The power of random quantum circuits

Bill Fefferman

(University of Chicago)

Based on "**On the complexity and verification of quantum random circuit sampling**" with A. Bouland, C. Nirkhe, U. Vazirani (Nature Physics, arXiv: 1803.04402)

"**Noise and the frontier of quantum supremacy**" with A. Bouland, Z. Landau, Y. Liu (arXiv: 2102.01738)

And **Efficient classical simulation of noisy random quantum circuits in one dimension** with K. Noh, L. Jiang (Quantum, vol. 4, arXiv:2003.13163)

Invited talk, APS March Meeting, March 16, 2021

The first “Quantum supremacy” claims have now been made…

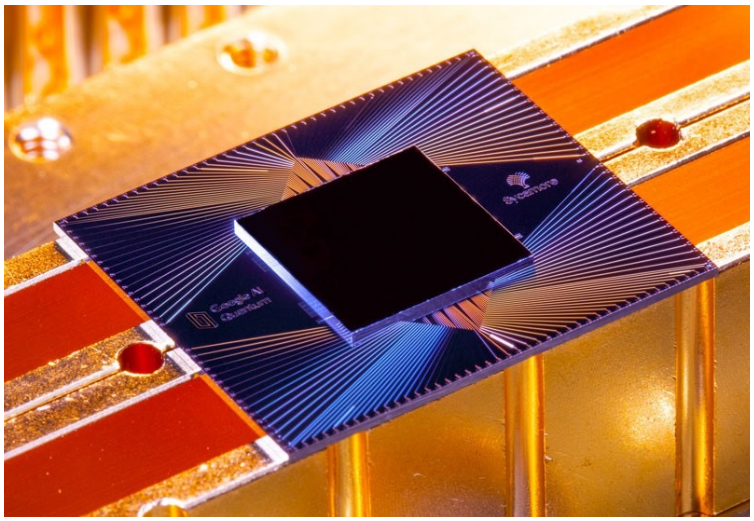

Google “Sycamore” in late 2019

USTC “Jiuzhang” in late 2020

Today: latest evidence that these random quantum circuit experiments are solving hard problems for classical computers

Why random circuits?

- Experimentally feasible
  - Hardness at comparatively low depth and system size
- Advantages for verification/benchmarking
  - e.g., Output distribution of Google's random circuits have "Porter-Thomas" property
    - For any outcome x, $\Pr_C\left[|\langle x|C|0^n\rangle|^2 = \frac{q}{N}\right] \sim e^{-q}$
  - We can use this property to calculate the ideal score of a random circuit on benchmarking tests (e.g., to understand the ideal "cross-entropy" score)

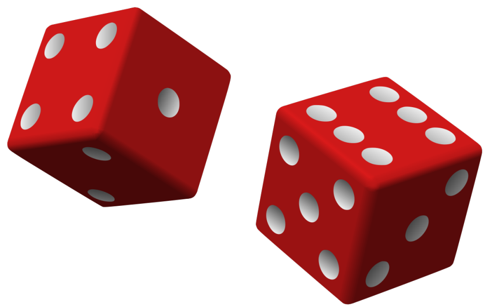

Agenda

1. Summarize the latest theoretical hardness results for *noiseless* random quantum circuits [BFNV'19][BFLL'21]
2. Show that these hardness results extend to certain models of *noisy* quantum experiments [BFLL'21]
3. Finally, give new heuristic classical algorithms for simulating noisy random circuits in 1D [NJF'20]

Part 1: Theoretical hardness results for *noiseless* random quantum circuits [BFNV'19][BFLL'21]

Random circuit sampling

- Current q.s proposals solve "random quantum circuit sampling"
  - i.e., the hard problem is to sample from the output distribution of a randomly chosen quantum circuit
- **Theory goal:** prove impossibility of an efficient *"classical Sampler"* algorithm that:
  - takes as input a random circuit C with output distribution D_C over $\{0,1\}^n$
  - outputs a sample from *any* distribution X so that:
    - $|X - D_C|_{TV} \leq \epsilon$ with high probability over choice of circuit C

