nature physics

ARTICLES

https://doi.org/10.1038/s41567-018-0124-x

Characterizing quantum supremacy in near-term devices

Sergio Boixo[1]*, Sergei V. Isakov[2], Vadim N. Smelyanskiy[1], Ryan Babbush[1], Nan Ding[1], Zhang Jiang[3,4], Michael J. Bremner[5], John M. Martinis[6,7] and Hartmut Neven[1]

A critical question for quantum computing in the near future is whether quantum devices without error correction can perform a well-defined computational task beyond the capabilities of supercomputers. Such a demonstration of what is referred to as quantum supremacy requires a reliable evaluation of the resources required to solve tasks with classical approaches. Here, we propose the task of sampling from the output distribution of random quantum circuits as a demonstration of quantum supremacy. We extend previous results in computational complexity to argue that this sampling task must take exponential time in a classical computer. We introduce cross-entropy benchmarking to obtain the experimental fidelity of complex multiqubit dynamics. This can be estimated and extrapolated to give a success metric for a quantum supremacy demonstration. We study the computational cost of relevant classical algorithms and conclude that quantum supremacy can be achieved with circuits in a two-dimensional lattice of 7 × 7 qubits and around 40 clock cycles. This requires an error rate of around 0.5% for two-qubit gates (0.05% for one-qubit gates), and it would demonstrate the basic building blocks for a fault-tolerant quantum computer.

The controlled evolution of ideal quantum systems offers computational resources more powerful than classical computers. On the basis of results in quantum chaos[1–11] and computational complexity theory[12–17], we propose an experiment for characterizing 'quantum supremacy'[18] in the presence of errors in the near term. Quantum supremacy is achieved when a formal computational task is performed with an existing quantum device, but the same task cannot be performed using any known algorithm running on an existing classical supercomputer in a reasonable amount of time.

Time-accurate simulations of classical dynamical systems with chaotic behaviour are among the hardest numerical tasks. Examples include turbulence and population dynamics, essential for the study of meteorology, biology, finance and so on. In all of these cases, a direct numerical simulation is required to get an accurate description of the system state after a finite time. Our minimal-resource demonstration of quantum supremacy is based on the implementation of (pseudo-)random quantum circuits with gates sampled from a universal gate set. These are examples of quantum chaotic evolutions that naturally lend themselves to the quantum computational framework[1,3,6]. A circuit, corresponding to a unitary transformation U, is a sequence of d clock cycles of one- and two-qubit gates, with gates applied to different qubits in the same cycle. With realistic superconducting hardware constraints[19,20], gates act in parallel on distinct sets of qubits restricted to a one-dimensional (1D) or 2D lattice. The cycle number t plays the role of time in the chaotic dynamics of the quantum state $|\psi_t\rangle$. In analogy with classical Lyapunov exponents, a signature of quantum chaos is the decrease of the overlap $|\langle\psi_t|\psi_t^{\varepsilon}\rangle|^2$ of the quantum state $|\psi_t\rangle$ with the state $|\psi_t^{\varepsilon}\rangle$ resulting from a small perturbation ε to the Hamiltonian that evolves $|\psi_t\rangle$ (refs [2,21–23]). The overlap decreases exponentially in the evolution time t and ε because chaotic evolutions give rise to delocalization of quantum states[1,24]. Such states are closely related to ensembles of random unitary matrices studied in random matrix theory[24,25], they possess no symmetries, and they are spread over Hilbert space. Therefore, as in the case of classical chaos, obtaining a description of $|\psi_t\rangle$ requires a high-fidelity classical simulation. This challenge is compounded by the exponential growth of the Hilbert space dimension $N = 2^n$ with the number of qubits n.

We introduce cross-entropy benchmarking as a method to estimate the fidelity of complex multi-qubit implementations. The cross-entropy difference is a measure of correspondence between experimentally obtained samples and the output distribution of the ideal circuit. We study numerically the convergence of the distribution of output probabilities to the exponential distribution, which in this context we call the Porter–Thomas distribution[26], characteristic of quantum chaos[27]. We find a good convergence for the first 10 moments and the entropy at depth ~20 with circuits of up to 7 × 6 qubits in a 2D lattice. Using chaos theory, the properties of the Porter–Thomas distribution, and numerical simulations, we argue that the cross-entropy is closely related to the circuit fidelity. State-of-the-art supercomputers cannot simulate universal random circuits of sufficient depth (approximately 40 clock cycles) in a 2D lattice of approximately 7 × 7 qubits with any known algorithm and significant fidelity.

Sampling random quantum circuits

Given a random quantum circuit U of depth d, we are interested in sampling from the distribution $p_U(x) \equiv |\langle x|\psi_d\rangle|^2$ of bit-strings in the computational basis $\{|x\rangle\}$. It has been argued that classically solving related sampling problems requires computational resources with asymptotic exponential scaling[12–15,17]. Examples include BosonSampling[14] and approximate simulation of commuting quantum computations, so-called instantaneous quantum polynomial-time (IQP) circuits[13,17].

