

ARTICLE OPEN



Boundaries of quantum supremacy via random circuit sampling

Alexander Zlokapa¹✉, Benjamin Villalonga², Sergio Boixo² and Daniel A. Lidar³

Google's quantum supremacy experiment heralded a transition point where quantum computers can evaluate a computational task, random circuit sampling, faster than classical supercomputers. We examine the constraints on the region of quantum advantage for quantum circuits with a larger number of qubits and gates than experimentally implemented. At near-term gate fidelities, we demonstrate that quantum supremacy is limited to circuits with a qubit count and circuit depth of a few hundred. Larger circuits encounter two distinct boundaries: a return of a classical advantage and practically infeasible quantum runtimes. Decreasing error rates cause the region of a quantum advantage to grow rapidly. At error rates required for early implementations of the surface code, the largest circuit size within the quantum supremacy regime coincides approximately with the smallest circuit size needed to implement error correction. Thus, the boundaries of quantum supremacy may fortuitously coincide with the advent of scalable, error-corrected quantum computing.

npj Quantum Information (2023)9:36; <https://doi.org/10.1038/s41534-023-00703-x>

INTRODUCTION

A recent seminal result^{1,2} by Google Quantum AI and collaborators claimed quantum supremacy^{3–11}, sampling pseudo-random quantum circuits on noisy intermediate-scale quantum (NISQ) hardware¹² beyond what can be done in practice by state-of-the-art supercomputers. The observed runtime advantage over classical simulation methods in cross-entropy benchmarking through random circuit sampling provides a critical achievement in establishing a performance benchmark between NISQ computers and classical supercomputers. It shows a quantum runtime speedup beyond the reach of implemented classical algorithms^{6,13–19}. However, it is inevitable that implemented classical algorithms will continue to improve, and there also exist proposals^{20,21} for faster classical implementations which have not been realized so far.

As the quantum circuit width and depth increase, circuit fidelity is observed to decrease exponentially in the absence of quantum error correction (QEC)^{22,23}. We describe the computational task of cross-entropy benchmarking in terms of this fidelity and its variance, with the concomitant increased effort of random circuit sampling on a quantum computer as the fidelity worsens. Our main finding is that currently known classical simulation algorithms can *recover* a classical advantage for cross-entropy benchmarking on sufficiently large quantum circuits. In other words, the quantum advantage is in fact a relatively limited region in the circuit width and depth plane. Using data from Refs. ^{1,2}, we explicitly characterize this region with respect to quantum circuit width and depth. At the current fidelity of quantum hardware established by the Google Sycamore chip, numerical estimates indicate that a quantum advantage may extend to circuits with a depth of a few hundred and a width of several hundred qubits, beyond which a classical advantage returns.

As NISQ hardware continues to improve gate and readout fidelity, we quantify the trajectory of the quantum runtime advantage in cross-entropy benchmarking and place it in the

context of current and future milestones for quantum computing in the NISQ era. This era is expected to last up to the point where QEC becomes pervasive and beneficial—the QEC era—estimated to be around a few thousand qubits for gate-model quantum computers, assuming a concurrent improvement in metrics of fidelity and coherence^{24–27}. We show that improving the gate fidelity enlarges the region of quantum supremacy inversely with gate error, although a classical advantage will still remain for sufficiently large quantum circuits. At the gate fidelity required for QEC, the boundary of quantum supremacy provided by cross-entropy benchmarking may naively be expected to be situated at circuit sizes distant from those in proposed QEC protocols. However, our results suggest that the boundary of advantage in quantum random circuit sampling will soon coincide with circuit sizes sufficiently large for early implementations of QEC. Once in the QEC regime, more significant benefits than are possible from random circuit sampling are anticipated from scaling advantages of quantum algorithms for important problems such as prime factorization²⁸, matrix inversion²⁹ or quantum simulation³⁰. Hence, the extrapolations provided below indicate that we may witness a smooth transition from beyond-classical computation in the NISQ era to applications in the QEC era.

RESULTS

We first briefly review random circuit sampling and the linear cross-entropy benchmarking metric \mathcal{F}_{XEB} to define the precise computational task of quantum supremacy for random circuit sampling. Using an empirical fidelity model described in the Methods section, we assess the expected performance of classical algorithms. We show an asymptotic scaling advantage above a threshold depth with the Schrödinger algorithm; for our main results about the boundaries of a quantum advantage, we also consider the Schrödinger-Feynman algorithm and tensor networks.

¹Division of Physics, Mathematics and Astronomy, Caltech, Pasadena 91125 CA, USA. ²Google AI Quantum, Venice 90291 CA, USA. ³Departments of Electrical and Computer Engineering, Chemistry, and Physics & Astronomy, and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles 90089 CA, USA. ✉email: azlokapa@caltech.edu

Random circuit sampling

In *circuit sampling*, we seek to sample from the probability distribution of outcomes $p_U(x) = |\langle x|U|0\rangle|^2$ for a given quantum circuit U and bitstrings $|x\rangle$, starting from the all-zero string $|0\rangle$. Historically, circuit sampling was first studied for quantum circuits from a family with certain structure: instantaneous quantum polynomial time (IQP) circuits were shown to be hard to simulate classically under reasonable average-case conjectures^{5,31}. In *random circuit sampling* (RCS)¹⁰ as experimentally realized by the Google experiment⁶, we sample from circuits $U \in \mathcal{U}$, defining the set \mathcal{U} to be n -qubit circuits with m cycles, where each cycle consists of a layer of randomly chosen single-qubit gates applied to all qubits followed by a layer of two-qubit gates.

RCS is classically hard for noiseless, fully randomly drawn quantum circuits subject to conjectures widely believed to hold^{10,11} (note that the set \mathcal{U} is not fully random). At the other extreme, if the circuits are completely noisy RCS is classically trivial since then the output bitstrings are uniformly distributed. RCS is also classically easy for very small m ³². Additionally, sufficiently noisy one-dimensional random circuits can be simulated efficiently classically using matrix product state²¹ or operator³³ methods.

To quantify the hardness of RCS, we use an estimator of circuit fidelity called linear cross-entropy benchmarking (XEB) \mathcal{F}_{XEB} ^{1,2,6,9}:

$$\mathcal{F}_{\text{XEB}} + 1 = 2^n \left\langle \sum_x q_U(x) p_U(x) \right\rangle = 2^n \langle \langle p_U(x) \rangle \rangle, \quad (1)$$

where $q_U: \{0,1\}^n \rightarrow [0,1]$ is the probability density function of the distribution obtained by measuring $U|0^n\rangle$, $p_U(x)$ is the probability of bitstring $x \in \{0,1\}^n$ computed classically for the ideal quantum circuit U , and the outer average $\langle \cdot \rangle$ is over random circuits U . For a fully random quantum circuit, the probability distribution over $U|0^n\rangle$ is given by the Porter-Thomas distribution, i.e., an exponential distribution over a random permutation of bitstrings. In practice, ideal unitary circuits will be replaced by noisy versions, generally trace-preserving completely-positive maps \mathcal{E}_U . Thus the sampling is from the noisy probability distribution of outcomes $\tilde{q}_U(x) = \langle x|\mathcal{E}_U(|0\rangle\langle 0|)|x\rangle$, which replaces $q_U(x)$ in Eq. (1). In an experimental setting, the fidelity may still be effectively measured due to the existence of “heavy” bitstrings with a higher probability of being measured by virtue of the Porter-Thomas distribution.