All distributions over $\{0,1\}^n$

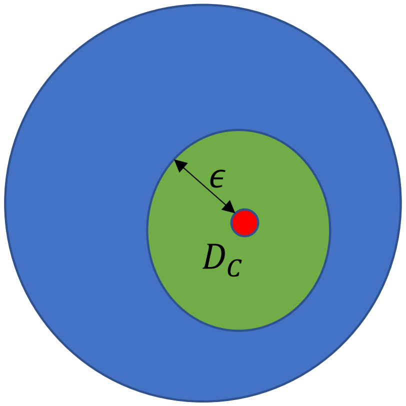

Proof first step: from *sampling* to *computing*

- Recall, our goal is to prove there ***does not*** exist a "classical Sampler" algorithm (under standard complexity theory assumptions)
- By well-known reductions [Stockmeyer '85], [Aaronson & Arkhipov '11] it suffices to prove that ***estimating*** the output probability of a ***random*** quantum circuit is **#P**-hard

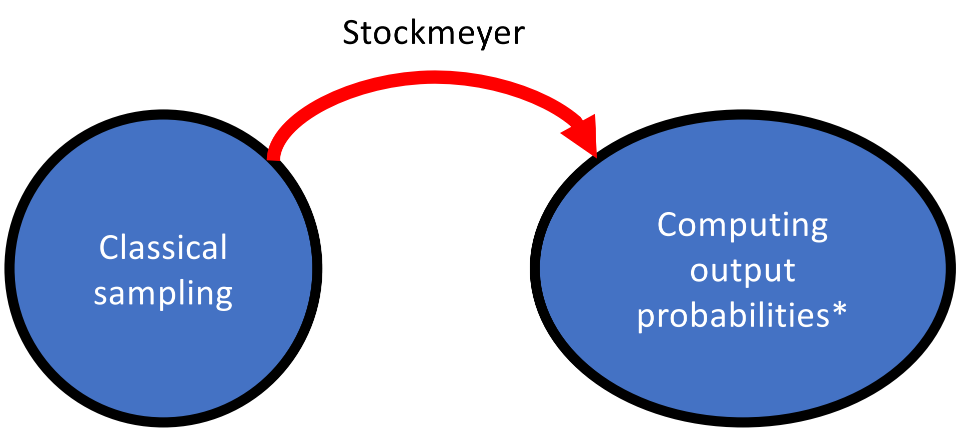

Formal statement of q. supremacy conjecture

- **Definition:** Let the "output probability", $p_0(C) = |\langle 0^n | C | 0^n \rangle|^2$
- Then consider the $\delta - \boldsymbol{Random\ Circuit\ Estimation}$ problem:

> Given as input circuit C, output q so that $|q - p_0(C)| \leq \boldsymbol{\delta}$ with probability 2/3 over C

- To prove goal, it suffices to show that the $\delta = O\left(\frac{1}{2^n}\right)$ problem is **#P**-hard
- **Our results** (as well as work by Movassagh '20 & Kondo et. al. '21): what we know is **#P**-hard for C on n qubits, size $m = O(n \cdot d)$

Hardness conjecture for BosonSampling

- In the case of BosonSampling, similar arguments take us “even closer” to the goal!
- With respect to BosonSampling with n photons, n^2 modes:

- So we’re only off by a factor of 6 in the exponent!
- But… we have hit a barrier, having to do with *noise resilience* of these techniques (more on this in a moment…)

Proof techniques

- The problem reduces to polynomial extrapolation on faulty data points [AA'11, BFNV'19, Mov'20]
  - i.e., Given $O(d^2)$ faulty evaluation points $\{(x_i, y_i)\}$ to a polynomial p(x) of degree d where:
    1. x_i are equally spaced in the interval $[0, \Delta \ll 1]$
    2. And we know ***at least*** *2/3* of y_i are δ-***close*** to $\mathrm{p}(x_i)$
  - Compute an estimate to $\mathrm{p}(1)$
- As compared with prior methods our new method [BFLL'21]:
  - tolerates more errors
  - reduces the extrapolation error
  - greatly simplified proof

