Eleanor G. Rieffel

NASA Ames Research Center

Moffett Field, California

Quantum supremacy using a programmable superconducting processor Google AI Quantum and collaboratorsy The tantalizing promise of quantum computers is that certain computational tasks might be executed exponentially faster on a quantum processor than on a classical processor. A fundamen- tal challenge is to build a high- delity processor capable of running quantum algorithms in an exponentially large computational space.

Here, we report using a processor with programmable superconducting qubits to create quantum states on 53 qubits, occupying a state space 253 ˇ1016. Measurements from repeated experiments sample the corresponding probability distribution, which we verify using classical simulations. While our processor takes about 200 seconds to sample one instance of the quantum circuit 1 million times, a state-of-the-art supercomputer would require approximately 10,000 years to perform the equivalent task.

This dramatic speedup relative to all known classical algorithms provides an experimental realization of quantum supremacy on a com- putational task and heralds the advent of a much-anticipated computing paradigm. In the early 1980s, Richard Feynman proposed that a quantum computer would be an e ective tool to solve problems in physics and chemistry, as it is exponentially costly to simulate large quantum systems with classical computers [1]. Realizing Feynman's vision poses signi – cant experimental and theoretical challenges. First, can a quantum system be engineered to perform a computa- tion in a large enough computational (Hilbert) space and with low enough errors to provide a quantum speedup? Second, can we formulate a problem that is hard for a classical computer but easy for a quantum computer? By computing a novel benchmark task on our superconduct- ing qubit processor[2{7], we tackle both questions. Our experiment marks a milestone towards full scale quantum computing: quantum supremacy[8].

In reaching this milestone, we show that quantum speedup is achievable in a real-world system and is not precluded by any hidden physical laws. Quantum supremacy also heralds the era of Noisy Intermediate- Scale Quantum (NISQ) technologies. The benchmark task we demonstrate has an immediate application in generating certi able random numbers[9]; other initial uses for this new computational capability may include optimization optimization [10{12], machine learning[13{ 15], materials science and chemistry [16{18]. However, realizing the full promise of quantum computing (e.g. Shor's algorithm for factoring) still requires technical leaps to engineer fault-tolerant logical qubits[19{23].

To achieve quantum supremacy, we made a number of technical advances which also pave the way towards er- ror correction. We developed fast, high- delity gates that can

be executed simultaneously across a two-dimensional qubit array. We calibrated and benchmarked the pro- cessor at both the component and system level using a powerful new tool: cross-entropy benchmarking (XEB).

Finally, we used component-level delities to accurately predict the performance of the whole system, further showing that quantum information behaves as expected when scaling to large systems. A COMPUTATIONAL TASK TO DEMONSTRATE QUANTUM SUPREMACY To demonstrate quantum supremacy, we compare our quantum processor against state-of-the-art classical com- puters in the task of sampling the output of a pseudo- random quantum circuit[24{26].

Random circuits are a suitable choice for benchmarking since they do not pos- sess structure and therefore allow for limited guarantees of computational hardness[24, 25, 27, 28]. We design the circuits to entangle a set of quantum bits (qubits) by re- peated application of single-qubit and two-qubit logical operations. Sampling the quantum circuit's output pro- duces a set of bitstrings, e.g. f0000101, 1011100, …g. Due to quantum interference, the probability distribution of the bitstrings resembles a speckled intensity pattern produced by light interference in laser scatter, such that some bitstrings are much more likely to occur than oth- ers. Classically computing this probability distribution becomes exponentially more di cult as the number of qubits (width) and number of gate cycles (depth) grows. We verify that the quantum processor is working prop- erly using a method called cross-entropy benchmarking (XEB) [24, 26], which compares how often each bitstring is observed experimentally with its corresponding ideal probability computed via simulation on a classical com- puter. For a given circuit, we collect the measured bit- strings fx igand compute the linear XEB delity [24{ 26, 29], which is the mean of the simulated probabilities of the bitstrings we measured: F XEB = 2 nhP(x i)i i 1 (1) where nis the number of qubits, P(x i) is the probability of bitstring x i computed for the ideal quantum circuit, and the average is over the observed bitstrings. Intu- itively, F XEB is correlated with how often we sample high probability bitstrings. When there are no errors in the quantum circuit, sampling the probability distribution will produce F XEB = 1. On the other hand, sampling from the uniform distribution will give hP(x i)i i = 1=2n and produce F XEB = 0. Values of F XEB between 0 and 2 Qubit Adjustable coupler a b 10 millimeters FIG. 1. The Sycamore processor. a, Layout of processor showing a rectangular array of 54 qubits (gray), each con- nected to its four nearest neighbors with couplers (blue). In- operable qubit is outlined. b, Optical image of the Sycamore chip. 1 correspond to the probability that no error has oc- curred while running the circuit.

