p_z and $\tilde{p}_z$ are method and ground-truth probabilities of state z , to corroborate the main metric ϵ_c at small sizes. Fig. 2C scatters p_z against $\tilde{p}_z$ based on 10^6 samples, again using QPU and MPS ($\chi=64$); resulting estimates of $\mathcal{F}$ are similar. In both Figs. 2B and 2C, it is evident that the approximate MPS yields data closely resembling that of the QPU, as manifested in both error measures ϵ_c and $\mathcal{F}$. In the following sections, the bond dimension that




FIG. 2. **QPU and MPS on quenched 2D spin glasses.** (A–C) Data for a typical realization of a 6×6 cylindrical spin glass. (A) Squared order parameter (top) and residual energy per spin (bottom) data exhibit close agreement between QPU and the ground truth across nearly three decades of t_a . For very fast quenches, correlations and therefore $\langle q^2 \rangle$ are small; for very slow quenches, the system evolves adiabatically and produces classical ground states with vanishing residual energy. (B) Spin-spin correlations c_{ij} compared against ground-truth values $\tilde{c}_{ij}$ for $t_a=7$ ns. Left, QPU. Right, MPS at a bond dimension producing output of similar quality. (C) Comparison of individual state probabilities for the same simulations as B; classical fidelities $\mathcal{F}$ estimated from 10^6 samples are comparable. (D–E) Ensembles of 20 $L \times L$ instances (disorder realizations). (D) Median correlation error ϵ_c remains roughly constant across system sizes for QPU (left), but matching the QPU error level in MPS requires an increasing QPU-equivalent bond dimension χ_Q . (E) χ_Q exhibits exponential growth in L for two quench rates. For $t_a=20$ ns, the superior processor, ADV2, shows lower errors, and therefore steeper scaling of χ_Q . All error bars represent 95% bootstrap confidence intervals of the ensemble median.

achieves comparable accuracy to the QPU, the “QPU-equivalent bond dimension”⁴, is denoted as χ_Q (≈ 64 in this example).

Next, we address scaling behavior of $L \times L$ lattices. Fig. 2D (left) shows flat scaling of ϵ_c for QPU in L (median over 20 random disorder realizations) using both ADV1 and ADV2 processors. Interpolating MPS χ -dependent error (right) allows us to estimate χ_Q . Fig. 2E shows χ_Q as a function of L and t_a . Crucially, we observe exponential dependence on L , which is consistent with area-law scaling⁵⁸. For $t_a=20$ ns, ADV2 has lower error, therefore higher χ_Q , than ADV1—henceforth all QPU results shown in the main text are from ADV2.

MPS can produce QPU-quality results—and better—for all annealing times of interest up to $L=8$. In contrast, we found that PEPS and NQS methods struggle for slower quenches ($t_a=20$ ns), even for those relatively small systems. Many classical methods, including PEPS, are especially effective in short-time regimes, where locality can be exploited with relative insensitivity to system size¹⁶.

We have exhaustively probed PEPS approaches to continuous time-evolution³⁷, starting with a simple update scheme⁴³, where only the immediate neighborhood is accounted for during the application of each gate of Trotterized evolution—a strategy that proved successful for the kicked Ising model^{16–18} despite its simplicity. We then systematically enlarged the size of the neighborhood—with quickly increasing numerical cost and heuristic compromises—to a scale approaching the entire lattice for an 8×8 example.




FIG. 3. **Breakdown of PEPS in 8×8 systems.** We show correlation error as a function of PEPS bond dimension D for three annealing times. The best-performing PEPS variant, of many evaluated³⁷, was “plqt-NN+”. Computation time increases sharply with bond dimension D . PEPS easily beats QPU error at $t_a=2$ ns, but degrades for slower quenches. At 20 ns, all PEPS experiments have higher error than QPU, and PEPS appears to be neighborhood-limited: increasing D does not significantly reduce error.

Fig. 3 highlights the results of the best-performing approach (“plaquette-nearest-neighbor+” [plqt-NN+], see SM³⁷), showing errors generated during the PEPS evolution. All procedures perform well for $t_a=2$ ns, even with small PEPS bond dimension $D=4$. A combination of beyond-nearest-neighbor environment and increasing D allows for systematic improvement at 7 ns. Finally, neighborhood-limited errors saturate well above the QPU level for $t_a=20$ ns. This is consistent with observations from spatial decomposition methods we eval-

uated³⁷, which suggest the presence of system-spanning correlations for the longer quenches.

We have also attempted PEPS simulations and spatial decomposition of cubic lattices³⁷. It is shown that for the largest $6 \times 6 \times 6$ problem considered in the next section, the quantum correlations are system spanning at $t_a = 20$ ns, preventing PEPS from exploiting locality. The biclique problems have no such locality to exploit.

While NQS ansätze do not rely explicitly on locality, we arrive at qualitatively similar conclusions³⁷. Accurate results are easily reproduced for fast quenches ($t_a = 2$ ns). However, modeling dynamics even at $t_a = 7$ ns requires numerical effort that grows prohibitively with system size, owing to the increase in required sampling precision at each step of the time evolution.

We conclude that these various PEPS and NQS methods, which we detail in the SM³⁷, do not match the utility of MPS as a competitive solver for the higher-dimensional topologies, which we now consider for the remainder of the paper.

HIGHER-DIMENSIONAL SYSTEMS

We now apply the approach outlined in the preceding section to evaluate sampling on more complex topologies: cubic, diamond and biclique lattices. The cubic and biclique lattices are dimerized, with each standard lattice site represented by a pair of spins rather than a single spin³⁷.

As in square lattices (Fig. 2D), QPU error appears nearly constant for cubic and diamond lattices, as seen in Fig. 4A. For biclique problems, ϵ_c increases with system size, which we believe to be a finite-size effect. Nonetheless, the increase is accounted for in the calculation of χ_Q , which is shown for all topologies in Fig. 4B. The x -axis is the MPS bipartition area, which scales to leading order as $N^{(d-1)/d}$. Just as we observed in the square lattice behavior shown in Fig. 2E, each of these higher-dimensional topologies exhibits exponential dependence on area, again compatible with area-law scaling.

This scaling of bond dimension is a direct consequence of entanglement in the system: In Fig. 4C we plot the maximum entanglement entropy $S_{\max}$ (maximum throughout the quench and over all MPS bipartitions) against $\log \chi_Q$, and observe a linear relation for multiple topologies, sizes, and quench rates. This underscores the intimate connection between MPS representation complexity and entanglement. The inset demonstrates close agreement between $S_{\max}$ at bond dimension χ_Q and the ground-truth entanglement $\tilde{S}_{\max}$. Together, this tells us two things: First, MPS methods must capture almost all the entanglement in the system to match QPU quality. Second, the MPS that achieves this is efficient, in the sense that its entanglement is close to the maximum possible value, $\log \chi_Q$, in an MPS of bond dimension χ_Q . We further note that for all topologies, χ_Q decreases as t_a increases from 7 ns to 20 ns—this trend, opposite from the




FIG. 4. **Entanglement and QPU-equivalent bond dimension.** (A) QPU median correlation error for the four topologies as color coded in the left panels. (B) For all four problem sets, the bond dimension required by MPS to match QPU simulation quality, χ_Q , exhibits exponential dependence on the bipartition area (asymptotically $N^{(d-1)/d}$, see inset example), consistent with the area-law scaling of entanglement entropy (two-point fit is trivial for diamond at $t_a = 7$ ns). (C) Median $\log \chi_Q$ against maximum entanglement entropy $S_{\max}(\chi_Q)$ for MPS with bond dimension χ_Q . All data points for the four topologies and two annealing times fall roughly on the same line. Inset: $S_{\max}(\chi_Q)$ versus ground-truth $\tilde{S}_{\max}$ estimates for each input.

expectation in the large-system limit, is a finite-size effect indicating system-spanning quantum correlations³⁰.

Our claim of supremacy requires evidence of QPU accuracy both within and beyond the classically simulable regime. For small systems, in addition to the small distributional distance quantified by ϵ_c (and $\mathcal{F}$), we provide evidence in Fig. 5A with an instance-wise comparison of the spin-glass order parameter $\langle q^2 \rangle$ on the largest lattice of each topology for which the MPS ground truth is converged at both $t_a = 7$ ns and $t_a = 20$ ns. The resulting quantities have a median deviation of only 1% between quantum and classical simulation.

Beyond the classically verifiable regime, we show agreement with quantum mechanics up to the largest size using the theory of critical phenomena. As shown in Fig. 5B, we verify a dynamic finite-size scaling ansatz^{24,59} in which the Binder cumulant

$$U = \frac{1}{2} \left(3 - \frac{[\langle q^4 \rangle]}{[\langle q^2 \rangle]^2} \right), \quad (5)$$

collapses as a function of $t_a L^{-\mu}$, where μ is the universal Kibble-Zurek (KZ) exponent corresponding to the topol-




FIG. 5. **Dynamical scaling and resource estimates for large-scale quantum simulation.** (A) We compare $\langle q^2 \rangle$ in QPU output and the ground truth for 20 instances of each topology at $t_a = 7$ ns and 20 ns, establishing close agreement (median relative error $< 1\%$). (B) Using a range of system sizes that we cannot simulate classically, we collapse Binder cumulants U using a dynamic finite-size scaling ansatz, yielding best-fit KZ exponents μ . Curve in collapsed data indicates crossover from power-law (KZ) to system-spanning correlations. Data points are computed from ensembles of 300 to 20,400 simulated spin-glass realizations (collapsed sizes N , indicated by different shades and markers, are 16–324 for square, 54–432 for cubic, 72–567 for diamond, and 32–96 for biclique). Inset: Estimates of U using simulations shown in A for $t_a = 7$ ns and 20 ns (dark and light stars respectively, also in main plot) show close agreement with the ground truth at simulable sizes, bridging the gap between the classically-simulable and large-scale regimes. (C) Extracted KZ exponents for one-, two-, three-, and infinite-dimensional systems agree with previous estimates for the corresponding QPT universality classes. (D) Area-law scaling of χ_Q (Fig. 4B) allows extrapolation of required classical time based on $N\chi_Q^3$ scaling of MPS methods. Filled markers indicate classically-simulated 20-instance ensembles; larger problems are extrapolations.

ogy’s universality class. Here, $[\cdot]$ indicates an average over 300 to 20,400 disorder realizations. The inset shows estimates of U for the 20-instance ensembles for which we have a ground truth (Fig. 5A), with excellent agreement between experiment and model; inset $\tilde{U}$ values are marked as stars in the main plot (no ground truth could be computed for the smallest instances in the diamond collapse, $6 \times 6 \times 8$). The crossover of the collapse functions from power-law to saturation near $U=1$ is an indication of system-spanning correlations, preventing PEPS or other approaches from exploiting locality. QPU estimates of μ for square, cubic, diamond, and biclique ensembles—extracted from experimental data collapse—are in agreement with previous estimates^{24,36,60} and theory⁶¹, as shown in Fig. 5C (agreement for 1D and “cubic no-dimer” lattices is also confirmed³⁷). These exponents, along with the stars in Fig. 5B, show that the QPU captures both universal quantum critical scaling and quantitative details of the non-universal collapse functions.

The area-law scaling observed in Fig. 4B provides a foundation for extrapolating χ_Q to larger sizes, as shown in Fig. 5D. Based on classically simulable systems (Fig. 4B), we can estimate the computational resources required to match QPU simulation quality with MPS, using leading scaling forms of $N\chi_Q^2$ for memory and $N\chi_Q^3$ for time (and energy) requirements. We generously assume perfect parallelization, and extrapolate GPU time required to generate MPS results³⁷, ignoring practical scaling challenges. Extrapolation (open symbols in Fig. 5D) indicates that MPS requires infeasible

resources to accurately reproduce the measurements in Fig. 5B: On the largest problems, MPS would take millions of years on the Frontier supercomputer per input to match QPU quality. Memory requirements would exceed its 700 PB storage, and electricity requirements would exceed annual global consumption³⁷.

SUMMARY AND CONCLUSIONS

We have demonstrated computational supremacy in the quantum simulation of nonequilibrium magnetic spin dynamics. We simulated square, cubic, diamond, and biclique topologies that are relevant to applications in materials science and AI, and are amenable to scaling analysis through the area law of entanglement and universal quantum critical scaling. More challenging and irregular topologies—for example those corresponding to deep neural networks³⁹—can also be simulated with the same QPUs.

Two distinct QPU generations yielded consistent solutions in agreement with quantum theory. At small scales this is verified by simulating the Schrödinger equation. At larger scales, where solving the Schrödinger equation is infeasible, QPU observables follow the anticipated universal quantum critical dynamic scaling, with critical exponents in close agreement with independent estimates.

To evaluate classical computation costs, we examined tensor networks and neural networks as the leading frameworks capable of providing approximate solu-

tions to the Schrödinger equation, seeking QPU-quality output. Among all algorithms, we could only extract a reliable scaling from MPS; PEPS and NQS approaches failed even on modestly-sized 2D problems.

For MPS, we demonstrated a stretched-exponential dependence of the required bond dimension χ_Q on the number of qubits, consistent with area-law entanglement scaling. This enabled a reliable extrapolation of χ_Q , thus computational resources, to larger systems. Extrapolation beyond the simulable scale to a few hundred qubits, far below the maximum experimental scale, indicates a hypothetical runtime on the Frontier supercomputer surpassing millions of years with infeasible memory and energy requirements.

The QPU has been shown to accurately sample from a variety of random models with correct predictions of

microscopic, macroscopic and scaling statistics without *a priori* knowledge of the ground truth. We hope and expect that our results will inspire novel numerical techniques for quantum simulation. However, competing classical methods not only have to match QPU-quality results for all models at the scales where ground truths were obtained, but also should reproduce both universal and non-universal details of quantum critical scaling up to the largest experimentally achievable sizes.

The nonequilibrium quantum dynamics studied here are responsible for a scaling advantage over classical dynamics in optimization²⁴. This impracticability of classical simulation opens the door to quantum advantage in optimization and AI, addressing scientific questions that may otherwise remain unanswered, and applications that may be classically impossible.

-
- [1] H. S. Zhong, Y. H. Deng, J. Qin, H. Wang, M. C. Chen, *et al.*, Phase-Programmable Gaussian Boson Sampling Using Stimulated Squeezed Light, *Physical Review Letters* **127**, 180502 (2021).
 - [2] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, *et al.*, Quantum computational advantage with a programmable photonic processor, *Nature* **606**, 75 (2022).
 - [3] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, *Science* **376**, 1209 (2022).
 - [4] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, *Benchmarking highly entangled states on a 60-atom analog quantum simulator* (2023), [arXiv:2308.07914](#).
 - [5] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
 - [6] Y. Wu, W. S. Bao, S. Cao, F. Chen, M. C. Chen, *et al.*, Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, *Physical Review Letters* **127**, 180501 (2021).
 - [7] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, *et al.*, *Phase transition in Random Circuit Sampling* (2023), [arXiv:2304.11119](#).
 - [8] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, *et al.*, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
 - [9] J. Preskill, *Quantum computing and the entanglement frontier* (2012), [arXiv:1203.5813](#).
 - [10] A. P. Lund, M. J. Bremner, and T. C. Ralph, Quantum sampling problems, bosonsampling and quantum supremacy, *NPJ Quantum Information* **3**, 15 (2017).
 - [11] A. Zlokapa, B. Villalonga, S. Boixo, and D. A. Lidar, Boundaries of quantum supremacy via random circuit sampling, *NPJ Quantum Information* **9**, 36 (2023).
 - [12] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, A polynomial-time classical algorithm for noisy random circuit sampling, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (ACM, 2023) pp. 945–957.
 - [13] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, *Leveraging Secondary Storage to Simulate Deep 54-qubit Sycamore Circuits* (2019), [arXiv:1910.09534](#).
 - [14] C. Oh, L. Jiang, and B. Fefferman, Spoofing Cross-Entropy Measure in Boson Sampling, *Physical Review Letters* **131**, 010401 (2023).
 - [15] K. Kechedzhi, S. Isakov, S. Mandrà, B. Villalonga, X. Mi, S. Boixo, and V. Smelyanskiy, Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments, *Future Generation Computer Systems* **153**, 431 (2024).
 - [16] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, Efficient Tensor Network Simulation of IBM’s Eagle Kicked Ising Experiment, *PRX Quantum* **5**, 010308 (2024).
 - [17] T. Begušić, J. Gray, and G. K.-L. Chan, Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance, *Science Advances* **10**, eadk4321 (2024).
 - [18] H.-J. Liao, K. Wang, Z.-S. Zhou, P. Zhang, and T. Xiang, *Simulation of IBM’s kicked Ising experiment with Projected Entangled Pair Operator* (2023), [arXiv:2308.03082](#).
 - [19] S. Anand, K. Temme, A. Kandala, and M. Zaletel, *Classical benchmarking of zero noise extrapolation beyond the exactly-verifiable regime* (2023), [arXiv:2306.17839](#).
 - [20] E. G. D. Torre and M. M. Roses, *Dissipative mean-field theory of IBM utility experiment* (2023), [arXiv:2308.01339](#).
 - [21] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Physical Review E* **58**, 5355 (1998).
 - [22] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, *et al.*, Phase transitions in a programmable quantum spin glass simulator, *Science* **361**, 162 (2018).
 - [23] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, *et al.*, Coherent quantum annealing in a programmable 2,000 qubit Ising chain, *Nature Physics* **18**, 1324 (2022).
 - [24] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, *et al.*, Quantum critical dynamics in a 5,000-

- qubit programmable spin glass, *Nature* **617**, 61 (2023).
- [25] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, *et al.*, Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator, *Nature* **568**, 207 (2019).
- [26] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, *et al.*, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature* **595**, 233 (2021).
- [27] W. H. Zurek, Cosmological experiments in superfluid helium?, *Nature* **317**, 505 (1985).
- [28] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, 2011).
- [29] P. Calabrese and J. Cardy, Time Dependence of Correlation Functions Following a Quantum Quench, *Physical Review Letters* **96**, 136801 (2006).
- [30] E. Canovi, E. Ercolessi, P. Naldesi, L. Taddia, and D. Vodola, Dynamics of entanglement entropy and entanglement spectrum crossing a quantum phase transition, *Physical Review B* **89**, 104303 (2014).
- [31] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore, Colloquium: Nonequilibrium dynamics of closed interacting quantum systems, *Reviews of Modern Physics* **83**, 863 (2011).
- [32] A. del Campo and W. H. Zurek, Universality of phase transition dynamics: Topological defects from symmetry breaking, *International Journal of Modern Physics A* **29**, 1430018 (2014).
- [33] D. Rossini and E. Vicari, Coherent and dissipative dynamics at quantum phase transitions, *Phys. Rep.* **936**, 1 (2021), Coherent and dissipative dynamics at quantum phase transitions.
- [34] J. Brooke, D. Bitko, T. F., T. F. Rosenbaum, and G. Aeppli, Quantum Annealing of a Disordered Magnet, *Science* **284**, 779 (1999).
- [35] A. Das and B. K. Chakrabarti, Colloquium: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* **80**, 1061 (2008).
- [36] M. Bernaschi, I. G.-A. Pemartín, V. Martín-Mayor, and G. Parisi, *The Quantum Transition of the Two-Dimensional Ising Spin Glass: A Tale of Two Gaps* (2023), [arXiv:2310.07486](https://arxiv.org/abs/2310.07486).
- [37] Materials and methods are available as supplementary materials.
- [38] J. Bertinshaw, Y. Kim, G. Khaliullin, and B. Kim, Square Lattice Iridates, *Annual Review of Condensed Matter Physics* **10**, 315 (2019).
- [39] J. D. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [40] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Annals of Physics* **326**, 96 (2011).
- [41] E. Stoudenmire and S. R. White, Studying Two-Dimensional Systems with the Density Matrix Renormalization Group, *Annual Review of Condensed Matter Physics* **3**, 111 (2012).
- [42] F. Verstraete, V. Murg, and J. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
- [43] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Annals of Physics* **349**, 117 (2014).
- [44] S.-J. Ran, E. Tirrito, C. Peng, X. Chen, L. Tagliacozzo, G. Su, and M. Lewenstein, *Tensor network contractions: methods and applications to quantum many-body systems* (Springer Cham, 2020).
- [45] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).
- [46] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [47] M. Schmitt, M. M. Rams, J. Dziarmaga, M. Heyl, and W. H. Zurek, Quantum phase transition dynamics in the two-dimensional transverse-field Ising model, *Science Advances* **8**, eabl6850 (2022).
- [48] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the area law, and the computational power of projected entangled pair states, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [49] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Reviews of Modern Physics* **82**, 277 (2010).
- [50] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [51] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational complexity of projected entangled pair states, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [52] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, Contracting projected entangled pair states is average-case hard, *Phys. Rev. Res.* **2**, 013010 (2020).
- [53] P. C. G. Vlaar and P. Corboz, Simulation of three-dimensional quantum systems with projected entangled-pair states, *Phys. Rev. B* **103**, 205137 (2021).
- [54] G. Carleo, F. Becca, M. Schiró, and M. Fabrizio, Localization and glassy dynamics of many-body quantum systems, *Scientific reports* **2**, 243 (2012).
- [55] M. Medvidović and G. Carleo, Classical variational simulation of the quantum approximate optimization algorithm, *npj Quantum Information* **7**, 101 (2021).
- [56] K. Donatella, Z. Denis, A. Le Boité, and C. Ciuti, Dynamics with autoregressive neural quantum states: Application to critical quench dynamics, *Physical Review A* **108**, 022210 (2023).
- [57] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature* **595**, 227 (2021).
- [58] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, *Physical Review Letters* **91**, 147902 (2003).
- [59] C.-W. Liu, A. Polkovnikov, A. W. Sandvik, and A. P. Young, Universal dynamic scaling in three-dimensional Ising spin glasses, *Physical Review E* **92**, 022128 (2015).
- [60] M. Guo, R. N. Bhatt, and D. A. Huse, Quantum critical behavior of a three-dimensional Ising spin glass in a transverse magnetic field, *Physical Review Letters* **72**, 4137 (1994).
- [61] N. Read, S. Sachdev, and J. Ye, Landau theory of quantum spin glasses of rotors and Ising spins, *Physical Review B* **52**, 384 (1995).
- [62] K. Boothby, A. D. King, and J. Raymond, *Zephyr Topology of D-Wave Quantum Processors*, Tech. Rep. (D-

