Lecture 25, 10 Dec 2020

quantumcryptopocalypse.com/quantum-moores-law/

Quantum Moore's Law

While small now, quantum computers are expected to experience exponential growth in size, speed, and accuracy over the coming years. The following plots chart this progress, updated and extrapolated from past performance, and prognosticate a time when quantum computers will be powerful enough to crack digital signatures.

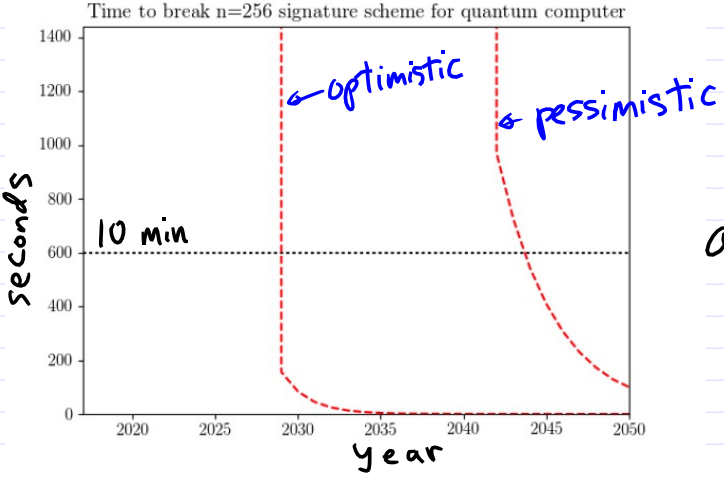

see arXiv: 1710.10377 for technical details

The last plot shows the predicted Shor time, as a function of development year, for a quantum computer to break the Elliptic Curve Digital Signature Algorithm based on the secp256k1 curve using Shor's discrete log algorithm (see Roetteler et al. 2017). The calculation includes overheads for a superconducting qubit quantum computer using surface code based quantum error correction. The left dashed line assumes optimistic predications for qubit number, gate frequency, and infidelities above, and the right dashed line assumes the pessimistic predictions. The horizontal grey dashed line is the average time to verify a transaction on the Bitcoin network, which would be compromised by a quantum attack on digital signatures.
For details about the assumptions of the model and how the overheads were calculated see: D. Aggarwal, G.K. Brennen, T. Lee, M. Santha, and M. Tomamichel, "Quantum attacks on Bitcoin, and how to protect against them," Ledger, [S.l.], v. 3, oct. (2018).

Quantum "advantage" -- "wave function sampling"

described in 1608.00263

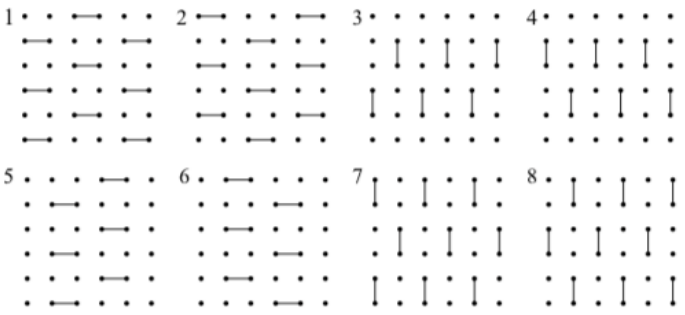

FIG. 6. Layouts of CZ gates in a 6×6 qubit lattice. It is currently not possible to perform two CZ gates simultaneously in two neighboring superconducting qubits [33, 34, 49, 52]. We iterate over these arrangements sequentially, from 1 to 8.

and really "hard":

Quantum Physics

[Submitted on 12 Mar 2018]

Quantum Supremacy and the Complexity of Random Circuit Sampling

Adam Bouland, Bill Fefferman, Chinmay Nirkhe, Umesh Vazirani

A critical milestone on the path to useful quantum computers is quantum supremacy – a demonstration of a quantum computation that is prohibitively hard for classical computers. A leading near-term candidate, put forth by the Google/UCSB team, is sampling from the probability distributions of randomly chosen quantum circuits, which we call Random Circuit Sampling (RCS).