The probabilities P(x i) must be obtained from classically simulating the quan- tum circuit, and thus computing F XEB is intractable in the regime of quantum supremacy. However, with certain circuit simpli cations, we can obtain quantitative delity estimates of a fully operating processor running wide and deep quantum circuits. Our goal is to achieve a high enough F XEB for a circuit with su cient width and depth such that the classical computing cost is prohibitively large. This is a di cult task because our logic gates are imperfect and the quan- tum states we intend to create are sensitive to errors.

A single bit or phase ip over the course of the algorithm will completely shu e the speckle pattern and result in close to 0 delity [24, 29].

Therefore, in order to claim quantum supremacy we need a quantum processor that executes the program with su ciently low error rates. BUILDING AND CHARACTERIZING A HIGH-FIDELITY PROCESSOR We designed a quantum processor named \Sycamore" which consists of a two-dimensional array of 54 transmon qubits, where each qubit is tunably coupled to four nearest-neighbors, in a rectangular lattice. The connec- tivity was chosen to be forward compatible with error-correction using the surface code [20].

A key systems- engineering advance of this device is achieving high- delity single- and two-qubit operations, not just in iso- lation but also while performing a realistic computation with simultaneous gate operations on many qubits. We discuss the highlights below; extended details can be found in the supplementary information. In a superconducting circuit, conduction electrons con- dense into a macroscopic quantum state, such that cur- rents and voltages behave quantum mechanically [2, 30]. Our processor uses transmon qubits [6], which can be thought of as nonlinear superconducting resonators at 5 to 7 GHz.

The qubit is encoded as the two lowest quan- tum eigenstates of the resonant circuit. Each transmon has two controls: a microwave drive to excite the qubit, and a magnetic ux control to tune the frequency. Each qubit is connected to a linear resonator used to read out the qubit state [5]. As shown in Fig. 1, each qubit is also connected to its neighboring qubits using a new ad- justable coupler [31, 32]. Our coupler design allows us to quickly tune the qubit-qubit coupling from completely o to 40 MHz. Since one qubit did not function properly the device uses 53 qubits and 86 couplers. The processor is fabricated using aluminum for metal- ization and Josephson junctions, and indium for bump- bonds between two silicon wafers. The chip is wire- bonded to a superconducting circuit board and cooled to below 20 mK in a dilution refrigerator to reduce am- bient thermal energy to well below the qubit energy. The processor is connected through lters and attenu- ators to room-temperature electronics, which synthesize the control signals.

The state of all qubits can be read simultaneously by using a frequency-multiplexing tech- nique[33, 34]. We use two stages of cryogenic ampli ers to boost the signal, which is digitized (8 bits at 1 GS/s) and demultiplexed digitally at room temperature. In to- tal, we orchestrate 277 digital-to-analog converters (14 bits at 1 GS/s) for complete control of the quantum pro- cessor. We execute single-qubit gates by driving 25 ns mi- crowave pulses resonant with the qubit frequency while the qubit-qubit coupling is turned o .