By definition, \mathcal{F}_{XEB} compares how often each bitstring is observed experimentally with its corresponding ideal probability computed via classical simulation, though it is also possible to avoid classical simulation to try to “spoof” \mathcal{F}_{XEB} ^{34,35}. It can also be understood as a test that checks that the observed samples tend to concentrate on the outputs that have higher probabilities under the ideal (Porter-Thomas) distribution for the given quantum circuit, or simply as the probability that no error has occurred while running the circuit. Thus, the cross-entropy benchmark fidelity \mathcal{F}_{XEB} is a good estimator of the experimental circuit fidelity F ^{1,2,6,9,15}. Note that a noiseless random circuit has $\mathcal{F}_{\text{XEB}} = 1$ on average for a perfect simulation, while a distribution p_U which is the uniform distribution or another distribution independent of U will give $\mathcal{F}_{\text{XEB}} = 0$ ³⁵. As described in the Methods section, we use an empirical fidelity model of a circuit with n qubits and m cycles given by

$$F = 2^{-\lambda m(3n - \sqrt{n})/2 - \gamma n}, \quad (2)$$

with parameters λ, γ fitted to data from Ref. 1 shown in Fig. 5 and found to correspond to average gate errors of 0.3% and readout errors of 3% (see Methods). The cross-entropy benchmark fidelity \mathcal{F}_{XEB} is a good estimator of the empirical fidelity F ^{1,6,9,15}, so henceforth we use the two interchangeably.

The quantum supremacy task

Under the conjecture XQUATH introduced in Ref. 34, XEB is shown to be classically hard: given a circuit U on n qubits, generate

N distinct samples x_1, \dots, x_N such that $\langle p_U(x) \rangle \geq b/2^n$ for $b > 1$. Phrased in terms of the fidelity estimator defined in Eq. (1), solving XEB (also referred to as XHOG) requires $N = O(\mathcal{F}_{\text{XEB}}^{-2})$ samples. For clarity, we can restate the task as generating a sample of bitstrings such that the random variable $\mathcal{F}_{\text{XEB}}(n, m)$ for random circuits with n qubits and m cycles can be estimated to within standard deviation $\sigma < \mathcal{F}_{\text{XEB}}$. By the central limit theorem $\sigma = N^{-1/2}$, i.e., this requires $N = O(\mathcal{F}_{\text{XEB}}^{-2})$ bitstrings to be generated². For a quantum computer, the task requires a random circuit to be run N times, while typically classical algorithms either simulate these circuits noiselessly or approximately^{15,17}. Critically, the factor of N samples required by a quantum computer causes the quantum runtime to scale poorly with fidelity, ultimately allowing a classical advantage for XEB at sufficiently large circuit sizes.

The latter is the computational problem solved in Google’s quantum supremacy work¹, which reported $\mathcal{F}_{\text{XEB}}(53, 20) \approx (2.24 \pm 0.21) \times 10^{-3}$, i.e., passed the $\sigma < \mathcal{F}_{\text{XEB}}$ test. As expected, this required $N \sim \mathcal{F}_{\text{XEB}}^{-2} \sim 10^6$ runs and was done in ~ 200 s, a time that has so far resisted attempts to be overtaken by classical algorithms^{17–19}.

We note that although XEB is known to be classically hard under certain plausible conjectures^{1,2,6,34}, it can be spoofed for circuits shallower than those in the Google experiment³⁵. In our analysis, we do not rule out alternative classical algorithms that completely bypass circuit simulation, but do not consider this possibility here given the conjectured exponential classical cost of RCS.

Asymptotic classical scaling advantage above a threshold depth

We address several classical simulation algorithms in this and the following sections for RCS with an n -qubit and m -cycle quantum circuit, including the Schrödinger algorithm^{15,20}, Schrödinger-Feynman algorithm^{7,15,16} and tensor networks^{17–19,36,37}. First, however, we address the scaling of a quantum computer for comparison. Given a time scaling of $T_Q \sim m/\mathcal{F}_{\text{XEB}}^2$ (assuming parallel readout) for the quantum computer to evaluate $1/\mathcal{F}_{\text{XEB}}^2$ samples for cross-entropy benchmarking, we evaluate the scaling for samples using our fidelity model:

$$T_Q \sim m 2^{\lambda m(3n - \sqrt{n}) + 2\gamma n}. \quad (3)$$

It should be noted that the depolarization model leading to this estimate [given by Eq. (2)] is subject to certain caveats noted in the Methods section; with this in mind, we proceed with this model to provide approximate numerical figures that help make the NISQ landscape more concrete.

Out of the different classical simulation algorithms we consider, we begin with the Schrödinger algorithm (SA)^{15,20} due to its favorable asymptotic scaling. Since it provides a full fidelity simulation, an optimal implementation of SA allows us to simulate only one randomly generated circuit and then repeatedly sample from its resulting amplitudes to solve XEB. From Eq. (107) of Ref. 2, this is completed in time

$$T_{\text{SA}} \sim mn 2^n = T_Q n 2^{n(1-2\gamma) - \lambda m(3n - \sqrt{n})}. \quad (4)$$

For sufficiently small constants λ and γ , XEB can be classically solved exponentially faster in m and n using SA for any m greater than a threshold value $m_{\text{th}}(n)$, corresponding to an asymptotic classical advantage for RCS for circuits deeper than

$$m_{\text{th}}(n) = \frac{n(1-2\gamma) + \log_2 n \xrightarrow{n \rightarrow \infty} 1 - 2\gamma}{\lambda(3n - \sqrt{n})} \frac{1}{3\lambda}. \quad (5)$$

For the Google Sycamore device with $n = 53$, this threshold occurs at $m \approx 87$. If the experimentally achieved values of λ, γ may be sustained for larger devices, an advantage for SA is achieved for all $m \gtrsim 71$ as $n \rightarrow \infty$. As these are relatively shallow depths, this

result may be compared to the algorithm of Bravyi, Gosset and Movassagh (BGM)³⁸ for simulating the circuit in time $T_{\text{BGM}} \sim n2^{O(m^2)}$, yielding a classical exponential speedup over quantum RCS for fixed m . Hence, a classical exponential advantage is achieved for RCS in all cases, above a certain width-dependent threshold circuit depth set by the quantum hardware's fidelity. However, classical hardware limitations constrain the experimental realization of such speedups, and thus we turn to the existence of a quantum runtime advantage at non-asymptotic widths and depths.