[1]Google Inc., Venice, CA, USA. [2]Google Inc., Zurich, Switzerland. [3]QuAIL, NASA Ames Research Center, Moffett Field, CA, USA. [4]SGT Inc., Greenbelt, MD, USA. [5]Centre for Quantum Computation and Communication Technology, Centre for Quantum Software and Information, Faculty of Engineering and Information Technology, University of Technology Sydney, Ultimo, New South Wales, Australia. [6]Google Inc., Santa Barbara, CA, USA. [7]Department of Physics, University of California, Santa Barbara, CA, USA. *e-mail: boixo@google.com

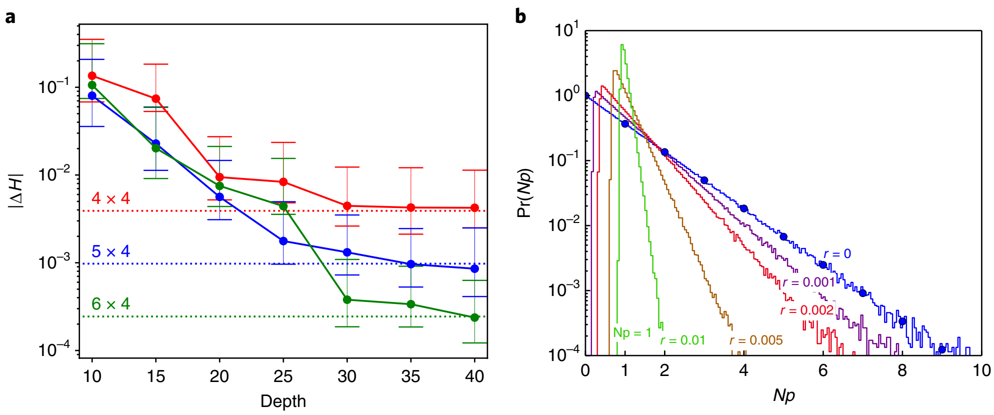

Fig. 1 | Sensitivity to errors of the output distribution of chaotic circuits. a, Cross-entropy difference (see text) as a function of depth between the ideal output distribution and the distribution after a single Z error (phase-flip) or X error (bit-flip) is applied. Different colours correspond to different sizes: 4 × 4, 5 × 4 or 6 × 4 qubits. Solid lines correspond to the median over all possible error positions, and error bars to the 0.25 and 0.75 quantiles. The dotted lines correspond to a cross-entropy difference of $2^{-n/2}$, for n qubits. (The Pearson's correlation coefficient results in a very similar plot.) **b**, Distribution of rescaled output bit-string probabilities Np for a typical random circuit U, with $N = 2^n$ and $\{p = p_U(x)\}$. The blue curve ($r = 0$) shows the distribution of $\{Np_U(x)\}$ obtained from numerical simulations of the ideal random circuit. This distribution is very close to the Porter–Thomas form $\Pr(Np) = e^{-Np}$ shown with blue dots. Curves with different colours show the distributions of probabilities obtained for different Pauli error rates r. More explicitly, each operation is followed by a depolarizing error channel with two-qubit error rate = measurement error = initialization error = r, and one-qubit error rate = $r/10$. The distributions with increasing error rates converge to the dashed line at $Np = 1$, which corresponds to the uniform distribution over bit-strings $\delta(p - 1/N)$. These numerics are obtained from simulations of a planar circuit with 5 × 4 qubits and gate depth of 40 ($n = 20$ and $N = 2^{20}$).

Hypersensitivity to perturbations is a signature of chaos[2]. We study the sensitivity of the distribution $p_U(x)$ to small perturbations numerically: the distribution obtained after a single random X or Z gate (error) is added to the circuit is almost uncorrelated with $p_U(x)$ (see Fig. 1a). Therefore, it is natural to conjecture that sampling from the distribution $p_U(x)$ requires a high-fidelity simulation of the quantum circuit, and that approximate classical simulations are very likely to fail (see Supplementary Information and ref. [28]). It follows that unless a classical algorithm uses resources that grow exponentially in n, its output would be almost statistically uncorrelated with the output distribution of a quantum circuit of enough depth.

States generated by pseudo-random circuits approximate aspects of the uniform distribution in Hilbert space with increasing depth[3,7,8]. A consequence of this is that the distribution of output bit-string probabilities $\{p = p_U(x)\}$ approaches the exponential or Porter–Thomas distribution Ne^{-pN} with mean $1/N$ (see Fig. 1b and Supplementary Information). We study numerically the convergence to the Porter–Thomas distribution entropy $-\sum_j p_U(x_j)\log p_U(x_j) \to \log N - 1 + \gamma$, where γ is Euler's constant, and moments $\sum_j p_U(x_j)^k \to k!/N^{k-1}$, up to $k = 10$. The specific ensemble we consider, which we find to converge rapidly to the Porter–Thomas distribution, starts with an initial layer of Hadamard gates to rotate to the X basis, because for this study we use only controlled-Z (CZ) two-qubit gates. The next d cycles alternate between eight configurations of CZ gates similar to Fig. 2a. We also place one-qubit gates from the set $\{X^{1/2}, Y^{1/2}, T\}$ at qubits not occupied by CZ gates at the same cycle. The gate $X^{1/2}$ ($Y^{1/2}$) is a $\pi/2$ rotation around the X (Y) axis of the Bloch sphere, and the non-Clifford T gate is the diagonal matrix $\{1, e^{i\pi/4}\}$. One-qubit gates are placed subject to the following rules: the first one-qubit gate for each qubit after the initial cycle of Hadamard gates is always a T gate; and we place a one-qubit gate only in the next cycle after a CZ gate in the same qubit. If this qubit has already seen a T gate, the gate is chosen with equal probability between the two gates different from the last one-qubit gate applied to this qubit. We obtain good convergence for circuits acting on as many as 7 × 6 qubits at depth ~20 (see Fig. 2b,c). We also numerically validate that the required depth grows sublinearly in the number of qubits, consistent with related results[8,29].