The pulses are shaped to minimize transitions to higher transmon states[35]. Gate performance varies strongly with fre- quency due to two-level-system (TLS) defects[36, 37], stray microwave modes, coupling to control lines and the readout resonator, residual stray coupling between qubits, ux noise, and pulse distortions. We therefore 3 Pauli and measurement errors CDF am, E ted histogr Integra e 1 e 2 e 2c e r a b

Average error Single-qubit (e 1) Two-qubit (e 2) Two-qubit, cycle (e 2c) Readout (e r) Isolated 0.15% 0.36% 0.65% 3.1% Simultaneous 0.16% 0.62% 0.93% 3.8% Simultaneous Pauli error e 1 , e 2 10 -2 10 -3 Isolated FIG. 2. System-wide Pauli and measurement errors. a, Integrated histogram (empirical cumulative distribution function, ECDF) of Pauli errors (black, green, blue) and readout errors (orange), measured on qubits in isolation (dotted lines) and when operating all qubits simultaneously (solid). The median of each distribution occurs at 0.50 on the vertical axis. Average (mean) values are shown below. b, Heatmap showing single- and two-qubit Pauli errors e 1 (crosses) and e 2 (bars) positioned in the layout of the processor. Values shown for all qubits operating simultaneously. optimize the single-qubit operation frequencies to mitigate these error mechanisms. We benchmark single-qubit gate performance by using the XEB protocol described above, reduced to the single- qubit level (n= 1), to measure the probability of an error occurring during a single-qubit gate. On each qubit, we apply a variable number mof randomly selected gates and measure F XEB averaged over many sequences; as m increases, errors accumulate and average F XEB decays.

We model this decay by [1 e 1=(1 1=D2)]m where e 1 is the Pauli error probability. The state (Hilbert) space di- mension term, D= 2n = 2, corrects for the depolarizing model where states with errors partially overlap with the ideal state. This procedure is similar to the more typical technique of randomized benchmarking [21, 38, 39], but supports non-Cli ord gatesets [40] and can separate out decoherence error from coherent control error. We then repeat the experiment with all qubits executing single- qubit gates simultaneously (Fig.2), which shows only a small increase in the error probabilities, demonstrating that our device has low microwave crosstalk. We perform two-qubit iSWAP-like entangling gates by bringing neighboring qubits on resonance and turning on a 20 MHz coupling for 12 ns, which allows the qubits to swap excitations. During this time, the qubits also ex- perience a controlled-phase (CZ) interaction, which orig- inates from the higher levels of the transmon. The two- qubit gate frequency trajectories of each pair of qubits are optimized to mitigate the same error mechanisms consid- ered in optimizing single-qubit operation frequencies. To characterize and benchmark the two-qubit gates, we run two-qubit circuits with mcycles, where each cy- cle contains a randomly chosen single-qubit gate on each of the two qubits followed by a xed two-qubit gate. We learn the parameters of the two-qubit unitary (e.g. the amount of iSWAP and CZ interaction) by using F XEB as a cost function. After this optimization, we extract the per-cycle error e 2c from the decay of F XEB with m, and isolate the two-qubit error e 2 by subtracting the two single-qubit errors e 1. We nd an average e 2 of 0:36%.

Additionally, we repeat the same procedure while simul- taneously running two-qubit circuits for the entire array. After updating the unitary parameters to account for ef- fects such as dispersive shifts and crosstalk, we nd an average e 2 of 0.62%. For the full experiment, we generate quantum circuits using the two-qubit unitaries measured for each pair dur- ing simultaneous operation, rather than a standard gate for all pairs. The typical two-qubit gate is a full iSWAP with 1=6 of a full CZ. In principle, our architecture could generate unitaries with arbitrary iSWAP and CZ inter- actions, but reliably generating a target unitary remains an active area of research. Finally, we benchmark

qubit readout using standard dispersive measurement [41]. Measurement errors aver- aged over the 0 and 1 states are shown in Fig 2a.