Limited-scale quantum runtime advantage

Due to limitations in random access memory, the Schrödinger algorithm is infeasible to run for a sufficiently large number of qubits, requiring storage of 2^n complex amplitudes. Similarly, the BGM algorithm has prohibitively large runtime with increasing depth since it scales as $2^{O(m^2)}$. In contrast, while achieving worse asymptotic performance than SA or BGM, the Schrödinger-Feynman algorithm (SFA)^{7,15,16} and tensor networks (TN)^{17–19,36,37} are more suitable to accommodate constraints of available classical hardware due to the use of Feynman paths and patching techniques that do not require keeping the entire quantum state in memory, as explained in more detail below. Both of these classical methods allow circuits to be simulated to partial fidelity F .

For SFA, we optimize the number of patches $p \geq 2$ and the number of paths simulated to satisfy XEB. We need to simulate $2^{kpBm\sqrt{n}}F$ permutations or paths of Schmidt decompositions of cross-gates between patches^{2,15}. After simulating each patch (at a time cost of $2^{n/p}$) we must compute the partial amplitudes of $1/F^2$ bitstrings (but at most 2^n). Assuming that both simulation within patches and in between patches have similar runtime prefactors, the time scaling is

$$T_{\text{SFA}} = 2^{kpBm\sqrt{n}}F \left(p2^{n/p} + \min(F^{-2}, 2^n) \right) \quad (6)$$

where $k = 1$ for $p = 2$ patches, $k = 3/4$ for $p = 4$ patches, and so on^{2,15,16}, approximated as $k = \frac{1}{2} + \frac{1}{p}$. The constant $B = 0.24$ is given by the grid layout of the Sycamore chip². Optimizing the runtime T_{SFA} as a function of the simulation fidelity F gives $F^{-2} = p2^{n/p}$ for $n > \log_2(p)/(1 - 1/p)$. In contrast to the SA memory usage of 2^n complex amplitudes, SFA with p patches only requires $2p2^{n/p}$ complex amplitudes at the optimal simulation fidelity. Although increasing the number of patches reduces memory requirements, larger p increases T_{SFA} runtime as well. In practice, SFA runtimes may be improved by taking checkpoints during simulation¹⁵, but the leading order in runtime scaling is given by Eq. (6).

To compare the expected runtime of a classical model to quantum hardware given our fidelity model, we define the total runtime of the computational task $R_x(n, m) = \tau_x T_x(n, m)$, with $x \in \{Q, \text{SA}, \text{SFA}\}$. The runtime constants τ_x are obtained from Ref. 2; for SFA this constant captures the fact that parallelization may reduce the runtime significantly. Since the total runtime is exponential in both m and n in all cases, and any scaling advantage is therefore due to a reduced exponent, we also define a more natural notion of an *effective polynomial quantum speedup* $\alpha_C > 0$ if for a classical method C (with, for our purposes, $C \in \{\text{SA}, \text{SFA}\}$) the time scalings are related via $T_C(n, m) = T_Q(n, m)^{1+\alpha}$, i.e.,

$$\alpha_C = \frac{\log T_C}{\log T_Q} - 1. \quad (7)$$

A value of $\alpha_C > 0$ ($\alpha_C < 0$) implies a quantum (classical) advantage, capturing the relative reduction of the exponent in $T_Q(n, m)$ ($T_C(n, m)$). We compare the runtime advantage boundaries created by the effective polynomial speedup analysis and a direct runtime analysis. The result is shown in Fig. 1, with runtime constants estimated from empirical data². The Google

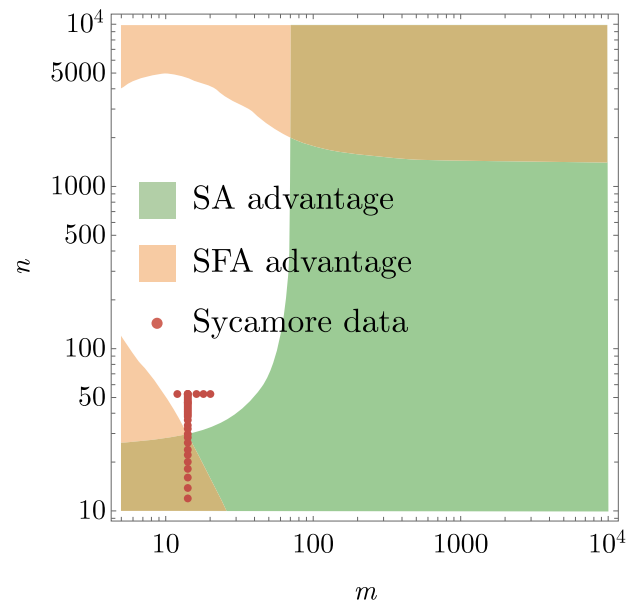


Fig. 1 Boundaries of an effective polynomial quantum speedup (Eq. (7)) for circuits with n qubits and m cycles (without accounting for memory constraints). Circuits that are too wide or deep have low fidelity, enabling a classical advantage. However, the Google Sycamore experiments (red dots) achieved a sufficiently high fidelity to enter the region of a quantum runtime advantage, denoted by the white region ($\alpha_C > 0$). Only depths larger than 5 cycles are shown due to recent polynomial-time simulation results for shallow 2D circuits³².

Sycamore experiments are inside the boundary of a quantum advantage, which extends to a maximum depth and width of approximately 70 cycles and 10^4 qubits, beyond which classical algorithms dominate.

The results shown in Fig. 1 do not account for memory constraints. To more accurately place bounds on a quantum runtime advantage, the memory limitations of classical hardware must be considered. An additional boundary is imposed by infeasible quantum runtimes, resulting in Fig. 2. Besides considering current quantum devices, we project superconducting qubit NISQ error rates and classical hardware to estimate the computational feasibility of XEB in the near term. Error rates of two-qubit gates of superconducting qubit NISQ devices are taken from isolated measurements of transmons at a given year^{1,39–44} to determine an exponential fit shown in the inset of Fig. 2(c), which is then applied to the fitted constants γ and λ in the error model of Eq. (2). Although not exact, this provides an estimate of a reasonable range of error rates to consider on a 5-year timescale.