Cross-entropy benchmarking

Consider a sample $S = \{x_1, \ldots, x_m\}$ of bit-strings x_j in the computational basis. For a typical sample S, and using that the $p_U(x)$ terms are approximately independent and identically distributed according to the Porter–Thomas distribution, the central limit theorem implies that $\log \Pr_U(S) = -m(\log N - 1 + \gamma) + O(m^{1/2})$ (where γ is Euler's constant, see Supplementary Information). As a baseline, consider now a sample $S_{\mathrm{unc}} = \{x_1^{\mathrm{unc}}, \ldots, x_m^{\mathrm{unc}}\}$ taken from a distribution uncorrelated with $p_U(x)$. We now focus on the probability $\Pr_U(S_{\mathrm{unc}})$ that this sample S_{unc} is observed from the output $|\psi_d\rangle$ of the circuit U. The central limit theorem implies that $\log \Pr_U(S_{\mathrm{unc}}) = -m\mathrm{H}(p_{\mathrm{unc}}, p_U) + O(m^{1/2})$, where $H(p_{\mathrm{unc}}, p_U) \equiv -\sum_{j=1}^{N} p_{\mathrm{unc}}(x_j)\log p_U(x_j)$ is the cross-entropy between $p_{\mathrm{unc}}(x)$ and $p_U(x)$. Averaging over the ensemble $\{U\}$ can be done independently for $p_{\mathrm{unc}}(x)$ and $\log p_U(x)$ to obtain $\mathbb{E}_U[H(p_{\mathrm{unc}}, p_U)] = -\sum_{j=1}^{N} \mathbb{E}_U[p_{\mathrm{unc}}(x_j)]\,\mathbb{E}_U[\log p_U(x_j)]$. For fixed x_j, the distribution of values $\{p_U(x_j)\}$ also converges towards the Porter–Thomas form if we use sufficiently deep random quantum circuits. This allows us to define the quantity $H_0 \equiv -\mathbb{E}_U[\log p_U(x_j)] = \log N + \gamma$. Then we obtain that the difference between the average of the log of the probabilities of an m-sample from the ideal circuit and from an uncorrelated distribution is $\mathbb{E}_U[\log \Pr_U(S) - \log \Pr_U(S_{\mathrm{unc}})] = m$. This equation reveals that a typical m-sample S from a random circuit U represents a unique signature of that circuit.

We can measure the quality of any algorithm A to sample bit-strings from the distribution $p_U(x)$ as the difference between its cross-entropy and H_0. We call this the cross-entropy difference $\Delta H(p_A) \equiv H_0 - H(p_A, p_U)$. This quantity is unity for the ideal random circuit and zero for an uncorrelated distribution over bit-strings when averaging over U. We refer to an experimental implementation of a quantum circuit U as $A_{\mathrm{exp}}(U)$ and associate with it the probability distribution $p_{\mathrm{exp}}(x_j|U)$ and samples S_{exp}. The experimental cross-entropy difference is $\alpha \equiv \mathbb{E}_U[\Delta H(p_{\mathrm{exp}})]$. Quantum supremacy is achieved, in practice, when $1 \geq \alpha > C$, where (a lower bound for) C is given by the performance of the best known classical algorithm A^*