We have also measured the error when operating all qubits simul- taneously, by randomly preparing each qubit in the 0 or 1 state and then measuring all qubits for the probability of the correct result. We nd that simultaneous readout incurs only a modest increase in per-qubit measurement errors. Having found the error rates of the individual gates and readout, we can model the delity of a quantum circuit as the product of the probabilities of error-free opera- 4 single-qubit gate: 25 ns qubit XY control two-qubit gate: 12 ns qubit 1 Z control qubit 2 Z control coupler cycle: 1 2 3 4 5 6 m time column row 7 8 A BC D A B D C a b FIG. 3. Control operations for the quantum supremacy circuits. a, Example quantum circuit instance used in our experiment. Every cycle includes a layer each of single- and two-qubit gates. The single-qubit gates are chosen randomly from f p X; p Y; p Wg. The sequence of two-qubit gates are chosen according to a tiling pattern, coupling each qubit sequentially to its four nearest-neighbor qubits.

The couplers are divided into four subsets (ABCD), each of which is executed simultaneously across the entire array corresponding to shaded colors. Here we show an intractable sequence (repeat ABCDCDAB); we also use di erent coupler subsets along with a simpli able sequence (repeat EFGHEFGH, not shown) that can be simulated on a classical computer. b, Waveform of control signals for single- and two-qubit gates. tion of all gates and measurements. Our largest random quantum circuits have 53 qubits, 1113 single-qubit gates, 430 two-qubit gates, and a measurement on each qubit, for which we predict a total delity of 0:2%.

This delity should be resolvable with a few million measurements, since the uncertainty on F XEB is 1= p N s, where N s is the number of samples. Our model assumes that entangling larger and larger systems does not introduce additional error sources beyond the errors we measure at the single- and two-qubit level | in the next section we will see how well this hypothesis holds. FIDELITY ESTIMATION IN THE SUPREMACY REGIME The gate sequence for our pseudo-random quantum circuit generation is shown in Fig.3. One cycle of the algorithm consists of applying single-qubit gates chosen randomly from f p X; p Y; p Wgon all qubits, followed by two-qubit gates on pairs of qubits. The sequences of gates which form the \supremacy circuits" are designed to minimize the circuit depth required to create a highly entangled state, which ensures computational complexity and classical hardness. While we cannot compute F XEB in the supremacy regime, we can estimate it using three variations to re- duce the complexity of the circuits.

In \patch circuits", we remove a slice of two-qubit gates (a small fraction of the total number of two-qubit gates), splitting the cir- cuit into two spatially isolated, non-interacting patches of qubits. We then compute the total delity as the product of the patch delities, each of which can be easily calcu- lated. In \elided circuits", we remove only a fraction of the initial two-qubit gates along the slice, allowing for entanglement between patches, which more closely mim- ics the full experiment while still maintaining simulation feasibility. Finally, we can also run full \veri cation cir- cuits" with the same

gate counts as our supremacy cir- cuits, but with a di erent pattern for the sequence of two- qubit gates which is much easier to simulate classically [29]. Comparison between these variations allows track- ing of the system delity as we approach the supremacy regime. We rst check that the patch and elided versions of the veri cation circuits produce the same delity as the full veri cation circuits up to 53 qubits, as shown in Fig.4a. For each data point, we typically collect N s = 5 106 total samples over ten circuit instances, where instances di er only in the choices of single-qubit gates in each cycle.

We also show predicted F XEB values computed by multiplying the no-error probabilities of single- and two-qubit gates and measurement [29]. Patch, elided, and predicted delities all show good agreement with the delities of the corresponding full circuits, despite the vast di erences in computational complexity and en- tanglement. This gives us con dence that elided circuits can be used to accurately estimate the delity of more complex circuits. We proceed now to benchmark our most computa- tionally di cult circuits. In Fig.4b, we show the mea- sured F XEB for 53-qubit patch and elided versions of the full supremacy circuits with increasing depth.

