

PHYSICS

Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance

Tomislav Begušić†, Johnnie Gray†, Garnet Kin-Lic Chan*

A recent quantum simulation of observables of the kicked Ising model on 127 qubits implemented circuits that exceed the capabilities of exact classical simulation. We show that several approximate classical methods, based on sparse Pauli dynamics and tensor network algorithms, can simulate these observables orders of magnitude faster than the quantum experiment and can also be systematically converged beyond the experimental accuracy. Our most accurate technique combines a mixed Schrödinger and Heisenberg tensor network representation with the Bethe free entropy relation of belief propagation to compute expectation values with an effective wave function–operator sandwich bond dimension $>16,000,000$, achieving an absolute accuracy, without extrapolation, in the observables of <0.01 , which is converged for many practical purposes. We thereby identify inaccuracies in the experimental extrapolations and suggest how future experiments can be implemented to increase the classical hardness.

INTRODUCTION

As quantum computers mature, it is critical to benchmark their performance against classical simulations. One approach is to compare against exact classical simulations, such as state-vector methods that store the wave function amplitudes of the n qubits as a 2^n -dimensional vector, or exact tensor network (TN) contractions (1–6), which have an exponential cost in circuit tree width (7). However, as qubit counts and circuit depths increase, quantum experiments are now beginning to exceed the capabilities of exact classical simulations.

A recent quantum experiment of this kind simulated the kicked Ising model on IBM's 127-qubit processor (8). Using zero-noise extrapolation (9) of the experimental data, Kim *et al.* (8) presented evidence for the utility of quantum computing before fault tolerance. Starting from the ferromagnetic state, the experiment implemented the circuit

$$U = \left(\prod_{(j,k)} e^{i\pi Z_j Z_k / 4} \prod_j e^{-i\theta_h X_j / 2} \right)^T \quad (1)$$

consisting of T steps of Clifford ZZ and non-Clifford X rotations. In Eq. 1, θ_h is a rotation angle that is varied, j and k run over all 127 qubits, and $\langle j, k \rangle$ denotes neighbors on the heavy hexagon lattice of the Eagle quantum processor `ibm_kyiv` (Fig. 1A). It was argued in (8) that the expectation values $\langle O \rangle = \langle 0 | U^\dagger O U | 0 \rangle$ of different Pauli observables O were not only beyond the reach of exact classical simulations but, furthermore, could not be computed to the experimental accuracy with certain approximate classical techniques, such as matrix product state (MPS) or isometric two-dimensional (2D) TN simulations (10).

However, in a recent unpublished, preliminary note by some of us (11) using an approximate classical method termed sparse Pauli dynamics (SPD), we showed that we could compute the experimental expectation values with an accuracy comparable to the experimental

extrapolation, in a time faster than the experiment. The same conclusion was reached in multiple preprints from other groups (12–16), including some of the authors of the original experiment (17). Nonetheless, as pointed out in (17), the remaining errors in the classical simulations, as well as in the extrapolated experimental data, still left considerable uncertainty in the precise values of the observables targeted in the original experiment.

Here, we show that multiple approximate classical methods can be converged for all examples studied in (8), including the 20-step dynamics involving 127 qubits, well beyond the results in the initial preprints (12–14, 17) and with a remaining uncertainty substantially less than the experimental extrapolation error bars. To this end, we provide a full account and further extend the application, of the Clifford-based SPD in (18) to larger-scale simulations. We further introduce new TN techniques that use lazy belief propagation (BP), in the Schrödinger picture [for projected entangled pair states (PEPSs)], in the Heisenberg picture [for projected entangled pair operators (PEPOs)], and in a mixed Schrödinger/Heisenberg picture (PEPS/PEPO expectation value). An important component in these improved TN simulations is the application of the Bethe free entropy BP formula, extending the ideas in (19), that allows us to completely avoid the expensive TN contraction and use large bond dimensions in the mixed Schrödinger/Heisenberg picture. We carefully benchmark the errors associated with these techniques and introduce a uniform norm metric to compare the quality of the different approximations. In our most accurate simulations of the largest circuits in the mixed picture (Fig. 1B) and using the free entropy formula, we effectively evaluate a TN representation of the expectation value with bond dimension of 16,777,216. Without resorting to any extrapolations, the associated absolute error that we achieve in the targeted observable is conservatively estimated as <0.01 . For many practical applications, this may be considered fully converged.

RESULTS

Approximate classical methods for quantum circuit expectation values

Various classical strategies have been developed to approximate quantum circuit expectation values with a cost that depends on

Copyright © 2024 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).

Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125, USA.

*Corresponding author. Email: gkc1000@gmail.com

†These authors contributed equally to this work.

characteristics of the quantum circuit. For example, Clifford-based approaches (20–24) can treat circuits with large numbers of qubits and highly entangled states as long as there is only a small number of non-Clifford gates. Similarly, approximate TN methods provide accurate results so long as the circuit generates a limited amount of entanglement. In practice, the goal with such approximate techniques is to numerically converge the results with respect to a parameter controlling the accuracy. Below, we describe approximate simulation methods in both the near-Clifford and TN categories and the innovations that we use to converge the expectation values of the kicked Ising experiment.

Sparse Pauli dynamics

The SPD approach (11) was recently introduced as a variant of Clifford perturbation theory (18). It is based on representing a target observable in the Heisenberg picture as a sum of Pauli operators $O = \sum_{P \in \mathcal{P}} a_P P$, where $\mathcal{P}$ is some subset of n -qubit Paulis and a_P are the corresponding complex coefficients. A Pauli rotation gate $U_\sigma(\theta) = e^{-i\theta\sigma/2}$, defined by a real rotation angle θ and a Hermitian n -qubit Pauli operator σ , transforms any n -qubit Pauli operator P according to

$$U_\sigma(\theta)^\dagger P U_\sigma(\theta) = \begin{cases} \cos(\theta)P + i\sin(\theta)\sigma P, & \{\sigma, P\} = 0 \\ P, & [\sigma, P] = 0 \end{cases} \quad (2)$$

Consequently, for each Pauli operator $P \in \mathcal{P}$ that anticommutes with σ of the rotation gate, the representation of the observable O must be expanded to $\mathcal{P}' = \mathcal{P} \cup \{\sigma P\}$. This will, in general, lead to an exponential growth of the number of terms in the sum. However, if

the rotation gate is Clifford, meaning that $\theta = k\pi/2$ for integer k , then only one of the terms in the right-hand side of Eq. 2 will remain and the number of Paulis will not increase. In SPD, we make use of this by initially transforming the non-Clifford Pauli rotation gates and the observable by the Clifford gates of the circuit. Similarly, we rewrite the Pauli rotation angles as $\theta = \theta' + k\pi/2$ with $\theta' \in (-\pi/4, \pi/4]$, where the term proportional to $\pi/2$ is treated as a separate Clifford gate. To truncate the exponentially growing Pauli series representing the Heisenberg-evolved observable, in (11), we used a perturbative criterion for adding or removing Pauli operators. Here, we formulate SPD using a threshold-based criterion, i.e., by truncating the Pauli representation of the observable to those Paulis whose coefficients are greater than the prescribed threshold (see Methods for full details). Then, the accuracy of the simulation is systematically improved by reducing this threshold until convergence. [We note that, recently, other related Pauli-based approaches have been proposed, with truncation based on Fourier expansion or Hamming weights of the Pauli operators (14, 16, 25, 26).]