Besides the return of a classical advantage for sufficiently large quantum circuits due to requiring $O(1/F^2)$ samples to solve XEB with a quantum computer, the numerical extrapolations above indicate additional boundaries on quantum supremacy caused by intractable quantum runtimes and classical memory limitations. We observe that the Google experiments have achieved a critical fidelity threshold to gain a runtime advantage over classical simulation. Had error rates been around $2.8 \times$ larger than Sycamore's (corresponding to an increased isolated two-qubit gate error rate of 1%), no quantum advantage would have been achieved in cross-entropy benchmarking within 100-year quantum runtimes [Fig. 2(a)]. However, even at the fidelity achieved by the Google experiment, the quantum runtime advantage for XEB stops at a few hundred qubits due to long quantum runtimes (Fig. 2(b)). Extrapolations suggest that even at achievable near-term fidelities below surface code thresholds, cross-entropy benchmarking will yield a runtime advantage up to at most a thousand

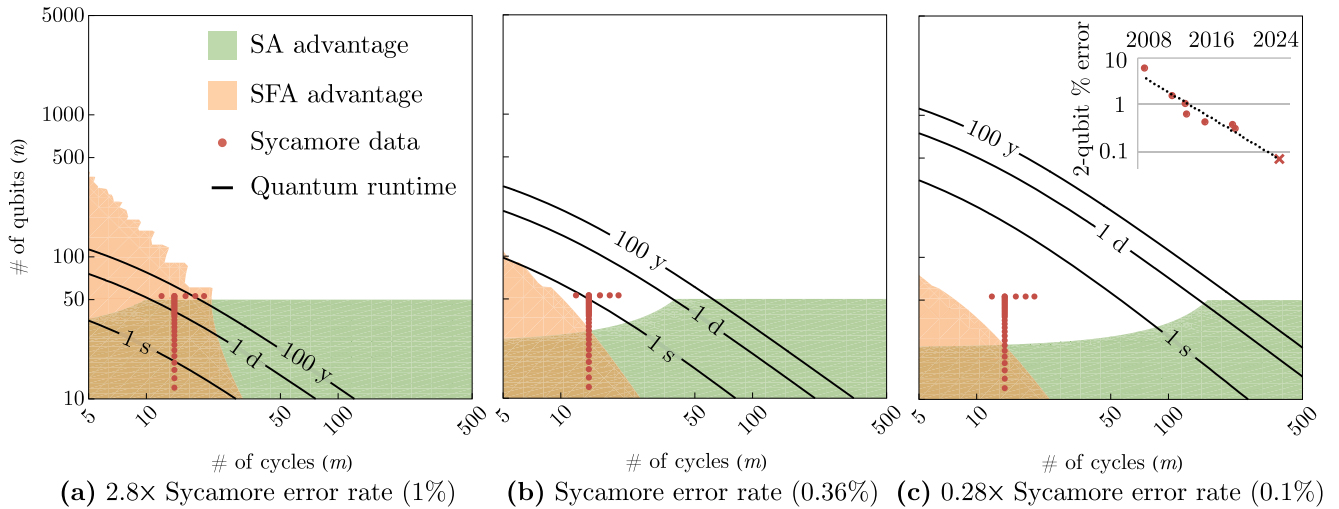


Fig. 2 Classical and quantum runtime advantage boundaries after accounting for memory constraints. Subfigures show different error rates relative to the Sycamore error model fitted in Fig. 5, with corresponding captions indicating isolated two-qubit gate error rates. Black contours indicate quantum device runtimes; colored regions indicate where a classical runtime advantage is expected according to supercomputer memory and performance; red dots show the circuit width (n) and number of cycles (m) of each Sycamore experiment reported in¹. Panel (a) shows that a mere $2.8\times$ factor increase in the error rate relative to the Sycamore chip would have required a quantum runtime of 100 years to break even with SA and SFA. Panel (b) shows the performance of SA and SFA relative to quantum runtimes at the error rate of the Sycamore device. Panel (c) uses the extrapolated fidelity of a state-of-the-art NISQ device in 2025 (inset, error rates decay by a factor of ~ 0.77 per year), and illustrates how even modest gains in error rates can significantly move the feasible quantum supremacy boundary. Extrapolated error rates are given by an exponential regression over transmon device two-qubit gate errors^{1,39–44}. All errors of the Sycamore device (single-qubit/two-qubit gates, readout errors) are scaled proportionally to the extrapolation. Runtimes are computed using Eqs. (3), (4), and (6) with p optimized within memory limitations, using performance and memory values reported in Ref. ².

qubits, beyond which quantum runtimes are computationally infeasible [Fig. 2(c)]. However, the regime of quantum advantage rapidly improves with lower error rates, underscoring the importance of achieving lower error rates for NISQ devices.

Sampling using tensor networks

In this section we consider the estimated runtimes of tensor network methods in the quantum supremacy sampling task. The simulation cost using tensor networks has been drastically reduced since the original experimental supremacy result (see, e.g., Refs. ^{45,46}), and in this section we provide an up-to-date analysis.

We consider random circuits of depth m similar to those of Ref. ¹ over square lattices of qubits of size $\sqrt{n}\times\sqrt{n}$. In line with Section the quantum supremacy task we require the number of bitstrings N to be large enough to resolve the variable \mathcal{F}_{XEB} with certainty comparable to that of Ref. ¹. In order to do so, we consider the parameters of the largest circuits ran in Ref. ¹, i.e., $n = 53$, $m = 20$, $N = 10^6$, and $\mathcal{F}_{\text{XEB}} = 2.24 \times 10^{-3}$. If we sample with $\mathcal{F}_{\text{XEB}} = 1$, then $N \propto O(\mathcal{F}_{\text{XEB}}^{-2})$ yields $N = 10^6 \left(\frac{2.24 \times 10^{-3}}{1} \right)^2 = 5.02 \approx 5$. Note that tensor network methods can sample at low fidelity with a speedup approximately equal to $\mathcal{F}_{\text{XEB}}^{-1}$. However, given that the number of bitstrings N necessary to achieve a constant $\mathcal{F}_{\text{XEB}}/\sigma$ ratio grows with $\mathcal{F}_{\text{XEB}}^{-2}$, it is advantageous for tensor network methods to sample with maximal $\mathcal{F}_{\text{XEB}} = 1$. As introduced in Ref. ⁴⁷, we assume that each contraction of the tensor network generated by the random circuit yields an uncorrelated sample from the output distribution. Ref. ⁴⁶ introduced a method to sample many uncorrelated bitstrings per contraction. Similarly, Ref. ⁴⁵ introduced a method to reuse intermediate computations across many independent tensor network contractions. However, these methods were heavily optimized for the $m = 20$, $n = 53$ circuits of Ref. ¹ and it is unclear how they scale asymptotically. We do not include them in our analysis. As discussed below, this does not affect the validity of our lowest time estimates.

As mentioned above, one can also define the sampling task with a fixed number of bitstrings $N \approx 10^6$ and a target fidelity \mathcal{F}_{XEB} ^{1,6,15,17–19,47}. In that case, tensor network methods can advantageously leverage the linear speedup introduced in Refs. ^{15,47}, i.e., sampling with fidelity \mathcal{F}_{XEB} with a reduction in computation time by a factor $\mathcal{F}_{\text{XEB}}^{-1}$. In this section we analyze the performance of tensor networks for both definitions of the task.