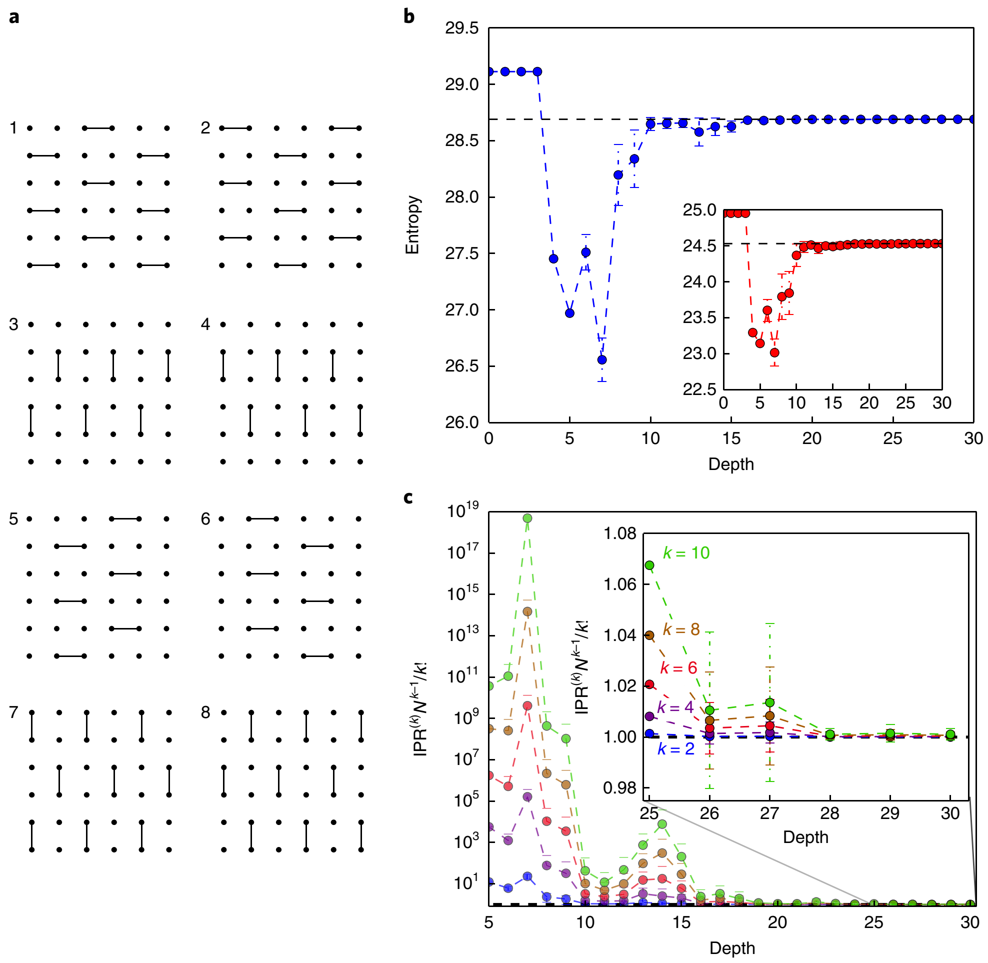

Fig. 2 | Convergence to the Porter–Thomas (or exponential) distribution. **a**, Layouts of CZ gates in a 6 × 6 qubit lattice. It is currently not possible to perform two CZ gates simultaneously in two neighbouring superconducting qubits[19,20,40]. We iterate over these arrangements sequentially, from 1 to 8. **b**, Mean entropy of the output distribution as a function of depth. The main figure pertains to circuits with 7 × 6 qubits, and the inset pertains to circuits with 6 × 6 qubits. The entropies converge to the dashed line, corresponding to the Porter–Thomas distribution $\sum_j p_U(x_j)\log p_U(x_j) \to \log N - 1 + \gamma$ (see Supplementary Information). Error bars are standard deviations among different circuit instances. **c**, Mean normalized inverse participation ratios $k \in [2, .., 10]$ of the output distribution ($\mathrm{IPR}^{(k)} = \sum_j p_U(x_j)^k \to k!/N^{k-1}$) as a function of depth for circuits with 7 × 6 qubits. They converge, at approximately the same depth, to the black dashed line, which corresponds to the Porter–Thomas distribution. Error bars show the standard deviation between different circuit instances.

executed on an existing classical computer, $C = \mathbb{E}_U[\Delta H(p^*)]$. Here p^* is the output distribution of A^*.

The space and time complexity of exactly calculating one output amplitude of a random circuit is exponential in the treewidth of the interaction graph of a corresponding Ising model (see below, Supplementary Information, the implementation in ref.[30] and also refs[31–33]). The treewidth is proportional to $\min(d, n)$ in a 1D lattice, and $\min(d\sqrt{n}, n)$ in a 2D lattice (see Fig. 3a). From the numerical estimation of the treewidth (and the simulation times reported in the Supplementary Information and ref.[30]), we estimate that the computation of an output amplitude for circuits with 7 × 7 qubits and depth of approximately 40 cycles is not currently viable. For large depth d, algorithms are limited by the memory required to store the wavefunction in random-access memory, which in single precision is $2^n \times 2 \times 4$ bytes. For $n = 48$ qubits, this requires at least 2.25 petabytes, which is approximately the limit of what can be done on the largest supercomputers of today. For example, Trinity, the sixth fastest supercomputer in TOP500 has about two petabytes of primary memory, which is one of the largest. For circuits of small depth or less than approximately 48 qubits, direct simulation is viable, so $C = 1$ and quantum supremacy is impossible.

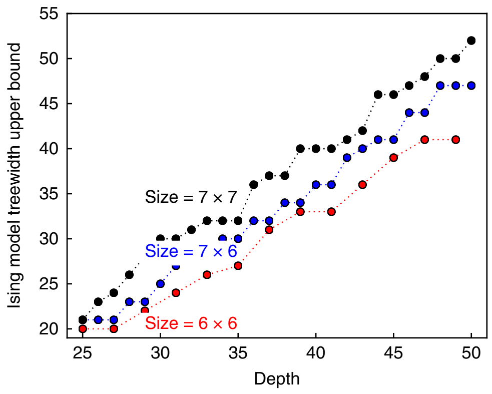

Fig. 3 | Numerical upper bound for the treewidth. We plot the treewidth of the interaction graph of the Ising model corresponding to circuits with 6 × 6, 7 × 6, and 7 × 7 qubits as a function of the circuit depth. The cost of calculating an output amplitude is exponential in the treewidth. We estimate that the computation of an output amplitude for a circuit with 7 × 7 qubits and depth 40 is not viable.