TN simulations

In our TN simulations, we allow ourselves to partition the unitary between the evolution of the states and that of the observable O . Thus, in general, we evolve both a PEPS representation forward to step τ , $|\psi(0)\rangle \rightarrow |\psi(\tau)\rangle$ and a PEPO backward $\Phi(T+1) \rightarrow \Phi(\tau+1)$ [note $\Phi(T+1) = O$]. In the limit that $\tau = T$, only the PEPS is evolved (Schrödinger picture) and the entanglement is entirely stored in the state. Conversely, if $\tau = 0$, then only the PEPO is evolved (Heisenberg picture), and the entanglement is stored in the operator. By allowing other values of τ (mixed picture), we can limit ourselves to moderate




Fig. 1. Schematic overview of the simulation. (A) Qubit layout of the quantum device used in (8). (B) Visualization of the full TN representation of the quantum circuit expectation value for 20 steps. Green tensors are entangling RZZ gates; orange tensors are RX rotations.

entanglement growth in both the state and operator. A schematic overview of the TN methods is shown in Fig. 2 (further details in Methods and the Supplementary Materials). Below, we refer to the case of purely Schrödinger evolution ($\tau = T$) as the “PEPS method,” the case of purely Heisenberg evolution ($\tau = 0$) as the “PEPO method,” and the case of half and half ($\tau = T/2$) as the “MIX method.” It is this MIX method that we find converges fastest and we show is the most accurate.

We use two key BP tools to perform these evolutions efficiently and, crucially, evaluate the final overlap $\langle \psi(\tau) | \Phi(\tau + 1) | \psi(\tau) \rangle$ at bond dimensions, χ , that are far beyond exact contraction.

The first is 2-norm BP compression. 2-norm BP compression is a version of simple update compression (27–29) where TN bonds are compressed by the singular value decomposition of the contraction of pairs of tensors, and the rest of the TN (the environment) is represented as a product of gauges. To obtain these gauges, the BP iteration is applied to the doubled TN used to compute the 2-norm (i.e., Fig. 2Bi for PEPS compression and Fig. 2Bii for PEPO compression). As well summarized in (30, 31), at the convergence of the BP procedure, the BP messages correspond to working in the “super orthogonal” gauge defined in (32). The key difference with the usual simple update formulation is that the BP message iteration allows for this gauge to be computed without modifying the original tensors. This then allows for the lazy computation of gauges without forming intermediate tensors of large bond dimension (33, 34). We refer to this lazy implementation of 2-norm BP compression as L2BP. On a lattice such as heavy hex with coordination number 3, it scales as $\mathcal{O}(\chi^4)$. The

second tool, which we refer to as lazy 1-norm BP (L1BP) contraction, is an extension to quantum expectation values of recent work connecting BP techniques and TNs (19). The essential idea is to approximate the final overlap contraction between states and operators as the exponential of the Bethe free entropy (35). To obtain this, we first compute converged BP messages of the sandwich TN of the operators and states (Fig. 2Biii), and the value of the contraction is approximated by a simple function of the contractions between tensors and converged BP messages (Fig. 2E). In the quantum setting, we apply this relation despite the TN being nonpositive. L1BP contraction can be used to compute local and nonlocal operators (such as high weight Paulis) as well as the norms of $|\psi\rangle$ and Φ . On the heavy-hex lattice, assuming both ψ and Φ have bonds of size χ (this refers to the MIX method), it scales as $\mathcal{O}(\chi^6)$. Using it, we here evaluate the expectation value of a PEPS/PEPO expectation value where the PEPS and PEPO each have bond dimension $\chi = 256$, for a combined effective TN expectation value bond dimension of $\chi = 256^3 = 16,777,216$. L1BP contraction is exact for tree geometries but is only a heuristic approximation away from such geometries. Consequently, we verify the accuracy of the L1BP contraction for our expectation values in multiple ways in the simulations below.

For clarity, we briefly note some differences to other approaches. Tindall *et al.* (12) perform PEPS evolution using 2-norm BP to compress after each gate application. This is subtly different from our PEPS evolution, where, because of our lazy approach, it is natural to evolve and compress one or more full layers of gates simultaneously, which we find improves accuracy. A second difference is that




Fig. 2. TN simulation overview. Full discussion in main text and Methods. (A) schematic of mixed PEPS/PEPO evolution on toy hexagonal lattice. The PEPS (red) is evolved using the next unitary layer and then compressed using L2BP. The PEPO (pink) is evolved and compressed similarly. Green bonds, “physical” indices connecting a single site. Once all unitary layers have been incorporated, the remaining PEPS/PEPO/PEPS sandwich is contracted using L1BP. (B) BP on the lattice. Blue circles, tensors on sites (which represent lazy contraction of tensors from multiple layers, represented by internal dots). Each bond has left and right tensor messages (yellow/red pairs of dots). We run BP algorithms on three types of TNs that differ in the definition of the site tensors (shown in insets). (i) Site tensor when compressing the PEPS and its unitary update, minimizing the 2-norm error (L2BP method). (ii) Site tensor when compressing the PEPO and its unitary update, minimizing the 2-norm error (L2BP method). (iii) Site tensor when using the Bethe free entropy formula to evaluate the PEPS-PEPO-PEPS expectation value (L1BP method). (C) Lazy BP contraction for updating a single message. (D) Lazy 2-norm belief propagation (L2BP) compression of a local site. From converged messages, pairs of projectors (purple rectangles), from the top right approximation, are found for each bond. Projectors adjacent to each site are then contracted in, yielding single tensors with maximum dimension χ . (E) Lazy 1-norm belief propagation (L1BP) TN contraction. Given converged messages, the global contracted value is estimated from contractions of tensors and messages.

Tindall *et al.* (12) evaluate the observables using the approximate local density matrices computed in the BP approach, rather than our L1BP expression. We discuss the implications of this for the accuracy of $\langle Z_{62} \rangle$ at $\theta_h > \pi/4$ later. Last, in another paper that appeared at the same time as our preprint, Liao *et al.* (15) perform PEPO Heisenberg evolution using simple update (27), before contracting the final PEPO exactly at $\tau = 0$. In our PEPO method, we apply entire layers of gates simultaneously, and the use of L1BP contraction enables us to directly contract Φ with $\log_2(\chi)/2$ layers remaining, avoiding the extra compression steps in (15).

Simulation of expectation values of the kicked Ising circuit

We start with a brief recapitulation of the results in the unpublished note (11), recomputed using the slightly modified SPD algorithm in the current work, as well as from fast but approximate TN simulations. In Fig. 3, we show these classical simulations of the experimental data of (8). The observables include the magnetization $M_Z = \sum_j \langle Z_j \rangle / n$ and two high-weight Pauli observables ($X_{13,29,31} Y_{9,30} Z_{8,12,17,28,32}$ and $X_{37,41,52,56,57,58,62,79} Y_{75} Z_{38,40,42,63,72,80,90,91}$) after five steps ($T = 5$ in Eq. 1), another weight-17 Pauli observable $X_{37,41,52,56,57,58,62,79} Y_{38,40,42,63,72,80,90,91} Z_{75}$ for a circuit consisting of five steps and an additional layer of X rotations, and $\langle Z_{62} \rangle$ after 20 steps of the circuit in Eq. 1. Exact benchmarks, computed with an exact TN simplification and contraction method (36), are shown for the first four examples. [We note that the exact result shown in Fig. 3D was not computed in (8).]