In this paper we study both the hardness and verification of RCS. While RCS was defined with experimental realization in mind, we show complexity theoretic evidence of hardness that is on par with the strongest theoretical proposals for supremacy. Specifically, we show that RCS satisfies an average-case hardness condition – computing output probabilities of typical quantum circuits is as hard as computing them in the worst-case, and therefore #P-hard. Our reduction exploits the polynomial structure in the output amplitudes of random quantum circuits, enabled by the Feynman path integral. In addition, it follows from known results that RCS satisfies an anti-concentration property, making it the first supremacy proposal with both average-case hardness and anti-concentration.

data from

github.com/sboixo/GRCS

(my visualizations)

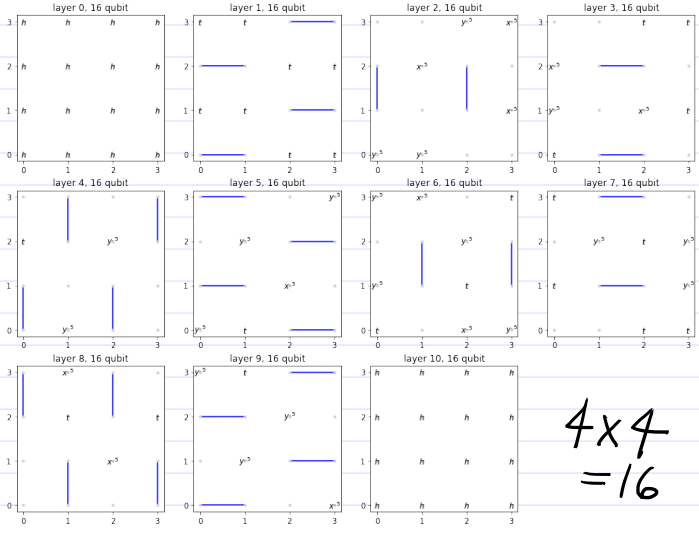

4x4 = 16

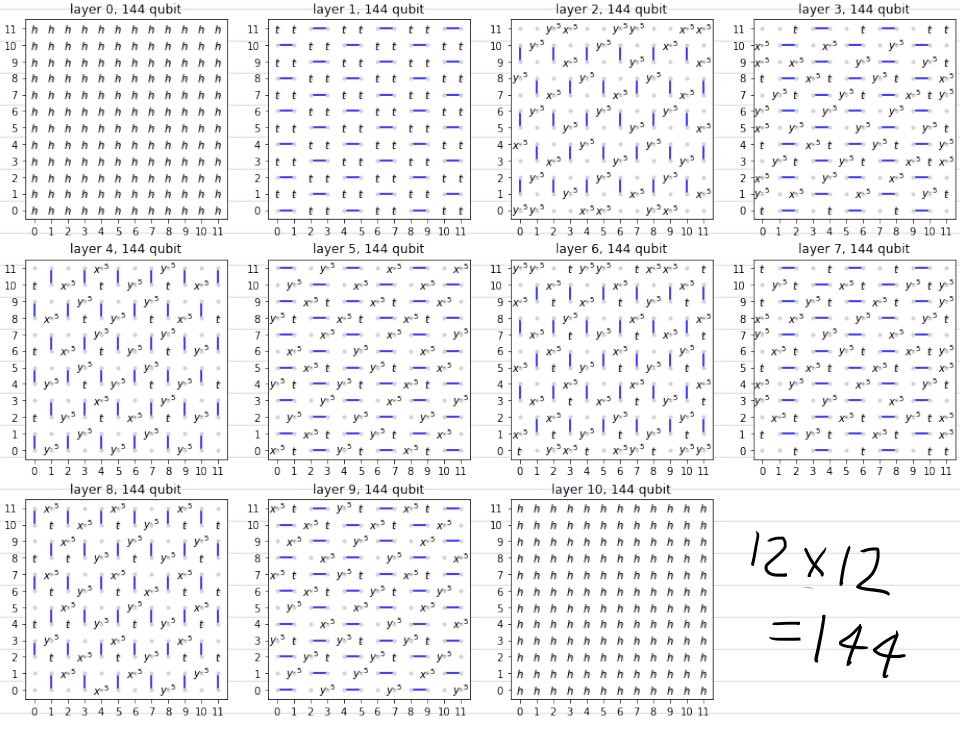

12x12 = 144

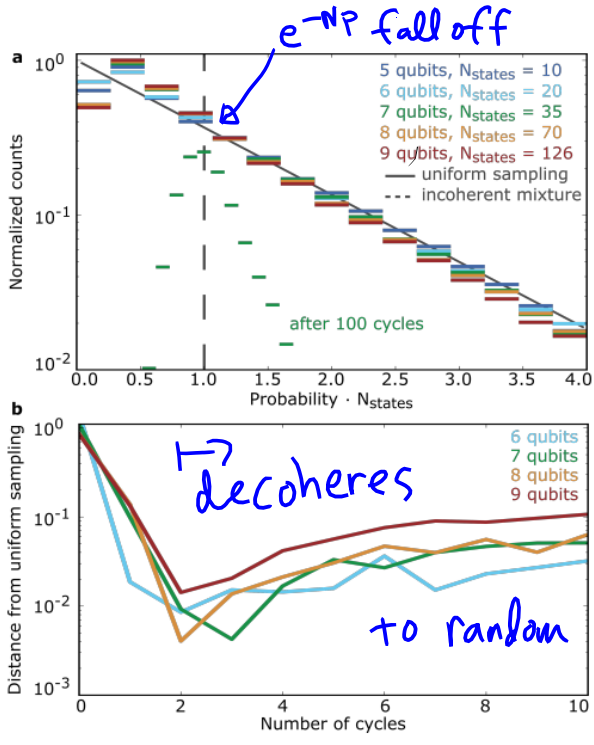

Figure 3. **Complexity: uniform sampling of an exponentially growing state-space. a** Histograms of the raw probabilities (see Fig. 1b) for 5 to 9 qubit experiments, after five cycles of evolution. Before making the histogram, probabilities are weighted by the number of states in the Hilbert-space, this places all of the curves onto a universal axis. The width of the bars represents the size of the bins used to construct the histogram. The data is taken from over 29.7 million experiments. For dynamics which uniformly explore all states, this histogram decays exponentially; an exponential decay is shown as a solid line for comparison. For contrast, we plot a histogram of the probabilities for 7 qubits after 100 cycles. Here, decoherence dominates and we observe a tall narrow peak around 1. **b** In order to measure convergence of the measured histogram to an exponential distribution, we compute their distance as a function of the number of cycles. Distance is measured using the KL-divergence (see Eq. 2). We find that a maximum overlap occurs after just two cycles, following which decoherence increases their distance.