- Wave, 2021).
- [63] K. Boothby, P. Bunyk, J. Raymond, and A. Roy, *Next-Generation Topology of D-Wave Quantum Processors* (2020), [arXiv:2003.00133](#).
 - [64] R. Harris, J. Johansson, A. J. Berkley, M. W. Johnson, T. Lanting, *et al.*, Experimental demonstration of a robust and scalable flux qubit, *Physical Review B* **81**, 134510 (2010).
 - [65] K. Chern, K. Boothby, J. Raymond, P. Farré, and A. D. King, Tutorial: calibration refinement in quantum annealing, *Frontiers in Computer Science* **5**, 10.3389/fcomp.2023.1238988 (2023).
 - [66] C. McCreesh, P. Prosser, and J. Trimble, The Glasgow Subgraph Solver: Using Constraint Programming to Tackle Hard Subgraph Isomorphism Problem Variants, in *Graph Transformation*, edited by F. Gadducci and T. Kehrler (Springer International Publishing, Cham, 2020) pp. 316–324.
 - [67] D. Venturelli, S. Mandrà, S. Knysh, B. O’Gorman, R. Biswas, and V. Smelyanskiy, Quantum Optimization of Fully Connected Spin Glasses, *Physical Review X* **5**, 031040 (2015).
 - [68] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
 - [69] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
 - [70] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
 - [71] M. B. Hastings, An area law for one-dimensional quantum systems, *Journal of statistical mechanics: theory and experiment* **2007**, P08024 (2007).
 - [72] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-dependent variational principle for quantum lattices, *Physical Review Letters* **107**, 1 (2011).
 - [73] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Annals of Physics* **411**, 167998 (2019).
 - [74] M. Fishman, S. White, and E. Stoudenmire, The ITensor Software Library for Tensor Network Calculations, *SciPost Physics Codebases* **4**, 4 (2022).
 - [75] G. Alvarez, The density matrix renormalization group for strongly correlated electron systems: A generic implementation, *Computer Physics Communications* **180**, 1572 (2009).
 - [76] M. M. Rams, G. Wójtowicz, A. Sinha, and J. Hasik, *Yet another symmetric tensor network* (2024).
 - [77] M. Yang and S. R. White, Time-dependent variational principle with ancillary krylov subspace, *Phys. Rev. B* **102**, 094315 (2020).
 - [78] J.-W. Li, A. Gleis, and J. Von Delft, Time-dependent variational principle with controlled bond expansion for matrix product states, *arXiv preprint arXiv:2208.10972* (2022).
 - [79] M. Ganahl, J. Beall, M. Hauru, A. G. Lewis, T. Wójcicki, J. H. Yoo, Y. Zou, and G. Vidal, Density matrix renormalization group with tensor processing units, *PRX Quantum* **4**, 010317 (2023).
 - [80] J. Unfried, J. Hauschild, and F. Pollmann, Fast time evolution of matrix product states using the qr decomposition, *Phys. Rev. B* **107**, 155133 (2023).
 - [81] P. Secular, N. Gourianov, M. Lubasch, S. Dolgov, S. R. Clark, and D. Jaksch, Parallel time-dependent variational principle algorithm for matrix product states, *Phys. Rev. B* **101**, 235123 (2020).
 - [82] Y. Da Liao, H. Li, Z. Yan, H.-T. Wei, W. Li, Y. Qi, and Z. Y. Meng, Phase diagram of the quantum Ising model on a triangular lattice under external field, *Physical Review B* **103**, 104416 (2021).
 - [83] B. O’Gorman, Parameterization of Tensor Network Contraction (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2019) pp. 10:1–10:19, 14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019).
 - [84] G. Cataldi, A. Abedi, G. Magnifico, S. Notarnicola, N. D. Pozza, V. Giovannetti, and S. Montangero, Hilbert curve vs hilbert space: exploiting fractal 2d covering to increase tensor network efficiency, *Quantum* **5**, 556 (2021).
 - [85] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
 - [86] V. Murg, F. Verstraete, Ö. Legeza, and R. M. Noack, Simulating strongly correlated quantum systems with tree tensor networks, *Phys. Rev. B* **82**, 205105 (2010).
 - [87] G. Vidal, Entanglement renormalization, *Phys. Rev. Lett.* **99**, 220405 (2007).
 - [88] M. Rizzi, S. Montangero, and G. Vidal, Simulation of time evolution with multiscale entanglement renormalization ansatz, *Phys. Rev. A* **77**, 052328 (2008).
 - [89] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, *arXiv preprint cond-mat/0407066* (2004).
 - [90] Y. Nishio, N. Maeshima, A. Gendiar, and T. Nishino, Tensor product variational formulation for quantum systems, *arXiv:cond-mat/0401115* (2004).
 - [91] Y. Ge and J. Eisert, Area laws and efficient descriptions of quantum many-body states, *New Journal of Physics* **18**, 083026 (2016).
 - [92] S. Patra, S. S. Jahromi, S. Singh, and R. Orus, *Efficient tensor network simulation of IBM’s largest quantum processors* (2023), [arXiv:2309.15642](#).
 - [93] I. L. Markov and Y. Shi, Simulating quantum computation by contracting tensor networks, *SIAM Journal on Computing* **38**, 963 (2008).
 - [94] J. Dziarmaga, Time evolution of an infinite projected entangled pair state: Neighborhood tensor update, *Phys. Rev. B* **104**, 094411 (2021).
 - [95] H. C. Jiang, Z. Y. Weng, and T. Xiang, Accurate determination of tensor network state of quantum lattice models in two dimensions, *Phys. Rev. Lett.* **101**, 090603 (2008).
 - [96] D. J. Harvey and D. R. Wood, *Journal of Combinatorial Theory, Series B*, volume 132, pages 157–179 **132**, 157 (2018).
 - [97] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).
 - [98] T. Vieijra, J. Haegeman, F. Verstraete, and L. Van derstraeten, Direct sampling of projected entangled-pair states, *Phys. Rev. B* **104**, 235141 (2021).
 - [99] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable programming tensor networks, *Phys. Rev. X* **9**,

- 031041 (2019).
- [100] W.-Y. Liu, J. Hasik, S.-S. Gong, D. Poilblanc, W.-Q. Chen, and Z.-C. Gu, Emergence of gapless quantum spin liquid from deconfined quantum critical point, *Phys. Rev. X* **12**, 031039 (2022).
 - [101] F. Pan, K. Chen, and P. Zhang, Solving the sampling problem of the sycamore quantum circuits, *Phys. Rev. Lett.* **129**, 090502 (2022).
 - [102] T. N. Ikeda, A. Abrar, I. L. Chuang, and S. Sug-iura, Minimum Trotterization Formulas for a Time-Dependent Hamiltonian, *Quantum* **7**, 1168 (2023).
 - [103] G. H. Low and I. L. Chuang, Hamiltonian Simulation by Qubitization, *Quantum* **3**, 163 (2019).
 - [104] J. Gonzalez-Conde, Z. Morrell, M. Vuffray, T. Albash, and C. Coffrin, The cost of emulating a small quantum annealing problem in the circuit-model (2024), [arXiv:2402.17667](https://arxiv.org/abs/2402.17667).
 - [105] M. P. Zaletel and F. Pollmann, Isometric tensor network states in two dimensions, *Phys. Rev. Lett.* **124**, 037201 (2020).
 - [106] G. Evenbly, Gauge fixing, canonical forms, and optimal truncations in tensor networks with closed loops, *Phys. Rev. B* **98**, 085155 (2018).
 - [107] R. Alkabetz and I. Arad, Tensor networks contraction and the belief propagation algorithm, *Phys. Rev. Res.* **3**, 023073 (2021).
 - [108] B. Dozier and J. Sapir, Simple versus nonsimple loops on random regular graphs, *Journal of Graph Theory* , 1 (2023).
 - [109] L. Cincio, J. Dziarmaga, M. M. Rams, and W. H. Zurek, Entropy of entanglement and correlations induced by a quench: Dynamics of a quantum phase transition in the quantum Ising model, *Physical Review A* **75**, 052321 (2007).
 - [110] J. Dziarmaga and J. M. Mazur, Tensor network simulation of the quantum Kibble-Zurek quench from the Mott to the superfluid phase in the two-dimensional Bose-Hubbard model, *Physical Review B* **107**, 144510 (2023).
 - [111] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nature Physics* **14**, 10.1038/s41567-018-0124-x (2018).
 - [112] P. Grassberger, Entropy estimates from insufficient samplings, [arXiv: Data Analysis, Statistics and Probability](https://arxiv.org/abs/2003.04437) (2003).
 - [113] P. Calabrese and J. Cardy, Evolution of entanglement entropy in one-dimensional systems, *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P04010 (2005).
 - [114] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Entanglement in many-body systems, *Rev. Mod. Phys.* **80**, 517 (2008).
 - [115] P. Calabrese and J. Cardy, Entanglement entropy and quantum field theory, *Journal of statistical mechanics: theory and experiment* **2004**, P06002 (2004).
 - [116] S. Czischek, M. Gärttner, and T. Gasenzer, Quenches near ising quantum criticality as a challenge for artificial neural networks, *Physical Review B* **98**, 024311 (2018).
 - [117] C. Roth and A. H. MacDonald, Group convolutional neural networks improve quantum state accuracy (2021), [2104.05085](https://arxiv.org/abs/2104.05085).
 - [118] M. Schmitt and M. Heyl, Quantum many-body dynamics in two dimensions with artificial neural networks, *Physical Review Letters* **125**, 100503 (2020).
 - [119] T. Schmale, M. Reh, and M. Gärttner, Efficient quantum state tomography with convolutional neural networks, *npj Quantum Information* **8**, 115 (2022).
 - [120] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Recurrent neural network wave functions, *Physical Review Research* **2**, 023358 (2020).
 - [121] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, Supplementing Recurrent Neural Network Wave Functions with Symmetry and Annealing to Improve Accuracy (2022) 35th Conference on Neural Information Processing Systems, [arXiv:2207.14314 \[cond-mat.dis-nn\]](https://arxiv.org/abs/2207.14314).
 - [122] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, Investigating topological order using recurrent neural networks, *Phys. Rev. B* **108**, 075152 (2023).
 - [123] H. Lange, F. Döschl, J. Carrasquilla, and A. Bohrdt, Neural network approach to quasiparticle dispersions in doped antiferromagnets, [arXiv preprint arXiv:2310.08578](https://arxiv.org/abs/2310.08578) [10.48550/arXiv:2310.08578](https://arxiv.org/abs/2310.08578) (2023).
 - [124] J. Carrasquilla, D. Luo, F. Pérez, A. Milsted, B. K. Clark, M. Volkovs, and L. Aolita, Probabilistic simulation of quantum circuits using a deep-learning architecture, *Phys. Rev. A* **104**, 032610 (2021).
 - [125] K. Sprague and S. Czischek, *Variational monte carlo with large patched transformers* (2023), [2306.03921](https://arxiv.org/abs/2306.03921).
 - [126] R. Rende, L. L. Viteritti, L. Bardone, F. Becca, and S. Goldt, A simple linear algebra identity to optimize large-scale neural network quantum states, [arXiv preprint arXiv:2310.05715](https://arxiv.org/abs/2310.05715) [10.48550/arXiv:2310.05715](https://arxiv.org/abs/2310.05715) (2023).
 - [127] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
 - [128] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, *et al.*, *Variational benchmarks for quantum many-body problems* (2023), [arXiv:2302.04919](https://arxiv.org/abs/2302.04919).
 - [129] M. Bortone, Y. Rath, and G. H. Booth, Impact of conditional modelling for universal autoregressive quantum states, [arXiv preprint arXiv:2306.05917](https://arxiv.org/abs/2306.05917) [10.48550/arXiv:2306.05917](https://arxiv.org/abs/2306.05917) (2023).
 - [130] S.-I. Amari, Natural gradient works efficiently in learning, *Neural computation* **10**, 251 (1998).
 - [131] L. Hackl, T. Guaita, T. Shi, J. Haegeman, E. Demler, and J. I. Cirac, Geometry of variational methods: dynamics of closed quantum systems, *SciPost Phys.* **9**, 048 (2020).
 - [132] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C (2nd ed.): the art of scientific computing* (Cambridge University Press, USA, 1992).
 - [133] S. Barison, F. Vicentini, and G. Carleo, An efficient quantum algorithm for the time evolution of parameterized circuits, *Quantum* **5**, 512 (2021).
 - [134] I. L. Gutiérrez and C. B. Mendl, Real time evolution with neural-network quantum states, *Quantum* **6**, 627 (2022).
 - [135] A. Sinibaldi, C. Giuliani, G. Carleo, and F. Vicentini, *Unbiasing time-dependent variational Monte Carlo by projected quantum evolution* (2023), [2305.14294](https://arxiv.org/abs/2305.14294).
 - [136] A. Gelman and D. B. Rubin, Inference from iterative simulation using multiple sequences, *Statistical Science* **7**, 457 (1992).

- [137] A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin, *Bayesian Data Analysis* (Chapman and Hall/CRC, 2013).
- [138] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, NetKet 3: Machine Learning Toolbox for Many-Body Quantum Systems, *SciPost Phys. Codebases*, **7** (2022).
- [139] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, *JAX: composable transformations of Python+NumPy programs* (2018).
- [140] T. W. B. Kibble, Topology of cosmic domains and strings, *J. Phys. A* **9**, 1387–1398 (1976).
- [141] J. Dziarmaga, Dynamics of a quantum phase transition and relaxation to a steady state, *Adv. Phys.* **59**, 1063 (2010).
- [142] P. Weinberg, M. Tylutki, J. M. Rönkkö, J. Westerholm, J. A. Åström, P. Manninen, P. Törmä, and A. W. Sandvik, Scaling and Diabatic Effects in Quantum Annealing with a D-Wave Device, *Physical Review Letters* **124**, 090502 (2020).
- [143] Y. Bando, Y. Susa, H. Oshiyama, N. Shibata, M. Ohzeki, F. J. Gómez-Ruiz, D. A. Lidar, S. Suzuki, A. del Campo, and H. Nishimori, Probing the universality of topological defect formation in a quantum annealer: Kibble-Zurek mechanism and beyond, *Physical Review Research* **2**, 033369 (2020).
- [144] H. Rieger and A. P. Young, Zero-temperature quantum phase transition of a two-dimensional Ising spin glass, *Physical Review Letters* **72**, 4141 (1994).
- [145] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by Simulated Annealing, *Science* **220**, 671 (1983).
- [146] S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, Experimental signature of programmable quantum annealing, *Nature Communications* **4**, 2067 (2013).
- [147] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Defining and detecting quantum speedup, *Science* **345**, 420 (2014).
- [148] J. A. Smolin and G. Smith, Classical signature of quantum annealing, *Frontiers in Physics* **2**, 8 (2014).
- [149] S. W. Shin, G. Smith, J. A. Smolin, and U. Vazirani, How “Quantum” is the D-Wave Machine? (2014), [arXiv:1401.7087](https://arxiv.org/abs/1401.7087).
- [150] T. Albash and J. Marshall, Comparing Relaxation Mechanisms in Quantum and Classical Transverse-Field Annealing, *Physical Review Applied* **15**, 014029 (2021).
- [151] T. Albash, W. Vinci, A. Mishra, P. A. Warburton, and D. A. Lidar, Consistency tests of classical and quantum models for a quantum annealer, *Physical Review A* **91**, 042314 (2015).
- [152] T. Albash, T. F. Rønnow, M. Troyer, and D. A. Lidar, Reexamining classical and quantum models for the D-Wave One processor, *The European Physical Journal Special Topics* **224**, 111 (2015).
- [153] S. Kumar and P. Majumdar, A travelling cluster approximation for lattice fermions strongly coupled to classical degrees of freedom, *The European Physical Journal B* **50**, 571 (2006).

DATA AVAILABILITY

A limited set of experimental and numerical data is available in the Zenodo online repository <https://doi.org/10.5281/zenodo.10709612>.

ACKNOWLEDGMENTS

We are grateful to the technical and non-technical staff at D-Wave, without whom this work would not have been possible. We thank L. Cincio, W. Zurek, V. Martin-Mayor, S. Boixo, A. Potter, D. Huerga for insightful and helpful discussions. We also thank the ITensor community—especially M. Fishman and K. Pierce—for providing flexible and efficient tensor-network methods with support for GPU acceleration. Work at UBC was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) Alliance Quantum Program (Grant ALLRP-578555), CIFAR and the Canada First Research Excellence Fund, Quantum Materials and Future Technologies Program. 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. G.A. acknowledges support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center. This research was supported in part by the National Science Centre (NCN), Poland under projects 2019/35/B/ST3/01028 (J.D.) and 2020/38/E/ST3/00150 (M.M.R.). R.G.M. was supported by NSERC. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade. J. Carrasquilla acknowledges support of NSERC, Compute Canada, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/#partners.

Supplementary Materials: Computational supremacy in quantum simulation

CONTENTS

I. Quantum annealing	13
A. Quantum annealing processors	13
B. Annealing schedules	13
1. Annealing time as a function of energy scale	13
C. Calibration refinement	14
1. Flux offsets	15
2. Anneal offsets	15
3. Embedding and automorphism averaging	15
D. Embedding in qubit graph	16
II. Experimental testbed	16
A. Definition of inputs	16
1. Square	17
2. Cubic dimer	17
3. Cubic no-dimer	17
4. Diamond	18
5. Biclique	19
III. Tensor networks	19
A. Matrix Product States (MPS)	20
1. Initial state preparation	21
2. Time-dependent variational principle (TDVP)	21
3. Classical computational platform limitations	22
4. MPS ordering	22
5. Tree tensor networks and multi-scale entanglement	23
B. Projected entangled pair states (PEPS)	23
1. Contraction of lattice structured PEPS	24
2. Circuit model and $d+1$ -dimensional PEPS contraction	25
3. Locally tree-like approximations	27
4. Simple update	28
5. Non-canonical MPS in 1D	28
6. 2D PEPS	31
7. 3D PEPS	32
IV. Measuring the quality of sampled distributions	34
V. Dynamics of entanglement entropy	36
VI. Comparing and evaluating QPU and MPS experiments	37
A. Estimating entanglement entropy	37
B. Evaluating ground-truth convergence	37
C. Extrapolation of bond dimension	38
D. Effects of time step and stopping point	39
E. Oblong 2D lattices	40
F. Area-law scaling for slower quenches	41
VII. Neural networks	41
VIII. Classical compute resources	43
IX. Critical exponents and dynamic finite-size scaling ansatz	44

X. Monte Carlo dynamics	47
XI. Spatial decomposition techniques	49
A. Cluster approximations with MPS	49
B. Cluster approximations with the QPU	50
1. Two dimensions	51
2. Three dimensions	53
C. Classical extrapolation of statistics from local to global	53

I. QUANTUM ANNEALING

A. Quantum annealing processors

Experiments were performed on two models of D-Wave quantum annealing processor: a prototype Advantage2™ processor (ADV2) and an Advantage™ processor (ADV1, online for general access as “`Advantage_system4.1`” during experiments). Both processors were maintained at a cryostat set point of 12 mK. ADV1 was used in a recent demonstration of coherent annealing of 3D spin glasses²⁴. ADV2 was fabricated using a lower-noise process and a different qubit connectivity family (Zephyr⁶²) than ADV1 (Pegasus⁶³). ADV2 and ADV1 had, respectively, 1222 and 5627 qubits; the qubit connectivity graphs, which dictate the available programmable qubits and couplers, are shown in Fig. S1.

B. Annealing schedules

The QPUs use radio-frequency superconducting quantum interference device (rf-SQUID) flux qubits to approximately realize ideal spin-1/2 particles in a TFIM⁶⁴. Each QPU has its own annealing schedule consisting of time-dependent transverse (x -basis) and longitudinal (z -basis) energy scales $\Gamma(s)$ and $\mathcal{J}(s)$. In a previous work²⁴ we performed detailed modeling of the effective many-body qubit Hamiltonian and mapped it to an effective TFIM schedule for ADV1, shown in Fig. S2A and Fig. 1. We use this ADV1 schedule as the model quench schedule, which defines the parameters of the classical methods used. We use both ADV1 and ADV2 to simulate this quench.

The ADV2 schedule (Fig. S2b) differs significantly, most importantly by having much higher energy scales than ADV1 in the region where $\Gamma(s) \approx \mathcal{J}(s)$. Simulating the ADV1 quench with the ADV2 QPU therefore requires parametric adjustments as described below.

1. Annealing time as a function of energy scale

The energy scales in the annealing schedule dictate the time scales relevant to Schrödinger evolution. QPUs are calibrated to perform an anneal on a requested timescale, $\hat{t}_a$ (an input to the control electronics), which is well controlled in the generally available operational regime of the device ($\hat{t}_a > 500$ ns). However, deviations from this nominal time are significant for the model quench times t_a studied here, and we make a refined mapping of requested time to model time as discussed in this section. Synchronization of s —the normalized time t/t_a —for all qubits also becomes challenging in the small anneal-time limit, for which correction methods are discussed in Appendix IC2. We remark that any future availability of $t_a < 500$ ns will have properly synchronized annealing lines and translated t_a , making this section unnecessary for future users of these QPUs.

We construct a mapping from a given model quench time, t_a , to a pair of QPU parameters: a nominal QPU quench time $\hat{t}_a$ and an energy scale κ (a multiplicative prefactor on $\mathcal{H}_P$ in Eq. 2). This approach is preferred over the more obvious approach of changing only $\hat{t}_a$ because of a technicality in this regime: due to the specifics of controlling waveforms at this speed, the effective quench time t_a is not a smooth or monotone function of the nominal (i.e., requested) quench time $\hat{t}_a$. It is, however, a smooth and monotone function of κ .

To calibrate time, we simulate an ensemble of 20 2D spin glasses on 4×6 qubits, which is small enough to routinely determine accurate ground-truth correlations for many quench times. We determine effective quench times t_a for a range of $\hat{t}_a$ and κ pairs by lining up a statistic that increases with t_a : the squared spin-glass order parameter $\langle q^2 \rangle$, Eq. 3.