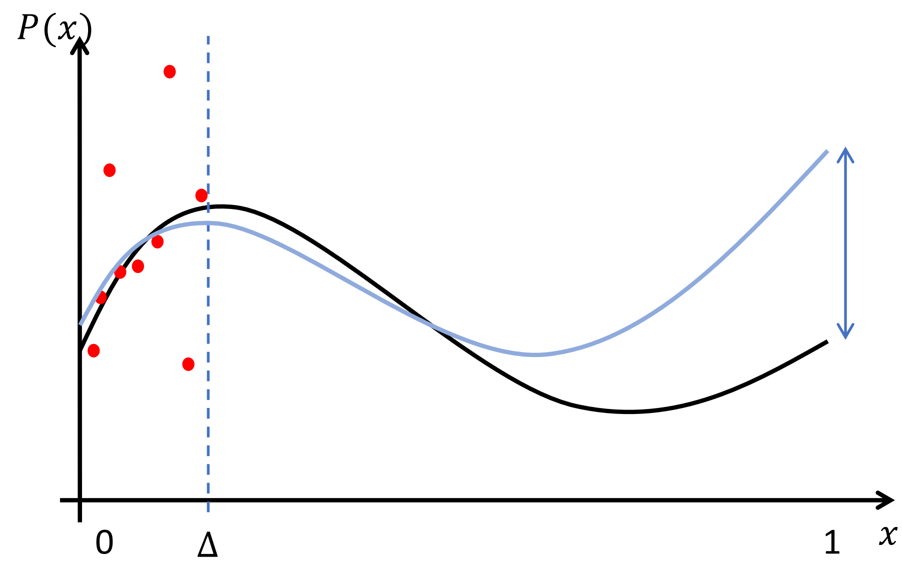

Part 2: Hardness arguments for *noisy* experiments [BFLL'21]

Hardness of noisy random circuits [BFLL'21]

- Without error mitigation noise eventually overwhelms
  - e.g., Google's RCS experiment ~0.2% fidelity and 99.8% noise
- How can we model this theoretically for random circuits?
  - Each random gate C_i is followed by two qubit depolarizing noise channel:
  - $\mathcal{E}_i = (1-\gamma)\rho + \frac{\gamma}{15}\sum_{\alpha,\beta \in \mathcal{P}\times\mathcal{P}-(I,I)}(\sigma_\alpha \otimes \sigma_\beta)\rho(\sigma_\alpha \otimes \sigma_\beta)$
- That is, we can think about choosing a noisy random circuit by:
  - First pick ideal circuit $C = C_m C_{m-1} \ldots C_1$ from the random circuit distribution
  - Then environment chooses operators N, from a distribution $\mathcal{N}$ (i.e., via $\mathcal{E}_i$)
  - We get a sample from output distribution of $N \cdot C$ without learning the noise operators

The same arguments work for the noisy case!

- By linearity, can write the output probability of the noisy circuit as:
  - $E_{N\sim\mathcal{N}}[|\langle 0^n|N\cdot C|0^n\rangle|^2] = E_{N\sim\mathcal{N}}[p_0(N\cdot C)]$
- The previous arguments can be easily extended to give a worst-to-average case reduction for estimating this noisy output probability to precision $2^{-O(m\log m)}$
- [Fujii '16] has shown that this quantity is also hard to compute in the worst-case if noise rate, γ, is a sufficiently small constant

But there's also a (trivial) algorithm here!

- We've now established that computing output probabilities of noisy random circuits of size $m = n \cdot d$ is **hard** to within precision $2^{-O(m \log m)}$
- **Issue:** uncorrected depolarization noise causes output distribution to rapidly converge to uniform as system size grows
  - And it's clearly not hard to output a probability from the uniform distribution!
- How fast is this convergence?
  - **Google's conjecture**: for random circuits $2^{-O(m)}$ [e.g., Boixo, Smelyansky, Neven '17]!
- If we believe this too, then our result is *essentially tight* in this setting!
- Moreover, this presents a **barrier** for the *noiseless* setting
  - i.e., if we want to improve our results we need to find proof techniques that *do not prove hardness in the presence of noise*!