e^{-Np} falloff

(semi-log)

1709.06678

5-9 qubits

Probability(Np) $\propto e^{-Np}$

"Porter Thomas" distribution, probability distribution over probabilities p of states (N = #states)

decoheres

to random

bottom line: large # of low probability states, small number of high probability states, with characteristic e^{-Np} falloff.

(Nature Article)

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5]*

The promise of quantum computers is that certain computational tasks might be executed exponentially faster on a quantum processor than on a classical processor[1]. A fundamental challenge is to build a high-fidelity processor capable of running quantum algorithms in an exponentially large computational space. Here we report the use of a processor with programmable superconducting qubits[2–7] to create quantum states on 53 qubits, corresponding to a computational state-space of dimension 2^{53} (about 10^{16}). Measurements from repeated experiments sample the resulting probability distribution, which we verify using classical simulations. Our Sycamore processor takes about 200 seconds to sample one instance of a quantum circuit a million times—our benchmarks currently indicate that the equivalent task for a state-of-the-art classical supercomputer would take approximately 10,000 years. This dramatic increase in speed compared to all known classical algorithms is an experimental realization of quantum supremacy[8–14] for this specific computational task, heralding a much-anticipated computing paradigm.

and arXiv: 1910.11333

1 broken, from 9×6 grid, so 54 − 1 = 53 qubits

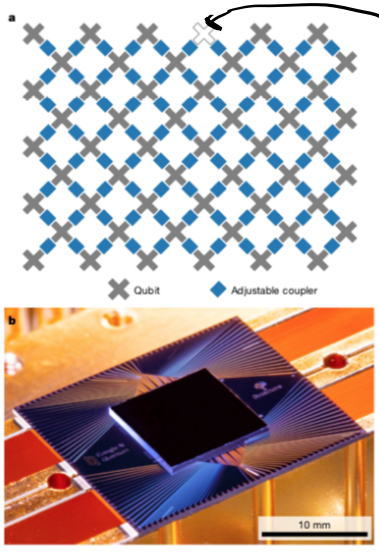

Fig. 1 | The Sycamore processor. a, Layout of processor, showing a rectangular array of 54 qubits (grey), each connected to its four nearest neighbours with couplers (blue). The inoperable qubit is outlined. **b**, Photograph of the Sycamore chip.