We measure $\langle q^2 \rangle$ for $\kappa \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ and a dense selection of nominal times $\hat{t}_a \in \mathcal{T}$, along with smaller values of κ down to 0.01 for the fastest $\hat{t}_a$, 3.6 ns. For each $\hat{t}_a \in \mathcal{T}$, we perform a polynomial fit between $\log \kappa$ and $\log t_a$. This gives us, for each $\hat{t}_a$, an invertible mapping $f_{\hat{t}_a}$ from κ to effective t_a . It only remains to choose,




FIG. S1. **Qubit connectivity graphs.** **A**, ADV1 and **B**, ADV2 qubit graphs. Gray points and lines indicate operable qubits and couplers respectively; red points and lines indicate inoperable qubits and couplers. Embeddings of example systems are shown (zoomed, **C** and **D**): cubic dimer, $4 \times 4 \times 4$ dimers in orange, biclique $K_{24,24}$ with 48 dimers in purple. Diagonal couplings are inter-dimer; horizontal and vertical couplings are intra-dimer.

for t_a , a canonical nominal QPU $\hat{t}_a$ to use. We choose the fastest $\hat{t}_a$ in $\mathcal{T}$ such that the appropriate energy scale $\kappa = f_{\hat{t}_a}^{-1}(t_a) \leq 1$. The resulting time-calibration mappings are shown in Fig. S3.

C. Calibration refinement

As mentioned, the QPUs used in this work are subject to a baseline calibration appropriate to the operational regime $\hat{t}_a > 500$ ns. Thus we do some additional work to refine the calibration for the relevant experimental regime




FIG. S2. **Annealing schedules for ADV1 and ADV2.** **A** and **B** show effective TFIM schedules for Advantage and Advantage2 systems, respectively. ADV1 annealing schedule is used to characterize the modeled quench, but ADV2 is used for main measurements, since the Advantage2 processor has lower noise and higher energy scale. The annealing parameter s is a normalized time: $s = t/t_a$.

$3.6 \text{ ns} \leq \hat{t}_a \leq 27 \text{ ns}$. All these methods have been applied effectively in previous works, especially Refs. 23,24, and as detailed in Ref. 65.

To correct the calibration we use only symmetries of the input models. For example, due to a lack of symmetry-breaking longitudinal field in the models studied, each qubit should have average magnetization 0 in all experiments. Further, the sampling distribution should be independent of the programming batch and the embedding; i.e., placement on the processor. This allows us to refine parameters (flux offsets and anneal offsets) by self-consistent criteria.

Calibration refinement of effective time t_a and anneal offsets are independent of the programmed spin-glass topology and coupling distribution. Calibrations of flux offsets restore trivial model symmetries. Hence these calibrations do not create any model-dependent computational advantage or overfitting relative to classical simulation techniques where such calibrations also exist by definition.

1. Flux offsets

All inputs studied in this work have no linear biases in the nominal Hamiltonian, and as a consequence $\langle \sigma_i^z \rangle = 0$ for well-calibrated experiments. This symmetry is respected in ideal flux qubits, but is broken by flux drift and cross-talks that can be corrected on a per-qubit basis with compensating applied fields (flux offsets). Magnetization is smooth and monotonic in the flux bias for our models and dynamics. We calibrate flux offsets to achieve approximately zero magnetization on every qubit by an iterative method for each experiment (ensemble, $\hat{t}_a, \kappa$)⁶⁵.

2. Anneal offsets

The QA Hamiltonian is controlled using several annealing lines (eight in ADV1, six in ADV2) that must be well synchronized. Calibrated synchronization has a typical tolerance of $\approx 1 \text{ ns}$, which is suited to anneals longer than 500 ns but too large to ignore in our experimental regime. We therefore offset the qubits on each line, either advancing or delaying them depending on observed spin-spin correlations in a 512-qubit chain, as in Ref. 23. We determine these per-line anneal offsets for each $\hat{t}_a$ and κ corresponding to the studied model quench times t_a , and use them throughout all experiments across all problem ensembles.

3. Embedding and automorphism averaging

When possible, we embed multiple copies (up to 200) of each problem into the QPU in parallel. Typically around half of the qubits are used, in total. For example, in ADV1 and ADV2 we use 87 and 23 parallel embeddings, respectively, of 6×6 square lattices.

The unweighted graphs corresponding to the lattices studied in this work have rich automorphism groups. On each programming of the QPU, we apply a random automorphism to the graph, which changes the mapping of model spins to physical qubits. This allows averaging out of calibration imperfections even when only one embedding is available.




FIG. S3. **Calibrating QPU parameters by quench time for ADV1 (top) and ADV2 (bottom) QPUs.** **A, C,** effective model quench time t_a is calculated for a range of energy scales κ and nominal quench times $\hat{t}_a$ (legends); t_a is taken as the median effective quench rate (matching $\langle q^2 \rangle$) over a 4×6 2D spin glass ensemble. **B, D,** for each t_a we choose an appropriate $\hat{t}_a$ and determine the appropriate κ based on a power-law fit in κ based on the data shown in the left-hand panels.

D. Embedding in qubit graph

Any problem must be mapped to the available coupling topology of the quantum annealing processor. The problem of mapping an input to the qubit graph is an instance of subgraph isomorphism, rather than minor embedding. We use the powerful Glasgow Subgraph Solver⁶⁶ for this mapping. Examples of embedding with Glasgow Subgraph Solver, with code, can be found in Ref. 65.

II. EXPERIMENTAL TESTBED

A. Definition of inputs

The Ising spin glasses studied are defined for 2D, 3D and high-dimensional graphs. These can be subdivided into models with dimers (biclique, cubic dimer), and those without. In the dimerized cases a subset of edges defined by a




FIG. S4. **Square lattice.** Largest cylinder embedded on ADV2, 18×18 , shown.

perfect matching are strongly ferromagnetically coupled ($J = -2$). To all non-dimer edges we assign uniformly and identically distributed random couplings J_{ij} , subject to a technicality in the cubic dimer lattice, described below. The inputs we study are “low precision” and “high precision” with respect to the random J values.

Low-precision inputs have random couplings drawn from $\{-1, 1\}$, while high-precision instances have couplings drawn from $\{-\frac{128}{128}, \dots, \frac{-1}{128}, \frac{1}{128}, \dots, \frac{128}{128}\}$. In the special case of bicliques, these couplings are further subject to a size-dependent rescaling, described below. All edges in the graph are assigned couplings from the same distribution.

Note that low- and high-precision ensembles differ not only in precision, but also effective energy scale—a lower energy scale alone (in the high-precision problem) can account in part for a delayed rate of correlation-length growth, but is not a sufficient explanation for the variation between the two ensembles.

Here we define five input topologies: Square, cubic dimer, cubic no-dimer, diamond, and biclique. All models, dimerized or not, are subgraphs of the QPU architecture and embed as a one-to-one mapping of lattice sites to qubits. Cubic no-dimer, as well as all low precision problems, are omitted from the main text for brevity.

1. Square

The 2D lattices we study are $L_x \times L_y$, periodic in the x -dimension (Fig. S4). There are no dimers.

2. Cubic dimer

Cubic dimer inputs (simply called “cubic” in the main text), which were studied extensively in previous work²⁴, are constructed starting with a spin glass on a simple cubic lattice of size $L_x \times L_y \times L_z$. Each cubic lattice site, with coordinates x, y, z , is expanded into a dimer of two spins $s_{x,y,z}$ and $s'_{x,y,z}$. Couplings in the x -dimension are placed incident to $s_{x,y,z}$, and couplings in the y -dimension are placed incident to $s'_{x,y,z}$. In the z -dimension, the coupling is divided equally across two couplers; for example, if the cubic lattice spin glass has a coupling of -1 between sites x, y, z and $x, y, z + 1$, the cubic dimer spin glass will accordingly have a coupling of $-1/2$ between $s_{x,y,z}$ and $s'_{x,y,z+1}$ and a coupling of $-1/2$ between $s'_{x,y,z}$ and $s_{x,y,z+1}$.

All cubic dimer inputs have open x - and y -boundaries, and are periodic in the z dimension if and only if $L_z > 2$. The $6 \times 6 \times 6$ -dimer system is shown (without periodic z -couplings) in Fig. S5 (left).

3. Cubic no-dimer

Cubic no-dimer inputs are derived from cubic dimer inputs by making two changes:

1. one coupling is omitted in the z -dimension: the one between $s'_{x,y,z}$ and $s_{x,y,z+1}$.
2. the coupling between $s_{x,y,z}$ and $s'_{x,y,z}$ is no longer a dimer.

The $6 \times 6 \times 6$ -dimer system is shown (without periodic z -couplings) in Fig. S5 (right).




FIG. S5. **Cubic dimer and no-dimer lattices.** Left, cubic dimer, and right, cubic no-dimer lattices with $(L_x, L_y, L_z) = (6, 6, 6)$ on 424 qubits. In cubic dimer problems, thick gray lines indicate dimer couplings ($J = -2$), dark lines indicate x - and y -dimension couplings, and light yellow lines indicate z -dimension couplings, in which a coupling term is distributed equally between two couplers joining two dimers. In cubic no-dimer inputs (right), all couplers are drawn from the same distribution, and half of the z -couplings from the cubic dimer connectivity are discarded. The disordered “no-dimer” couplings prevent the spins from being paired into approximately two-level objects. Periodic couplings between top and bottom not shown.




FIG. S6. **Diamond lattice.** The largest input studied in this work is a diamond lattice with dimensional parameters $(12, 12, 16)$, in which 567 of 576 qubits (98.4%) are operable. Inoperable qubits not shown. Periodic couplings between top and bottom not shown.

4. Diamond

The diamond lattice, like the square lattice, is 4-regular. We parameterize these inputs with dimensions L_x , L_y , and L_z , which do not correspond exactly to Euclidean coordinates of diamond lattice sites. Rather, assuming an edge length of 1 and in the limit of large systems, the Euclidean dimensions of an input of size $L_x \times L_y \times L_z$ are $\frac{L_x}{\sqrt{3/2}}$, $\frac{L_y}{\sqrt{3/2}}$, $\frac{L_z}{\sqrt{3}}$, and such an instance has only $N = L_x L_y L_z / 4$ spins, not $L_x L_y L_z$ spins. For the purposes of finite-size scaling, we take L to be $N^{1/3}$. The inputs studied here are open in the x - and y -dimensions and periodic in the z -dimension; this requires L_z to be divisible by 4. The maximum values of L_z for inputs embeddable in ADV1 and ADV2 are 8 and 16, respectively. The largest system we study has $(L_x, L_y, L_z) = (12, 12, 16)$ (Fig. S6), embedded in ADV2 with 9 defects (site vacancies) out of 576. The largest diamond lattice is the only case for which we permit defects; due to their small number these are expected to have no significant impact on algorithmic complexity of simulation.




FIG. S7. **Biclique lattice.** Shown is a $K_{24,24}$ topology on 96 spins, where gray lines indicate $J = -2$ coupled dimers, and purple lines indicate inter-dimer couplings, forming a complete bipartite structure.

5. Biclique

Biclique inputs in this work are based on complete bipartite graphs $K_{M,M}$ embedded using two-spin dimers for each biclique site. We may call the bipartition of dimers A and B , and the dimers in this bipartition $A_1, \dots, A_M$ and $B_1, \dots, B_M$. Between any A_i and B_j there exists exactly one coupler. Since the number of random couplings scales as M^2 it is necessary to rescale the coupling strength to ensure extensive energy (correct scaling behavior)⁶⁷; we rescale couplers as $J_{ij} \rightarrow \frac{J_{ij}}{\sqrt{M/4}}$ for every i in A and j in B . The largest input studied is $K_{24,24}$ (Fig. S7); an embedding is shown in Fig. S1.

III. TENSOR NETWORKS




FIG. S8. **MPS and PEPS for a small cylindrical TFIM.** A 4×4 cylindrical transverse-field Ising model may be represented diagrammatically with spins or qubits depicted by circles and couplers depicted by lines (left). To represent a quantum state associated with such a lattice we may use a tensor network (MPS or PEPS). The above diagrams depict contracted networks for purposes of amplitude calculation, as in Eq. (S1). Tensors are represented by squares, lines adjoining two tensors imply a contraction, and free lines depict unsummed indices (in this example, the qubit states). In an MPS (center), the tensors form a path with three indices (lines) per tensor (except on the first and last tensor where there are two), one index or line being of dimension 2 in the depicted case (matching the state-space size of a qubit). In our canonical ordering on a square lattice the tensor contraction circles around the lattice in the periodic dimension, minimizing pathwidth on the target TFIM. Local entanglement is captured indirectly through paths with longer-range connections, but contraction is efficient. In a PEPS (right), one dimension (bond) can be assigned for each non-zero coupling in the TFIM, capturing local full entanglement, but exact contraction is not efficient. In practice, tensor contractions for purposes of time evolution and estimation can be performed iteratively over local subsets, or effective approximate tensor environments, avoiding creation of large objects (many unpaired dimensions).

Tensor networks are considered the state of the art for simulation of quantum states and dynamics in low-dimensional quantum systems^{26,40–44,47}. Several forms of tensor network state exist, suited to different applications. Sections III A and III B respectively describe matrix product states (MPS)—and certain generalizations—and projected entangled pair states (PEPS), and discuss details of our implementations and their applicability to our problems.

The area law of entanglement and Kibble-Zurek scaling of the correlation length limit the performance of tensor network methods on lattices. In MPS, the area law of entanglement manifests as exponential scaling of the bond dimension $\chi \propto \exp(N^{(d-1)/d})$, where N is the number of qubits and d is the dimension. Memory, time, and energy requirements scale polynomially in this parameter. In PEPS, certain types of area law are simulable with finite bond dimension, but the numerical requirements can grow prohibitively with an increasing range of correlations.

A. Matrix Product States (MPS)

MPS or “tensor train” networks are ansätze for efficient factorization of high-dimensional tensors, including quantum states of many-qubit systems^{40,68–70}.

For an N -qubit system, an MPS may be defined in terms of N tensors of the form $T_{l_i, z_i, l_{i+1}}^i$. The indices $\{l_i\}$ label virtual indices (bonds), and z_i physical ones, with dimensions $\dim(l_i) \leq \chi \ \forall i$ and $\dim(z_i) = p$. In general p scales exponentially in the number of qubits associated with the tensor, which in all our experiments is constant; in the standard representation with one tensor per qubit, $p = 2$. For a wave function $|\psi\rangle = \sum_z a_z |z\rangle$, amplitudes in the computational basis can be represented as

$$a_{z_1, \dots, z_N} = \sum_{\{l\}} \prod_{i=1}^N T_{l_{p(i)}, z_i, l_{\pi(i)+1}}^{\pi(i)} \quad (\text{S1})$$

where bond indices $\{l_*\}$ are summed over the corresponding bond dimension^I and $\pi(i)$ is the ordering of the qubits in the MPS. A graphical representation of the MPS contraction, Eq. (S1), is shown in Fig. S8. It is important to point out the efficiency of the MPS representation: while the space of amplitudes is of size 2^N , the space of tensors is of size $O(Np\chi^2)$, controlled by the maximum bond-dimension parameter χ .

The quality of an MPS (ordering) can be understood in terms of entanglement entropy. Consider a system divided into two subsystems, A and B , at some bond. If ρ_A and ρ_B denote the reduced density matrices for these subsystems, the entanglement entropy is defined as

$$S_A = -\text{Tr}_A(\rho_A \log \rho_A) \quad (\text{S2})$$

The reduced density matrix ρ_A possesses 2^{N_A} diagonal elements, where N_A is the number of qubits in subsystem A . The maximum possible entropy corresponds to the extreme case of a uniform probability distribution: $\max(S_A) = N_A \log 2$. This means it scales with volume of the subsystem. For a random state in the Hilbert space this scaling is also expected for any subsystem and is sometimes referred to as the volume law of entanglement entropy. However, in systems with local connectivity and a gapped ground state, the entanglement entropy of the ground state exhibits a linear scaling with the number of qubits at the boundary between the two subsystems. This phenomenon is commonly referred to as the area law of entanglement⁴⁹: $S_A \propto$ bipartition area. As discussed in Section V, and in agreement with previous observations³⁰, we show that quenched quantum dynamics deviating from the ground state also show area-law scaling.

In open-boundary one-dimensional systems, optimal MPS ordering (for purposes of minimizing bond dimension) is typically aligned with the 1D topology. Such MPS states naturally encode the area law of entanglement⁷¹: ground states of gapped Hamiltonians of N qubits can be approximated to arbitrarily high precision with an MPS of bond dimension χ , which does not grow with the size of the system. Therefore, the entanglement entropy of a block of L qubits constructed out of an MPS state, $S(L) = -\text{Tr}[\rho_L \log(\rho_L)] = O(\log \chi)$, with ρ_L being the density matrix of the block, is bounded by a constant $S(L) \propto \text{const}$ for $L \gg 1$.

Choosing an optimal MPS ordering for higher-dimensional systems is ambiguous. In spite of this, geometrically intuitive orderings can be defined that are optimal from a scaling perspective, as discussed in Section III A 4. For these orderings, area-law compliant states of d -dimensional systems have $S(N) \propto N^{(d-1)/d}$, requiring the MPS bond-dimension $\chi \sim \exp(S(N))$ to grow exponentially in the area of the lattice. Our experiments present evidence that coherently-annealed quantum states obey an area law of entanglement scaling. Hence, although the MPS representation is more concise than the full representation, it cannot be reduced further without incurring systematic errors.

^I States near the boundary can be represented with reduced bond

dimension, for example $\chi_1 = \chi_N = 1$ without loss of generality.

MPS representations can also be modified by representing multiple physical qubits per tensor. Qubits composing dimers in our models exhibit stronger correlations than other qubit pairs, and it can be efficient to represent them in a single tensor—in this case the prescription Eq. (S1) applies but with $p = 2^2$, reflecting the expanded state space. By extension, we might collapse a d -dimensional system to a $d - 1$ -dimensional system by collecting together $O(N^{1/d})$ qubits into a single state. Since p grows exponentially in the number of qubits collected together, such a method cannot be used to bypass area-law scaling constraints, and typically leads to scaling inefficiencies. These collection methods may lead to some modest practical speedups beyond the benefit of simply processing dimer pairs consecutively in the tensor train. Aside from dimers, a further example of multi-qubit tensor efficiency is examined in Section III B 6 for PEPS.

1. Initial state preparation

If the dynamics were initialized with the ground state of only the driver Hamiltonian $\mathcal{H}_D$, this would trivially be a product state with $\chi = 1$. Time-evolving an MPS with $\chi = 1$ with the time-dependent variational principle (TDVP, see next subsection for more details) is ill-defined, as many of the long-range interactions are lost during the projection step and the resulting time evolution would be wrong^{50,72,73}. To avoid this problem, we run a ground-state Density Matrix Renormalization Group (DMRG) algorithm on the full system Hamiltonian at $s = 0$ to estimate the initial state⁷⁴. This is performed by setting the truncation error threshold to zero and keeping the state at a bond dimension χ larger than one.

For initialization and dynamics of MPS using the TDVP method we rely primarily on the `ITensor`⁷⁴ library. We also tested the dynamics using the local-Krylov approach for time-evolution as implemented in the `DMRG++`⁷⁵ library, but found the TDVP superior in terms of running time and precision for larger time steps. For the GPU implementation of the TDVP method, we have mainly used the `ITensor` library and adapted the GPU backend to both Nvidia and AMD GPUs as used in the Summit and Frontier clusters, respectively. Our primary platform for PEPS experiments is the `yastn` library⁷⁶.

2. Time-dependent variational principle (TDVP)

For MPS time-evolution we employed a standard time-dependent variational principle algorithm with a two-site update, with second-order Trotterization error^{50,74,75}. Unless stated otherwise, our presented results rely upon a time-step of $dt = 0.01$ ns, and a discarded-weight threshold for the singular value decomposition (SVD) of 10^{-12} . We have checked that this value of the time-step gives converged ground-truth results for all the different topologies studied.

An important comment is in order about our use of the two-site versus one-site update for the TDVP algorithm. Refs.^{77,78} proposed a scheme to dynamically enlarge the bond dimension within a one-site update, which is shown to perform better than a standard two-site TDVP algorithm. We expect a factor of 2-3 speedup when this refinement is implemented. Nevertheless, it is clear that these improvements won't change the exponential scaling of computing costs, which will eventually become classically prohibitive for the system sizes currently achievable with the QPU (see also Section III A 3).

Continuing the discussion about details of our implementation of the TDVP method: for the solution of the forward and backward time-dependent Schrödinger equations for the local MPS and center matrix tensors, we keep up to 100 Krylov vectors, allowing a Krylov truncation error of 10^{-12} . Computational time also scales in proportion to $n_t = t_a/dt$, which is for practical purposes significant, since we consider evolutions on time scales of order 10 ns; values of dt larger than 0.01 ns quickly lead to large systematic errors in many simulations. The truncation threshold has little impact on the statistics or running time in the small-value limit (small χ and large N) in which we operate; increasing the truncation threshold can allow efficiencies through the early-time evolution of states where tensors need not saturate the bond-dimension bound, but this phase of simulation is not a computational bottleneck. We found that more careful optimization of these quantities did not impact our conclusions qualitatively—in other words, their impact is insignificant compared to the choice of χ .

The dominant contribution to computing time is the choice of maximum bond dimension χ , which is applied uniformly to all bonds and all times, and is saturated for a large finite fraction of tensors in space and time, during our simulation. Memory requirements scale as $O(N\chi^2)$, and time requirements scale as $O(Nn_t\chi^3)$, owing to SVD and Krylov iterations that must be applied to all bonds and at every time step of the algorithm.

MPS errors are systematically controlled by bond dimension χ and dynamical discretization dt . Our statistics converge smoothly at large χ and small enough dt , and it is possible to infer a ground truth. By reducing χ , it is possible to obtain a computationally less expensive, but less accurate distribution.

3. Classical computational platform limitations

Computational platforms used for MPS simulations include CPU, GPU and tensor processing unit (TPU)⁷⁹ implementations. Our walltimes and energy consumption estimates of Figures 4C–4D of the main text are based on an implementation of TDVP algorithm on a single GPU node with a two-site update, which by definition necessitates an SVD of joined consecutive MPS tensor pairs to implement the MPS tensor truncations.

We performed this SVD using the divide and conquer algorithm, which is notoriously difficult to implement efficiently on GPU architectures. In this regard, Li et al.⁸⁰ recently proposed a QR decomposition to eliminate the SVD step of the time-evolving block decimation (TEBD) algorithm (not for the TDVP method), demonstrating substantial speedups on single GPU architectures. Furthermore, a real-space parallelized version of the TDVP algorithm has been recently proposed in Ref. 81, which might offer further performance speedups. We believe that combining all these approaches together with a multi-GPU or TPU⁷⁹ implementation of the two-site (or better one-site update with bond controlled expansion) TDVP scheme will give the possibility of benchmarking the QPU against even larger system sizes both in 3D and biclique lattices, and that our time and energy estimates could likely be reduced by one order of magnitude. Nevertheless, even including this renormalization factor, we stress that current TPUs are optimized to fit at most a $\chi = 2^{16} = 65536$ bond dimension⁷⁹, while for a 3D or biclique lattice our work estimates that a bond dimension of $\chi_Q \simeq 10^7$ is required to match the QPU error. For diamond lattices, the estimated requirement grows to $\chi_Q \simeq 10^{11}$.