The data generated using SPD are simulated within on average 10 s per point on a single core of a laptop computer and are compatible with the experimental zero-noise extrapolated data. As observed in

(11), this is three orders of magnitude faster than the reported quantum wall-clock run time for Fig. 3 (D and E) (4 and 9.5 hours, respectively) and faster than the hypothesized run time of IBM's Eagle processor without any classical processing steps (estimated at about 5 min for Fig. 3E). We also show data from the MIX TN method with $\chi = 64$. The bond dimension here is chosen such that the most expensive computations (Fig. 3, A and E) complete in less than 3 min of computation time on a single consumer graphics processing unit (GPU) [Fig. 3 (B to D); only require about 10 s per point]. We see again that the TN simulations generate data compatible with the experimental zero-noise extrapolated data faster than the run time of the quantum processor.

We next turn to computing more accurate data using the approximate SPD and TN approaches. To do so, in Fig. 4, we first show the absolute errors of these approximate simulation methods for the examples in Fig. 3 (A to D) corresponding to shallow circuits, where we have access to an exact result. Our fast SPD results already match the exact values very well, with a maximum error over all points on the order of 10^{-3} . The magnetization can be computed numerically exactly with SPD (zero threshold) within ~ 15 min per point on a single core, while the error of the other three observables can be evaluated with a maximum error around 10^{-4} (Fig. 4, B to D, dark blue line), where the most expensive points take up to 45 min, 5 hours, and 6 hours (see table S2), respectively, on four cores in our Python implementation. For these shallow-depth circuits, the TN calculations with $\chi = 64$ (as discussed, taking a maximum of 3 min on the consumer GPU) already result in no truncation; thus, the error is solely due to the 1-norm BP contraction approximation. These errors are shown




Fig. 3. Expectation values of the quantum circuit in Eq. 1. (A) Magnetization; (B) $X_{13,29,31} Y_{9,30} Z_{8,12,17,28,32}$; (C) $X_{37,41,52,56,57,58,62,79} Y_{75} Z_{38,40,42,63,72,80,90,91}$, all after five steps; (D) $X_{37,41,52,56,57,58,62,79} Y_{38,40,42,63,72,80,90,91} Z_{75}$ for a five-step circuit with an additional X -rotation layer, and (E) $\langle Z_{62} \rangle$ after 20 steps. The results presented include the exact benchmarks calculated using TN contraction [“Exact,” (A) to (D)], SPD results obtained within on average 10 s per point on a single core of a laptop computer (“SPD 10s”), MIX TN results on a single GPU under 3 min (“MIX 3 min”), and the zero-noise extrapolated quantum experiment (“Quantum”) of (8). All data apart from the quantum simulations were computed at θ_h between zero and $\pi/2$ in increments of $\pi/32$. Thresholds used for the SPD simulations are reported in table S1.




Fig. 4. Absolute errors of approximate methods for the observables presented in Fig. 3. (A) Magnetization; (B) $X_{[3]}Y_{[2]}Z_{[5]} \equiv X_{13,29,31}Y_{9,30}Z_{8,12,17,28,32}$; (C) $X_{[8]}Y_{[1]}Z_{[8]} \equiv X_{37,41,52,56,57,58,62,79}Y_{75}Z_{38,40,42,63,72,80,90,91}$, all after five steps; (D) $X_{[8]}Y_{[8]}Z_{[11]} \equiv X_{37,41,52,56,57,58,62,79}Y_{38,40,42,63,72,80,90,91}Z_{75}$ for a five-step circuit with an additional X-rotation layer. We show the SPD results obtained within on average 10 s per point (SPD 10 s), converged SPD results ("SPD"), and the three TN methods (B to D): PEPS, PEPO, and MIX. Thresholds used for the converged SPD simulations are reported in table S1.

for the PEPS method, the PEPO method, and the MIX method for the observables computed in Fig. 4 (B to D); in all cases, they are less than 10^{-4} . We now discuss converging the most difficult observable in (8), namely, $\langle Z_{62} \rangle$ with a 20-step circuit. For orientation, we first give results for a simpler nine-step circuit for this observable where we can still compute an exact benchmark. The data are presented in Fig. 5. For the nine-step expectation value (Fig. 5A), all results agree to within 0.01 with the exact TN benchmark. At our used parameters (see table S3 for SPD thresholds and table S4 for TN χ values), the largest errors appear at the highly non-Clifford points in the middle of the studied θ_h range. For the 20-step circuit (Fig. 5B; SPD thresholds in table S3, up to $\chi = 320$; for TN methods, see table S4), we no longer have exact data, but we can still examine the maximum spread across all the methods. We find this to be less than 0.045. Notably, at each computed θ_h , the spread in the classical results is substantially smaller than the error bar in the experimental extrapolation at a similar θ_h .

In Fig. 6, we show the convergence with respect to threshold (for SPD) and bond dimension (for the TN methods). At intermediate θ_h , the SPD data appears to be less converged than the TN data at the smallest accessible values of the threshold. Both Heisenberg picture methods, namely, PEPO and SPD, exhibit a non-monotonic convergence behavior (see fig. S3), similar to the Heisenberg picture

matrix-product operator (MPO) simulations of (17). In turn, PEPS and (to some extent) MIX simulations exhibit a monotonic convergence with respect to χ . Figure 6C shows that the MIX results are the most converged among the three TN methods. In the Supplementary Materials, we provide three estimates of the residual error of the MIX method with respect to χ : These estimates range from well less than 10^{-2} to perhaps as small as 10^{-3} .

It is useful to define a metric that allows us to compare the results of different approximate methods on a more equal footing. For this, we use the norm as a common convergence parameter and an approximate measure of fidelity. For PEPS, we use the wave function norm $N_\psi = \|\psi\|$; for operator-based methods (PEPO and SPD), we use the Frobenius norm $N_O = \sqrt{\text{Tr}(O^\dagger O) / 2^n}$; and for the mixed-picture approach, we use $N_{\text{MIX}} = N_\psi N_O$. These are shown in Fig. 7A as functions of the rotation angle θ_h in the 20-step simulation. Note that achieving an accurate fidelity is expected to be much harder than computing an expectation value; thus, this represents a stringent test of the accuracy of the classical methods. We find that the operator-based approaches (SPD and PEPO) follow a similar trend, where the norm is close to 1 for Clifford and near-Clifford angles and low or even near zero for intermediate values of θ_h . Conversely, the PEPS approach yields near-zero norms at large θ_h , even near the Clifford




Fig. 5. Expectation value of $\langle Z_{62} \rangle$. (A) After 9 steps; (B) after 20 steps of the circuit in Eq. 1, as calculated with SPD and PEPS, PEPO, and MIX TN methods. For the 9-step result, we compare these to the exact benchmark; for the 20-step result, we compare the spread of these classically simulated expectation values with the confidence intervals (CI) of the extrapolation procedure used in the quantum simulation experiment (8). All data were computed at θ_h between zero and $\pi/2$ in increments of $\pi/32$. Thresholds used for SPD are reported in table S3, and bond dimensions used for TN simulations are in table S4.




Fig. 6. Convergence of $\langle Z_{62} \rangle$ after 20 steps. (A) SPD with respect to the threshold; (B) the PEPS, PEPO, and MIX methods with respect to bond dimension χ . The results are presented for a range of θ_h values indicated by different markers. (C) Zoom-in on the convergence of the MIX method at $\theta_h = 6\pi/32, 8\pi/32$, and $10\pi/32$ (see fig. S1 for other values of θ_h).