We now address the question of how the cross-entropy difference α can be estimated from an experimental sample of bit-strings S_{exp} obtained by measuring the output of $A_{\text{exp}}(U)$ after m realizations of the circuit. For a typical sample S_{exp}, the central limit theorem implies that $\alpha \simeq H_0 + \frac{1}{m}\sum_{j=1}^{m} \log p_U(x_j^{\text{exp}})$, with statistical error $\kappa / \sqrt{m}$ and n-independent constant $\kappa \simeq 1$ (see Supplementary Information). Therefore, if we can compute the quantities $\log p_U(x_j^{\text{exp}})$ with the aid of a sufficiently powerful classical computer, we can estimate α (see Supplementary Information and ref. [30]). We note that, for most quantum states, the outcome of any fixed measurement concentrates exponentially on the average value over Hilbert space[34]. We have circumvented this limitation by designing a state-specific global measurement. A close correspondence between experiment, numerics and theory provides a reliable foundation from which to extrapolate α to larger circuits where the quantities $p_U(x_j)$ can no longer be obtained numerically. After this point, on the basis of the results of our numerical studies and aforementioned insights from quantum chaos, we assume that the output of a classical algorithm quickly becomes statistically uncorrelated with $p_U(x)$, and quantum supremacy can be achieved. The value of α can be extrapolated from circuits that can be simulated because they have either fewer qubits, mostly Clifford gates (stabilizer simulations)[35] or less depth[30].

We now present a theoretical error model for α that can be compared with experiment. The output ρ of the experimental realization of a random circuit U is $\rho = \alpha_{\text{f}} U\,|\psi_0\rangle\langle\psi_0|\,U^{\dagger} + (1-\alpha_{\text{f}})\sigma_U$, where $\langle\psi_0|U^{\dagger}\sigma_U U|\psi_0\rangle = 0$ and α_{f} is the circuit fidelity. As U is a random circuit implementing a chaotic evolution, we see in numerical simulations that the probabilities $p_U(x)$ and $\langle x|\sigma_U|x\rangle$ are almost uncorrelated (see Fig. 1a). Under this ansatz, by the same derivation used above, we have $\sum_j \langle x_j|\sigma_U|x_j\rangle \log p_U(x_j) = H_0$ (see Supplementary Information). Therefore, the circuit fidelity α_{f} is approximately equal to the cross-entropy difference; that is, $\alpha \approx \alpha_{\text{f}}$. This introduces a fundamentally new way to estimate the fidelity of complex quantum circuits, which we call cross-entropy benchmarking. A common approximation for studying circuit fidelities is a digital error model where each gate is followed by a depolarizing error channel[19,36–40]. Within this idealized model, a simple estimate of the fidelity is $\alpha \approx \exp(-r_1 g_1 - r_2 g_2 - r_{\text{init}} n - r_{\text{res}} n)$, where $r_1, r_2 \ll 1$ are the Pauli error rates for one- and two-qubit gates, $r_{\text{init}}, r_{\text{res}} \gg 1$ are the initialization and measurement error rates and $g_1, g_2 \gg 1$ are the numbers of one- and two-qubits gates, respectively. We compare numerically cross-entropy benchmarking with this idealized estimate of the fidelity and observe a good fit between these two quantities (see Fig. 4a). Furthermore, because random Porter–Thomas distributed states are near-maximally entangled[5,9–11], we expect that even one Pauli error completely destroys correlations with the ideal sampling[41]. We therefore make the ansatz $\rho_{\mathcal{K}} = \alpha\,|\psi_d\rangle\langle\psi_d| + (1-\alpha)\frac{1}{N}$ for the output state of a circuit implementation with fidelity α. The cross-entropy difference ΔH defined above is given by the probability distribution of $z = \log(p_U(x))$ where the bit-strings x are sampled from the output $\rho_{\mathcal{K}}$ of a circuit implementation with fidelity α. Using the ansatz for $\rho_{\mathcal{K}}$ and the Porter–Thomas distribution for $p_U(x)$, we obtain $\Pr_{\alpha}(z) = \mathrm{e}^{z-\mathrm{e}^{z}}(1+\alpha(\mathrm{e}^{z}-1))$. We observe an excellent correspondence between this equation and simulations using the digital error model (see Fig. 4b).