We conclude that quantum simulations of coherent anneal quenches with sufficiently large annealing time t_a on systems with several hundreds or thousands of qubits are prohibitive for current MPS-TDVP approaches on existing GPU or TPU architectures.

4. MPS ordering

The efficiency of an MPS representation, and its manipulation in TDVP, depends critically upon tensor ordering. For one-dimensional systems with open boundaries, the optimal path matches the topology. In higher dimensions, it is generally hard to find the optimal MPS ordering to represent a generic quantum state. In this work, we resort instead to geometric intuition, and have tested this intuition against some heuristic methods. Standard orderings that have proven successful for 2D align with our heuristic method^{26,41,47,82}.

Ordering in an MPS becomes more efficient when strongly-correlated sites are close in ordering. Directly coupled qubits are more entangled (on average), as are qubits coupled by short coupling paths (short geometric distance), in line with the intuition that there will be a path length beyond which one might assume weak correlation (a correlation length).

We can formalize this intuition with a heuristic objective function to be minimized, a quadratic assignment problem with a cost

$$Q(p) = \sum_{i,j} F_{ij} R(\pi(i), \pi(j)) , \quad (\text{S3})$$

where R is a distance between qubits under a particular ordering choice, and F is a flow (‘correlation strength’) between those qubits.

Minimizing the cost under various intuitive definitions of flow and distance remains hard, but an approximate solution can be efficiently obtained and add value relative to an uninformed MPS ordering. This problem can be mapped to an unconstrained quadratic model, and naturally, we find this is amenable to heuristic solution by the QPU through a LeapTM hybrid solver as one option.

A special intuitive case of the measure takes $R(x, y) = \|x - y\|_\infty$ (maximum distance between tensors) and $F_{ij} = 1 - \delta(J_{ij}, 0)$ (a penalty for coupled qubits, no penalty for uncoupled qubits). I.e., we minimize how far apart coupled qubits are in the tensor train. The objective is then equivalent to the graph-theoretic problem of pathwidth minimization, which is also relevant to tensor decomposition in other contexts⁸³.

For results presented in this paper we have used orderings that aim to minimize pathwidth^{II}. The MPS ordering used for the square lattice covers first the periodic dimension and then the open dimension, creating a spiral path along the cylinder (an example is shown in Fig. S8). Similar “one-dimension at a time” patterns apply to the diamond

^{II} A formal proof of minimal pathwidth is challenging in some topologies; we settle for simple scalable patterns that are prov-

ably optimal to within $O(1)$ corrections (as N becomes large), in some cases provably optimal.

and cubic lattice cases, proceeding first over dimer pairs where relevant. For bicliques we proceed first over half of one partition, then the entire other partition, then the remaining half of the first partition. A feature of these orderings is that the partitions define lattice interfaces that scale in worst case as $N^{(d-1)/d}$. I.e., bonds define partitions, and these partitions typically define $d-1$ -dimensional planes with respect to the lattice topology. As such, we do the best job possible in minimizing the entanglement entropy that a bond must support (in a scaling sense), since we cannot do better than $d-1$ -dimensional partitions (but can do worse).

It is expected that we might obtain efficiencies by working harder to determine an efficient MPS ordering. To test this we have used several forms of R (1,2, and infinity norms) combined with several definitions of F ($1 - \delta(J_{ij})$, $|J_{ij}|^2$, and mutual information based on simulated states) incorporating disorder-realization specific information. The computational savings made by such optimizations were not reliable, and did not show clear evidence of scaling advantage. We also investigated the use of a Hilbert-curve ordering on the square lattice. This ordering has been used to improve MPS efficiency to solve the quantum Ising model on square lattices, but while the ordering better preserves locality and has a lower average neighbor distance than the pathwidth-minimizing order, it comes at the cost of a maximum $O(N)$ neighbor distance (as opposed to the pathwidth-minimizing case of $N^{1/2}$)⁸⁴. We found the Hilbert-curve ordering used with TDVP produced higher correlation error across all disorder realizations tested, as well as worse scaling. In our final analysis of time we used the simpler pathwidth-minimizing (lattice specific, but disorder-realization independent) MPS orderings.

5. Tree tensor networks and multi-scale entanglement

Tree tensor networks (TTN) are a natural generalization of MPS, maintaining qualitatively similar properties in our application^{85,86}. Virtual bonds define a tree rather than a path. Canonical forms can be constructed by iterating from the leaves, at a cost scaling as χ^{k+1} , where k is the maximum connectivity of any node ($k \geq 3$ in contrast to $k = 2$ for MPS). The space of TTN is a superset of the space of MPS (as defined by orderings), thus we might expect some gain in optimization of the TTN to the problem (lattice and disorder realization) along the lines explored in Sec. III A 4.

It is possible that the optimal TTN choice is not an MPS and that performance gains are sufficient at $k > 2$ to overcome the polynomial disadvantage in χ scaling (plus any implementation-specific complications). However, given that there is no obvious tree-like substructure to a homogeneous spin-glass lattice problem, there is no obvious choice for a TTN. As in MPS, bonds in TTN define partitions, and it is not possible to define trees with finite k such that the worst partition scales better than $N^{(d-1)/d}$. Thus given area-law entanglement scaling, a tree tensor network cannot provide a scalable advantage relative to an MPS. Empirical investigations fall beyond the scope of this work.

The multi-scale entanglement renormalization ansatz (MERA) is a means to expand the scale of quantum entanglement captured without exponential growth in the contraction complexity⁸⁷. This is achieved by combining a TTN with disentanglers constrained to an isometric (canonical) form for efficient contraction. This has been shown to effectively capture scalable correlation structures that describe one-dimensional critical states, in line with renormalization-group scale invariance that is present by construction. MERAs have been specialized for use in the analysis of ground states for 1D systems with translational invariance in the large-system limit. Although a time-dependent MERA algorithm has been proposed and successfully applied to translationally invariant 1D systems⁸⁸, its unfavorable computational scaling $O(\chi^9)$ makes a generalization to our finite, higher dimensional disordered models impractical.

B. Projected entangled pair states (PEPS)

Projected entangled pair states (PEPS), also called tensor product states, are a generalization of matrix product states to arbitrary graphical structure^{42-44,89,90}. For some states subject to geometrically constrained patterns of entanglement and correlation, in particular a subset of those exhibiting area-law entanglement, PEPS can provide a memory-efficient representation relative to MPS when virtual indices (bonds) are chosen to match the target topology^{48,91}. Unfortunately the contraction of PEPS in 2D and higher is $\#P$ -complete⁵¹, and in practically interesting cases it relies on heuristic contraction schemes (lattice contraction is described in Sec. III B 1). Contraction is required for all inference purposes: the calculation of bond environments for purposes of (polynomial-scaling) Trotterized evolution, the evaluation of local expectations, and the calculation of state probabilities.

Efficient contractions have been achieved for some interesting lattice problems appearing in the context of quantum supremacy claims that are practically inaccessible to conventional MPS methods^{16,92}. Despite being a $\#P$ -complete problem in worst case, it is understood that structured states, in particular those controlled by a finite correlation

length, ξ , can allow efficiencies in scaling^{III}. Supposing a weak dependence of some environment (local observables) on the state (tensors) beyond some neighborhood scale (ξ), we can propose to contract a neighborhood of volume $O(\xi^d)$ surrounding the environment. The vertex congestion (and tree width) associated with contraction on such a neighborhood scales as ξ^{d-1} , and computational complexity scales exponentially in this quantity. In so far as the contraction neighborhood scale is linked to the Kibble-Zurek scaling of correlation length we might estimate scalable bounds on such approximations. Thus we can anticipate the failure of PEPS neighborhood approximations owing to a combination of dimensionality and correlation length.

In line with this intuition we show that PEPS is accurate for lattice evolution on very short time scales, where short correlation lengths apply. At longer times and for larger problem energy scales (correlated with connectivity and reduced precision), ξ grows and larger neighborhoods are required. We investigate a variety of practical heuristics and conclude that the systematic errors for PEPS are larger than those of the QPU in our more challenging experiments.

We present the fundamental problem of PEPS contraction for lattices in Section III B 1. A $d + 1$ dimensional PEPS representation, compatible with circuit models, is examined in Section III B 2. In Sections III B 3 and III B 4 we examine locally tree-like approximations and an associated simple update (and simple observable) algorithm. In Sections III B 6-III B 7 we consider more complex neighborhood update methods tailored to 2D square and 3D cubic lattices. The interaction of correlation length and neighborhood contraction scale is examined in the well-controlled context of 1D models in Section III B 5.

1. Contraction of lattice structured PEPS

Lattice PEPS are a generalization of MPS in which tensor virtual bonds are in alignment with lattice edges; an example is shown in Figure S8 for a square lattice. In this section we consider primarily the standard form of PEPS with one tensor per qubit (state dimension $p = 2$) fully entangled with neighboring tensors, and with a uniform bound on bond dimension (D). For the square and cubic lattices, we consider also tensors in alignment with plaquette and cubic states ($p = 2^4, 2^8$ respectively), as shown in Figure S14G. Use of finite spatially-local (plaquette or cubic base) tensors does not change qualitative dimensionality considerations, but can be practically useful as shown later. All tensors are of size $\leq pD^k$, k being the lattice connectivity.

The process of contraction involves taking inner products of tensors (contracting tensor-pairs to single tensors), creating sequences of intermediate tensors. Contractions are required for controlled truncation (compression) of PEPS, and for the calculation of observables. Successive multiplications can be represented by a binary tree in which one bond (adjoining two tensors) is eliminated to create a new tensor at each node proceeding inwards from the leaves to a root. Assuming inner products are the dominant computational operation, the complexity (memory and time) of the contraction is tied to the (bond-dimension-weighted) vertex and edge congestion of the best possible binary tree⁸³. If bond dimension is homogeneous, the scaling of floating-point operations (and memory) is equivalently understood as being exponential in the treewidth of the line graph⁹³. Line graphs of d -dimensional lattices are also (different) d -dimensional lattices, and have treewidth that scales as $N^{(d-1)/d}$. Therefore, even at very modest scales, heuristics are required for contraction.

A heuristic of primary interest to us involves a neighborhood approximation: rather than contracting over the full volume we might consider a symmetrical neighborhood of volume $O(\xi^d)$, so that contraction only scales as $\xi^{(d-1)/d}$ ⁹⁴. Another heuristic of interest is boundary-MPS (or boundary-PEPS), where a $d-1$ -dimensional PEPS (MPS in the case $d = 2$), defined at some boundary, is approximately contracted by a sequence of matrix-product operations with truncation⁴³. Whether for full, neighborhood or boundary-MPS/PEPS contractions, the details of the contraction process can be limited not just by dimensionality, but also the specifics of the lattice. It should be noted that practical (small finite D , χ_B and ξ) implementations can be dominated by costs other than the inner products, and many methods require iterative procedures that do not guarantee convergence and have a high propensity toward numerical instability.

Mirroring our MPS method, our PEPS methods separate the time evolution of a lattice-structured PEPS from the calculation of observables on the time-evolved PEPS. As we are interested in z -basis sampling, our more-costly numerical experiments are terminated at $s = 0.6$, where the transverse field becomes negligible (see Figs. S2A and S23).

^{III} The Kibble-Zurek mechanism anticipates a correlation length $\hat{\xi}$ growing with quench duration, and as a function of the problem class and energy scale; we have demonstrated various experimental results aligned with this intuition. Here the correlation length

might be associated with the Kibble-Zurek mechanism, but other phenomena can contribute to correlation length; the use is more generic.

Method	Lattice type and dimensions	Vertex congestion (complexity)
Tree-like (SU,SVDU)	single bond environment	$O(pD^{\text{connectivity}+1})$ (Ref. 95)
NTU(NN, NN+)	2×3 core, on square lattice; Fig. S14A,B	$O(D^8 + pD^7)$ (NN in Ref. 94)
NTU(NNN, NNN+)	3×4 core, on square lattice; Fig. S14C,D	$O(pD^{10})$
NTU(3DNN)	4 nearest loops, on simple cubic lattice; Fig. S17C	$O(D^{12} + pD^{11})$
NTU(3Dcube)	3x3x4 core, on simple cubic lattice	$O(pD^{29})$
Normalization	square	$O(pD^{2L+6}) [O(p\chi_B^2 D^8 + \chi_B^3 D^6)]$
MCMC	square	$O(D^{L+4}) [O(\chi_B^2 D^5 + \chi_B^3 D^3)]$
MCMC	cubic	$O(D^{L^2+L+5})$
MCMC	biclique $K_{h,v}$	$O(D^{hv/2+1})$ (Ref. 96)

TABLE I. **Vertex congestion (time complexity) scaling as a function of bond dimension D at large L .** Scaling is enumerated for: local environment approximations (NTU, SU) for time evolution; full contraction for amplitude estimation (MCMC sampling); and full bra-ket contraction (normalization, exact environment estimation or FU). Cubic and square lattice are assumed periodic in one dimension per main-matter experiments. Scaling of boundary-MPS methods for the square lattice are shown, for some choice of heuristic MPS bond dimension χ_B , in square brackets. Prefactors that scale as volume (and time) also apply to a practical usage, just as with MPS. Neighborhood-update costs, beyond simple update, have severe scaling as local loops are incorporated in higher-dimensional lattices⁵³. Tree-like decorations, or one-dimensional NTU-extensions⁹⁷, on neighborhoods can be added, raising only the prefactor scaling of contractions; the limited power of such decorations is explored in experiments.

Section III B 2 discusses an interpretation of our PEPS contraction in terms of a $d+1$ (space plus time) dimensional PEPS problem; from this perspective, the time-evolution element of the contraction can be interpreted as a boundary-PEPS method ($\chi_B = D$), with the boundary proceeding from time $t = 0$ to our final state at t_a . An initial state is efficiently prepared (as discussed in Section III A 1, the ground state of a driver Hamiltonian is already sufficient) and Trotterized time evolution applied. Coupled tensors are updated via action of a rank-2 (or 4) matrix product operator, as illustrated in Figure S9. This enlarged tensor pair is then compressed via heuristic methods to prevent exponential growth of the rank (memory requirements). Estimation of observables is required following the time evolution. We can again contract on some neighborhood for this purpose. The neighborhood might match that used in time evolution, for example in Section III B 4 a single-edge neighborhood is used (simple observable matching the simple update). In most places we consider a larger neighborhood—the full lattice, exploiting efficiencies (such as MCMC) for approximate contraction. Errors can arise either in the dynamics or in the final contraction of the time-evolved state. We find that the former results in more practically significant computational constraints.

We consider two types of heuristic truncation for time evolution: simple update (SU) and neighborhood tensor updates (NTU) at various scales. For final state measurement of observables, we consider boundary-MPS, exploiting where necessary the efficiencies of Markov-chain Monte Carlo sampling methods, as well as contractions limited to small-neighborhood scales.

A summary of contraction complexity for several lattices and contraction schemes is provided in Table I. Calculation of environments for purposes of heuristic truncation require a bra-ket contraction, and the complexity of these operations is a function of the neighborhood structure. For our neighborhoods, we can determine an optimal ordering, and denote the vertex congestion assuming homogeneous bond dimension. Whereas many methods for calculation of observables also require bra-ket local (or global) contraction, we rely primarily upon Markov-chain Monte Carlo (MCMC) sampling⁹⁸. An amplitude can be calculated with respect to only a bra (or ket) network by fixing the state legs; taking the absolute square of the amplitude yields an unnormalized probability. This achieves approximately quadratic speedup in contraction as shown in Table I. Provided correlation strengths are not too large, approximated probability ratios can be used as rejection criteria for an MCMC method for z -basis sampling, sufficient for estimation of the observables of interest.

While methods are well developed for square-lattice models, particularly ground-state estimation in translationally invariant models^{99,100}, the performance of practical approximations in three and higher dimensions are poorly understood in application to disordered models and diabatic time evolution. Even on the 3D simple cubic lattice, scaling of basic neighborhood bra-ket contractions (beyond locally tree-like approximations) can be as severe as $O(D^{29})$, thereby limiting bond dimension to $D \leq 2$ and hence reducing the support for entanglement⁵³.

2. Circuit model and $d+1$ -dimensional PEPS contraction

We formulate the 2nd-order Trotterized model of dynamics as a circuit model ($d+1$ -dimensional PEPS). We show that the scale of such circuits is well beyond the capability of any current circuit (gate)-model computer. Given a




FIG. S9. **A PEPS (circuit-model equivalent) representing a 2nd-order Trotterized approximation to time evolution of duration dt .** Repeated action of a dt -operator on a PEPS enables a discretized simulation of dynamics with error controlled by dt . The dt operator can be decomposed into $2 + \kappa$ layers of rank-2 unitaries (gates), where κ is the edge chromatic number of the lattice being simulated ($\kappa = 2 =$ site connectivity for the case of a 1D chain shown). 2-site operators are $\exp(-i2\pi dt \mathcal{J}(t/t_a) J_{ij} \sigma_i^z \sigma_j^z)$ and single-site operators are $\exp(-i\pi dt \Gamma(t/t_a) \sigma_i^x)$. The depth of 2-site operators is controlled by κ , and the total number of 2-gates matches the number of edges in the target lattice. To compress the circuit, 1-site operators between time steps can be combined. With care in the order of operations (to ensure 2nd-order Trotterization error), depth can be reduced to κ per unit dt by use of 2-site (rank 4) gates of the form $\exp\left(-i2\pi dt [\mathcal{J}(t/t_a) J_{ij} \sigma_i^z \sigma_j^z - \Gamma(t/t_a) (\frac{\sigma_i^x}{k_i} + \frac{\sigma_j^x}{k_j})]\right)$, k_* being site connectivity, eliminating the need for 1-site gates.

circuit model, we might also evaluate the complexity of contraction by circuit-model-inspired PEPS methods that have proved successful in undermining some supremacy claims^{15,16,18,92,101}. The PEPS formulation motivates boundary-PEPS (time-evolved lattice-PEPS) formulations, which we evaluate in detail in other sections.

Simulation of the time-dependent Schrödinger equation is a challenging problem that has been tackled by a variety of methods from the fields of physics and ordinary differential equations^{102–104}. Trotterization methods are commonly employed, and allow discretization errors that are controlled polynomially in dt , with a number of unitary operations that scales linearly in the number of Hamiltonian terms and simulation time^{IV}. Trotterized dynamics can be formulated as a PEPS with one tensor per unitary operator (gate); these gates are coupled in a discretized time dimension, so that a lattice-structured PEPS analysis becomes a problem of $d + 1$ dimension. The challenging task is to contract to the final time layer and to measure z -basis statistics. The PEPS can be interpreted as a circuit model. A high-fidelity circuit-model computer with an amenable structure, programmable gates, and z -basis measurement also suffices for dynamical simulation. Our classical methods, based upon the mid-point rule method in PEPS (2nd-order Trotterization) and 2-site 2nd-order TDVP methods in MPS, show that Trotterization with $dt \sim 0.01$ ns is sufficient for accurate simulation. It is therefore interesting to evaluate the circuit-model formulation with a view to either implementation on a hypothetical circuit-model quantum computer, or by PEPS heuristics.

Figure S9 demonstrates a PEPS formulation of 2nd-order Trotterized dynamics applied to a 1D model. Higher-dimensional models require additional 2-site unitary layers aligned with the edge chromatic number κ (in our lattices, matches the lattice connectivity). To perform a single time step requires at least κ layers of $|E|$ 2-qubit gates, where $|E|$ is the number of non-zero couplers in the problem Hamiltonian. Repeating the action of these operations $n_t = t_a/dt$ times creates a PEPS sufficient for simulation of dynamics.

Implementation of this PEPS, assuming suitable programmability of 2-gates matched to the target lattice architecture, could require as few as $|E|t_a/dt$ high-fidelity 2-site gates with depth kt_a/dt . Assuming a measurement at $s = 0.5$ and a step size $dt = 0.1$ ns are sufficient, as shown in the context of second-order TDVP dynamics (Figs. S23 and S24)^V, this would modify our lower-bounds on depth/gates by a factor 20 compared to our analysis with $dt = 0.01$ ns evolution to $s = 1$; it does not change the conclusions.

^{IV} State-of-the-art (in scaling) qubitization approximations to time-dependent evolution, based on 2-qubit unitaries (gates), have been proposed that allow for exponential reduction in the error as a function of dt , with an overhead relative to Trotterization that does not scale with simulation time t_a ¹⁰³. However, no such circuits have been tested in practice. In our view the most

practical means to simulate our results, assuming availability of a large digital universal quantum computer, would be a 2nd order Trotterization.

^V Limited experiments indicate that 2nd-order Trotterization (Fig. S9, PEPS dynamics) may be more sensitive to time discretization than 2nd-order TDVP (MPS dynamics), requiring smaller dt ; $dt \sim 0.01$ ns appears to provide sufficient accuracy in both cases.

In the context of simulating to the end of the anneal ($s = 1$) with $dt = 0.01$ ns, per our work, lower bounds on gate-model resources can be considered, assuming no constraints on coherence, gate tuning or processor connectivity. For our smallest experimental results—a 4×4 cylindrical lattice (28 couplers, edge chromatic number 4) at 7 ns—this would require depth of 2800 and 39,200 2-gates. Bounding the resources for our more challenging experiments, the Advantage processor used in this study has an edge-chromatic number of 15, and 40279 programmable couplers, so that a 2nd-order Trotterized circuit simulation at 20 ns requires depth 30,000 and 80,558,000 2-gates. Recent state-of-the-art simulations with time-independent coarsely-grained Trotterization on $N = 127$ qubits ($|E| < 1.25N$, mean connectivity $\lesssim 2.5$) heavy-hex lattices have depth 60 and 2880 CNOT gates by contrast. Simulation by state-of-the-art circuit-model quantum computers is limited by gate type, scale, topology and fidelity. Since we see many orders of magnitude separation between the scale of digital simulation required for QPU emulation and published state-of-the-art performance, we can conclude that near-term circuit-model quantum computers are impractical for this purpose. Similar conclusions have been drawn in more detailed analyses of related problems¹⁰⁴.