Fig. 7. Convergence analysis of norm and $\langle Z_{62} \rangle$ after 20 steps. Norm as a function of angle θ_h (A) and the convergence of $\langle Z_{62} \rangle$ after 20 steps for SPD and PEPO (B), and PEPS and MIX (C) methods with respect to the norm. θ_h values are indicated by different markers. Horizontal lines in (B) correspond to the highest- χ MIX results at respective θ_h .

angle $\pi/2$. The MIX approach achieves norms ~ 1 for $\theta_h < 10\pi/32$ and outperforms all other approaches for $\theta_h < 13\pi/32$, although it still yields near-zero norm near $\pi/2$. However, both the PEPS and MIX approaches correctly find that the expectation values vanish at the high- θ_h points, indicating that any complicated entanglement generated by the circuit is not reflected in the expectation value. In Fig. 7 (B and C), we plot the convergence of $\langle Z_{62} \rangle$ against norm. Similar to that in fig. S3, we observe non-monotonic convergence in the expectation value for the Heisenberg-based methods, while PEPS and MIX show smooth and monotonic convergence with norm. Notably, the MIX method achieves a large norm for all nonzero $\langle Z_{62} \rangle$.

The above studies support the convergence of our simulation methods with respect to the threshold/bond dimension and demonstrate, in particular, that estimates from the MIX TN method are particularly well converged with respect to χ . However, as discussed earlier, the TN methods implemented here include an additional approximation in the form of the 1-norm BP (L1BP) contraction that

introduces an uncontrolled error into the expectation value. We now verify the accuracy of this approximation in three separate ways. The first verification is that, for exactly accessible reference quantities such as the depth 5 and depth 9 observables or the 31 qubit circuits in (13), we can directly see that the total error (i.e., including truncation) is $< 10^{-3}$ in all cases. The second verification is that, in the case of the PEPS method, when the bond dimension is too large for exact contraction of the observable but still small enough that we can contract amplitudes, i.e., $\langle x|\psi \rangle$ for a given bit string x , we can compute an unbiased Monte Carlo estimation of the observable by sampling from the state. In all our checks, we find that L1BP contraction is consistent with such estimates, which themselves are accurate to about 0.1 to 1% limited by shot noise (see Fig. 8A). The final check that we can perform is that, for PEPO evolution when the final bond dimension is small enough, $\chi \lesssim 128$, we can exactly contract $\langle 0|\Phi|0 \rangle$. In this regime, we find a discrepancy with L1BP $\lesssim 5 \times 10^{-4}$, as shown in Fig. 8 (B and C). From these checks of the L1BP approximation, as well as




Fig. 8. Verifying L1BP contraction for PEPS and PEPO evolution. (A) Comparison of $\langle Z_{62} \rangle$ computed by L1BP contraction (lines) and via Monte Carlo sampling (markers with error bars) for the PEPS method at 20 steps at a range of χ and θ_h . The sampled points are computed from 10,000 shots each. **(B)** Expectation values $\langle Z_{62} \rangle$ for a 20-step circuit computed with L1BP contraction (y axis) and exact (x axis) contraction at different values of θ_h and for different bond dimensions χ . **(C)** histogram of the absolute errors of all points shown. Note that, to enable exact contraction, we evolve the PEPO all the way to time 0, and the final layer is not compressed leading to a contribution of $\times 4$ to χ . This is slightly different to the PEPO method elsewhere in the text.

the convergence studies with respect to χ and norm, we can conservatively estimate that the MIX TN method, without extrapolation, provides results for the 20-step circuit $\langle Z_{62} \rangle$ to an absolute accuracy of better than 0.01. The most expensive (i.e., largest χ) MIX calculations required to reach this accuracy took approximately 1 day of computation time on a NVidia A100 GPU. A more complete set of timing data is provided in the Supplementary Materials.

We now use our converged estimates of $\langle Z_{62} \rangle$ from the 20-step circuit using the MIX method to assess the accuracy of the extrapolated experimental results as well as results from other classical simulations. In Fig. 9A, we see that most of the experimental data points agree with our results, but the experimental extrapolation at $6\pi/32$ has a significant error outside of its error bars, and that the data at $7\pi/32$ are also at the limit of its error bars. It is interesting to see that these highly non-Clifford circuits, which are in some ways more challenging classically, are also challenging for the experimental extrapolation.

In Fig. 9B, we show the MPS and MPO results from (8) and (17). As already argued in those works, these results are not well converged, and this is confirmed by our accurate data.

In Fig. 9C, we show results from the BP-TN approach of (12). The unextrapolated data ($\chi = 500$) agree well with our data up to $8\pi/32$ but deviate slightly at larger angles, where they are slightly above our MIX results. In (12), $\langle Z_{62} \rangle$ is computed using the approximate local density matrices associated with the 2-norm BP messages, rather than the L1BP contraction. This implicitly renormalizes the PEPS wave function after compression, due to how the local density matrices are defined in terms of the BP messages. We find that such a renormalization leads to convergence behavior from above, as observed in their data around $\theta_h \sim 1$. On the other hand, the MIX method (where we do not renormalize the state) converges from below. As argued from Fig. 7, the unextrapolated MIX results are likely of higher fidelity than the PEPS results near $\theta_h \sim 1$. We note that the extrapolated data from (12) is in good agreement with our unextrapolated MIX data for all θ_h . This is illustrated in the inset that compares convergence at $\theta_h = 0.7$, showing extrapolation for the BP-PEPS method in (12), our own

PEPS method (where the state is not renormalized after compression) and the reference MIX result.

In Fig. 9D, we see that the reduced 31-qubit model of (13) is in good agreement with our results at small and large angles but shows substantial deviation at intermediate θ_h , with a discrepancy of ~ 0.1 . (Figure S7 shows that our MIX approach accurately reproduces the exact 31-qubit result, indicating that the discrepancy in Fig. 9D arises due to the model reduction to 31 qubits.) Next, in Fig. 9E, the observable's back-propagation on Pauli paths (OBPPP) method of (14), which is itself closely related to SPD, shows good agreement with our data for larger angles of $\theta_h > 5\pi/32$ but exhibits larger errors at low θ_h , where all other classical simulations and the quantum simulations agree very well. In the absence of additional convergence data for OBPPP, it is difficult to say whether the agreement at the other points is fortuitous or not. Last, in Fig. 9F, we compare against the PEPO method described in the recent preprint (15), which used simple update style Heisenberg evolution along with exact contraction of the final TN. We show both the highest χ result of that paper and their extrapolated result. We see that, although the largest χ data are in good agreement with our reference data, the extrapolated data in (15) overshoot considerably. We can trace this to non-monotonic convergence of Heisenberg picture methods (clearly seen in the SPD and PEPO data in Fig. 7B; see also inset of Fig. 9F), which renders such extrapolations misleading.

DISCUSSION

Our work demonstrates that classical algorithms can simulate the expectation values of the quantum circuits corresponding to the kicked Ising dynamics experiment on 127 qubits, not only faster than but also well-beyond the accuracy of the current quantum experiments. For this specific experiment, the classical data may be considered for practical purposes to be converged, and this brings to a close the question of quantum utility. However, we also identify that, for certain parameters of the quantum circuits, the accuracy in the classical simulation of expectation values is achieved despite a low estimated




Fig. 9. Comparison of various simulation results for $\langle Z_{62} \rangle$. Expectation values $\langle Z_{62} \rangle$ for a 20-step circuit computed with the reference, MIX method (shown in all panels, labeled “This work”); (A) the zero-noise extrapolated quantum simulation (“Quantum”); (B) MPS and MPO approaches (8, 17); (C) BP-TN method of (12) with $\chi = 500$ and the extrapolated $\chi \rightarrow \infty$ result, the inset compares convergence at $\theta_h = 0.7$ with our MIX and PEPS methods; (D) exact state-vector simulation on a reduced 31-qubit model (13); (E) OBPPP approach of (14) (data used with permission from Z. Liu); (F) PEPO method of (15) with $\chi = 256$ and the extrapolated result, the inset compares convergence at $\theta_h = 7\pi/32$ with that of our MIX and PEPO methods. The estimated error bars of the MIX method (less than 0.01) are smaller than the plot markers. Horizontal black lines in the insets of (C) and (F) correspond to the MIX results at highest available χ , while the dotted cyan (C) and green (F) lines correspond to the extrapolation fits of (12) and (15), respectively.

global fidelity (as measured by the error in the global norm). Identifying physically relevant observables with a stronger sensitivity to the global fidelity, for example, through analyses similar to that in (13), will therefore be key to designing future experiments to demonstrate quantum utility. Last, the advances in classical simulation methods reported here hold promise for many areas of quantum simulation and highlight the rich landscape of approximate classical algorithms that have yet to be explored for the simulation of quantum circuits and quantum dynamics.