In order to get tensor network runtime estimates for the sampling tasks we proceed as follows. First, we generate random circuits similar to those run on Sycamore over square lattices of size $n = 3 \times 3$ through $n = 20 \times 20$, and depths $m = 4$ through $m = 60$. We then optimize the contraction ordering of the tensor network induced by these circuits using CoTenGra¹⁸ for a maximum contraction width of 29, i.e., the largest tensor in the contraction cannot exceed 2^{29} in size^{18,19,48}. This width is chosen so that the memory requirement of a single contraction is smaller than the memory available on each GPU in Summit, as discussed in more length in Refs. ^{18,19}. We use the KaHyPar driver in the optimization and make use of the subtree reconfiguration functionality implemented in CoTenGra for subtrees of size 10^{19} . For each circuit, we run the optimization 500 times. This yields an upper bound on the number of floating point operations required for each circuit of parameters (m, n) , for a discrete set of (m, n) points. We then linearly interpolate the logarithm of this function, $\log_2 \text{FLOP}(m, n)$, which lets us estimate the FLOP cost of a contraction over the (m, n) plane. The estimated cost for the $(m = 20, n = 53)$ circuits is larger than the one found in Ref. ¹⁹, presumably due to the hardness of simulation of square lattices as compared to the rectangular $n = 53$ Sycamore topology. (Sycamore's layout is not really a rectangle, but a rotated surface code layout with a slight asymmetry in the length of both axes. In addition, there is a missing qubit on one edge, which makes it easier to simulate.) In order to be consistent with the estimate of Ref. ¹⁹, we define a new function

$$\log_2 \text{FLOP}'(m, n) = \log_2 \text{FLOP}(m, n) - D, \quad (8)$$

where $D = \log_2 \text{FLOP}(20, 53) - \log_2 \text{FLOP}_{\text{Syc}}(20, 53)$ and $\log_2 \text{FLOP}_{\text{Syc}}(20, 53)$ is the estimate given by Ref. ¹⁹. In other words, we shift $\log_2 \text{FLOP}(m, n)$ in order for it to match the best known ($m = 20, n = 53$) estimate for the Sycamore topology. Finally, the time estimate $T_{\text{Summit}}(m, n)$ to finish the sampling task on the Summit supercomputer using single precision floating point complex numbers is

$$T_{\text{Summit}}(m, n) = 8E_{\text{Summit}} \frac{2^{\log_2 \text{FLOP}'(m, n)} \times N \times \mathcal{F}_{\text{XEB}}}{\text{FLOPs}_{\text{Summit}}}, \quad (9)$$

where N is the number of bitstrings sampled, the \mathcal{F}_{XEB} factor accounts for the linear speedup due to low fidelity sampling, $\text{FLOPs}_{\text{Summit}} \approx 4 \times 10^{17} \text{s}^{-1}$ is the peak single precision FLOP count per second delivered by Summit, the factor of 8 accounts for the use of single precision multiply-add operations in the matrix multiplications involved in the tensor contractions¹⁷, and $E_{\text{Summit}} \approx 0.15$ is the empirical efficiency of the tensor contractions run on Summit's NVIDIA V100 GPUs¹⁹.

Figure 3 shows the estimated compute time needed to complete the sampling tasks over n qubits and m cycles using the Summit supercomputer, $T_{\text{Summit}}(m, n)$. As expected, for constant, sufficiently small n , increasing depth m does not add complexity to the contraction beyond a linear factor. This is not true at constant m and increasing n on a square lattice.

Estimates for parameters $\mathcal{F}_{\text{XEB}} = 1$ and $N = 5$ are shown with black, solid lines. Estimates for $\mathcal{F}_{\text{XEB}} = 2.24 \times 10^{-3}$ and $N = 10^6$ are shown with blue, dotted lines. As discussed above, the first set of parameters is targeted at achieving $\mathcal{F}_{\text{XEB}} > \sigma$, while the second is targeted at sampling $N = 10^6$ output bitstrings with fidelity similar to that one of the experiment of Ref. ¹. Sampling 5 bitstrings from the ($m = 20, n = 53$) circuit with fidelity $\mathcal{F}_{\text{XEB}} = 1$ is estimated to take about 1.23 h on Summit. Sampling 10^6 bitstrings from the same circuit with fidelity $\mathcal{F}_{\text{XEB}} \approx 2.24 \times 10^{-3}$ is estimated to take about 23 days on Summit, which is consistent with Ref. ¹⁹, which estimated 20 days with $\mathcal{F}_{\text{XEB}} \approx 2 \times 10^{-3}$. Note that we have not included the methods studied in Refs. ^{45,46} in our analysis, given that it is not clear how they would scale at larger values of m and n . However, the low number of bitstrings needed when sampling with $\mathcal{F}_{\text{XEB}} = 1$ makes our estimate for the ($m = 20, n = 53$) circuits

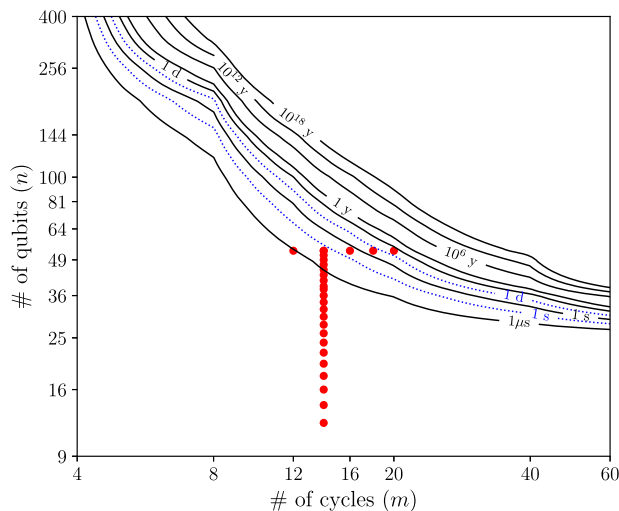


Fig. 3 Time estimates for the quantum supremacy sampling task using tensor networks and the Summit supercomputer. Black, solid contours give estimates for perfect fidelity simulations, which require only 5 bitstrings to achieve $\mathcal{F}_{\text{XEB}} > 0$ with similar confidence to the experiment of Ref. ¹. Blue, dotted contours consider a fidelity and a number of bitstrings similar to the one achieved by the experiment of Ref. ¹ for the largest circuits, i.e., $n = 53$ qubits and $m = 20$ cycles. See the main text for a discussion of both sampling tasks.

consistent with the results of Ref. ⁴⁶, which manages to sample a large number of uncorrelated bitstrings at a cost not much larger than that of computing a single amplitude. Ref. ⁴⁶ uses 512 NVIDIA A100 GPUs and completes the sampling task in 15 h. Summit, on the other hand, has 27648 NVIDIA V100 GPUs, which have less computational power and memory than the A100 model. Extrapolating time estimates for this method to Summit is therefore not straightforward. Similarly, our time estimate for these circuits is much lower than the 7.5 days of Ref. ⁴⁵. Note that Ref. ⁴⁹ claimed that it solved the sampling task in about 300 seconds. However, this work generated a single set of correlated bitstrings, unlike Refs. ^{45,46} and our own estimates, which sample uncorrelated bitstrings, similar to experimental RCS.