We can also consider the implications for PEPS and MPS simulation in the context of the $d + 1$ -dimensional PEPS that is the circuit model. Full contraction can be interpreted as a problem of $d + 1$ -dimensional PEPS contraction, but special care must be taken, relative to Section III B 1, owing to the asymmetry of the time dimension (and reduced bond dimension approaching $t = 0$). Complexity arises in the contraction of the network, in line with standard vertex congestion considerations⁸³. In the regime of deep circuits where $t_a/dt > N$ (our experimental regime), an exact contraction requires resources scaling exponentially in N , worse than for contraction of the d -dimensional PEPS models. Our time-evolved PEPS methods (and also time-evolved MPS methods) are attempts to leverage the improved scaling understood for PEPS-lattice models at fixed time. They can be interpreted as boundary-PEPS heuristics for the circuit model: we approximate the boundary $t = 0$ as an (effectively factorized) PEPS and contract in time with compression after each operation to maintain a maximum bond dimension. The shortcomings of various standard and bespoke boundary-PEPS heuristics are evaluated in the following sections.

Many methods have been developed to accelerate the contraction of circuit-model PEPS. These can exploit inhomogeneity in the circuit structure, small circuit depth, rank-reduction of gates, Clifford-gate methods, an assumed fidelity threshold, or locally tree-like structure (or many of the above)^{15,16,18,92,101}. Considering our ensemble of rank-2 gates, $\exp(-i\pi/2[4dt\mathcal{J}(t/t_a)J_{ij}]\sigma_i^z\sigma_j^z)$ and $\exp(-i\pi/2[4dt\Gamma(t/t_a)]\sigma_i^x)$, we can note that these are close to Clifford gates when angles $\theta = [\cdot]$ are integer: if dt could be chosen to satisfy this criterion for a significant subset of gates throughout the anneal, efficiencies become possible. Since small dt is required and $\{\mathcal{J}, \Gamma\}$ vary across a broad range of values, we view discretization using Clifford gates as infeasible. To achieve a sufficiently refined time evolution (faithful high-fidelity J_{ij} representation, and small Trotterization error $[dt]$, at the evolved energy scales throughout the anneal) a large variety of angles θ is required; ideally a continuous range, which is impractical for calibration purposes in current circuit-model quantum computer platforms. We evaluate locally tree-like methods in sections III B 3 and III B 4 without success. The demonstrated growth of correlations and entanglement through the quantum phase transitions lead us to anticipate, more generally, a failure of perturbative expansion methods. Homogeneity of gates in space in time, a large number of short loops, and large depth, leaves few opportunities to leverage slicing and contraction-order optimization beyond those exploited in our boundary-MPS and boundary-PEPS methods.

3. Locally tree-like approximations

The success of MPS methods is contingent on the ability to efficiently manipulate isometries and canonical forms, which allow efficient full-scale contraction. Heuristic methods are required for the manipulation of isometries in PEPS and have only been well studied, with limitations, in square lattices¹⁰⁵. Computationally useful canonical forms have not been defined for disordered lattice models. Foregoing global consistency of isometries and gauges, a local consistency of gauges and/or isometries can be sought, and such approximations have had much success. Uniqueness of a locally self-consistent gauge, in agreement with full contraction on environments up to small corrections, is expected for “locally tree-like lattices” with short correlation lengths¹⁰⁶.

When such gauges are stable, two methods are commonly used: belief propagation (BP) and simple update (SU)^{16,95,107}. The gauges resulting from a locally consistent criterion can be defined based only on tensors (tensor pairs) and their neighbors, matching those on a tree tensor network. These can be iterated from some initial condition, but unlike a tree tensor network it is not possible to systematically contract from leaves to produce a canonical form. Convergence of the iterated local criteria is not guaranteed, and systematic errors associated with converged gauges are typically difficult to quantify. However, if correlation length is short compared to the length of loops, it follows intuitively that such methods can converge on a locally tree-like environment with reasonable accuracy.

For these types of approximations to be efficient, a tensor network needs to look “tree-like”, which means that the network can locally be represented by a tree with a diameter bounded by the smallest loop. For example, in the

aforementioned simulation of the kicked Ising experiment on the heavy-hex lattice, the shortest loops are length 12, and tree-tensor networks proved to be very fast and economical in computational resources since correlation lengths are relatively short¹⁶.

For a random k -regular graph, it is easy to show that the expected average loop size goes to infinity with the number of nodes¹⁰⁸, and hence a tree-tensor approach is expected to efficiently perform quantum annealing protocols on such graphs for any finite correlation length. Our k -regular lattice models, in contrast, have many short loops incident on every site: length 2 for dimerized bicliques, length 4 for dimerized-cubic and square lattices, and length 6 in diamond lattices and non-dimerized 3D lattices. As shown in Section III B 4, we can use a locally tree-like approximation efficiently in all these lattices for both time updates and estimation of observables—agnostic to the requirement of gauge convergence—but find that for even the shortest experimental time scales and simplest lattices, large systematic errors result.

4. Simple update

Simple update assumes a locally tree-like approximation for purposes of truncation during time evolution⁹⁵. The initial condition of factorized spins meets the locally tree-like gauge conditions. In our annealing context, SU with fine-grained Trotterization ought to track a gauge consistent with BP where such a fixed point is stable¹⁰⁷, and so we omit an independent evaluation of BP.

The impact of an unconverged, or converged (but poor), gauge during time evolution is that inferred truncations are uncontrolled. The systematic errors associated with uncontrolled truncation can nevertheless represent a good trade-off when errors in the final state are permissible. SU methods have been successfully applied for the simulation of a kicked Ising model on a heavy-hex lattice for example, which despite finite dimensionality (2) has a paucity of short loops^{16,92}.

Simple update is based upon an approximation of the bond environment. This environment can be used not only for PEPS truncation in time evolution, but also for estimation of expectation values. We call this approximation SO (simple observable). The self-consistent combination of SO with SU has been successfully applied in related PEPS applications⁹², and this method is feasible in all of our lattices.

In application of SU and SO to the experimental regimes many results were well converged, with a weak dependence of ϵ_c on bond dimension D (for $D \geq 2$). A weak dependence of errors on D might result from the insufficient neighborhood scale. In some more strongly correlated regimes SU was found to be unstable, with sensitivity to dt and the minimum singular value threshold. This is not unexpected since inversion of small singular values necessary in the algorithm can amplify errors quickly, and locally-tree-like assumptions fail as correlation length grows. We present in Fig. S10 results with $dt = \{2^{-7}, 2^{-8}\}$ ns and $D = \{2, 3, 4\}$ optimized for minimal error on a per instance basis, and with a small (10^{-10}) cutoff in the singular values supported. Correlation error ϵ_c (NN) is taken only over nearest-neighbor pairs; this quantity correlates strongly with ϵ_c in (4), but SO environments are local and hence more suited for estimation only of nearest neighbor correlations, so we use the nearest-neighbor variant for convenience. Errors are significantly larger than those obtained from the QPU in our experimental regime (lattices and time scales).

Sources of errors might be attributed to SU or SO. A spin-glass ($\pm J$, low precision) ensemble for small square open-boundary lattices provides a controlled context for this evaluation. SU might be combined with one of: SO; contractions on 2×3 neighborhoods (combined with Frobenius norm outside the sublattice, matching the standard NTU update for time-evolution); or exact full contraction. As shown in Fig. S11, SU at small bond dimension combined with SO is accurate for short time scales ($t_a < 1$ ns). 3×2 and full contractions improve slightly on errors in these regimes, in line with perturbative loop corrections. SO degrades at longer time scales in a manner that is almost independent of bond dimension, whereas 3×2 and full contraction respond to variation in bond dimension, with full contraction superior. For contraction methods, on time scales $\gtrsim 7$ ns, both full and 3×2 contractions fail in a manner that cannot be reduced systematically by increasing bond dimension. In this regime the accumulation of errors in SU prevents attainment of small errors by any contraction method for observables. This aligns with our intuition that as correlation length grows more complex, neighborhoods must be contracted^{94,97}. We examine this in the context of square and cubic lattices in Sections III B 6 and III B 7.

5. Non-canonical MPS in 1D

In this work we use the standard PEPS techniques that do not refer to canonical forms. Potential advantages of the isometric PEPS¹⁰⁵ seem to be offset by the errors of Moses moves that attempt to change the canonical structure of the network. In order to better understand problems arising from a non-canonical structure of PEPS in 2D, we simulated the Kibble-Zurek (KZ) quench in 1D for an infinite pure-ferromagnetic quantum Ising chain with an infinite




FIG. S10. **SU (evolution) plus SO (correlation estimation) versus QPU.** (A) Ensemble-median 8×8 square-lattice (high-precision) correlation error, with numerator and denominator taken only over nearest-neighbor pairs, for SU and QPU. (B) Nearest-neighbor correlation error for high-precision square (8×8), cubic ($3 \times 3 \times 3$), diamond ($4 \times 4 \times 8$), and biclique ($K_{6,6}$) ensembles at $t_a = 7$ ns (colored markers) and 20 ns (white markers). PEPS SU+SO errors are smallest on shorter time scales and in models with fewer short loops (like the diamond ensemble), as expected.




FIG. S11. **Absolute correlation error on the (unique) central bond in an open-boundary 3×4 $J = \pm 1$ square lattice as a function of quench time.** Exact diagonalization methods are applicable at this scale. We present a root mean square error on SU dynamics with respect to 17 disorder realizations. Three different contraction methods for calculation of observables are used, indicated by color, and are combined with SU at various bond dimensions: SO ($D = 2, \dots, 6$), 3×2 neighborhood contraction ($D = 2, \dots, 4$), and exact contraction ($D = 2, \dots, 6$). As simulation time approaches zero, all methods approach zero error for bond dimension ≥ 2 . At the shortest time shown, error decreases monotonically with bond dimension for all methods. For time $\ll 1$ ns all methods converge to zero error at small bond dimension. Approaching the short-time limit, errors go to zero much quicker (in absolute terms shown, and also relative to c_{ij}), and SO is sufficient. At intermediate time scales estimation quality improves both as a function of neighborhood scale and bond dimension (albeit SO has a very weak dependence on bond dimension). For time $\gtrsim 4$ ns all methods approach larger errors, and without a monotonic dependence of error on bond dimension and neighborhood size. Even with no error in observable calculations we see the systematic errors at experimental time scales—larger problems with additional loops can be anticipated to have larger systematic errors, explaining the trends in Figure S10.

MPS ansatz, but without enforcing any canonical gauge. The ansatz is a tensor train with two sublattice tensors: $\dots - A - B - A - B - \dots$. The time evolution is performed in a similar manner as for an infinite PEPS in 2D. Not counting local gates with a transverse field, in each time step there are two 2-site Suzuki-Trotter gates with a ferromagnetic interaction on bonds AB and BA that commute with each other. After the application of a gate, the bond dimension is doubled and has to be truncated back to the original bond dimension D in a way that minimizes truncation error (see Fig. S12B). The error is measured on a cluster of $2 + 2M$ sites including the two central sites with the gate and two blocks of M spins adjacent to the central sites to the left and the right. The error $\delta_{t_i}^g$ is the Frobenius norm of the difference between the clusters before and after the truncation (see Fig. S12B), following application of the gate at time step t_i to a bond $g = AB$ or BA . In the worst-case scenario, the individual errors sum up to a (upper estimate) total truncation error $\Delta_g = \sum_i \delta_{t_i}^g$. We define the error estimate as

$$\Delta = \max_g \Delta_g = \max_g \sum_i \delta_{t_i}^g. \quad (\text{S4})$$




FIG. S12. **Optimizing bond dimension truncation in 2D PEPS and non-canonical 1D MPS.** In **A**, for non-canonical PEPS in 2D, the two central tensors are applied with a Trotter gate g with SVD rank $r = 2$. When the gate tensors are absorbed into the central ones, the bond dimension increases from D to Dr . It is truncated back to D by approximating the central tensors with a pair of new tensors with bond dimension D . The new tensors minimize the Frobenius norm $\delta_{t_i}^g$ of the difference between the left and the right diagram, where t_i index time steps of the Trotterized evolution. Here, in addition to the two central sites, the 2D cluster includes their six nearest neighbor sites. The total cluster size during such NTU algorithm⁹⁴ update is 8 sites. In **B**, we consider a non-canonical MPS in 1D. Here, in addition to the two central sites, the diagrams include $M = 2$ sites to the left and $M = 2$ sites to the right. A total cluster size is $2 + 2M$ sites.

For lack of canonical structure, (S4) serves as a rough estimate for a relative error inflicted on observables by the truncations. We will show below that it correlates well with the errors of the expectation values calculated at the end of the quench.

The total error after the quench with annealing times 7, 14, 28 ns, simulated with 10 different cluster sizes M , is shown in Fig. S13. Overall, a longer annealing time results in a final state with a longer KZ correlation length $\hat{\xi} \propto \tau_Q^{1/2}$ and requires a higher bond dimension to reach the same level of accuracy^{VI}. Here $\tau_Q = t_a/t_0$ is a dimensionless KZ quench time with t_0 being a characteristic system's time scale. As we can see in Figure S13, for shortest values of M




FIG. S13. **Truncation errors in non-canonical MPS.** In **A**, **B**, and **C**, we show the truncation error $\Delta = \max(\Delta^{AB}, \Delta^{BA})$, according to Eq. (S4), as a function of the bond dimension D for annealing times $t_a = 7, 14$, and 28 ns, respectively, for a range of cluster sizes M . For each t_a , the function $\Delta(D)$ saturates above a critical value of M . In **D**, we plot the correlation error, ϵ_c in Eq. (4), as a function of the truncation error. The two are roughly linearly correlated. The correlation error is an order of magnitude lower than the truncation error—the dashed line indicating $\Delta/10$ is a guide to the eye. In this panel, $t_a = 28$ ns and $D = 4, \dots, 60$.

^{VI} We use $\hat{\xi}$ to denote the Kibble-Zurek correlation length in line with literature notation. This is to be contrasted with the usage

of a more generic correlation length (ξ) elsewhere in our paper (which is not assumed to follow KZ scaling).




FIG. S14. **2D PEPS: environments for truncating enlarged bonds in time-evolution simulations.** The two-layer diagrams correspond to the calculation of the Frobenius norms depicted in Fig. S12A. In panel A, we show the neighborhood tensor update environment and, in B–D, its direct extensions that are contracted numerically exactly. In B, SVD₁ indicates a product approximation, where we retain a single leading singular value for tensor decomposition along black dashed cuts. Larger environments, such as the one in E, require approximate contraction, which, however, might disrupt the positivity of the depicted bond metric. Here, to that end, we employ boundary MPS. In F, we indicate a standard PEPS representation with one tensor per spin and, in G, a plaquette PEPS with one tensor for four spins. While a 2×2 plaquette of a standard PEPS with bond dimension D can be contracted into a single plaquette PEPS tensor of dimension D^2 , such bonds might not be used optimally. Indeed, we observe that a larger effective environment in plaquette PEPS might result in the compression of the required bond dimension, $D' < D^2$, for reaching similar accuracy.

the error improves slowly, if at all, with increasing D . We can also see that for a given D , the error decreases with the cluster size until it saturates for long enough M . The minimal M needed for convergence, M_c , depends on the annealing time. It might be natural to expect $M_c \propto \hat{\xi} \propto \sqrt{t_a}$, but it grows faster than $\sqrt{t_a}$, as t_a . For example at $t_a = 7$ ns we can estimate $M_c \approx 16$ while for $t_a = 28$ ns it is four times longer, $M_c \approx 64$. As such, $M_c \propto t_a$ seems to be more consistent with the data. This steeper scaling is not necessarily at odds with the KZ mechanism. After correlated pairs of quasiparticles with opposite quasimomenta $(k, -k)$ are excited near the critical point, during the following quench across the ferromagnetic phase, they are free to move along the chain and spread their pairwise entanglement. In this way, a second length scale, $l = \sqrt{\tau_Q} \ln \tau_Q$, develops¹⁰⁹ that is longer than $\hat{\xi}$. In the limited range of τ_Q where we are operating, what might at first be dismissed as a mere logarithmic correction to $\hat{\xi}$, is actually substantial and makes l appear to grow nearly linearly with $\tau_Q \propto t_a$.

In order to see how the total truncation error affects observables we compare the results with the ground-truth correlations. Fig. S13D shows the correlation error, Eq. (4) of the main text, as a function of the total truncation error (S4). The log-log plot is not inconsistent with a linear dependence. It is noteworthy that the correlation error is an order of magnitude lower than the truncation error. We observe that $\epsilon_c \approx \Delta/10$ is a good approximation to the data, allowing prediction of the correlation error based on the truncation error.

The 1D toy model shows how important it is for the cluster size to be at least comparable to the range of correlations in order to make efficient use of the limited bond dimension. The issue becomes more important in 2D and 3D where the limitations are more stringent.

6. 2D PEPS

In 2D, we start with the current state-of-the-art algorithm that can match QPU accuracy only for the fastest annealing times considered. We then systematically examine a family of extensions, which reduce errors below QPU level at intermediate time scales (though not for the slowest experimentally relevant quenches).

As our baseline method, we employ the neighborhood tensor update (NTU)^{94,110} with one PEPS tensor assigned to each lattice site. In NTU, the cluster used for truncation includes the two central sites with the gate plus their six nearest neighbor sites, including two loops, as shown in Fig. S12A. Such a neighborhood is contracted numerically-exactly at a cost of $\mathcal{O}(D^8 + pD^7)$ operations, where p is the physical dimension of PEPS tensor. This exactness preserves the desired positivity of the bond metric tensor in Fig. S14A that gets inverted during optimization of the truncation.

We next systematically probe several methods that expand the neighborhood scale considered in the truncation. Within the paradigm of local bond environments that get contracted exactly (preserving positivity), we expand the nearest neighbor (NN) environment in Fig. S14A by adding partial information from the next layer of tensors, treated tree-like (see Fig. S14B). Calculating this metric has the same $\mathcal{O}(D^8 + pD^7)$ asymptotic numerical cost. Further

extensions involve including extra loops in the exactly-contracted environment, with both next nearest neighbor (NNN) and NNN+ in Fig. S14C,D coming at $\mathcal{O}(pD^{10})$ numerical cost. Finally, larger environments can be considered, going toward full update. They rely on approximate contraction schemes, such as boundary MPS shown in Fig. S14E. Unlike exact contraction, such approximations do not preserve the non-negativity of the metric tensor.

Another strategy still uses the NTU cluster for truncation, but each PEPS tensor covers a plaquette of 2×2 sites, as shown in Fig. S14G. In this way, the NTU cluster of plaquettes covers four times as many qubits as in the standard setup of Fig. S14F. Surprisingly, for many considered protocols, the bond dimension required for the same truncation error is less than a square of the bond dimension in the one-site NTU. Intuitively, two bond indices with dimension D extending in a given direction out of a 2×2 plaquette in the one-site PEPS should fuse into one bond index with dimension D^2 . We observe, however, that the dimension needed to achieve equivalent accuracy is lower than D^2 . Such compression is possible because a plaquette represented by a contraction of 4 one-site PEPS tensors is more constrained than the same plaquette covered by a single plaquette tensor (see Fig. S14F and G). This compression also enables a numerically exact contraction of the (four-times bigger) plaquette NTU cluster. This larger effective cluster allows making more efficient use of the truncated bond dimension, contributing to the compression.

In all the above cases, expectation values follow from the approximate contraction of the entire network. We employ a boundary MPS approach adapted to a cylinder. Larger bond dimensions require exploiting Monte Carlo sampling operating in the z -basis. Here, we follow the approach of Ref. 98, but with a local trial probability distribution. As compared to the deterministic evaluation of correlators, which requires a double-layer PEPS (PEPS contracted with its conjugate), Monte Carlo sampling requires only the contraction of a single (projected) PEPS layer. It is worth noting, that we produce on the order of one MC sample per second for an 8×8 system at $D = 10$ (running on a single CPU core), which is roughly three orders of magnitude slower than the QPU sampling rate.

We find the systematic errors associated with the calculation of the expectation values to be much smaller than those resulting from approximate truncation in the time evolution part of the simulation. In Fig. S15, where we collect simulations done for 8×8 cylinder lattice, both for high and low inputs precision studied on the QPU, we show the relation between total truncation error in Eq. (S4) and the correlation error ϵ_c . The latter is roughly an order of magnitude smaller than the corresponding truncation error Δ among many tested environments for standard and plaquette PEPS and a range of bond dimensions used to simulate time evolution for several annealing times t_a . Interestingly, a rough estimate of the correlation error as $\Delta/10$ is consistent with a similar estimate in the 1D toy example in Fig. S13. This corroborates that time evolution is the main source of errors in the PEPS results and truncation error is a valid metric to estimate the quality of the approximation. In those simulations, we terminate the dynamics at $s = 0.6$. This has insignificant impact on statistics in the case of exact simulation (see Fig. S23A) but prevents PEPS from accumulating extra errors during the last stage of the evolution.

We use the truncation error metric in Fig. S16 to compare the simulation errors with the simulation times, where all times are for the time-evolution only; i.e., without sampling from the final state, executed on a single CPU core. While for $t_a = 2$ ns all tested simulation methods produce errors below QPU precision, already for $D = 4$, longer annealing times require larger environments and bond dimensions. In particular, we identify a plaquette PEPS with NN+ environments as a leading method for 7 ns quenches. However, all tested methods give worse than QPU results for 20 ns quenches. In particular, low-precision inputs prove to be harder to simulate using PEPS (note that the data in Fig. 3 of the main text shows high-precision results). Finally, we note that the simulated problem has a global $\mathbb{Z}_2$ symmetry. We observe that imposing this global symmetry directly in PEPS ansatz does not affect the quality of the results compared to using dense tensors with no symmetry imposed.

7. 3D PEPS

We performed simulations on a cubic $L \times L \times L$ lattice with open boundary conditions and random nearest-neighbor couplings J_{ij} from a uniform distribution in $[-1, 1]$. This is a simplification of the dimerized cubic model used in experiments which has twice as many qubits, but the simplification is already challenging enough to highlight limitations in simulation. Four truncation schemes, shown in Fig. S17, were applied. The simplest one is the SVD truncation in **A** that ignores any tensor environment and requires only $\mathcal{O}(pD^7)$ operations. Scheme **B** includes NN tensors but still ignores contractions between them, keeping the numerical cost to calculate the Gramm-Schmidt metric tensor at the same level of $\mathcal{O}(pD^7)$. Scheme **C** includes the contractions, raising the cost to $\mathcal{O}(D^{12} + pD^{11})$ operations. Finally, scheme **D** includes up to third nearest neighbors. In order to keep the cost at the tolerable level of $\mathcal{O}(pD^8)$ a sequence of SVD truncations to the single leading singular value (SVD₁) is applied that renders an approximate tree-like diagram. First the eight double-PEPS tensors at the corners of the diagram are SVD₁-truncated and then, after almost all contractions are completed, the four vertical double bonds are cut. As in Section III B 6 we are able to truncate dynamics at $s = 0.6$ with negligible impact on results.