Computational complexity

The intuitive arguments about the classical cost of sampling from $pU(x)$ can be made more rigorous in the asymptotic limit of large n using computational complexity theory. Note that approaching this limit experimentally requires error correction. It has been shown that for IQP circuits, the function $pU(x) = |\langle x|\psi_d\rangle|^2$ maps directly to the partition function of a random complex Ising model[13,17] $\langle x|\psi_d\rangle = \lambda\sum_{\mathbf{s}} e^{i\frac{\pi}{4}H_x(\mathbf{s})}$, where $H_x(\mathbf{s}) = \mathbf{h}_x \cdot \mathbf{s} + \mathbf{s}\cdot\hat{\mathbf{J}}\cdot\mathbf{s}$ is a classical energy, $\mathbf{s}$ is a vector of classical spins ± 1, $\mathbf{h}_x$ is a vector of local fields, $\hat{\mathbf{J}}$ is the coupling matrix, $i\frac{\pi}{4}$ is the inverse imaginary temperature and λ is a scaling constant. It is a strongly held conjecture in computational complexity theory that probabilistically approximating partition functions with purely imaginary temperatures is much harder, in the worst case, than any problem that can be solved with an NP oracle[13,15,16]. Ref. [17] further conjectures that this applies to any sufficiently large fraction of partition functions where $\mathbf{h}_x$ and $\hat{\mathbf{J}}$ are chosen uniformly at random. Finally, ref. [17] proves that if a classical algorithm could efficiently approximately (in the ℓ_1 norm) sample from the output of an IQP circuit, then this would imply there exists an algorithm that could use a non-deterministic polynomial time (NP) oracle to probabilistically approximate a large fraction of random complex Ising models, contradicting their conjecture (see also the Supplementary Information).

We are interested in circuits on a 2D lattice with gates in the set {CZ, $\mathrm{X}^{1/2}$, $\mathrm{Y}^{1/2}$, T}. The gates {CZ, T} are diagonal in the computational basis, while the gates {$\mathrm{X}^{1/2}$, $\mathrm{Y}^{1/2}$} are two-sparse. We show that for these circuits $\langle x|\psi_d\rangle = 2^{-G/2}\sum_s \exp\left(\frac{i\pi}{4}H_s(x)\right)$, where G is the number of two-sparse gates, $s \in \{+1,-1\}^G$ and $H_s(x)$ is an Ising model with a quasi-3D interaction graph (see Supplementary Information). We give the explicit form of the probability distribution of the couplings and local fields for the Ising models corresponding to the quantum circuits above in the Supplementary Information. It follows that we can write $\langle x|\psi_d\rangle = 2^{-G/2}Z$, where $Z = \sum_{j=0}^{7} M_j e^{i\frac{2\pi}{8}E_j}$ is a partition function, the E_j terms are different energies of the Ising model (mod 8) and $M_j \sim 2^G$. As $|\langle x|\psi_d\rangle| \sim 2^{-n/2}$, the partition function $|Z| \sim 2^{(G-n)/2}$ is exponentially smaller in G than the individual terms M_j in its sum. This very strong cancellation seems to prevent any efficient algorithm from being able to accurately estimate the quantity $\langle x|\psi_d\rangle$. Furthermore, except for the restriction to a quasi-3D interaction, these Ising models are similar to the ones obtained from IQP circuits. Therefore, similarly to ref. [17], one might also conjecture that, if the treewidth of the interaction graph is at least linear in n (corresponding to a depth approximately $\sqrt{n}$), any sufficiently large fraction of partition functions can not be probabilistically approximated using an NP oracle. More precisely, if the second moment of the output distribution is close to the value of the Porter–Thomas distribution (verified numerically for depth approximately $\sqrt{n}$), the proof of ref. [17] applies. That is, if a classical algorithm could approximately (in the ℓ_1 norm) sample from the output of these circuits, this would contradict the conjecture on the complexity of approximating partition functions, as above. A related recent conjecture states directly that no polynomial classical algorithm can estimate if $pU(x)$ is above the median with bias better than ~2^{-n} (ref. [32]).

We can use the mapping of $pU(x)$ to the partition function of a complex Ising model to define approximate sampling algorithms for this distribution. We analyse a Bayesian probabilistic algorithm that uses a prior for $pU(x)$ given by the Porter–Thomas distribution in the Supplementary Information. The posterior is updated after sampling random spin-strings s to approximate the partition function. We show that the classical cost required to obtain an appreciable improvement scales as $N^{1/2}2^{G/4}$. There exists also a classical algorithm for sampling so-called extended stabilizer states[35], which scales as $2^{0.23t}$, where t is the number of T gates. Both algorithms are prohibitive for circuits with 7×7 qubits and depth $\gtrsim$40.

Discussion

A crucial aspect of a near-term quantum supremacy proposal is that the computational task can be performed classically only through a direct simulation with cost exponential in the number of qubits. Direct simulations are required for chaotic systems, such as random quantum circuits[1,2]. State-of-the-art supercomputers fail to simulate universal random circuits with more than approximately 48 qubits