METHODS

Sparse Pauli dynamics

Let us consider an observable $O = \sum_{P \in \mathcal{P}} a_P P$ that is a sum of a subset $\mathcal{P}$ of n -qubit Pauli operators P with complex coefficients a_P . To compute its expectation value

$$\langle O \rangle = \langle 0 | U^\dagger O U | 0 \rangle \quad (3)$$

for a quantum circuit U , in SPD, we apply the circuit gates to the observable in the Heisenberg picture

$$U^\dagger O U = \mathcal{U} O \quad (4)$$

where we introduced the Liouville-space unitary operator $\mathcal{U}$. Specifically, we consider circuits composed of Clifford gates and, possibly, non-Clifford, Pauli rotation gates $U_j \equiv U_{\sigma_j}(\theta_j) = \exp(-i\theta_j \sigma_j/2)$. For the rotation gates, we first transform the angles $\theta_j = \theta'_j + k\pi/2$, where

the integer multiple of $\pi/2$ forms a separate Clifford operator and $|\theta'_j| < \pi/4$. We can now write the Heisenberg-evolved observable as

$$\mathcal{U} O = C_N \mathcal{U}_N \dots C_1 \mathcal{U}_1 O \quad (5)$$

where C_j corresponds to Clifford operators and $\mathcal{U}_j$ corresponds to Pauli rotation gates. Because Clifford operators transform any Pauli operator into another Pauli operator, we can apply them to the Pauli rotation gates and observable to arrive at

$$\mathcal{U} O = \tilde{\mathcal{U}}_N \dots \tilde{\mathcal{U}}_1 \tilde{O} \quad (6)$$

where $\tilde{\mathcal{U}}_j$ corresponds to a rotation gate $U_{\tilde{\sigma}_j}(\tilde{\theta}_j)$ with $\tilde{\sigma}_j = \tilde{C}_N \dots \tilde{C}_j \sigma_j$ and, similarly, $\tilde{O} = \tilde{C}_N \dots \tilde{C}_1 O$. This is also known as Clifford recompilation and has been used in the context of stabilizer-state simulations (37). In the remainder, we assume that the observable and Pauli rotation gates have been transformed as above. Applying a non-Clifford Pauli rotation gate $U_\sigma(\theta)$ to the observable $O = \sum_{P \in \mathcal{P}} a_P P$ yields

$$U_\sigma(\theta)^\dagger O U_\sigma(\theta) = \sum_{P \in \mathcal{P}'} a'_P P \quad (7)$$

where the new set of Pauli operators is $\mathcal{P}' = \mathcal{P} \cup \{\sigma P | P \in \mathcal{P}, \{\sigma, P\} = 0\}$ and the updated coefficients are

$$a'_P = \begin{cases} a_P \cos(\theta) + i a_{\sigma P} \sin(\theta), & \{\sigma, P\} = 0 \\ a_P, & [\sigma, P] = 0 \end{cases} \quad (8)$$

In general, $2|\mathcal{P}| \geq |\mathcal{P}'| \geq |\mathcal{P}|$, i.e., the number of Pauli operators representing the observable grows with each applied gate, leading to the worst-case scaling of $\mathcal{O}(2^N)$. As an approximation to the exact Heisenberg evolution, we replace the quantum circuit $\mathcal{U}$ by

$$\mathcal{U}_{SPD} = \Pi_\delta \mathcal{U}_N \dots \Pi_\delta \mathcal{U}_1 \quad (9)$$

where $\Pi_\delta(\sum_{P \in \mathcal{P}} a_P P) = \sum_{P \in \mathcal{P}'} a_P P$ so that $|a_P| \geq \delta$ for all $P \in \mathcal{P}'$. In words, the exact representation of the observable is truncated after each gate to only include the terms whose coefficients are greater than the prescribed threshold δ . This can be considered an alternative to methods that truncate the sum based on the Hamming weight of Pauli operators, i.e., the number of nonidentity Pauli matrices appearing in a Pauli operator (14, 38, 39).

We store Pauli operators as $2n$ bit strings $(z_1, \dots, z_n, x_1, \dots, x_n) \in \{0, 1\}^{2n}$ and integer phases q so that $P = (-i)^q \prod_{i=1}^n Z_i^{z_i} X_i^{x_i}$. In the implementation, the bit strings are encoded into arrays of 64-bit unsigned integers of length $n_{64} = 2\lceil n/64 \rceil$. Therefore, we can represent a sum of N_{Pauli} Pauli operators as a $N_{\text{Pauli}} \times n_{64}$ array p of Pauli operators, a N_{Pauli} integer array of phases q , and an array a of N_{Pauli} complex coefficients. The Pauli operator array is sorted so that finding a Pauli in p takes $\mathcal{O}[n_{64} \log_2(N_{\text{Pauli}})]$ time. Insertion and deletion of Pauli operators are implemented in a way that preserves this ordering. Overall, the algorithm to apply a single Pauli rotation gate $U_\sigma(\theta)$ can be described as follows (see Fig. 10):

1. Find Pauli operators representing the observable (i.e., rows of the array p) that anticommute with σ , compute new Pauli operators $\sigma \cdot P$ for all P that anticommute with σ , and store them in a separate array p_{new} .
2. Check which rows of p_{new} are already in p .
3. For new Pauli operators, compute the coefficients as $a'_{\sigma P} = i \sin(\theta) a_P$ and store these in a new array a_{new} . For all Paulis that are already in the representation of the observable, update the coefficients according to Eq. 8.
4. Delete rows of p and elements of the phase (q) and coefficient (a) arrays where $|a| < \delta$.
5. Insert rows of p_{new} into p (and analogously for q and a arrays) where $|a_{\text{new}}| \geq \delta$.

Last, the expectation value is the sum of elements of a that correspond to Pauli operators with $x_{i \in \{1, \dots, n\}} = 0$ (i.e., those consisting only of identity and Z matrices) and the Frobenius norm of an operator in

Pauli representation is the 2-norm of its vector of coefficients, $N_O = \|a\|$. Figure S6 shows how the computational cost of the method and the Frobenius norm of the evolved operator scale with the threshold.

TN simulations

Our TN method involves three basic stages: (i) Schrödinger evolution of a PEPS forward in time, $|\psi(0)\rangle \rightarrow |\psi(\tau)\rangle$; (ii) Heisenberg evolution of a PEPO backward in time, $\Phi(T+1) \rightarrow \Phi(\tau+1)$; and (iii) contraction of the expectation (a TN “sandwich”) between the two, $\langle \psi(\tau) | \Phi(\tau+1) | \psi(\tau) \rangle$. At the limit $\tau = T$, we call this the “PEPS method”; when $\tau = 0$, we call this the “PEPO method”; and, lastly, we also study the case $\tau = T/2$, which we dub the “MIX method.” This approach allows us to exploit low entanglement in both the Schrödinger and Heisenberg pictures (40), alongside the device geometry, yielding a very efficient description of the overall dynamics.