In 2D it was practical to combine neighborhood updates with approximate full contraction, the former being a




FIG. S15. **2D PEPS: Correlation errors vs truncation error.** We compare the correlation errors ϵ_c with the truncation error in Eq. (S4). The two quantities are roughly linearly correlated for a broad range of simulation parameters, with ϵ_c being an order of magnitude smaller than Δ . The dashed line showing $\epsilon_c = \Delta/10$ is a guide for an eye. Each point corresponds to a different PEPS bond dimension (color), the environment used for truncation (symbols, see Fig. S14), annealing time, and coupling precision (the latter two are not distinguished by color nor symbol). Data for 8×8 cylindric lattice.




FIG. S16. **2D PEPS: errors vs numerical costs.** We show estimates of PEPS accuracy, employing truncation errors $\Delta/10$, as a function of simulation times for problems defined on 8×8 cylinders. Each panel corresponds to a fixed annealing time and type of random couplings (high- or low-precision, see Sec. II). Points in each panel represent PEPS attempts to simulate a given problem for a selection of PEPS bond dimensions and type of environment used in truncation (see Fig. S14). Those plots complement Fig. 3 of the main text, which shows the correlation errors for $t_a = 2, 7$, and 20 ns for high-precision couplings. We see systematic accuracy improvements with the bond dimension for shorter annealing times. This trend systematically stagnates for slower evolutions, where the necessity to employ larger environments becomes apparent. The low-precision couplings prove more challenging for PEPS to simulate than the high-precision. Indeed, energy degeneracies related to discrete $J = \pm 1$ couplings are expected to facilitate entanglement spreading as compared to high-precision inputs, where such degeneracies are much less common. As CPU times, we report simulation wall times executed on a single CPU core.

computational bottleneck. By contrast, the calculation of observables by comparable methods in 3D is a computational bottleneck. The deterministic scheme employing double-layer PEPS⁵³ is not practical beyond $D = 4$. We also tested code for a 3D generalization of the single-layer Monte Carlo⁹⁸, which partly trades the contraction cost for sampling,




FIG. S17. **3D PEPS**. The left panel shows four truncation clusters used in 3D PEPS time evolution. The diagrams represent Gramm-Schmidt metric tensors that measure the truncation error. With cluster **A** the truncation is just an SVD, cluster **B** includes whole nearest-neighbor tensors (more than the simple update), **C** includes vertical bonds between the nearest-neighbor tensors (like the 2D NTU, called 3DNN). The biggest cluster **D** reaches out to third nearest neighbors. It is made tree-like by SVD truncations to just the single leading singular value (SVD_1) that is marked by the black dashed cuts. The right panel, **E**, shows total truncation error as a function of the annealing time t_a for clusters **A**, **B**, **C**, **D** and bond dimension $D = 4$ and 6 . Here $L = 5$, $dt = 0.01$ ns, and the error bars result from averaging over 8 realizations of the random couplings. A dashed line indicates a 2% level of error obtained in QPU simulations for 3D problems with dimers.

but it still does not extend much beyond $D = 4$. We estimate that $D = 6$ is the upper limit, if allowed a few months of computations, but D is expected to grow exponentially with the Kibble-Zurek length and so does the computation cost.

The right panel of Fig. S17 shows the total truncation error as a function of the annealing time for $D = 4, 6$. For a given bond dimension, the errors of schemes **B**, **C** are practically the same, while those of **A**/**D** are much larger/smaller, suggesting that the cluster size matters more than the contractions between the nearest-neighbor tensors. For $D = 4, 6$ the error of scheme **D** crosses the 2% QPU benchmark around 4 ns. In the regime of parameters of Fig. S17, the accuracy of the evolution is limited mainly by the cluster size: there is little improvement with increasing D but much with the size. With approximate contraction, the size could be further increased but, for long-enough annealing time, a crossover to a D -limited regime would inevitably take place and the cost of time evolution would start to grow exponentially with t_a .

As a means to increase the exactly contractible cluster size, we also attempted a 3D cubic PEPS, similar to the plaquette one in 2D, where a single PEPS tensor covers a cube of 2^3 sites. In 3D, a one-site PEPS has 4 bond indices with dimension D extending in each direction out of the 2^3 cube. In a cube PEPS tensor, the 4 indices can be fused into one index with dimension D^4 . Therefore, the minimum required to simulate the evolution with only the classical nearest-neighbor Hamiltonian (the transverse field set to zero) is $D = 2$ for the one-site PEPS, $D^2 = 4$ for the plaquette PEPS, and $D^4 = 16$ for the cubic one. The cubic PEPS has a much higher minimal bond dimension than the plaquette one. In 3D, $D \geq 16$ may still be feasible for time evolution, despite an overhead due to the physical dimension, $2^8 = 256$, but the evaluation of expectation values within a controlled approximation seems out of reach.

IV. MEASURING THE QUALITY OF SAMPLED DISTRIBUTIONS

The computational task we demonstrate is approximately independent and identically distributed sampling from a quenched quantum state projected in the computational basis. The quality of approximation can be captured by various distributional divergence measures: a correlation-error metric ϵ_c and classical fidelity (F , also called Bhattacharyya distance) are presented in the main text.

A popular measure of sampling quality relevant to random-circuit sampling is the cross entropy¹¹¹; in that application, the target distribution approaches a Porter-Thomas distribution and a cross-entropy-based measure is known to approach 1 from below as a function of gate fidelity. The chaotic nature of the distribution makes spoofing of large values challenging.

In this work, we might define the cross entropy as $-\langle \log \tilde{p}(z) \rangle_p$, where $\tilde{p}$ is the ground-truth probability of states projected in the z -basis and angled brackets denote an expectation taken by Monte-Carlo sampling of some approximating distribution p (the QPU). Unfortunately such a measure appears not to be readily bounded in a manner informative for fidelity. Furthermore, and unlike random-circuit distributions, state probabilities in our application




FIG. S18. **Estimating QPU-equivalent bond dimension from classical fidelity.** (A) Fig. 2C (right) is extended, with data for four MPS bond dimensions, up to and including the ground-truth bond dimension $\tilde{\chi}$, for which only sampling error remains. (B) Fig. 2C (left) is shown for comparison, qualitatively similar in spread to $\chi = 64$. (C) We can estimate the QPU-equivalent bond dimension from $\mathcal{F}$ instead of from ϵ_c , by using 10^6 samples or by extrapolating to the limit of infinitely many samples drawn. We call these quantities χ_Q' and χ_Q'' , similar to the estimate χ_Q derived from correlation error. Analogously to χ_Q , error ($1 - \mathcal{F}$ in this case) is plotted for a range of MPS bond dimensions and the crossing point of QPU error (solid horizontal lines) gives the QPU-equivalent bond dimension. (D) χ_Q' is shown for varying resampled sets of size 10^5 to 10^6 , and plotted against the inverse sample size along with χ_Q'' .




FIG. S19. **Comparing correlation-based and fidelity-based QPU-equivalent bond dimensions.** Shown are χ_Q and χ_Q'' (as in Fig. 2D and Fig. S18C respectively). At a fixed maximum sample size 10^6 , we expect χ_Q'' to deviate downwards from χ_Q for systems of increasing entropy, since fidelity estimation requires exponentially many samples.

are correlated with efficiently calculated statistics (such as $\mathcal{H}_P$); the strength of correlation with $\mathcal{H}_P$ is controlled monotonically by t_a and noise. It is easy to raise (lower) the cross-entropy measure by preferential sampling of states with low (high) classical energy. Thus the cross-entropy measure is not a useful quality measure.

The interpretation of a cross-entropy measure is made simpler by comparison with (subtraction of) the model entropy, $KL[p, \tilde{p}] = \langle \log \frac{p(x)}{\tilde{p}} \rangle_p$, yielding a Kullback-Leibler divergence measure that approaches zero only when the distributions are in agreement. The problem with such a measure is twofold: first, it measures the ability of the ground truth to predict model samples; second, the calculation of entropy from sampling requires a number of samples that scales with the model entropy¹¹², and the estimate can be subject to strong variance if the ground-truth distribution is narrower than the model's one. To solve the first issue we might consider the opposite form, $KL[\tilde{p}, p]$, but the problem of scalability persists with this definition. Due to entropy scaling with system size, we could not formulate an unbiased estimator for either form of KL-divergence that was informative on larger problems. We found that several estimators exhibit related issues, including the classical fidelity. Because the standard techniques for dealing with scaling—coarse graining, binning, and projection of samples into a lower-dimensional space—are challenging to implement without simplifying the computational task, we selected the scalable measure ϵ_c as our primary metric. As we show, this correlates well with other measures of distribution divergence.

Alongside ϵ_c we present classical fidelity, $\sum_x \sqrt{p(x)\tilde{p}(x)}$, as a complementary measure to test distribution divergence beyond that captured by two-body correlations. Fidelity is upper bounded by 1 (when distributions agree) and lower bounded by 0. We present data with 2^N bins (one per state), but limit usage to smaller scales where the measure is not dominated by sampling bias since, like the entropy-based measures already discussed, the number of samples must scale with entropy to avoid a dominant sampling bias.

Distribution divergence measures, such as ϵ_c and $\mathcal{F}$, compress high-dimensional information to a number. These numbers may be skewed in some non-trivial manner; for example, our measure ϵ_c is not sensitive to sign-symmetry breaking ($p(s) \neq p(-s)$). A divergence of this type would yield a strong signal in classical fidelity by contrast. For this reason it is useful to consider several numbers and also contextualize them with additional distributional information where possible. It is also important to demonstrate that our core results, based upon ϵ_c , are robust to a change in divergence measure. With this in mind, we provide visualizations of the full probability distribution to provide confidence in the measures, and show that if our QPU-equivalent bond-dimension analysis, based on ϵ_c , is replaced by an analysis based on $\mathcal{F}$, no practically significant difference in results is apparent.

Fig. 2C scatters p_z against $\tilde{p}_z$ for 10^6 samples drawn from QPU and from approximate MPS at bond dimension $\chi = 64$. We first extend this data in Fig. S18A, showing the spread of p_z versus $\tilde{p}_z$ for various values of χ . When $\chi = \tilde{\chi} = 256$, the only error in the distribution is sampling error. We compare this with QPU samples in Fig. S18B, noting, as in the main text, that results look roughly equivalent in quality to MPS at $\chi = 64$.

To provide a more precise estimate of QPU-equivalent bond dimension, we first define χ'_Q as the intersection in infidelity $1 - \mathcal{F}$ with the interpolating function of many evaluated bond dimensions—this is exactly how we define χ_Q , based on correlation error. This infidelity is a function of the sample-set size used, since we use an empirical estimate to the distribution p based on sampling. We can suppress the effect of finite sample sizes by extrapolating infidelity to the infinite-sample limit; assuming that sampling error decreases as the square root of the number of samples is sufficient in lower-entropy use cases. Using the infinite-sample limit gives another measure, χ''_Q , which we analyze instead of χ'_Q . Sampling error in measurement of χ'_Q for one input is shown in Fig. S18D.

We compare χ_Q with χ''_Q in Fig. S19, for ensembles of 2D problems. As previously noted, sampling error in estimating $\mathcal{F}$ is severe for distributions with high entropy (i.e., larger problems and faster anneals), which will affect not only direct estimates of χ'_Q , but also extrapolations for χ''_Q . As seen in Fig. S18D, we can expect small sample sizes to result in underestimates of χ''_Q , and therefore we should see a deviation downwards—compared to χ_Q —for larger problems and faster anneals. Indeed, in Fig. S19 this is what we see, with strong quantitative agreement between χ_Q and χ''_Q beginning to decay for larger problems at $t_a = 7$ ns. Nonetheless, these results give confidence that χ_Q —derived from pairwise correlation errors—is closely linked to the more expressive fidelity measure, and thus the low distributional distance seen for the QPU is not just limited to a plurality of two-body correlations, but rather is based on accurate reproduction of the amplitudes in the final wave function.

V. DYNAMICS OF ENTANGLEMENT ENTROPY

The dynamics of entanglement in a quantum quench have been widely studied in quantum many-body systems, both in the context of instantaneous quench and in the context, relevant to this work, of a continuous quench across a QPT^{30,109,113,114}. In “anneal” quenches, the state of an N -qubit system, prepared initially in the ground state, is evolved in time with *slowly* varying time-dependent Hamiltonian parameters. In such a finite system, the evolution across a continuous critical point can remain adiabatic for sufficiently slow quenches due to a finite-size gap at the




FIG. S20. **Convergence of peak entanglement entropy.** Entanglement $S_{\max}$ grows as a function of MPS bond dimension χ , converging in the $\chi \rightarrow \infty$ limit. **A** and **B** show this convergence in 2D models for 7 ns and 20 ns quenches, respectively. The line $\log(\chi)$ indicates the upper bound of entanglement that can be expressed in an MPS of bond dimension χ .

critical point. Consequently, at least in one dimension the entanglement entropy of bipartition follows instantaneous ground-state behavior in this scenario. With faster parameter variation, the system gets excited upon crossing the critical point. Entanglement dynamics in such a case has been studied in a homogeneous 1D TFIM in Refs. 30,109, showing that at the end of the anneal it grows logarithmically as a function of annealing time t_a , $S(t = t_a) \propto \log(t_a)$. Focusing on a finite system³⁰, there is a crossover between those two regimes, where, in the adiabatic regime, the entanglement entropy at the critical point is $S \propto \log(N)$ ¹¹⁵. Much less is known about entanglement scaling as a function of number of qubits and quench time t_a in anneal quenches in dimensions $d > 1$.

Accordingly, this work relies mostly on MPS to reveal the scaling behavior of entanglement entropy in anneal quenches for lattice topologies with dimensions $d > 1$. Section VI analyzes the behavior of the QPU-equivalent bond dimension, χ_Q , in terms of the correlation error, ϵ_c , and ground-truth bond dimension. We further show how we extrapolated our MPS results in bond dimension in order to establish the ground truths.

VI. COMPARING AND EVALUATING QPU AND MPS EXPERIMENTS

In our experiments, we estimate bipartite von Neumann entanglement entropy S_A using MPS, with our main practical constraint being the exponential growth of χ_Q as a function of MPS area. The main quantity of interest is the peak entanglement $S_{\max}$, which is the maximum both over MPS cuts (A, B) and over annealing parameter $s = t/t_a$.

A. Estimating entanglement entropy

For both $S_{\max}$ and χ_Q , our approach is to grow χ in factors of $\sqrt{2}$ until quantities are converged. Fig. S20 shows convergence of $S_{\max}$ as a function of MPS bond dimension χ for 2D models with $t_a = 7$ ns and $t_a = 20$ ns. In general we find that $S_{\max}$ converges to the $\chi \rightarrow \infty$ limit faster than χ_Q does.

B. Evaluating ground-truth convergence

The ground truth is ideally obtained in our experiment by taking the limit $\chi \rightarrow \infty$ in our MPS simulations. Our distribution accuracy measures— ϵ_c , $\mathcal{F}$, and χ_Q —can be shown to converge as a function of increasing χ approaching this limit much in the same way as entanglement entropy (see Fig. S21).

Statistics converge smoothly to an asymptote as shown, consistent with an exponential tail of truncated singular values⁵⁸. Ideally we might observe a clear and sustained asymptote across a large range of χ , indicating that singular value truncation is a negligible source of error. For the hardest problems, this is not practical due to constraints on computational resources. We therefore consider the possibility of extending the reach of MPS simulations by extrapolating in χ .




FIG. S21. **Convergence of error and equivalent bond dimension as a function of ground-truth bond dimension.** As in Fig. S20, we study square lattices at $t_a = 7$ ns and $t_a = 20$ ns. QPU error ϵ_c (A–B) and QPU-equivalent bond dimension χ_Q (C–D) are shown as a function of MPS bond dimension χ . The MPS used for the ground truth must have sufficient bond dimension such that these quantities are converged. Gray lines indicate the upper bound $\chi_Q \leq \chi$.




FIG. S22. **Effect of extrapolation in MPS bond dimension on correlation errors.** MPS experiments of varying χ are evaluated under three types of two-point extrapolation of spin-spin correlations, in terms of their error against a well-converged ground truth with $t_a = 7$ ns.

C. Extrapolation of bond dimension

Since observables in MPS simulations converge as a function of χ , it is often helpful to extrapolate to the limit of unbounded bond dimension ($\chi = \infty$). In this work, we study problems that are near the edge of our computational capacity; we are therefore sufficiently far from convergence that extrapolation must be approached carefully. Extrapolating using too-small bond dimensions is unreliable, so we restrict extrapolation to linear extrapolation of spin-spin correlations in $1/\chi$ based on the largest two bond dimensions available. With this restriction, extrapolation




FIG. S23. **MPS error from stopping at $s < 1$, compared with growth of entanglement entropy.** (A) (Top) We plot median correlation error of MPS correlations throughout the quench, using a reference ground truth at $s = 1$. Square (8×8 , left), cubic ($3 \times 3 \times 3$, middle), and biclique ($K_{8,8}$, right) ensembles ($t_a = 7$ ns, $\chi = 1024$) approach the $s \rightarrow 1$ limit in varying manner. (Bottom) Growth of ensemble-median $S_{\max}$ (maximum over MPS cuts for each instance) saturates when ϵ_c is near 0.1, higher than QPU error, consistent with Fig. 4C (inset). (B), Bipartite entanglement entropy in a single 8×8 spin-glass instance, taken across the central MPS cut (32 spins on each side) as a function of s for varying t_a . Even $t_a = 1$ ns is far slower than the speed-limited regime in which entanglement grows monotonically as a function of s .

to $\chi = \infty$ (“full extrapolation”) tends to increase errors, but we find that extrapolating from $\chi_{\max}/\sqrt{2}$ and $\chi_{\max}$ to $\sqrt{2}\chi_{\max}$ (“partial extrapolation”) generally reduces errors, particularly in the important region $\epsilon_c \approx 0.03$, where errors are the same order of magnitude as QA errors. Fig. S22 shows examples of 2D, 3D, and biclique ensemble errors for no extrapolation, full extrapolation, and partial extrapolation. Errors are evaluated against the ground truth of maximum available bond dimension $\chi_{\max}$ with no extrapolation. In all cases, partial extrapolation gains at best a factor $\sqrt{2}$ in bond dimension. For this reason, although partial extrapolation is useful for refining ground-truth statistical estimates, it does not change the competitiveness of approximate MPS methods; we therefore restrict the use of extrapolation to ground-truth estimation.

D. Effects of time step and stopping point

The main parameter we have tuned to trade speed against accuracy in MPS methods is χ . We found MPS incapable of matching QPU quality for the largest simulations due to both time and memory requirements, and showed that nearly all entanglement entropy must be captured in the MPS to match QPU quality. In Fig. S23A we show that the time evolution of the quench cannot be truncated early—before nearly full entanglement is reached, requiring




FIG. S24. **MPS error by time step dt .** (A) $t_a = 7$ ns and (B) $t_a = 20$ ns. Correlation error is plotted as a function of dt , against extrapolation to $dt \rightarrow 0$. System sizes are 8×8 square, $3 \times 3 \times 3$ cubic, and $K_{8,8}$ biclique.




FIG. S25. **QPU error and entanglement in oblong cylinders.** We study $L_x \times L_y$ systems with $L_x = 6$. Data are analogous to Fig. 4; filled and empty markers indicate $t_a = 7$ ns and 20 ns respectively. **A**, QPU correlation error ϵ_c is flat as a function of $N = 6L_y$. **B**, χ_Q is nearly flat for $t_a = 7$ ns, but shows a slight scaling for $t_a = 20$ ns. **C**, $S_{\max}$ and $\log(\chi_Q)$ are similarly related as in other ensembles, with $S_{\max}$ increasing gently with L_y . In the limit $L_y \rightarrow \infty$, we expect $S_{\max}$ to increase as a function of t_a ; larger systems are required in this range of t_a .

explosion of the bond dimension—without increasing correlation errors above QPU levels. Therefore some time can be saved by stopping the simulation early, sometimes at a small cost to accuracy, but it will not be useful in reducing memory requirements. Fig. S23B shows entanglement dynamics for a single MPS cut in a single instance as a function of s . Oscillations in S after the critical point are hidden when taking the ensemble average and maximum over all cuts.

Throughout the work we have generally used $dt = 0.01$ ns; in Fig. S24 we verify that this time step is not a leading source of error.

E. Oblong 2D lattices

We are interested in the scaling of QPU error and entanglement as a function of system size, but for the ensembles studied in the main text, this quickly becomes intractable. One compromise is to study quasi-1D models, which should have favorable scaling in MPS bond dimension requirements. Here we provide data on $6 \times L_y$ cylinders for $6 \leq L_y \leq 15$, whose MPS area remains constant at 6.

As in the main text, we once again see flat scaling of QPU error ϵ_c as a function of system size (Fig. S25A); the bounded MPS width allows extending to larger systems than before. As expected, QPU-equivalent bond dimension χ_Q remains almost constant, with a mild scaling evident at $t_a = 20$ ns. In all other systems we saw χ_Q (and $S_{\max}$) decreasing as a function of t_a , suggesting that we are near the adiabatic, finite-size regime, as opposed to a large-system Kibble-Zurek regime. The data in Fig. S25B hint that we may see a crossover for larger systems, as was previously observed in 1D chains³⁰.




FIG. S26. **Correlation error and area-law scaling of χ_Q at $t_a = 20$ ns.** (A) and (B) are analogous to Fig. 4A and B respectively.

F. Area-law scaling for slower quenches

Fig. S26 shows data analogous to Fig. 4 for $t_a = 20$ ns instead of 7 ns. We note that correlation errors ϵ_c are slightly higher for all topologies, and show less size-dependent growth for biclique problems than at 7 ns. In all topologies the longer quenches have lower χ_Q , which allows us to evaluate larger lattices in some cases.