For the largest circuit with 53 qubits and 20 cycles, we collected N s = 30 106 samples over 10 circuit instances, obtaining F XEB = (2:24 0:21) 10 3 for the elided circuits. With 5˙con dence, we assert that the average delity of run- ning these circuits on the quantum processor is greater than at least 0.1%. The full data for Fig.4b should have similar delities, but are only archived since the simula- tion times (red numbers) take too long. It is thus in the quantum supremacy regime. 5 number of qubits, n number of cycles, m n = 53 qubits a Classically veriable b Supremacy regime idelity, XEB F

XEB m = 14 cycles Prediction from gate and measurement errors Full circuit Elided circuit Patch circuit Prediction Patch E F G H A B C D C D A B Elided (±5 error bars) 10 millennia 100 years 600 years 4 years 4 years 2 weeks 1 week 2 hour sC la ic mp ng @ Sycamore 5 hours Classical verication Sycamore sampling (N s = 1M): 200 seconds 10 15 20 25 30 35 40 45 50 55 12 14 16 18 20 10 -3 10 -2 10 -1 10 0 FIG. 4. Demonstrating quantum supremacy. a, Veri cation of benchmarking methods. F XEB values for patch, elided, and full veri cation circuits are calculated from measured bitstrings and the corresponding probabilities predicted by classical simulation. Here, the two-qubit gates are applied in a simpli able tiling and sequence such that the full circuits can be simulated out to n= 53;m= 14 in a reasonable amount of time. Each data point is an average over 10 distinct quantum circuit instances that di er in their single-qubit gates (for n= 39;42;43 only 2 instances were simulated). For each n, each instance is sampled with N s between 0:5M and 2:5M.

The black line shows predicted F XEB based on single- and two-qubit gate and measurement errors. The close correspondence between all four curves, despite their vast di erences in complexity, justi es the use of elided circuits to estimate delity in the supremacy regime. b, Estimating F XEB in the quantum supremacy regime. Here, the two-qubit gates are applied in a non-simpli able tiling and sequence for which it is much harder to simulate. For the largest elided data (n= 53, m= 20, total N s = 30M), we nd an

average F XEB >0.1% with 5˙con dence, where ˙includes both systematic and statistical uncertainties.

The corresponding full circuit data, not simulated but archived, is expected to show similarly signi cant delity. For m= 20, obtaining 1M samples on the quantum processor takes 200 seconds, while an equal delity classical sampling would take 10,000 years on 1M cores, and verifying the delity would take millions of years. DETERMINING THE CLASSICAL COMPUTATIONAL COST We simulate the quantum circuits used in the exper- iment on classical computers for two purposes: verify- ing our quantum processor and benchmarking methods by computing F XEB where possible using simpli able circuits (Fig.4a), and estimating F XEB as well as the classical cost of sampling our hardest circuits (Fig.4b).

Up to 43 qubits, we use a Schr• odinger algorithm (SA) which simulates the evolution of the full quantum state; the Ju• lich supercomputer(100k cores, 250TB) runs the largest cases. Above this size, there is not enough RAM to store the quantum state [42]. For larger qubit num- bers, we use a hybrid Schr• odinger-Feynman algorithm (SFA)[43] running on Google data centers to compute the amplitudes of individual bitstrings. This algorithm breaks the circuit up into two patches of qubits and e - ciently simulates each patch using a Schr• odinger method, before connecting them using an approach reminiscent of the Feynman path-integral.

While it is more memory- e cient, SFA becomes exponentially more computation- ally expensive with increasing circuit depth due to the exponential growth of paths with the number of gates connecting the patches. To estimate the classical computational cost of the supremacy circuits (gray numbers, Fig.4b), we ran por- tions of the quantum circuit simulation on both the Sum- mit supercomputer as well as on Google clusters and ex- trapolated to the full cost. In this extrapolation, we account for the computational cost scaling with F XEB, e.g. the 0.1% delity decreases the cost by 1000[43, 44]. On the Summit supercomputer, which is currently the most powerful in the world, we used a method inspired by Feynman path-integrals that is most e cient at low depth[44{47].