We use two particular techniques that enable the steps above: (i) L2BP compression of both the PEPS and the PEPO as they evolve and (ii) L1BP contraction of the final expectation. The first is equivalent to methods already demonstrated for wave functions (27, 31, 32) and operators (41) but with an efficient “lazy” implementation that allows us to globally compress large and complex TNs. The second is an extension of an idea that has only recently been introduced in the context of TNs (19) and allows us to (approximately) compute nonlocal quantities such as high weight observables and the norm of the PEPS and PEPO very efficiently.

We first set up a TN representation of

$$\langle O \rangle = \langle 0 | \prod_{t=1..T} U_t^\dagger O \prod_{t=T..1} U_t | 0 \rangle$$

where O is the target observable, with $|0\rangle$ we denote $\otimes_{i=1..N} |0_i\rangle$ for sites $\{i\}$, and U_t contains all gates in layer t of a depth T circuit. We decompose any entangling gates spatially so that every tensor is associated to a single site, i , and layer, t . We take as our initial PEPS $|\psi(0)\rangle = |0\rangle$ and initial PEPO as $\Phi(T+1) = O$. For each PEPS evolution step, we take the object $U_{t+1}|\psi(t)\rangle$ and use L2BP compression with maximum bond dimension χ to yield the next PEPS $|\psi(t+1)\rangle$. Similarly, for each PEPO evolution step, we take the object $U_{t-1}^\dagger \Phi(t) U_{t-1}$ and use L2BP compression to yield $\Phi(t-1)$. This is sketched in Fig. 2A. Last, we evaluate the resulting TN with L1BP contraction.

To define L2BP compression and L1BP contraction, we first introduce “lazy” BP for TNs, an iterative message passing algorithm.




Fig. 10. Scheme representing the computation of $U_\sigma^\dagger (\sum_P a_P P) U_\sigma$ within SPD. (A) The observable is represented by an array of Pauli operators P with corresponding coefficients a_P . (B) In SPD, we identify which P commute with σ (white) and which do not (gray). For P that anticommute with σ , we compute $\sigma \cdot P$ (orange). (C) Coefficients are updated according to Eq. 8. This requires that we find Pauli operators $\sigma \cdot P$ that existed in the initial observable, which is accomplished by keeping the Pauli arrays ordered and performing a binary search. New Pauli operators that were not in the representation of the initial observable are shown in orange and not crossed out. (D) After all coefficients have been updated, we find Pauli operators whose coefficients are below threshold (additional two crossed out elements in the scheme). (E) Last, the remaining Pauli operators are merged into a new array representing the updated observable.

Consider first a TN with a single effective tensor, $T^{[i]}$, at each site i , with bonds between pairs $\{(i, j)\}$. We define vector messages $m_{i \rightarrow j}$ and $m_{j \rightarrow i}$ for each bond (Fig. 2B) and then iterate all messages according to the update contraction in Fig. 2C until convergence (30, 35, 42–44). Because the update is linear, we can allow $T^{[i]}$ to itself be a network of tensors per site, with multi-bond connections, as long as we can still lazily perform the contraction of the combined collection with input messages, thereby avoiding forming the dense version of $T^{[i]}$ (see Fig. 2B, insets). Exact contraction of these unstructured networks then becomes the dominant computational step, and we make use of recent advances to do so efficiently and with controlled memory (2, 45). The lazy approach improves both the memory and cost scaling of BP, which we explicitly detail in the Supplementary Materials. Other recent works (33, 34) have used a similar approach to BP where the sites are themselves networks, and (12, 31) can be thought of as using a lazy network of two tensors per site.

Now, we describe L2BP compression. To run BP and compress a TN, say, $|\psi\rangle$ or Φ , with dangling indices, we first have to form the 2-norm object $\langle\psi|\psi\rangle$ or $\langle\Phi|\Phi\rangle$ by tracing each dangling index with its dual to form a scalar network, shown in Fig. 2B (i and ii). In our case, we always have the PEPS or PEPO as well as the full next layer of gates. The messages in this case factorize as positive semidefinite matrices of size $d \times d$, where d is the total size of shared bonds between i and j in our original TN. Once converged, we can take each pair $m_{i \rightarrow j}$ and $m_{j \rightarrow i}$ and compute R_A and R_B as the “reduced factors” of an approximate orthogonalization around that bond. This is done by eigen-decomposition, with $m_{i \rightarrow j} = W_A \lambda_A W_A^\dagger$ giving $R_A = \sqrt{\lambda_A} W_A^\dagger$ and, similarly, $m_{j \rightarrow i} = W_B \lambda_B W_B^\dagger$ giving $R_B = (\sqrt{\lambda_B} W_B^\dagger)^T$. Taking the truncated SVD $R_A R_B \approx U \sigma V^\dagger$, we can define the optimal projectors $P_A = R_B V \sigma^{-1/2}$ and $P_B = \sigma^{-1/2} U^\dagger R_A$ and insert them into our original TN. Once this has been done for every bond, we can group and contract each site’s tensors now with all adjacent projectors, yielding a single tensor per site with maximum single bond dimension χ (31, 46, 47). This process is depicted in Fig. 2D.

Next, we describe L1BP contraction. Here, we again consider that we have multiple tensors per site but arranged in a TN that contracts to a scalar such as $\langle\psi(\tau)|\Phi(\tau+1)|\psi(\tau)\rangle$ and that we can perform the message updates lazily. Note that this is a real but not strictly positive quantity. To estimate the contracted value (or “1-norm”), we make use of an interpretation of the TN as the exponential of the Bethe free entropy $\exp(F) = Z$ (19, 35). Given converged messages, we can estimate

$$\langle O \rangle \approx Z = \left(\prod_i T^{[i]} \cdot \bigotimes_k m_{k \rightarrow i} \right) / \left(\prod_{\{(i,j)\}} m_{i \rightarrow j} \cdot m_{j \rightarrow i} \right)$$

where the index k runs over neighbors of i and $\square \cdot \square$ denotes contraction. This quantity is shown graphically in Fig. 2E. Note that each message appears exactly once in the numerator and once in the denominator, so its overall normalization does not matter. Given the nonpositive structure of our TN, it is not a priori clear that the messages should converge and Z be accurate in this case, but, as demonstrated in the Results section, for this problem, we find it to perform well. Notably, it allows us to compute not just high weight observables but also do so without relying on the renormalization of marginals/local environments. For both the PEPS and PEPO methods (with the exception of the PEPO benchmark in Fig. 8), we start performing the L1BP contraction once the cost of doing so is similar to continuing compression. For PEPS, this is at $\tau = T - 2$, and, for PEPO, this is

$\tau = \log_2(\chi)/2$ [which likely results in the faster convergence compared to (15) seen in Fig. 9F]. For the MIX method with $\chi = 256$, we note the final overlap sandwich contraction, if flattened to a single 2D TN, would require a bond dimension of $256^3 \sim 16,000,000$ (or if flattened to a 1D TN, a bond dimension of $256^{3 \times 4} \sim 10^{29}$).

There are two sources of error with this TN method. The first is the bond truncation performed during the PEPS and PEPO evolution. This is a well-understood source of error, and we know that, as $\chi \rightarrow \mathcal{O}(2^T)$, the TN becomes exact, with the L2BP aspect functioning simply as a gauge in which to perform the truncation when $\chi < \mathcal{O}(2^T)$. Heuristically, if a computation converges in χ before this point, then it is usually assumed to be close to the exact value. We can also assess how much of the TN has been truncated away by computing $\langle\psi|\psi\rangle$ or $\langle\langle\Phi|\Phi\rangle\rangle/2^n$, i.e., the norm, after all compressions. Because the action of the projectors P_A and P_B can only decrease this quantity from the exact value of 1, it serves as a reasonable proxy of fidelity. In our case, using L1BP contraction, we can compute this norm and find that, in all exactly accessible contexts, as it approaches 1, it also heralds high accuracy. The second source of error comes from the L1BP contraction of the final scalar TN. This method also has an exact limit when the underlying network becomes a tree. Away from this regime, it becomes only a heuristic but can be accurate, for example, in cases with sparse connectivity without strongly correlated loops.