VII. NEURAL NETWORKS

Neural network quantum states (NQS) offer a powerful approach to constructing computationally efficient ansätze wave functions. Several types of neural networks, such as complex restricted Boltzmann machines (cRBMs)^{46,116}, feed-forward neural networks (FNNs)^{117–119}, recurrent neural networks (RNNs)^{120–123}, and transformers^{124–126}, have been employed to approximate ground states and out-of-equilibrium dynamics of quantum many-body systems. For the ground-state problem, the variational Monte Carlo method¹²⁷ provides a numerical method for ground-state estimation that is competitive with other state-of-the-art approaches such as tensor networks or quantum Monte Carlo¹²⁸. Crucial for the success of this method is the ability to sample from the NQS. Whereas cRBMs and FNNs parameterize unnormalized quantum states and are sampled using Markov chain Monte Carlo, autoregressive models such as RNNs and transformers can be sampled using direct sampling. The latter come with a drawback, since recent results indicate that the autoregressive property of the models may restrict the expressiveness of the model^{123,129}.

Beyond ground states, it is possible to simulate the Schrödinger equation on the variational manifold induced by the neural network ansatz. There are currently two approaches to simulate real-time dynamics. The first method is called time-dependent variational Monte Carlo (tVMC)⁵⁴. For an NQS $|\psi_{\theta_t}\rangle$ with p parameters θ , we calculate the force vector F at time t , which depends on the variational energy of the Hamiltonian driving the dynamics. The parameter update to the network for a small time step is then given by

$$\dot{\theta} = S^{-1}F, \quad (\text{S5})$$

which is closely related to the second-order optimization method of natural gradient descent¹³⁰. Here, S is the metric on the variational manifold induced by the Fubini-Study metric, which is the natural metric on a Hilbert space¹³¹. As S is a $p \times p$ matrix, its inversion takes time $O(p^3)$, which significantly limits the number of variational parameters that can be used. In addition, both S and F must be accurately estimated by sampling from the variational state. For a stable simulation, one typically requires $N_s \sim 10^4$ samples per time step. In practice, the matrix S is typically ill-conditioned and needs to be regularized, for which there exist several approaches^{46,55,118}. The stability of the regularization depends on many factors, including the model architecture, the current variational state and the Hamiltonian. Regularization problems are particularly acute for the RNN architecture. The authors of Ref. 56 argue that RNNs fail for tVMC due to problems with regularizing the S matrix and recent results indicate that the autoregressive property of the model may restrict its expressiveness¹²⁹. Once the S matrix is successfully regularized, we can solve the differential equation in Eq. S5 for a small time step τ with a second-order differential-equation solver (e.g., Heun's method or Runge-Kutta-Fehlberg¹³²).

The second approach for time evolution of variational quantum states is to learn the effect of the time-evolution propagator $\mathcal{U}$ on our variational state^{55,56,133–135}. We refer to this method as projected time-dependent variational

Monte Carlo (ptVMC). Here, instead of solving a differential equation at each time step, we explicitly minimize the distance, or infidelity, between the variational state at time t and time $t + \tau$:

$$\min_{\theta_{t+\tau}} d(|\psi_{\theta_{t+\tau}}\rangle, e^{i\tau H} |\psi_{\theta_t}\rangle). \quad (\text{S6})$$

We see that there are two variational states: the state ψ_{θ_t} at time t and the state $\psi_{\theta_{t+\tau}}$. We “freeze” the parameters θ_t and optimize the parameters $\theta_{t+\tau}$. Since $e^{i\tau H}$ is not a local operator that we can easily evaluate on our variational wave function, in practice we have to approximate the full unitary propagator, which can be done in several ways^{56,134,135}. In order to efficiently estimate the infidelity, the authors of Ref. 135 propose a control-variate method to reduce the variance of the estimator similar to Ref. 120. This estimator is quite efficient and, in our experiments, requires only order $N_s = 10^2$ samples per optimization step. However, this comes at the cost of having to minimize the above distance, which could take on the order of 10^3 steps in our experiments. Additionally, to achieve stable results at long times, a second-order approximation to the unitary propagator is required, which results in a quadratic increase in complexity that becomes prohibitive at larger system size. In ptVMC, we have traded the sampling problem of tVMC for an optimization problem. The advantage of ptVMC over tVMC is that the former is unconstrained in terms of parameters, which opens up the possibility of using highly expressive models for the simulation of real-time quantum dynamics.

For the problem at hand, which is the simulation of non-adiabatic critical dynamics of the TFIM, the Kibble-Zurek mechanism has been studied with an NQS approach that uses tVMC⁴⁷. Additionally, Ref. 56 studied critical dynamics with ptVMC. There are two major differences between previous investigations of critical dynamics and the work presented here. First, the Ising model considered has no symmetries (apart from $\mathbb{Z}_2$), hence we cannot make use of translation or lattice symmetries, and we have to use a general variational ansatz. Second, the time scales considered here are beyond what has previously been investigated in the literature. In both Ref. 47 and Ref. 56, NQS approaches are only used close to the critical point. In particular, Ref. 47 notes the NQS method was not stable in the ferromagnetic regime at long times. Finally, a recent work on the simulation of Rydberg arrays was unable to achieve stable results with NQS⁴.




FIG. S27. **Monte Carlo convergence statistics for an 8×8 instance at 7 ns.** (A) Gelman-Rubin statistic ($\hat{R}$) plotted versus time¹³⁶. A value of < 1.1 indicates that Markov chains are converging¹³⁷. (B) Autocorrelation time (τ) of the Markov chain. Longer autocorrelation times mean that more sweeps are required to obtain uncorrelated samples. Both statistics indicate that as we pass through the critical point, the Monte Carlo sampling becomes more difficult, resulting in a higher number of samples required to obtain the correct dynamics.

We note that before we start the dynamics, the state of the neural network is brought to the ground state of $\mathcal{H}(0)$ via variational Monte Carlo¹²⁷, which converges rapidly to the desired paramagnetic state. In our simulations, we consider only the high-precision square lattices (see Section II). Once the ground state is found, we can attempt to simulate the dynamics with either tVMC or ptVMC. With tVMC, we have been able to simulate systems up to 8×8 at a time scale of 7 ns at significant computational cost (3-4 days of compute time with multiple Nvidia T4 GPUs per instance). We chose a fixed time step of $\tau = 10^{-2}$ in units of $\Gamma_c \cdot \pi$ (~ 4.32 ps), where Γ_c is set to the point in the schedule where $\Gamma(s) = \mathcal{J}(s)$, which is where the dynamics are critical. Going beyond 8×8 is possible for shorter

annealing times $t_a < 7$ ns but would require a large number of GPUs to obtain the necessary samples in a reasonable time. We experimented with adaptive Runge-Kutta-Fehlberg schemes to adaptively set the time step τ . This could significantly reduce the computational cost, since we could speed up the integration at the beginning and slow down as we cross the critical point. However, we were not able to reach stable results for large system sizes with such an approach. The most successful architecture was a cRBM model with the number of hidden units equal to the number of spins (see Figure S28). Due to the $O(p^3)$ complexity of inverting S , we limited the number of hidden units to N ($\alpha = 1$), which was sufficient for the system sizes considered here.

While a Jastrow ansatz¹²⁷ proved to be much more numerically stable than the cRBM, it was not expressive enough to capture the dynamics of the system, resulting in larger errors than the cRBM. Feed-forward neural networks with multiple layers showed promise for small system sizes ($L < 5$), but failed at larger system sizes. Similarly, convolutional neural networks were ineffective, most likely due to the lack of translational symmetry of the spin-glass Hamiltonian considered here. Autoregressive architectures like the RNN and transformers came with S -matrix spectra that were hard to regularize, which resulted in the wrong dynamics.

The main bottleneck in the simulation of the dynamics is the estimation of the S matrix in Eq. S5. Here, we used a Metropolis-Hasting sampling method with local updates for the Markov chain, which can be parallelized over multiple GPU machines via Message Passing Interface (MPI). As we cross the critical point, the sampling becomes very challenging due to the glassy nature of the system (see Figure S27). The long auto-correlation times result in an increase in variance of the estimators of the forces and S matrix, necessitating a larger sampling budget to keep errors under control. For regularization of the S matrix, we explored both the method by Schmitt et al.¹¹⁸ and Medvidović⁵⁵, and found that the latter was the most stable at $r_{\text{cond}} = 10^{-5}$.

Autoregressive models such as the RNN and transformer fail for tVMC. Attempts to perform dynamics with these models with ptVMC indicate that this method has large (but practically feasible) costs, since the number of steps required to minimize the infidelity for each τ can be of order 10^3 . However, a major difficulty with this approach is that the infidelity optimization of Equation S6 can fail to converge, which leads to numerical instabilities. As a result of these convergence issues, we were not able to simulate system sizes larger than $L = 5$. Experiments with the cRBM using ptVMC were much more stable, but not competitive with the tVMC simulations in terms of accuracy, especially at larger system sizes.

Our studies have been limited to square lattices at limited space and quench times; square lattice experiments at 20 ns proved too costly to simulate. If we assume that simulation costs scale up as a function of the correlated volume (consistent with square lattice data from 1 ns to 7 ns in Figure S28), an NQS approach to more-connected topologies on experimental time scales may exhaust GPU resources at even shorter time scales. All numerical simulations were performed with NetKet¹³⁸, a Python framework built on top of JAX¹³⁹.

VIII. CLASSICAL COMPUTE RESOURCES

This work includes experiments run on a heterogeneous collection of classical CPU and GPU compute resources. These include the Summit and Frontier supercomputers at the Oak Ridge Leadership Computing Facility, as well as clusters at the Vector Institute, University of British Columbia, Jagiellonian University, and D-Wave.

The main MPS experiments were run on Summit and Frontier, whose workhorse GPUs are, respectively, Nvidia Tesla V100 and AMD Radeon Instinct MI250X. The NQS experiments were performed on multiple Nvidia T4 cards.

Our extrapolation of compute resources is based on measurements from Summit for a 54-qubit system, estimated to require 12 GB of memory and 7×10^4 seconds at a bond dimension $\chi = 724$ and $t_a = 7$ ns. These figures are extrapolated to larger systems and bond dimensions with the leading scaling forms

$$\text{memory} \propto N\chi^2 \quad (\text{S7})$$

and

$$\text{time} \propto Nt_a\chi^3. \quad (\text{S8})$$

since we use a constant time step $dt = 0.01$ ns throughout (see Fig. S24).

The V100 GPUs advertise 14 teraFLOPS in single precision; we compare this with a peak performance of 1.7 exaFLOPS for Frontier to produce the comparison in Fig. 5D stating that matching QPU quality on the largest systems would exceed millions of years using all of Frontier.

This is of course an oversimplification in many respects, most importantly ignoring the practical constraints on such computations, which are far too large to perform naively on a GPU. Each Frontier compute node has 4 TB of memory, and the total storage is 700 PB. These figures are marked in Fig. S29. We note that the largest bond dimension simulated to date⁷⁹ is $2^{16} = 6.5 \times 10^4$.




FIG. S28. **Average correlation error for a system of size $L \times L$ using the cRBM.** The correlations are estimated with a total of 10^6 samples (equivalent to QPU). Error bars represent the standard deviation across 5 different disorder realizations. We focus on correlations at $s = 0.6$ due to the smallest correlation error observed at this point (see also Figure S23). Depending on the system size and number of samples N_s , we increased the number of GPUs from 1 to 8 and implemented a chunking scheme that splits up the total number of samples into batches. This reduces memory requirements at the expense of increased compute time. (A-C) For times up to 4 ns, matching the QPU error is achievable for larger system sizes by increasing N_s . (D) As system sizes grow, matching the QPU error becomes more challenging. For all time scales, we observe that the error plateaus after reaching a certain sample count.

IX. CRITICAL EXPONENTS AND DYNAMIC FINITE-SIZE SCALING ANSATZ

In this section we briefly explain the dynamic finite-size scaling analysis of Fig. 5B, which generalizes the approach of Ref. 24 from cubic to multiple topologies.

A crucial element of any claim to computational supremacy is evidence that experimental results scale correctly beyond the classically verifiable regime. Since we simulate traversals of quantum phase transitions (QPTs), we can appeal to the existing theoretical foundation of critical phenomena to provide this evidence. For large problems, i.e., in the thermodynamic limit, the Kibble-Zurek mechanism^{27,31–33,140,141} allows us to predict the scaling of correlations




FIG. S29. **Resource requirements for MPS to match QPU quality.** (A) Area-law extrapolation is used to estimate the QPU-equivalent bond dimension χ_Q required by MPS methods to achieve QPU-level accuracy. Values far exceed the largest reported MPS bond dimension to date⁷⁹. (B) Dominant scaling of memory is $\propto N\chi_Q^2$; assuming this form we compute MPS memory requirements, which exceed both Frontier node flash memory (4 TB) and total storage of the Frontier filesystem (700 PB). (C) Power, like computation time, scales $\propto t_a N\chi_Q^3$. Here, GPU-time requirements (Fig. 5D) are converted to power requirements based on 300 W draw of an Nvidia Tesla V100. These requirements far exceed the total global electricity consumption throughout human history. All plots use $t_a = 7$ ns for biclique problems and $t_a = 20$ ns for diamond and cubic problems.

	Square	Cubic	Biclique
z , literature	2.46(17) (Ref. 36)	1.3 (Ref. 60)	2 (Ref. 61)
$1/\nu$, literature	0.71(24) (Ref. 36)	1.55(11) (Ref. 24)	4 (Ref. 61)
μ , literature	3.17(41)	2.85(21)	6
μ , experiment, all sizes, low precision	2.67(30)	2.65(12)	6.1(11)
μ , experiment, all sizes, high precision	2.54(13)	2.77(12)	6.5(7)
μ , experiment, largest sizes, low precision	2.99(19)	2.65(7)	6.9(79)
μ , experiment, largest sizes, high precision	2.67(29)	2.79(24)	6.9(49)
μ , experiment	2.83(24)	2.72(16)	6.3(9)

TABLE II. **Critical exponents.** We compare experimental (QPU) measurements of μ based on collapse of the Binder cumulant U against predictions from independent estimates based on classical Monte Carlo (square and cubic) and theoretical (biclique) for the three relevant universality classes. We perform collapses using all sizes available (square: $L = 4$ to 18; cubic: $L = 3$ to 6; biclique: $M = 8$ to 24), and the three largest sizes available for each topology, to take into account small-size distortions. For square and cubic our final estimate is the large-size average between high and low precision. For biclique our final estimate is the all-size average between high and low precision, since large-size estimates have very large uncertainties.

and related quantities—including the Binder cumulant—as a function of quench time. More specifically, quantities will be related by the universal critical exponents reflecting the universality class of the QPT. This approach has

previously been used to compare quantum simulations with theoretical expectations in superconducting quantum annealers^{23,24,142,143} and annealed Rydberg arrays^{25,57}.

Fig. 5B shows data collapses for the Binder cumulant

$$U = \frac{1}{2} \left(3 - \frac{[\langle q^4 \rangle]}{[\langle q^2 \rangle]^2} \right), \quad (\text{S9})$$

where $[\cdot]$ indicates an average over disorder realizations. The Binder cumulant gives a statistical signature of phase transitions. In this work, we use U to estimate—via collapse—the Kibble-Zurek exponent

$$\mu = z + 1/\nu, \quad (\text{S10})$$

where z and ν are the dynamic exponent and the equilibrium correlation-length exponent pertaining to the QPT at hand²⁸. That is, the paramagnetic-to-spin-glass QPT in the universality class for each topology. These are 2D Ising, 3D Ising, 3D Ising, 3D Ising, and mean-field (fully connected) for square, cubic, cubic no-dimer, diamond, and biclique, respectively.

At the critical point in the quench, it is expected⁵⁹ that both $\langle q^2 \rangle$ and U should collapse (separately) as a functions of L and t_a . While $\langle q^2 \rangle$ is less statistically sensitive, collapse of U is preferable here for two reasons. First, $\langle q^2 \rangle$ was found²⁴ to be susceptible to distortion from post-critical dynamics between the QPT and the end of the quench, where the output state is read. Second, $\langle q^2 \rangle$ must be collapsed as a function of both μ and the anomalous dimension η ; the collapse of U is only in one parameter μ , making the fit more straightforward in the face of limited system sizes and uncertainty in η .

Thus we measure U for multiple linear system sizes L and quench times t_a . For each topology (and treating high-precision and low-precision ensembles separately), we expect data collapse when plotting U against $t_a L^{-\mu}$. We therefore minimize deviation between the collapsed data and a common collapse curve of the parametric form

$$f(x) = a_0 + a_1 \log(1 + \exp(a_2(x - a_3))) \quad (\text{S11})$$

in the $(\log(U), \log(t_a L^{-\mu}))$ plane (as in Ref. 24, Fig. 3c–d), where for bicliques we take $L = M^{1/8} = (N/4)^{1/8}$ based on the upper-critical dimension 8 (Ref. 61). Estimates of μ can be impacted by the effect of small inputs in the best-fit collapse; we therefore estimate μ using the three largest simulated sizes of each topology, as well as estimating using all sizes. The results, along with related literature values, are summarized in Table II.

Precise estimates of U require averaging over many spin-glass realizations (often called “samples” in the spin-glass community). Therefore instead of focusing on 20-instance ensembles, which allow us to keep classical compute resources within reason, we simulate many realizations. For each of 300 QPU calls, we generate a distinct random realization for each individual parallel embedding into the qubit lattice. Thus the number of spin-glass instances simulated in QA for this analysis ranges from 300 (for the largest square and cubic instances) to 20,400 (for the smallest instances, where many random realizations are run in parallel in each QPU call).

Previous studies have provided independent estimates of z and $1/\nu$, allowing us to validate the scaling of correlations in the system beyond the classically verifiable regime. We compare our experimental estimates of μ against the results of classical Monte Carlo simulations (for square and cubic) and mean-field theory (for biclique).

For 2D and 3D, there has been some controversy over whether or not z is finite. For 2D, this was settled very recently³⁶ through the consideration of distinct parity sectors in the state space, using MC simulations of low-precision $L \times L$ square lattices (with fully periodic boundaries) with L ranging from 8 to 24. The authors extracted distinct estimates for the finite-size measurement (fixing their fitting parameter $\eta_e^0 = 0$) and the large-system limit of, respectively,

$$\mu = z + 1/\nu = 1.7(2) + 0.9(1) = 2.6(3) \quad (\text{S12})$$

and

$$\mu = z + 1/\nu = 2.46(17) + 0.71(24) = 3.27(41), \quad (\text{S13})$$

both of which are in remarkably good agreement with our all-size and largest-size estimates of μ from low-precision glasses (2.67(30) and 2.99(19) respectively). Future experiments in a larger QPU might clarify finite-size effects in these estimates. The classical estimates of Ref. 36 updated the previous literature estimates¹⁴⁴ of $z = 1.50(5)$ and $\nu = 1.0(1)$ —which are by now outdated in terms of computational reach—with extensive state-of-the-art GPU simulations. When considering small-system effects, the three estimates of μ (2.6(3), 2.5(2) and 2.67(30) from, respectively, Ref. 36, Ref. 144, and this work) concur extremely well.

For 3D, the finite- z scenario is supported by considerable evidence, including the above result for 2D and large-scale experiments in quantum annealers (this work itself and Ref. 24), and we may reasonably take the position that z

	Cubic no-dimer	Diamond
μ , literature	2.85	2.85
μ , experiment, largest sizes, low precision	2.61(23)	2.95(62)
μ , experiment, largest sizes, high precision	2.98(12)	2.96(15)
μ , experiment	2.80(18)	2.96(39)

TABLE III. **Critical exponents for cubic no-dimer and diamond glasses.** Estimates of μ are expected to be consistent with the 3D Ising value.




FIG. S30. **Dynamic finite-size scaling collapse for diamond and cubic no-dimer instances.** Curves of collapsed Binder cumulant U for diamond (green) and cubic no-dimer instances (orange).

is indeed finite. The prevailing classical estimate of $z = 1.3$ is, like the previous 2D estimate¹⁴⁴, computationally outdated and could benefit from an extension of the methods and computational resources brought to bear in Ref. 36. The estimate of $1/\nu = 1.55(11)$ has recently been updated for the purposes of evaluating QA²⁴. Our estimates of μ , summarized as $\mu = 2.72(16)$, are again in line with previous estimates. The fact that this estimate is slightly lower than that reported by Ref. 24 may be due to reduced noise in the ADV2 processor as compared to the ADV1 processor used previously, with possible influence of smaller system sizes used in the collapse in this work.

For biclique, we expect behavior in line with mean-field critical exponents $z = 2$ and $1/\nu = 4$, which would mean a collapse upon rescaling t_a by $M^{-\mu/8} \approx M^{-3/4}$ (Ref. 61). There are large per-size fluctuations relative to the other topologies; this may be related to per-size configurations of inter-dimer coupling structure. The estimates based on $M = 8$ to 24 are nonetheless very close to mean-field expectations of an infinite-dimensional topology.

In addition to the problem classes considered in the main text, we apply the same analysis to cubic no-dimer lattices, with collapses shown in Fig. S30 and compared with diamond lattices. We expect both of these to exhibit behavior consistent with the 3D quantum Ising spin glass universality class. For cubic no-dimer, we consider the same problem sizes as for cubic dimer. Both cubic no-dimer and diamond inputs admit estimates of μ consistent with the 3D quantum Ising spin-glass universality class, as summarized in Table III. The slightly higher estimate compared to cubic dimer—while still within uncertainties—may be related to the larger effective system sizes.

For 1D systems, we estimate μ by measuring kink density $\bar{n}$ as a function of t_a for a large system ($L = 512$, Fig. S31). As in previous works^{23,25}, the measured scaling exponent is consistent with the expected KZ defect scaling $\bar{n} \propto t_a^{-1/\mu}$ with $z = 1$ and $\nu = 1$.

Taken as a whole, the agreement in critical scaling, which relates time and correlation length via the Kibble-Zurek mechanism, provides indirect validation of QA results beyond the classically simulable regime. This provides a complement to the precise quantitative agreement shown for small problems.

X. MONTE CARLO DYNAMICS

Tensor-network approaches provide faithful approximations to quantum dynamics, and generate accurate sample distributions in the appropriate parametric limits. There are other methods, especially various Monte Carlo (MC) dynamics, which have been proposed to mimic the output of quantum annealers. Here we discuss several such MC dynamics and show that they cannot emulate coherent out-of-equilibrium quantum dynamics. The dynamics we discuss here are simulated annealing (SA)¹⁴⁵, simulated quantum annealing (SQA) with path-integral Monte Carlo^{146,147}, spin-vector MC (SVMC)^{148,149}, and SVMC with transverse-field freeze-out (SVMC-TF)¹⁵⁰.




FIG. S31. **Kink density in pure 1D quantum Ising chain.** Density of kinks (frustrated bonds) are shown as a function of t_a for a pure antiferromagnetic 1D chain with length $L = 512$. Assuming a kink density proportional to $t_a^{-1/\mu}$ yields an estimate $\mu = 2.04(5)$, in agreement with theory.