In our estimates we have targeted a contraction width of 29 in order to decrease memory requirements to fit the memory of NVIDIA V100 GPUs. One could wonder how much time estimates decrease when memory constraints are loosened. It turns out that in the unrealistic case where we do not constrain the memory requirements (as in Fig. 1), compute times go down by about a factor of 4¹⁹. This would place the time estimate for sampling 5 bitstrings from the ($m = 20, n = 53$) circuit at fidelity $\mathcal{F}_{\text{XEB}} = 1$ at about 18.5 min. Sampling 10^6 bitstrings at fidelity $\mathcal{F}_{\text{XEB}} = 2.24 \times 10^{-3}$ from the same circuit would take about 5.75 days. This is to be compared to 200 seconds for the experiment of Ref. ¹ with $N = 10^6$ and $\mathcal{F}_{\text{XEB}} = 2.24 \times 10^{-3}$.

DISCUSSION

Due to an asymptotic classical advantage for circuits with sufficiently small n and large m for random circuit sampling at fixed quantum fidelity, we find that the projected quantum runtime advantages for the next five years in solving the XEB problem underlying Google's quantum supremacy demonstration¹ are limited to the very early NISQ regime. However, reducing the component error rate increases the quantum advantage regime rapidly, which underscores the importance of a continued emphasis on error rate reduction. We observe that while our work can be interpreted as placing a practical upper bound on circuit width and depth for which RCS-based quantum supremacy holds, a rigorous lower bound based on complexity theory conjectures ruling out all possibility of competitive classical simulation algorithms, both known and unknown, was presented in Ref. ⁵⁰ for other supremacy proposals in the noiseless setting.

For narrow circuits within classical memory constraints of up to 50 qubits, a classical advantage is expected to remain in the near-term beyond depths of a few hundred [see Fig. 2]. As given by Eq. (2), these results assume depolarization error with the additional simplification into independent cycle and qubit errors. While the model provides a reasonable heuristic, the possible superexponential decay of fidelity discussed in the Methods section implies that this is an upper bound on the expected fidelity for larger circuits. Hence, the presence of the upper boundary on the quantum supremacy regime caused by requiring $O(1/F^2)$ samples with a decaying fidelity may occur at even shallower depths. In the absence of quantum data⁵¹ or an oracle⁵², we observe that the lack of structure in random circuit sampling provides a lower bound on hardness: any circuit with additional structure (such as quantum simulation) must lie within the boundaries of quantum RCS advantage to achieve a quantum speedup.

While RCS provides a purposeful milestone for measuring the progress of quantum devices, we have shown that the regime of quantum advantage for a standard approach to XEB in the near term is upper-bounded by about a thousand qubits, rapidly approaching circuit sizes sufficient for early QEC. Given that fidelities comparable to those achieved in the Google experiment are close to establishing surface codes at a few hundred to a few thousand physical qubits^{24–27}, we anticipate that the disappearance of the quantum runtime advantage in RCS shown here

approximately coincides with the onset of error-corrected quantum computing. Hence, while the point of comparison in XEB and other metrics are either established through or motivated by classical simulation of random quantum circuits, this transition point into QEC motivates the adoption of problem-specific or QEC-specific metrics of progress for the field. By focusing on problems that do not require direct circuit simulation — similar to the benchmarking of quantum annealers vs state-of-the-art classical optimization algorithms^{53–56} — we may obtain a more informative view of the usefulness of a quantum device.

In combination with the above limitations on intractable quantum runtimes for wide circuits, this result naively suggests a non-smooth transition to QEC characterized by a classical advantage, as circuits may be too large to achieve a quantum advantage via XEB but too small to perform computational tasks aided by QEC. However, the numerical results suggest that the onset of QEC will approximately coincide with the boundary of a feasible quantum advantage in random circuit sampling. To be precise, performing XEB on 1000 qubits within reasonable quantum runtime and depth $m = \sqrt{n}$ requires error rates that are an order of magnitude reduced from Sycamore; this coincides with error rates at the fault tolerance threshold. Many of the most appealing results for quantum computers are far more transformative in the presence of QEC. With this in mind we anticipate that given the limitations of the XEB problem in maintaining a quantum advantage upon its intersection with early QEC circuit sizes, it will be natural to transition from the NISQ era to the QEC era along with metrics designed to capture performance under QEC, such as logical error rate.

METHODS

Depolarization error model

We describe the fidelity model given in Eq. (2) and fitted to the Google experiment (Fig. 5).

We assume a depolarization error model, which yields a fidelity that decays like a power law with the addition of gates and qubits. For a gate set G , gate errors e_g , qubit set Q , and qubit errors (measurement and state preparation) e_q , we approximate the fidelity in the absence of QEC by $F = \prod_{g \in G} (1 - e_g) \prod_{q \in Q} (1 - e_q)^2$.

We simplify the approximation of fidelity into cycle errors, assigning a constant error per cycle and per readout. A single cycle consists of single-qubit gates applied to all n qubits followed by a pattern of two-qubit gates along edges in the planar graph, denoted by a single color in Fig. 4. For a square lattice of n qubits, each row and each column includes $\sqrt{n} - 1$ two-qubit gates. Each cycle contains two-qubit gates of just one color. For odd n , all four colors appear $\sqrt{n} \times (\sqrt{n} - 1) / 2 = (n - \sqrt{n}) / 2$ times. For even n , blue and yellow each appear $\sqrt{n} \times (\sqrt{n} / 2)$ times, whereas red and green appear $\sqrt{n} \times (\sqrt{n} / 2 - 1)$ times. Thus, also for even n on average $(n - \sqrt{n}) / 2$ gates are applied per cycle. Hence, a total of $m(n + (n - \sqrt{n}) / 2)$ gates are applied for each circuit of m cycles. Including a separate factor for the readout error, we thus approximate the fidelity as in Eq. (2).

Due to this error model, we may observe the noise dependence of the outer boundary of the quantum supremacy regime defined by setting T_Q to a given constant (e.g. 1 year). Since $T_Q \sim m/F^2$ and $F \sim (1 - e)^{-mn}$ to leading order for characteristic gate error rate e , fixing m (or n) and Taylor expanding $\log(1 - e)$ in small e shows that the number of gates (i.e., the boundary in m or n) grows like $O(1/e)$. This provides the rapid growth of the quantum supremacy regime with improved error rates, as discussed in the abstract and main text.