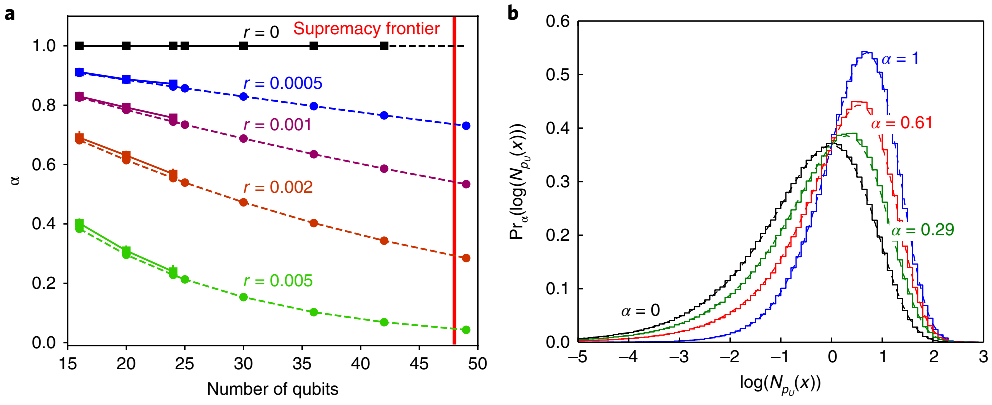

Fig. 4 | Cross-entropy benchmarking and fidelity. a, The circuit fidelity α as a function of the number of qubits. Different colours correspond to different Pauli error rates $r_2 = r_{\mathrm{init}} = r_{\mathrm{res}} = r$ and $r_1 = r/10$. The circular markers correspond to an idealized estimate of fidelities (see the main text). The square markers correspond to the average cross-entropy benchmarking among ten instances. The circuit depth in these simulations is 40. The red line, at 48 qubits and depth 40, is a reasonable estimate of the largest size that can be simulated with state-of-the-art classical computers in practice. The error bars correspond to the standard deviation among instances. **b**, Probability distribution of $\log(Np_U(x))$ where bit-strings x are sampled from a circuit of estimated fidelity α. The continuous-step histograms are obtained from numerical simulations with different Pauli error rates $r_2 = r_{\mathrm{init}} = r_{\mathrm{res}} = r$ and $r_1 = r/10$. The values of r are $r = 0$ for $\alpha = 1$ (blue), $r = 0.002$ for $\alpha = 0.61$ (red), $r = 0.005$ for $\alpha = 0.29$ (green) and uniform sampling of bit-strings for $\alpha = 0$. The superimposed dashed lines correspond to the theoretical distribution derived in the text. We chose a circuit of 5 × 4 qubits and depth 40.

and depth ~40 (see Fig. 3a and Supplementary Information). The evaluation of effective error models for large-scale universal quantum circuits is a difficult theoretical and experimental problem due to their complex nature. Existing proposals involve an expensive additional unitary transformation to the initial state[36] or are restricted to non-universal circuits[42]. We propose cross-entropy benchmarking as a novel way of characterizing and validating error models, and open quantum system theory in general. This method can also be applied to other systems, such as continuous chaotic Hamiltonian evolutions[43]. A successful implementation of the experimental proposal outlined in this paper would require an error rate of around 0.5% for two-qubit gates and 0.05% for one-qubit gates (see Fig. 4a) in a 2D arrangement of 7×7 qubits. This would demonstrate the basic building blocks for a large-scale quantum computer within the operational range of the surface code[19,39].

Data availability. The data that support the plots within this paper and other findings of this study are available from the corresponding author upon reasonable request.

Received: 18 October 2017; Accepted: 21 March 2018;
Published online: 23 April 2018

References

1. Emerson, J., Weinstein, Y. S., Saraceno, M., Lloyd, S. & Cory, D. G. Pseudo-random unitary operators for quantum information processing. *Science* **302**, 2098–2100 (2003).
2. Scott, A. J., Brun, T. A., Caves, C. M. & Schack, R. Hypersensitivity and chaos signatures in the quantum baker's maps. *J. Phys. A* **39**, 13405–13433 (2006).
3. Oliveira, R., Dahlsten, O. & Plenio, M. Generic entanglement can be generated efficiently. *Phys. Rev. Lett.* **98**, 130502 (2007).
4. Arnaud, L. & Braun, D. Efficiency of producing random unitary matrices with quantum circuits. *Phys. Rev. A* **78**, 062329 (2008).
5. Trail, C. M., Madhok, V. & Deutsch, I. H. Entanglement and the generation of random states in the quantum chaotic dynamics of kicked coupled tops. *Phys. Rev. E* **78**, 046211 (2008).
6. Harrow, A. W. & Low, R. A. Random quantum circuits are approximate 2-designs. *Comm. Math. Phys.* **291**, 257–302 (2009).
7. Weinstein, Y. S., Brown, W. G. & Viola, L. Parameters of pseudo-random quantum circuits. *Phys. Rev. A* **78**, 052332 (2008).
8. Brown, W. & Fawzi, O. Scrambling speed of random quantum circuits. Preprint at https://arxiv.org/abs/1210.6644 (2012).
9. Kim, H. & Huse, D. A. Ballistic spreading of entanglement in a diffusive nonintegrable system. *Phys. Rev. Lett.* **111**, 127205 (2013).
10. Hosur, P., Qi, X.-L., Roberts, D. A. & Yoshida, B. Chaos in quantum channels. *J. High. Energy Phys.* **2016**, 4 (2016).
11. Nahum, A., Ruhman, J., Vijay, S. & Haah, J. Quantum entanglement growth under random unitary dynamics. *Phys. Rev. X* **7**, 031016 (2017).
12. Aaronson, S. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. R. Soc. A* **461**, 3473–3482 (2005).
13. Bremner, M. J., Jozsa, R. & Shepherd, D. J. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. R. Soc. A* **467**, 459–472 (2011).
14. Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *STOC '11 Proc. Forty-Third Annual ACM Symp. Theory of Computing* 333–342 (ACM, New York, NY, 2011).
15. Fujii, K. & Morimae, T. Commuting quantum circuits and complexity of Ising partition functions. *New J. Phys.* **19**, 033003 (2017).
16. Goldberg, L. A. & Guo, H. The complexity of approximating complex-valued Ising and Tutte partition functions. *Comput. Complex.* **26**, 765–833 (2017).
17. Bremner, M. J., Montanaro, A. & Shepherd, D. J. Average-case complexity versus approximate simulation of commuting quantum computations. *Phys. Rev. Lett.* **117**, 080501 (2016).
18. Preskill, J. Quantum computing and the entanglement frontier. Preprint at https://arxiv.org/abs/1203.5813 (2012).
19. Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500–503 (2014).
20. Kelly, J. et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature* **519**, 66–69 (2015).
21. Peres, A. Stability of quantum motion in chaotic and regular systems. *Phys. Rev. A* **30**, 1610 (1984).
22. Schack, R. & Caves, C. M. Hypersensitivity to perturbations in the quantum baker's map. *Phys. Rev. Lett.* **71**, 525 (1993).
23. Gorin, T., Prosen, T., Seligman, T. H. & Žnidarič, M. Dynamics of Loschmidt echoes and fidelity decay. *Phys. Rep.* **435**, 33–156 (2006).
24. Beenakker, C. W. Random-matrix theory of quantum transport. *Rev. Mod. Phys.* **69**, 731 (1997).
25. Mehta, M. L. *Random Matrices* Vol. 142 (Academic, San Diego, CA, 2004).
26. Porter, C. & Thomas, R. Fluctuations of nuclear reaction widths. *Phys. Rev.* **104**, 483 (1956).
27. Haake, F. *Signatures of Quantum Chaos* (Springer, Berlin, 1991).
28. Boixo, S., Smelyanskiy, V. N. & Neven, H. Fourier analysis of sampling from noisy chaotic quantum circuits. Preprint at https://arxiv.org/abs/1708.01875 (2017).
29. Bremner, M. J., Montanaro, A. & Shepherd, D. J. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum* **1**, 8 (2017).
30. Boixo, S., Isakov, S. V., Smelyanskiy, V. N. & Neven, H. Simulation of low-depth quantum circuits as complex undirected graphical models. Preprint at https://arxiv.org/abs/1712.05384 (2017).
31. Markov, I. L. & Shi, Y. Simulating quantum computation by contracting tensor networks. *SICOMP* **38**, 963–981 (2008).