FIG. S32. **Dynamic finite-size scaling for Monte Carlo dynamics.** (A) Binder cumulant collapses for high-precision instances as in Fig. 5B for SA, SVMC, and SVMC-TF dynamics. Blue, yellow, green, and purple data are from square, cubic, diamond, and biclique lattices respectively. (B) Extracted KZ exponents μ , with quantum estimates indicated in black. For biclique problems we use the quantum upper-critical dimension of 8 rather than the classical value of 6, since we are testing consistency with quantum dynamics. Dark, medium, and white shaded markers indicate SA, SVMC, and SVMC-TF values respectively. Estimates are taken from high-precision instances only.

On one hand, these methods are computationally inexpensive compared to tensor-network approaches, and therefore can be run on all system sizes used in our QA experiments. On the other hand, they have already been shown to fail various emulation tests in toy problems^{151,152}. In particular, all of these dynamics fail in one way or another to qualitatively reproduce quantum dynamics even in a pure one-dimensional system²³ (see SM for Ref. 23, Figs. S4 and S5).

Since these methods can be run at large scale, we can collapse the Binder cumulant for these dynamics, as in Fig. 5B, and compare against quantum mechanical values. This was already done in Ref.²⁴, which showed clearly distinct Kibble-Zurek exponents for QA, SQA, and SA in (dimerized) cubic problems. As a consequence, QA was shown to have a dynamical advantage in optimization for these spin glasses. We now repeat a similar test for SVMC and SVMC-TF to further solidify the deviation between these classical dynamics and the Schrödinger dynamics realized within QA. For SA we use a geometric cooling schedule as in Ref.²⁴. For SVMC and SVMC-TF we use the model quench schedule with a low temperature equal to $\Gamma(s^*)/64$, where s^* is the crossing point with $\Gamma(s^*) = \mathcal{J}(s^*)$. SA is run up to $t_a = 2^{15}$ Monte Carlo sweeps, and SVMC and SVMC-TF are run up to 2^{17} Monte Carlo sweeps. For such slow anneals we do not expect sensitivity to the shape of the annealing schedule.

Results are shown in Fig. S32 for SA, SVMC, and SVMC-TF for high-precision problems only. SQA is far slower, so we disregard it since its consistency with quantum quench dynamics has already been refuted in 1D and 3D systems. For the three dynamics, we find consistent Binder cumulant estimates that are much higher than quantum values.

XI. SPATIAL DECOMPOSITION TECHNIQUES

Experimental evidence and theory presented throughout the paper are consistent with the notion that in some of our experimental systems the correlation length, ξ^{VII} , is not system spanning, particularly in higher precision (lower-energy scale) models at shorter time scales, and in lower dimensions. Intuition suggests that correlation length, more than system size, should dictate the complexity of approximate simulation.

This is exploited in our PEPS methods, which find success in square lattice models on short time scales. PEPS provides a globally consistent ansatz for the quantum state, which, given efficient contraction allows for estimation of both local and global observables. For calculation of local observables, some neighborhood (proportioned to ξ) contraction is expected to be sufficient.

For purposes of estimating local observables, even simpler approximations (than neighborhood-limited PEPS contraction) may suffice. If correlation length ξ is short throughout some dynamical process, we might truncate the system to a cluster of scale proportional to ξ and only simulate dynamics on such a cluster, with statistics at its center well approximated. Unlike a PEPS neighborhood contraction, there needn't be an approximation of quantum phenomena beyond the cluster boundary, or at the boundary (in the following we assume that the clusters have open boundary conditions). Such subsystems might be simulated with an MPS or PEPS method, or by the QPU itself. In the context of MPS, the complexity of simulation can then be reduced in proportion to the area-law of the cluster with the error being controlled by a reduced bond dimension due to the smaller cluster size. In the context of the QPU, we might aim to estimate statistics associated with systems that are beyond a programmable scale. A strength relative to PEPS is potentially the simple nature of such a *cluster* approximation.

In both QPU and MPS cases, we might locate clusters centered at various points on a lattice, and then estimate correlations at the center of such clusters. We establish this for MPS sampling in [XIA](#) and for QPU subsampling in [XIA](#) and [XIB](#). In [XIC](#) we consider a simple extrapolation method, based on such statistics, that allows accurate estimates of global statistics in some circumstances.

A. Cluster approximations with MPS

We already showed in the main text (Figs. 2 and 4) that the bond dimension required for MPS simulations to achieve the QPU accuracy follows area-law growth. Hence, naively it is expected that with increasing number of qubits the cost of MPS simulations attaining QPU accuracy should eventually go beyond the capability of classical computational resources. In this section we use MPS simulations on multiple open-boundary clusters around a chosen site inside an $L_x \times L_y$ square lattice cylinder, periodic in y , at a scale $L_x = L_y = 8$; see Fig. S33A. Panels B ($t_a = 7$ ns) and C ($t_a = 20$ ns) compare a subset of ground-truth spin-spin correlation estimates against those obtained with square clusters of sizes $L_c \times L_c$ with $L_c = 4, 5, 6$. For $t_a = 7$ ns, increasing accuracy is apparent as L_c grows, with a good quantitative match in the majority of nearest and next-nearest neighbor correlations already at $L_c = 5$. Accuracy is reduced as t_a increases towards 20 ns (Fig. S33C). This gives a qualitative understanding that for higher t_a we need larger clusters to reproduce accurate short-range correlations, in line with increasing correlation lengths.

The cluster position on the full lattice can be displaced to attain all short-range correlations. The quantum dynamics of all these small clusters can be trivially parallelized leading to only linear growth in the numerical cost with the system size (N for square clusters, L for boundary-spanning stripe clusters). A similar technique has been used extensively for performing Monte Carlo simulations on quantum-classical hybrid models¹⁵³ under the name of traveling cluster approximation (TCA). We study the accuracy of TCA for quantum annealing by calculating ϵ_c attained by the TCA against MPS simulation of the full system (Fig. S34). We tried three different cluster shapes: square clusters ($L_c \times L_c$), open stripes ($L_c \times L_y$), and periodic stripes ($L_x \times L_c$). We noticed all three cases perform well for small t_a . For example, for $t_a = 0.3$ ns, all three TCAs for the 6×6 system with cluster of linear length 4 give QPU-level accuracy, and the clusters of length 5 can produce trustworthy results up to $t_a = 2$ ns. The correlation error grows as we increase t_a . This is expected because the correlation length increases with t_a and the longer-range correlations, inaccessible to the small cluster, also becomes relevant. We also found periodic stripes perform best, followed by the open stripes and then square clusters. However, even at equivalent size, the cost of simulation of obtaining high accuracy results is a non-trivial function of the topology that must be considered. Periodic stripes require support for boundary spanning entanglement that must be accommodated with higher bond dimension. Similarly, the open stripes are costlier than square clusters of equivalent width not only in proportion to their additional length, but

^{VII} As discussed elsewhere, a correlation length scale might originate in the Kibble-Zurek mechanism, we call the Kibble Zurek correlation length $\hat{\xi}$. For purposes of this section we need not assume

the Kibble-Zurek mechanism to be the source of correlations, or assume a scaling form.




FIG. S33. **Short-range correlations.** Panel **A** shows the 6×6 (red), 5×5 (green), 4×4 (gray) clusters embedded in the 8×8 system, around site number 37. Correlations with respect to site=37, $c_{37,i}$, are shown for various cluster sizes in **B** ($t_a = 7$ ns) and **C** (20 ns). Correlations are only shown with respect to the central sites compatible with all clusters.

due to the higher bond dimension required to accommodate higher entanglement entropy (anticipated logarithmic corrections in the entropy as the system is passing through the QCP in the annealing process).




FIG. S34. **Correlation matrix error.** Panels **A**, **B**, and **C** show ϵ_c for various t_a as a function of cluster size using the square cluster, open stripes, and periodic stripes, respectively. All results are for the full system of 6×6 , except red dashed lines, which depict the ϵ_c of TCA on the 8×8 system. Legend values are provided in nanoseconds.

B. Cluster approximations with the QPU

As shown in the previous section, MPS methods become computationally expensive for large systems because even though correlation lengths are short (for the annealing times of our annealing experiments), they can still encompass many qubits. When MPS simulations become unfeasible, one can then think of using the QPU as a *solver* for each the traveling cluster. It is therefore possible to use QPU subsampling to explore the relationship between large systems




FIG. S35. **Rotating clusters in 2D.** For fast quenches, where final states have only short-ranged correlations, spoofing approaches can be based on simulating cheaper clusters—with one dimension having width W smaller than the linear system size L —within the full lattice, then recombining output. We evaluate this approach for square, cubic, and cubic no-dimer lattices. Shown are examples of cluster shapes simulated in an effort to spoof 18×18 square lattice results.




FIG. S36. **Error from rotating clusters in 2D.** (A) For 8×8 square lattices, we can compute MPS ground truths and measure ϵ_c directly for varying values of W . For $W = 8$, where each cluster is the same size as the full lattice, significant deviations from the full-lattice error level are seen for $t_a > 4$ ns. (B) Here, error in 8×8 systems is computed not from a ground truth, but from a full-lattice quantum simulation. An independently-run full-lattice simulation gives a sampling error pseudo-floor (which can be beaten due to averaging of clusters and smoothing of statistical noise). Error against QPU appears to lower-bound error against ground truth. (C) 18×18 system for which ground truth cannot be computed. At $W = 18$, ϵ_c is indistinguishable from the full lattice simulation even for $t_a = 20$ ns.

and the larger clusters required for their approximation. Moreover, we can compare the QPU statistics with and without clustering approximations. For small systems, we have established that the cluster scale required for accurate estimation in this context is a good proxy for the cluster scale required for emulation of ground-truth statistics.

In this section we show that clustered approximations at these larger scales are indeed sufficient when the cluster volume is large compared to cluster-boundary, and when the correlation length (controlled primarily by annealing time t_a) is relatively small. In some scenarios we show that even clusters at a scale comparable to the full system are insufficient for estimation of local correlations, indicating the presence of significant system-spanning correlations.

1. Two dimensions

For large 2D systems of size $L \times L$ with one periodic dimension, we use the following spatial decomposition to determine spin-spin correlations:

- Choose a block width $W \leq L$.
- In the periodic dimension, delete every W th row of couplers, leaving behind a set of $W \times L$ blocks that span from one open boundary to the other. Do this for every possible offset (0 to W , ignoring divisibility issues which should be handled in the obvious way).




FIG. S37. **Error from rotating clusters in 3D.** (A) Cubic dimer and (B) cubic no-dimer at the largest simulated size (432 qubits), and (C) diamond inputs at the largest defect-free simulated size (243 qubits). As with 2D systems, correlations from full-lattice simulation are compared against composed correlations from multiple simulations with varying cluster size. For $W = L$, the deviation of ϵ_c from the statistical floor (solid lines) indicates that a moving cluster containing all qubits, but with a single row of couplers deleted, is insufficient to capture correlations in the system for all three lattice types beyond short time scales.

- Simulate all the $W \times L$ blocks and determine spin-spin correlations within the blocks.
- Combine the correlations. For two spins that are in a block together, we take the weighted average of blocks that maximize the minimum distance of the two spins to the block cuts. Correlations with spins near the block boundaries can be discarded, or not. Long-distance correlations can be inferred from short-ranged correlations, or simply set to zero.

The cluster shapes are shown in S35, we might call this approximation a rotating-cluster approximation.

In Fig. S36 we show experimental results for $L = 18$ and $W = 3, 4, 6, 9, 18$. Note when $W = 18$ we are performing 18 times as many simulations on the same number of spins, and combining the correlations determined from all rotations of the single cut. Since we do not have a ground truth for $L > 8$, we instead take our reference distribution (and correlations) from the full-lattice simulation. The statistical noise floor, which indicates that correlations are indistinguishable from the full simulation, is given by the “self-error” of two independently generated reference distributions.

As expected this approach yields correlation errors that decrease as W increases. We also see that the signal deviates upwards from the noise floor as t_a increases. This is compatible with correlation length growth. The Kibble-Zurek mechanism anticipates a correlation length scaling as $\xi \propto t_a^{1/\mu}$, where μ is the Kibble-Zurek exponent. This is loosely corroborated by the fact that we require $W \approx 6$ to approach the noise floor at $t_a = 7$ ns, and $W \approx 9$ to approach it at $t_a = 20$ ns, given that $\mu \approx 3$ in 2D systems. These experiments are based on 20 inputs and 50,000 samples each. The sampling error floor in ϵ_c grows with system size quickly to dominate the signal of interest in our largest experiments, particularly for short correlation lengths (small t_a). Correlations at large distance (large compared to the correlation length) are noise-dominated, and the number of contributions scales as N^2 . These are integrated with equal weight in the numerator and denominator of ϵ_c , which is the origin of the sampling error.

Figure S36A (full lattice) indicates an error threshold that describes QPU quality estimation (of the ground truth). Error appears to be weakly dependent on system size (Fig. 4A), so we might anticipate that a similar scale describes accuracy of estimation also in 18×18 lattice experiments where the ground truth is unknown to us. Fig. S36C shows that, with respect to a time scale $t_a \approx 20$ ns it is possible for 9×18 strip-clusters to obtain correlations sufficient to match this threshold (in estimation of the full QPU distribution). We believe this is a good proxy for the quality a 9×18 -cluster MPS (or similarly high-fidelity) simulation might achieve with respect to a ground truth distribution at 18×18 scale. We cannot verify behavior at smaller t_a owing to sampling error, but trends in ϵ_c as a function of strip width and t_a established for smaller systems (8×8 , without the limitation of sampling error) indicate that strip width $W = 9$ may also be successful at shorter times. Thus a distribution of small ϵ_c , comparable to the QPU, can be established without a full-system simulation. The area-law scaling for strips of fixed width is anticipated to be at most logarithmic (in strip length), and so only modest growth in MPS resources might be required to simulate local correlations accurately for $L \gtrsim 9$, implying a deviation from the extrapolation of Fig. 5D for square lattices.

The estimation of local correlations sufficient for small ϵ_C does not in itself solve the task of fair sampling in the z -basis, but we might anticipate that it opens avenues to do so. One simple approach that shows promise is explored in Section [XIC](#). Fortunately for our demonstration, and as already established in the context of PEPS, square lattice problems are relatively easier than other lattices we simulate. In Section [XIB 2](#) we establish that for cubic lattices at the maximum experimental scale, such clustering approaches are not successful.

2. Three dimensions

One can extend the square lattice approach to three dimensions in a straightforward way, decomposing a 3D system with planar cuts as we decomposed a 2D system with linear cuts. Here we do so, and show results in Fig. [S37A](#) for an $L \times L \times L$ dimerized cubic system with $L = 6$. As in 2D, we see that deviation from the statistical floor increases with t_a . Indeed, it appears possible to spoof correlations fairly accurately with $W = 4$ up to $t_a = 7$ ns. However, at 20 ns there is significant deviation even at $W = 6$, where any two spins are inside the bulk of a cluster.

We repeat the approach for cubic no-dimer inputs, which—exactly like the diamond lattice—are four-regular and have girth (shortest cycle) six. As in the cubic dimer case, we see in Fig. [S37B](#) a separation for $W = L = 6$ at $t_a = 20$ ns, although it is less pronounced. For diamond lattices, we apply this approach to the largest defect-free lattice size, and find in Fig. [S37C](#) a very strong effect even at $W = L = 12$ —this is likely in part due to the nature of the cluster boundary, which sees each spin lose two of four incident couplers, rather than one of four in cubic no-dimer lattices, and one of six in the dimerized cubic lattice. Unlike in the square lattice case, errors established by clustering approximations are in all cases well above the threshold we anticipate to characterize the QPU (based on extrapolation of smaller system data, Fig. [4A](#)).

These results suggest that spoofing results for cubic, cubic no-dimer, and diamond systems by spatial decomposition is unpromising in this parametric regime. The simplest explanation for this is the presence of significant system-spanning correlations; the presence of which place limitations on a variety of related approximation methods.

C. Classical extrapolation of statistics from local to global

Spatial decomposition techniques typically allow estimation of local observables such as correlations and marginal probability distributions. However, inference of global properties is technically challenging on the basis of this information, particularly in the absence of a simple parametric model for the global distribution. In this section, we show that simple heuristics can, in some circumstances, extrapolate from (sufficiently accurate and numerous) local statistics of the projected distribution to macroscopic properties, with some tolerable error. In particular, we demonstrated examples in which fully visible Boltzmann distributions trained on accurate local pairwise correlations can be used to infer long-range correlations and approximately sample full states with an accuracy comparable to the QPU (in sampling a ground truth).

A fully visible classical Boltzmann machine, also called an exponential model, takes the form of $P(x) = \frac{1}{Z} \exp(\theta_i f_i(x))$, where Z is a normalization constant also called the partition function. The parameterization of this model can be inferred from expectations $\langle f_i(x) \rangle$. Supposing the statistics are geometrically local (e.g. $f_i(x) = x_i x_j$, for variables nearby on some lattice), then only local statistics are necessary. For smooth probability distribution described by essentially-classical correlation decay at larger geometric distances, a BM might provide good approximations (provided suitable choices are made for the set $f_i(x)$). Whilst it is technically NP-complete to train from sufficient statistics, training of a lattice model at the scales we present is practically straightforward, and methods such as maximum pseudolikelihood inference offer alternative efficient routes to model inference.

Assuming long-range correlations are weak and well described classically, and that we have access to accurate local observables, then we can train our Boltzmann machine from local observables. A complex pattern of local correlations in the projected distribution, perhaps induced by entanglement, might be reproduced by a sufficiently complex set of $f_i(x)$. The full distribution, and statistics thereof, are then smooth extrapolations consistent with the local observables. The Boltzmann distribution can, for practical purposes, be sampled by Markov chain Monte Carlo (short correlation length implies short autocorrelation time) and so we can reproduce any statistic amenable to Monte-Carlo sampling with some systematic error reflecting the restriction in the choice of f_i and quality of the training data.

Sections [XIB 1-XI A](#) demonstrate MPS and QPU methods for obtaining local correlation estimates. PEPS methods based on the neighborhood contraction might also provide a basis for extrapolation in principle.

We found in our analysis of $L \times L$ square high and low-precision systems up to a scale 8×8 that achieving small ϵ_c , for emulation of either the QPU or ground-truth distribution, was equally challenging (the error threshold obtained in training was similar for a fixed bound on the scale of local observables). Furthermore we were able to establish a




FIG. S38. **Error by extrapolation from local correlation estimates on 8×8 cylinders.** (A) $t_a = 7$ ns and (B) $t_a = 20$ ns. The divergence of the QPU from the ground-truth distribution is shown by the blue horizontal line. If the QPU distribution can be reconstructed from a rotating-cluster approximation with similar accuracy, we anticipate an equal-scale clustering approach can also be used for ground-truth estimation, as verified at smaller scale. Correlations on 8×8 cylinder models can be estimated by rotating clusters with one cut (Fig. S35 right); correlations are more accurate furthest from the cut. For each correlation, we can select a rotating cluster with the cut maximally separated from the site coordinates, yielding an error shown by the black horizontal line. We can select the best (spatially local) correlations from either the approximation (QPU[W=8]), or the target distribution (QPU), and train a Boltzmann distribution. As we incorporate correlations of increasing geometric distance, we can improve the quality of estimation relative to the raw correlation data. However, approximation data becomes increasingly errant, and so this only works for QPU[W=8] up to a limited scale. At 7 ns inclusion of ideal data up to geometric distance 3 is sufficient to reproduce the distribution with comparable accuracy to the ground truth, the effects of the boundary in the rotating cluster lead to a deviation upwards, so that ϵ_c remains just above the target threshold. At 20 ns, longer distance is required for the ideal data [QPU] to achieve low error, compatible with a longer correlation length, and the approximating data QPU[W=8] is relatively poor leading to weak approximation.

strong correlation between ϵ_c , classical fidelity and other distributional measures—measuring long range correlations as captured by ϵ_c was established as a good proxy for accurate reproduction of the full distribution. With this assumed, we present an analysis based on accurate estimation of ϵ_c only. Since observables approximated by large clusters are less expensive to create for QPU distributions, relative to MPS-based data, we present the approximation for the QPU generated approximations (XIB 1), understanding this to hold qualitatively also for approximation of the ground truth.

Figure S38 shows the result of training an Boltzmann machine to infer all correlations (and state probabilities) from local correlations within some geometrically constrained distance.^{VIII} For the 8×8 cylinder, the quality of the QPU distribution is captured by ϵ_c , measuring the distance from the ground truth to the QPU. For our approximation, we measure the gap between the cluster-based approximation and the QPU (as the ground truth). When the latter error is smaller than the former, our cluster method can be considered to have succeeded in emulating the global distribution. At 7 ns we show that sufficiently accurate calculation of pairwise correlations up to Euclidean distance $\lesssim 3$ on square lattice at 8×8 scale is sufficient to train a Boltzmann machine with high accuracy. At 20 ns, we do not see such a success, presumably due to longer correlation lengths. We present the training result based on both ideal data (from the full QPU distribution) and the estimates based on the rotating-cluster approximation. When incorporating data from sufficiently large geometric distance we expect the boundary effect to degrade performance of Boltzmann models trained with the latter set of correlations; this is found to be the case. At 7 ns the accuracy of local statistics is excellent, and we see little degradation in the performance out to distance 3. By contrast at 20 ns the correlation estimates are weaker, and we see an immediate deviation between the trained estimates.

As shown in the main text, our distribution can become less entangled at larger t_a , particularly as we cross over to the adiabatic regime with system-spanning correlations. Less-entangled distributions favor this classical approximation method—assuming local observables may be accurately estimated. However, correlation lengths increase with t_a , degrading the quality of local estimates attainable by cluster approximations. The quality of local estimations is also impacted by dimensionality: for fixed correlation length the boundary effect (relative to the cluster bulk) has a more favorable ratio in low dimensions, in addition to the fact that correlation lengths can grow more quickly. In this section we have shown that at 7 ns, 2D high-precision spin-glass problems at scale 8×8 and larger may be susceptible

^{VIII} Geometric distance being the Euclidean distance in the standard square lattice representation. For geometric distance 3 we in-

clude only pairs separated by at most 2 couplers (graph distance 2); in other cases all correlations are included.

to clustering approximations, even where global statistics are concerned. For other experimental systems with larger correlation length relative to linear scale (20 ns experiments, lower-precision, and/or higher-dimensional models) the method is understood to fail.