We perform an empirical fit of the fidelity model to data from Ref. ¹ shown in Fig. 5. The parameters λ , γ are constants that result from the regular application of cycles and readout errors, respectively, as random gate selection in RCS allows us to adopt

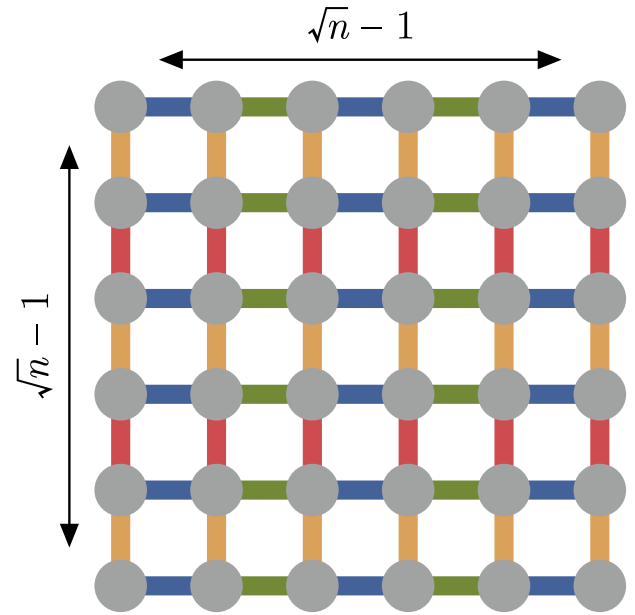


Fig. 4 Cycles of a random circuit sampling experiment on an n -qubit planar architecture. Each cycle consists of single-qubit gates followed by a pattern of two-qubit gates, such as those highlighted by a single color in the diagram.

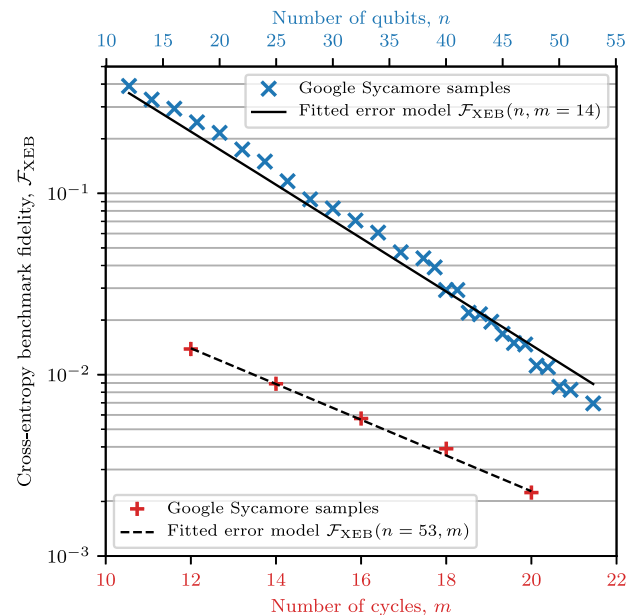


Fig. 5 Empirical fidelity model Eq. (2) for $\lambda = 0.0043 \pm 0.0008$ and $\gamma = 0.042 \pm 0.017$ (to two standard deviations) showing quality of fit for elided verifiable circuits of fixed depth (blue) and elided supremacy circuits of fixed width (red). The deviation from the fit is largely caused by the proportion of two-qubit gates used in larger circuits. Source: Fig. 4 of Ref. ¹.

an effective error rate per cycle or qubit. As noted in the Results section, the cross-entropy benchmark fidelity \mathcal{F}_{XEB} is a good estimator of the empirical fidelity $F^{1,6,9,15}$. The values of λ and γ may also be related directly to gate fidelities: rearranging the depolarization error model for F gives an average gate error of $e_g = 1 - 2^{-\lambda} \approx 0.3\%$ and an average readout error of $e_q = 1 - 2^{-\gamma} \approx 3\%$, assuming equal errors for all gates and qubits. These numbers closely match the observed gate and readout errors reported in the original quantum supremacy experiment of

0.1–1% (depending on the type of gate and presence neighboring gates) and 3.1–3.8% (depending on isolated or simultaneous measurement) respectively¹.

The superexponential decay in fidelity for larger circuits visible in Fig. 5 largely occurs due to the increased proportion of two-qubit gates for circuits with a relatively smaller boundary, as seen by the better fit in Fig. 4 of Ref. ¹ that considers gate-specific noise and accounts for a non-square geometry. If other sources of noise such as $1/f$ noise appear on longer timescales⁵⁷ or cross-talk between qubits increases with wider circuits, the quantum advantage region may be further constrained. Hence, the depolarization model provides an upper bound on the expected fidelity of large quantum circuits. Nevertheless, our results above showing a return of classical advantage for sufficiently large quantum circuits are robust to this, although the precise location of the boundary may shift if the depolarization error assumption is weakened.

DATA AVAILABILITY

We provide all data for computing quantum advantage boundaries in a supplemental GitHub repository⁵⁸.

Received: 13 February 2021; Accepted: 29 March 2023;