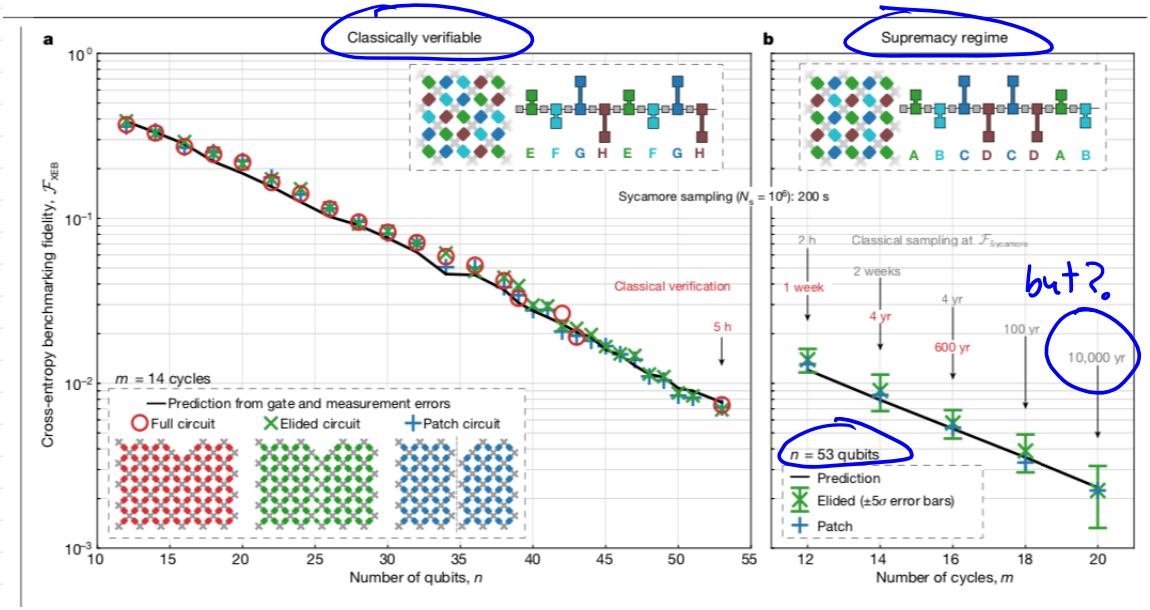

LC circuit = harmonic oscillator

Josephson junction ⇒ anharmonic oscillator, so $\omega_{12} \neq \omega_{01}$

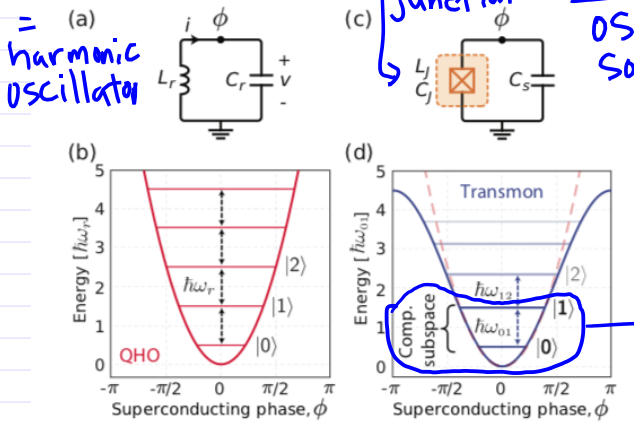

1904.06560

2-state quantum system = qubit

FIG. 1. **(a)** Circuit for a parallel LC-oscillator (quantum harmonic oscillator, QHO), with inductance L in parallel with capacitance, C. The superconducting phase on the island is denoted ϕ, referencing ground as zero. **(b)** Energy potential for the QHO, where energy levels are equidistantly spaced $\hbar\omega_r$ apart. **(c)** Josephson qubit circuit, where the nonlinear inductance L_J (represented with the Josephson-subcircuit in the dashed orange box) is shunted by a capacitance, C_s. **(d)** The Josephson inductance reshapes the quadratic energy potential (dashed red) into sinusoidal (solid blue), which yields non-equidistant energy levels. This allows us to isolate the two lowest energy levels $|0\rangle$ and $|1\rangle$, forming a computational subspace with an energy separation $\hbar\omega_{01}$, which is different than $\hbar\omega_{12}$.

Next time: how to Construct CNOT + complexity Z00

Other qubit possiblities:

NMR

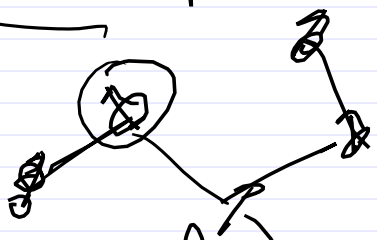

(doesn't scale)

Photons (hard to hold onto, can't reprogram)

Trapped ions

use magnetic interactions and vibrational modes, very long coherence times ~1sec, but scalable to 1000's or millions?

⇒ currently superconducting transmon qubits are best bet