To verify the accuracy of L1BP, one approach in the main text for the PEPS method is to use Monte Carlo sampling of the expectation value (Fig. 8A). We do so efficiently by sampling from the approximate distribution $|\langle x|\psi\rangle|^2 \approx \omega(x)$ where ω is the probability of sampling configuration x via BP decimation sampling. This works as follows: We compute all approximate qubit reduced density matrices, $\{\rho_i\}$, using BP; we then sample a value for the qubit with $\max_i |\text{Tr}(Z\rho_i)|$ according to $\text{diag}(\rho_i)$. We then condition on that value and restart BP on the remaining qubits. When all qubits have been fixed, we have a full string x , approximate probability $\omega(x)$, and, from exact contraction, can obtain the true probability $p(x) = |\langle x|\psi\rangle|^2$ (which is much easier than the exact observable due to only having a single layer). This is sufficient to perform efficient importance sampling allowing us to compute unbiased Monte Carlo estimates of observables, to verify the accuracy of L1BP contraction.

Overall, we note that the methods presented here are easily implemented for arbitrary geometries and gates and thus should be a useful tool in a wide variety of both real and imaginary time evolution settings. There are also various obvious pathways to improving the fidelity, for example, by matching the PEPS/PEPO norms to dynamically set the τ balance in the MIX method. One could also compress much larger blocks of gates or use MPS as the messages as in (34). Given the equivalence of BP compression to simple update, itself equivalent to cluster update (28, 29) with “zero cluster size,” an interesting question is whether the environment correlations can be systematically increased in a similar manner. However, we find that we can reach sufficient fidelities across the relevant range of observables with the MIX method without invoking such improvements. The maximum values of χ used are listed in table S4.

All TN simulations were performed using quimb (48) and cotengra (45). As well as controlling complexity through the bond dimension χ , we use a truncation cutoff $\kappa = 5 \times 10^{-6}$ throughout such that each local truncation error is $|AB - AP_A P_B B| < \kappa$. We also use light-cone cancellation (although this is also easily found in the PEPO picture by the truncation with a nonzero κ). We performed the PEPS and MIX simulations in single precision but did find that the PEPO

simulation requires double precision. To give an approximate sense of timings, all simulations up to 9 steps or 20 steps with $\chi \lesssim 64$ take on the order of min or less, including the cost of contraction path optimization. For large χ , we run calculations on a NVidia A100 GPU, with the most expensive points (MIX method with $\chi = 192$, $\theta_i > 8\pi/32$, for example) taking about a day. The “Mix 3 min” method shown in Fig. 3 was executed on a NVidia 4070 Ti. A more detailed runtime analysis on a subset of the data can be found in the Supplementary Materials.

Supplementary Materials

This PDF file includes:

Supplementary Text

Figs. S1 to S10

Tables S1 to S4

REFERENCES AND NOTES

1. S. Boixo, S. V. Isakov, V. N. Smelyanskiy, H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models. *arXiv:1712.05384 [quant-ph]* [preprint] (2017).
2. J. Chen, F. Zhang, C. Huang, M. Newman, Y. Shi, Classical simulation of intermediate-size quantum circuits. *arXiv:1805.01450 [quant-ph]* [preprint] (2018).
3. C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, M. Szegedy, Y. Shi, J. Chen, Classical simulation of quantum supremacy circuits. *arXiv:2005.06787 [quant-ph]* [preprint] (2020).
4. Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, J.-W. Pan, Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.* **127**, 180501 (2021).
5. F. Pan, K. Chen, P. Zhang, Solving the sampling problem of the sycamore quantum circuits. *Phys. Rev. Lett.* **129**, 090502 (2022).
6. A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov, J. Chau, G. Laun, R. Movassagh, A. Asfaw, L. T. A. N. Brandão, R. Peralta, D. Abanin, R. Acharya, R. Allen, T. I. Andersen, K. Anderson, M. Ansmann, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, A. Bيلمes, G. Bortoli, A. Bourassa, J. Bovaid, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, J. Campero, H. S. Chang, B. Chiaro, D. Chik, C. Chou, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. D. T. Barba, S. Demura, A. Di Paolo, A. Dunsborough, L. Faoro, E. Farhi, R. Fatemi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, G. Garcia, E. Genois, W. Gidney, C. Gidney, D. Gilboa, M. Giustina, R. Gosula, A. G. Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, P. Heu, M. R. Hoffmann, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, A. R. Klotz, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. W. Lee, Y. D. Lensky, B. J. Lester, A. T. Lill, W. Liu, A. Locharla, F. D. Malone, O. Martin, S. Martin, J. R. McClean, M. McEwen, K. C. Miao, A. Mieszala, S. Montazeri, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, S. Omonije, A. Opremcak, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, D. M. Rhodes, C. Rocque, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, V. Sivak, J. Skrzynny, W. C. Smith, R. D. Somma, G. Sterling, D. Strain, M. Szalay, D. Thor, A. Torres, G. Vidal, C. V. Heidweiller, T. White, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, N. Zobrist, E. G. Rieffel, R. Biswas, R. Babbush, D. Bacon, J. Hilton, E. Lucero, H. Neven, A. Megrant, J. Kelly, I. Aleiner, V. Smelyanskiy, K. Kechedzhi, Y. Chen, S. Boixo, Phase transition in random circuit sampling. *arXiv:2304.11119 [quant-ph]* [preprint] (2023).
7. I. L. Markov, Y. Shi, Simulating quantum computation by contracting tensor networks. *SIAM J. Comput.* **38**, 963–981 (2008).
8. Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, A. Kandala, Evidence for the utility of quantum computing before fault tolerance. *Nature* **618**, 500–505 (2023).
9. K. Temme, S. Bravyi, J. M. Gambetta, Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017).
10. M. P. Zaletel, F. Pollmann, Isometric tensor network states in two dimensions. *Phys. Rev. Lett.* **124**, 037201 (2020).
11. T. Begušić, G. K.-L. Chan, Fast classical simulation of evidence for the utility of quantum computing before fault tolerance. *arXiv:2306.16372 [quant-ph]* [preprint] (2023).
12. J. Tindall, M. Fishman, M. Stoudenmire, D. Sels, Efficient tensor network simulation of IBM's kicked Ising experiment. *arXiv:2306.14887 [quant-ph]* [preprint] (2023).
13. K. Kechedzhi, S. V. Isakov, S. Mandrà, B. Villalonga, X. Mi, S. Boixo, V. Smelyanskiy, Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments. *arXiv:2306.15970 [quant-ph]* [preprint] (2023).
14. Y. Shao, F. Wei, S. Cheng, Z. Liu, Simulating quantum mean values in noisy variational quantum algorithms: A polynomial-scale approach. *arXiv:2306.05804 [quant-ph]* [preprint] (2023).
15. H.-J. Liao, K. Wang, Z.-S. Zhou, P. Zhang, T. Xiang, Simulation of ibm's kicked ising experiment with projected entangled pair operator. *arXiv:2308.03082 [quant-ph]* [preprint] (2023).
16. M. S. Rudolph, E. Fontana, Z. Holmes, L. Cincio, Classical surrogate simulation of quantum systems with lowesa. *arXiv:2308.09109 [quant-ph]* [preprint] (2023).
17. S. Anand, K. Temme, A. Kandala, M. Zaletel, Classical benchmarking of zero noise extrapolation beyond the exactly-verifiable regime. *arXiv:2306.17839 [quant-ph]* [preprint] (2023).
18. T. Begušić, K. Hejazi, G. K.-L. Chan, Simulating quantum circuit expectation values by Clifford perturbation theory. *arXiv:2306.04797 [quant-ph]* [preprint] (2023).
19. N. Pancotti, J. Gray, One-step replica symmetry breaking in the language of tensor networks. *arXiv:2306.15004 [quant-ph]* [preprint] (2023).
20. S. Aaronson, D. Gottesman, Improved simulation of stabilizer circuits. *Phys. Rev. A* **70**, 052328 (2004).
21. S. Bravyi, D. Gosset, Improved classical simulation of quantum circuits dominated by clifford gates. *Phys. Rev. Lett.* **116**, 250501 (2016).
22. R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, R. C. Pooser, Unbiased simulation of near-Clifford quantum circuits. *Phys. Rev. A* **95**, 062337 (2017).
23. S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum* **3**, 181 (2019).
24. P. Rall, D. Liang, J. Cook, W. Kretschmer, Simulation of qubit quantum circuits via pauli propagation. *Phys. Rev. A* **99**, 062337 (2019).
25. N. A. Nemkov, E. O. Kiktenko, A. K. Fedorov, Fourier expansion in variational quantum algorithms. *Phys. Rev. A* **108**, 032406 (2023).
26. E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, C. Cirstoiu, Classical simulations of noisy variational quantum circuits. *arXiv:2306.05400 [quant-ph]* [preprint] (2023).
27. H. C. Jiang, Z. Y. Weng, T. Xiang, Accurate determination of tensor network state of quantum lattice models in two dimensions. *Phys. Rev. Lett.* **101**, 090603 (2008).
28. M. Lubasch, J. I. Cirac, M.-C. Banuls, Algorithms for finite projected entangled pair states. *Physical Review B* **90**, 064425 (2014).
29. M. Lubasch, J. I. Cirac, M.-C. Banuls, Unifying projected entangled pair state contractions. *New J. Phys.* **16**, 033014 (2014).
30. R. Alkabetz, I. Arad, Tensor networks contraction and the belief propagation algorithm. *Phys. Rev. Research* **3**, 023073 (2021).
31. J. Tindall, M. Fishman, Gauging tensor networks with belief propagation. *arXiv:2306.17837 [quant-ph]* [preprint] (2023).
32. S.-J. Ran, W. Li, B. Xi, Z. Zhang, G. Su, Optimized decimation of tensor networks with super-orthogonalization for two-dimensional quantum lattice models. *Phys. Rev. B* **86**, 134429 (2012).
33. Y. Wang, Y. E. Zhang, F. Pan, P. Zhang, Tensor network message passing. *arXiv:2305.01874 [cond-mat.stat-mech]* (2023).
34. C. Guo, D. Poletti, I. Arad, Block belief propagation algorithm for 2D tensor networks. *arXiv:2301.05844 [quant-ph]* [preprint] (2023).
35. M. Mézard, A. Montanari, *Information, Physics, and Computation* (Oxford University Press, 2009).
36. J. Gray, G. K.-L. Chan, Hyper-optimized compressed contraction of tensor networks with arbitrary geometry. *arXiv:2206.07044 [quant-ph]* [preprint] (2022).
37. H. Qassim, J. J. Wallman, J. Emerson, Clifford recompilation for faster classical simulation of quantum circuits. *Quantum Phys.* **3**, 170 (2019).
38. I. Kuprov, N. Wagner-Rundell, P. Hore, Polynomially scaling spin dynamics simulation algorithm based on adaptive state-space restriction. *J. Magn. Reson.* **189**, 241–250 (2007).
39. T. Rakovsky, C. von Keyserlingk, F. Pollmann, Dissipation-assisted operator evolution method for capturing hydrodynamic transport. *Phys. Rev. B* **105**, 075131 (2022).
40. D. M. Kennes, C. Karrasch, Extending the range of real time density matrix renormalization group simulations. *Comput. Phys. Commun.* **200**, 37–43 (2016).
41. A. Kshetrimayum, H. Weimer, R. Orús, A simple tensor network algorithm for two-dimensional steady states. *Nat. Commun.* **8**, 1291 (2017).
42. M. B. Hastings, Quantum belief propagation: An algorithm for thermal quantum systems. *Phys. Rev. B* **76**, 201102 (2007).
43. M. S. Leifer, D. Poulin, Quantum graphical models and belief propagation. *Ann. Phys.* **323**, 1899–1946 (2008).

44. S. Sahu, B. Swingle, Efficient tensor network simulation of quantum many-body physics on sparse graphs. arXiv:2206.04701 [quant-ph] [preprint] (2022).
45. J. Gray, S. Kourtis, Hyper-optimized tensor network contraction. *Quantum Phys.* **5**, 410 (2021).
46. Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, T. Xiang, Coarse-graining renormalization by higher-order singular value decomposition. *Phys. Rev. B* **86**, 045139 (2012).
47. S. Iino, S. Morita, N. Kawashima, Boundary tensor renormalization group. *Phys. Rev. B* **100**, 035449 (2019).
48. J. Gray, Quimb: A python package for quantum information and many-body calculations. *J. Open Source Softw.* **3**, 819 (2018).

Acknowledgments: J.G. thanks N. Pancotti for many fruitful discussions of BP and TNs. We thank the authors of (12–15, 17) for sharing data presented in Fig. 9. **Funding:** We were supported by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of Basic Energy Sciences, Scientific Discovery through Advanced Computing (SciDAC) program under award number DE-SC0022088. T.B. acknowledges financial support from the Swiss National Science Foundation through the Postdoc Mobility Fellowship (grant number P500PN-214214). G.K.-L.C. is a Simons Investigator in Physics. Computations presented here were partly conducted in the Resnick High Performance Computing Center, a facility supported by Resnick Sustainability Institute at the California Institute of Technology. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science

User Facility located at Lawrence Berkeley National Laboratory, operated under contract no. DE-AC02-05CH11231 using NERSC award BES-ERCAP0024087. **Author contributions:** G.K.-L.C. conceptualized and supervised the work. J.G. and T.B. participated in conceptualization, implemented the computational methods, carried out simulations, and prepared the data and figures. All authors contributed to the analysis of the results and writing of the manuscript. **Competing interests:** G.K.-L.C. is part owner of QSimulate Inc. and has served as a scientific consultant to the Flatiron Institute. Between 2019 and 2021, he was on the advisory board of the Qingdao Institute of Theoretical and Computational Chemistry. He is the Chief Scientific Advisor to QSimulate and an Associate Researcher of the HK Quantum AI lab. The other authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Data presented in Figs. 3 to 5 are available at www.github.com/tbegusic/arxiv-2308.05077-data (Zenodo access code: [www.doi.org/10.5281/zenodo.10223349](https://doi.org/10.5281/zenodo.10223349)). Code for running the SPD and TN simulations is available on GitHub and Zenodo (SPD, www.github.com/tbegusic/spd and <https://doi.org/10.5281/zenodo.10456459> and TN, www.github.com/jcmgray/mixpic-bp-quantum-dynamics and doi.org/10.5281/zenodo.10458595).

Submitted 22 August 2023

Accepted 19 December 2023

Published 17 January 2024

10.1126/sciadv.adk4321