Published online: 11 April 2023

REFERENCES

- Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
- Arute, F. et al. Supplementary information for “quantum supremacy using a programmable superconducting processor”. Preprint at <https://arxiv.org/abs/1910.11333> (2019).
- Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In: Proc. Annu. ACM Symp. Theory Comput. STOC '11, pp. 333–342. Association for Computing Machinery, New York, NY, USA <https://doi.org/10.1145/1993636.1993682> (2011).
- Bremner, M. J., Jozsa, R. & Shepherd, D. J. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **467**, 459–472 (2011).
- Bremner, M. J., Montanaro, A. & Shepherd, D. J. Average-case complexity versus approximate simulation of commuting quantum computations. *Phys. Rev. Lett.* **117**, 080501 (2016).
- Boixo, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
- Aaronson, S. & Chen, L. Complexity-Theoretic Foundations of Quantum Supremacy Experiments. In: O’Donnell, R. (ed.) Ann. IEEE Conf. Comput. Leibniz International Proceedings in Informatics (LIPIcs), vol. 79, pp. 22–12267. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany <https://doi.org/10.4230/LIPIcs.CCC.2017.22> (2017).
- Harrow, A. W. & Montanaro, A. Quantum computational supremacy. *Nature* **549**, 203 (2017).
- Neill, C. et al. A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science* **360**, 195 (2018).
- Bouland, A., Fefferman, B., Nirkhe, C. & Vazirani, U. On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* **15**, 159–163 (2018).
- Movassagh, R. Cayley path and quantum computational supremacy: A proof of average-case #P-hardness of random circuit sampling with quantified robustness. Preprint at <https://arxiv.org/abs/1909.06210> (2019).
- Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- Pednault, E. et al. Breaking the 49-qubit barrier in the simulation of quantum circuits. Preprint at <http://arxiv.org/abs/1710.05867> (2017).
- De Raedt, H. et al. Massively parallel quantum computer simulator, eleven years later. *Comput. Phys. Commun.* **237**, 47–61 (2019).
- Markov, I. L., Fatima, A., Isakov, S. V. & Boixo, S. Quantum Supremacy Is Both Closer and Farther than It Appears. Preprint at <https://arxiv.org/abs/1807.10749> (2018).
- Chen, Z.-Y. et al. 64-qubit quantum circuit simulation. *Sci. Bull.* **63**, 964–971 (2018).
- Villalonga, B. et al. Establishing the quantum supremacy frontier with a 281 pflop/s simulation. *Quantum Sci. Technol.* **5**, 034003 (2020).
- Gray, J. & Kourtis, S. Hyper-optimized tensor network contraction. *Quantum* **5**, 410 (2021).
- Huang, C. et al. Classical simulation of quantum supremacy circuits. Preprint at <https://arxiv.org/abs/2005.06787> (2020).
- Pednault, E., Gunnels, J. A., Nannicini, G., Hoesesh, L. & Wisnieff, R. Leveraging secondary storage to simulate deep 54-qubit sycamore circuits. Preprint at <https://arxiv.org/abs/1910.09534> (2019).
- Zhou, Y., Stoudenmire, E. M. & Waintal, X. What limits the simulation of quantum computers? *Phys. Rev.* **10**, 041038 (2020).
- Gottesman, D. An introduction to quantum error correction and fault-tolerant quantum computation. In: Lomonaco, S.J. (ed.) Quantum Information Science and Its Contributions to Mathematics. Proceedings of Symposia in Applied Mathematics, vol. 68, p. 13. Amer. Math. Soc., Rhode Island (2010).
- Lidar, D. A. & Brun, T. A. (eds.) Quantum Error Correction. Cambridge University Press, Cambridge, UK <http://www.cambridge.org/9780521897877> (2013).
- Raussendorf, R. & Harrington, J. Fault-tolerant quantum computation with high threshold in two dimensions. *Phys. Rev. Lett.* **98**, 190504 (2007).
- Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev.* **86**, 032324 (2012).
- Yoder, T. J. & Kim, I. H. The surface code with a twist. *Quantum* **1**, 2 (2017).
- Bravyi, S., Englbrecht, M., König, R. & Peard, N. Correcting coherent errors with surface codes. *npj Quantum Inf.* **4**, 55 (2018).
- Shor, P. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484–1509 (1997).
- Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009).
- Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
- Bremner, M. J., Montanaro, A. & Shepherd, D. J. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum* **1**, 8 (2017).
- Napp, J. C., La Placa, R. L., Dalzell, A. M., Brandão, F. G. S. L. & Harrow, A. W. Efficient classical simulation of random shallow 2d quantum circuits. *Phys. Rev.* **12**, 021021 (2022).
- Noh, K., Jiang, L. & Fefferman, B. Efficient classical simulation of noisy random quantum circuits in one dimension. *Quantum* **4**, 318 (2020).
- Aaronson, S. & Gunn, S. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory Comput.* **16**, 1–8 (2020).
- Barak, B., Chou, C. & Gao, X. Spoofing linear cross-entropy benchmarking in shallow quantum circuits. Preprint at <https://arxiv.org/abs/2005.02421> (2020).
- Markov, I. L. & Shi, Y. Simulating quantum computation by contracting tensor networks. *SIAM J. Comput.* **38**, 963–981 (2008).
- 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).
- Bravyi, S., Gosset, D. & Movassagh, R. Classical algorithms for quantum mean values. *Nat. Phys.* **17**, 337–341 (2021).
- DiCarlo, L. et al. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature* **460**, 240–244 (2009).
- Chow, J. M. et al. Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Phys. Rev. Lett.* **109**, 060501 (2012).
- Chen, Y. et al. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.* **113**, 220502 (2014).
- Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* **508**, 500–503 (2014).
- Sheldon, S., Magesan, E., Chow, J. M. & Gambetta, J. M. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Phys. Rev.* **93**, 060302 (2016).
- Kjaergaard, M. et al. A quantum instruction set implemented on a superconducting quantum processor. Preprint at <https://arxiv.org/abs/2001.08838> (2020).
- Kalachev, G., Panteleev, P. & Yung, M.-H. Multi-Tensor Contraction for XEB Verification of Quantum Circuits. Preprint at <https://arxiv.org/abs/2108.05665> (2021).
- Pan, F., Chen, K. & Zhang, P. Solving the sampling problem of the sycamore quantum circuits. *Phys. Rev. Lett.* **129**, 090502 (2022).
- Villalonga, B. et al. A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware. *npj Quantum Inf.* **5**, 86 (2019).
- Chen, J., Zhang, F., Huang, C., Newman, M. & Shi, Y. Classical simulation of intermediate-size quantum circuits. Preprint at <https://arxiv.org/abs/1805.01450> (2018).
- Liu, Y. A. et al. Closing the “quantum supremacy” gap: Achieving real-time simulation of a random quantum circuit using a new runway supercomputer. In: Int. Conf. High. Perfor. SC '21. Association for Computing Machinery, New York, NY, USA <https://doi.org/10.1145/3458817.3487399> (2021).
- Dalzell, A. M., Harrow, A. W., Koh, D. E. & La Placa, R. L. How many qubits are needed for quantum computational supremacy? *Quantum* **4**, 264 (2020).
- Huang, H.-Y. et al. Quantum advantage in learning from experiments. *Science* **376**, 1182–1186 (2022).

52. Pokharel, B. & Lidar, D. A. Demonstration of algorithmic quantum speedup. Preprint at <https://arxiv.org/abs/2207.07647> (2022).
53. Rønnow, T. F. et al. Defining and detecting quantum speedup. *Science* **345**, 420–424 (2014).
54. Albash, T. & Lidar, D. A. Demonstration of a scaling advantage for a quantum annealer over simulated annealing. *Phys. Rev.* **8**, 031016 (2018).
55. Denchev, V. S. et al. What is the computational value of finite-range tunneling? *Phys. Rev.* **6**, 031015 (2016).
56. Mandrà, S. & Katzgraber, H. G. A deceptive step towards quantum speedup detection. *Quantum Sci. Technol.* **3**, 04–01 (2018).
57. Quintana, C. M. et al. Observation of classical-quantum crossover of $1/f$ flux noise and its paramagnetic temperature dependence. *Phys. Rev. Lett.* **118**, 057702 (2017).
58. Zlokapa, A., Villalonga, B., Boixo, S., Lidar, D.: Boundaries of quantum supremacy via random circuit sampling – Supplementary Information. GitHub <https://github.com/quantummind/quantum-rqs-boundaries/> (2021).

ACKNOWLEDGEMENTS

The authors thank Tameem Albash, Fernando Brandão, Elizabeth Crosson, and Maria Spiropulu for discussions, and Johnnie Gray for tensor network simulations of Sycamore circuits. A.Z. and D.A.L. acknowledge support from Caltech's Intelligent Quantum Networks and Technologies (INQNET) research program and by the DOE/HEP QuantISED program grant, Quantum Machine Learning and Quantum Computation Frameworks (QMLQCF) for HEP, award number DE-SC0019227. D.A.L. further acknowledges support from the Oracle Corporation.

AUTHOR CONTRIBUTIONS

D.A.L. conceived the project. A.Z., B.V., S.B., and D.A.L. wrote the manuscript. A.Z., S.B., and D.A.L. performed the analytic analysis of runtimes. A.Z. implemented the simulations of S.A., and S.F.A., B.V. implemented the simulations of tensor networks.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Alexander Zlokapa.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023