At m= 20 the tensors do not reasonably t in node memory, so we can only measure runtimes up to m= 14, for which we estimate that sampling 3M bitstrings with 1% delity would require 1 year. 6 On Google Cloud servers, we estimate that perform- ing the same task for m= 20 with 0:1% delity using the SFA algorithm would cost 50 trillion core-hours and consume 1 petawatt hour of energy. To put this in per- spective, it took 600 seconds to sample the circuit on the quantum processor 3 million times, where sampling time is limited by control hardware communications; in fact, the net quantum processor time is only about 30 seconds. The bitstring samples from this largest circuit are archived online. One may wonder to what extent algorithmic innova- tion can enhance classical simulations. Our assumption, based on insights from complexity theory, is that the cost of this algorithmic task is exponential in nas well as m. Indeed, simulation methods have improved steadily over the past few years[42{50].

We expect that lower simula- tion costs than reported here will eventually be achieved, but we also expect they will be consistently outpaced by hardware improvements on larger quantum processors. VERIFYING THE DIGITAL ERROR MODEL A key assumption underlying the theory of quantum error correction is that quantum state errors may be con- sidered digitized and localized [38, 51]. Under such a dig- ital model, all errors in the evolving quantum state may be characterized by a set of localized Pauli errors (bit- and/or phase- ips) interspersed into the circuit. Since continuous amplitudes are fundamental to quantum me- chanics, it needs to be tested whether errors in a quantum system could be treated as discrete and probabilistic. In- deed, our experimental observations support the validity of this model for our processor. Our system delity is well predicted by a simple model in which the individ- ually characterized delities of each gate are multiplied together (Fig

4). To be successfully described by a digitized error model, a system should be low in correlated errors. We achieve this in our experiment by choosing circuits that ran- domize and decorrelate errors, by optimizing control to minimize systematic errors and leakage, and by design- ing gates that operate much faster than correlated noise sources, such as 1=f ux noise [37]. Demonstrating a pre- dictive uncorrelated error model up to a Hilbert space of size 253 shows that we can build a system where quantum resources, such as entanglement, are not prohibitively fragile. WHAT DOES THE FUTURE HOLD?

Quantum processors based on superconducting qubits can now perform computations in a Hilbert space of di- mension 253 ˇ9 1015, beyond the reach of the fastest classical supercomputers available today. To our knowl- edge, this experiment marks the rst computation that can only be performed on a quantum processor. Quan- tum processors have thus reached the regime of quantum supremacy. We expect their computational power will continue to grow at a double exponential rate: the clas- sical cost of simulating a quantum circuit increases expo- nentially with computational volume, and hardware im- provements will likely follow a quantum-processor equiv- alent of Moore’s law [52, 53], doubling this computational volume every few years. To sustain the double exponen- tial growth rate and to eventually o er the computational volume needed to run well-known quantum algorithms, such as the Shor or Grover algorithms [19, 54], the engi- neering of quantum error correction will have to become a focus of attention. The \Extended Church-Turing Thesis” formulated by Bernstein and Vazirani [55] asserts that any \reasonable” model of computation can be e ciently simulated by a Turing machine. Our experiment suggests that a model of computation may now be available that violates this assertion.

We have performed random quantum circuit sampling in polynomial time with a physically realized quantum processor (with su ciently low error rates), yet no e cient method is known to exist for classical comput- ing machinery. As a result of these developments, quan- tum computing is transitioning from a research topic to a technology that unlocks new computational capabilities. We are only one creative algorithm away from valuable near-term applications. Acknowledgments We are grateful to Eric Schmidt, Sergey Brin, Je Dean, and Jay Yagnik for their executive sponsorship

of the Google AI Quantum team, and for their continued engagement and support. We thank Peter Norvig for reviewing a draft of the manuscript, and Sergey Knysh for useful discussions.

We thank Kevin Kissel, Joey Raso, Davinci Yonge-Mallo, Orion Martin, and Niranjan Sridhar for their help with simulations. We thank Gina Bortoli and Lily Laws for keeping our team organized. This research used resources from the Oak Ridge Leadership Computing Facility, which is a DOE O ce of Science User Facility supported un- der Contract DE-AC05-00OR22725.

A portion of this work was performed in the UCSB Nanofabrication Facility, an open access laboratory. Author contributions The Google AI Quantum team conceived of the experiment. The applications and algorithms team provided the theoretical foundation and the speci cs of the algorithm. The hardware team carried out the experiment and collected the data. The data analysis was done jointly with outside collaborators.