32. Aaronson, S. & Chen, L. Complexity-theoretic foundations of quantum supremacy experiments. *CCC'* **17**, 22 (2017).
33. Pednault, E. et al. Breaking the 49-qubit barrier in the simulation of quantum circuits. Preprint at https://arxiv.org/abs/1710.05867 (2017).
34. Bremner, M. J., Mora, C. & Winter, A. Are random pure states useful for quantum computation? *Phys. Rev. Lett.* **102**, 190502 (2009).
35. Bravyi, S. & Gosset, D. Improved classical simulation of quantum circuits dominated by Clifford gates. *Phys. Rev. Lett.* **116**, 250501 (2016).
36. Emerson, J., Alicki, R. & Zyczkowski, K. Scalable noise estimation with random unitary operators. *J. Opt. B* **7**, S347–S352 (2005).
37. Knill, E. et al. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307 (2008).
38. Magesan, E., Gambetta, J. M. & Emerson, J. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.* **106**, 180504 (2011).
39. Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* **86**, 032324 (2012).
40. Barends, R. et al. Digital quantum simulation of fermionic models with a superconducting circuit. *Nat. Commun.* **6**, 7654 (2015).
41. Boixo, S. & Monras, A. Operational interpretation for global multipartite entanglement. *Phys. Rev. Lett.* **100**, 100503–100504 (2008).
42. Flammia, S. T. & Liu, Y.-K. Direct fidelity estimation from few Pauli measurements. *Phys. Rev. Lett.* **106**, 230501 (2011).
43. Neill, C. et al. A blueprint for demonstrating quantum supremacy with superconducting qubits. Preprint at https://arxiv.org/abs/1709.06678 (2017).

Acknowledgements

We especially acknowledge M. Smelyanskiy, from the Parallel Computing Laboratory, Intel Corporation, who performed the simulations of circuits with 6×6 and 7×6 qubits and wrote the corresponding section in the Supplementary Information. We would like to acknowledge A. Montanaro for multiple suggestions, especially regarding IQP circuits. We would like to thank S. Aaronson, A. Fowler, I. Markov, M. Mohseni and E. Rieffel for discussions. The authors also thank J. Hammond, from the Parallel Computing Laboratory, Intel Corporation, for his useful insights into MPI run-time performance and scalability. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the US Department of Energy under contract no. DEAC02-05CH11231. M.J.B. has received financial support from the Australian Research Council via the Future Fellowship scheme (project no. FT110101044) and as a member of the ARC Centre of Excellence for Quantum Computation and Communication Technology (project no. CE170100012).

Author contributions

S.B. designed the project. S.B. and V.N.S. developed most of the theory. S.V.I. performed numerical studies and designed the specific quantum circuits. All authors contributed to several tasks, such as analysis of theory and results and discussions of the draft.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41567-018-0124-x.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to S.B.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.