[1]Feynman, R. P. Simulating physics with computers. Int. J. Theor. Phys. 21, 467{488 (1982).

[2]Devoret, M. H., Martinis, J. M. & Clarke, J. Mea- surements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased josephson junction. Phys. Rev. Lett 55, 1908 (1985).

[3]Nakamura, Y., Chen, C. D. & Tsai, J. S. Spectroscopy of energy-level splitting between two macroscopic quan- tum states of charge coherently superposed by josephson coupling. Phys. Rev. Lett. 79, 2328 (1997).

[4]Mooij, J. et al. Josephson persistent-current qubit. Sci- ence 285, 1036 (1999).

[5]Wallra , A. et al. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrody- namics. Nature 431, 162 (2004).

[6]Koch, J. et al. Charge-insensitive qubit design derived from the cooper pair box. Phys. Rev. A 76, 042319 (2007).

[7]You, J. Q. & Nori, F. Atomic physics and quantum optics using superconducting circuits. Nature 474, 589 (2011).

[8]Preskill, J. Quantum computing and the entanglement frontier. Rapporteur talk at the 25th Solvay Conference on Physics, Brussels (2012).

[9]Aaronson, S. Certi ed randomness from quantum supremacy. In preparation .

[10]Hastings, M. B. Classical and Quantum Bounded Depth Approximation Algorithms. arXiv e-prints arXiv:1905.07047 (2019). 1905.07047.

[11]Kechedzhi, K. et al. E cient population transfer via non- ergodic extended states in quantum spin glass. arXiv e-prints arXiv:1807.04792 (2018). 1807.04792.

[12]Somma, R. D., Boixo, S., Barnum, H. & Knill, E. Quan- tum simulations of classical annealing processes. Phys. Rev. Lett. letters 101, 130504 (2008).

[13]McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural net- work training landscapes. Nat. Comm. 9, 4812 (2018).

[14]Cong, I., Choi, S. & Lukin, M. D. Quantum convolutional neural networks. arXiv:1810.03787 (2018).

[15]Bravyi, S., Gosset, D. & K• onig, R. Quantum advantage with shallow circuits. Science 362, 308{311 (2018).

[16]Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head- Gordon, M. Simulated quantum computation of molecu- lar energies. Science 309, 1704{1707 (2005).

[17]Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. Nat. Commun. 5, 4213 (2014). [18]Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. Phys. Rev. X 8, 031022 (2018).

[19]Shor, P. W. Algorithms for quantum computation: dis- crete logarithms and factoring proceedings. Proceedings 35th Annual Symposium on Foundations of Computer Science (1994).

[20]Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cle- land, A. N. Surface codes: Towards practical large-scale quantum computation. Phys. Rev. A 86, 032324 (2012).

[21]Barends, R. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. Nature 508, 500{503 (2014).

[22]C orcoles, A. D. et al. Demonstration of a quantum error detection code using a square lattice of four supercon- ducting qubits. Nat. Commun. 6, 6979 (2015).

[23]Ofek, N. et al. Extending the lifetime of a quantum bit with error correction in superconducting circuits. Nature 536, 441 (2016).

[24]Boixo, S. et al. Characterizing quantum supremacy in near-term devices. Nat. Phys. 14, 595 (2018).

[25]Aaronson, S. & Chen, L. Complexity-theoretic founda- tions of quantum supremacy experiments. In 32nd Com- putational Complexity Conference (CCC 2017) (2017).

[26]Neill, C. et al. A blueprint for demonstrating quantum supremacy with superconducting qubits. Science 360, 195{199 (2018).

[27]Bremner, M. J., Montanaro, A. & Shepherd, D. J. Average-case complexity versus approximate simulation of commuting quantum computations. Phys. Rev. Lett. 117, 080501 (2016).

[28]Bouland, A., Fe erman, B., Nirkhe, C. & Vazi- rani, U. Quantum supremacy and the com- plexity of random circuit sampling. Preprint at **arxiv.org/...** (2018).

[29]See supplementary information .

[30]Vool, U. & Devoret, M. Introduction to quantum electro- magnetic circuits. Int. J. Circ. Theor. Appl. 45, 897{934 (2017). 8

[31]Chen, Y. et al. Qubit architecture with high coherence and fast tunable coupling circuits. Phys. Rev. Lett. 113, 220502 (2014).

[32]Yan, F. et al. A tunable coupling scheme for implement- ing high- delity two-qubit gates. Phys. Rev. Applied 10, 054062 (2018).

[33]Schuster, D. I. et al. Resolving photon number states in a superconducting circuit. Nature 445, 515 (2007).

[34]Je rey, E. et al. Fast accurate state measurement with superconducting qubits. Phys. Rev. Lett. 112, 190504 (2014).

[35]Chen, Z. et al. Measuring and suppressing quantum state leakage in a superconducting qubit. Phys. Rev. Lett. 116, 020501 (2016).

[36]Klimov, P. V. et al. Fluctuations of energy-relaxation times in superconducting qubits. Phys. Rev. Lett. 121, 090502 (2018).

[37]Yan, F. et al. The ux qubit revisited to enhance coher- ence and reproducibility. Nat. Commun. 7, 12964 (2016).

[38]Knill, E. et al. Randomized benchmarking of quantum gates. Phys. Rev. A 77, 012307 (2008).

[39]Magesan, E., Gambetta, J. M. & Emerson, J. Scalable and robust randomized benchmarking of quantum pro- cesses. Phys. Rev. Lett. 106, 180504 (2011).

[40]Cross, A. W., Magesan, E., Bishop, L. S., Smolin, J. A. & Gambetta, J. M. Scalable randomised benchmarking of non-cli ord gates. NPJ Quantum Information 2, 16012 (2016).

[41]Wallra , A. et al. Approaching unit visibility for control of a superconducting qubit with dispersive readout. Phys. Rev. Lett. 95, 060501 (2005).

[42]De Raedt, H. et al. Massively parallel quantum computer simulator, eleven years later. Comput. Phys. Commun. 237, 47 { 61 (2019).

[43]Markov, I. L., Fatima, A., Isakov, S. V. & Boixo, S. Quantum supremacy is both closer and farther than it appears. Preprint at **arxiv.org/...** (2018).

[44]Villalonga, B. et al. A exible high-performance sim- ulator for the veri cation and benchmarking of quan- tum circuits implemented on real hardware. Preprint at **arxiv.org/...** (2018).

[45]Boixo, S., Isakov, S. V., Smelyanskiy, V. N. & Neven, H. Simulation of low-depth quantum circuits as complex undirected graphical models. Preprint at **arxiv.org/...** (2017).

[46]Chen, J., Zhang, F., Huang, C., Newman, M. & Shi, Y. Classical simulation of intermediate-size quantum circuits. Preprint at **arxiv.org/...** (2018).

[47]Villalonga, B. et al. Establishing the quantum supremacy frontier with a 281 p op/s simulation. Preprint at **arxiv.org/...** (2019).

[48]Pednault, E. et al. Breaking the 49-qubit barrier in the simulation of quantum circuits. Preprint at **arxiv.org/...** (2017).

[49]Chen, Z. Y. et al. 64-qubit quantum circuit simulation. Sci. Bull. 63, 964{971 (2018).

[50]Chen, M.-C. et al. Quantum teleportation-inspired al- gorithm for sampling large random quantum circuits. Preprint at **arxiv.org/...** (2019).

[51]Shor, P. W. Scheme for reducing decoherence in quan- tum computer memory. Phys. Rev. A 52, R2493{R2496 (1995).

[52]Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: An outlook. Science 339, 1169{1174 (2013).

[53]Mohseni, M. et al. Commercialize quantum technologies in ve years. Nature 543, 171 (2017).

[54]Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. letters 79, 325 (1997).

[55]Bernstein, E. & Vazirani, U. Quantum complexity the- ory. Proc. 25th Annual ACM Symposium on Theory of Computing (1993).