Part 3: Simulation algorithms for noisy quantum circuits [NJF ’20]

New easiness results

- There are many simulation algorithms for restricted classes of random quantum circuits (e.g., [Napp et. al. ’20], [Pan & Zhang’21])
- **Our focus:** 1D random circuits with Haar random two-qubit gates and ***local depolarizing noise***

Numerical results for noisy 1D RCS [NJF'20]

- We consider the "*MPO entanglement entropy*" of the output mixed state
  - A measure of quantum correlations between two disjoint subsystems of qubits $[1, \dots, \ell], [\ell + 1, \dots, n]$
  - Reduces to standard entanglement entropy in case of pure states
- *Motivation for this quantity:* determines the cost of classical simulation
  - Can compute the output probability in time $\sim 2^{S_{max-MPO-EE}(\rho)}$
    - Because "*Maximum MPO entanglement entropy*" can be used to bound the required bond dimension, χ, needed to accurately describe a mixed state
    - Running time is $poly(n, d, \chi)$ and so exponential in $S_{max-MPO-EE}(\rho)$

Growth of MPO EE with depth [NJF’20]

- Each plot has different fixed two-qubit noise rate p
- For each system size $n = 4 \dots 18$ we compute the *Max MPO Entanglement Entropy* measure, averaged over N_s =24 different random circuits

1. We see that for each noise rate, there’s a peak depth
2. Moreover at this peak depth, after sufficiently large system size, adding more qubits doesn’t change the *Max MPO Entanglement Entropy*

- So from the perspective of this particular algorithm, once we fix the noise rate, hardness “saturates” at fixed system size.

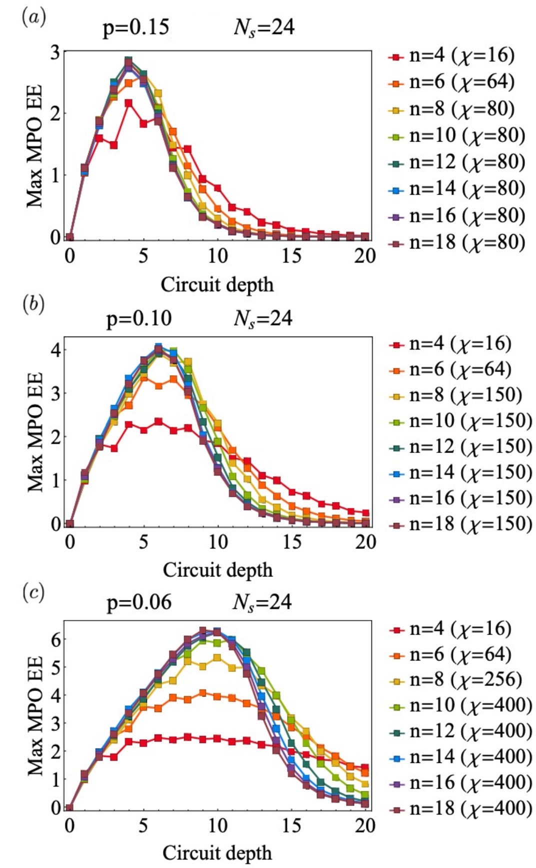

FIG. 5: Maximum MPO entanglement entropy $\mathcal{S}_{\max}$ (averaged over $N_s = 24$ circuit realizations) as a function of the circuit depth D for various number of qubits $4 \leq n \leq 18$ and two-qubit gate error rates (a) $p = 0.15$, (b) $p = 0.1$, and (c) $p = 0.06$. In all cases, we numerically confirm that the chosen bond dimensions are large enough to account for at least 99.1% of the total probability on average.

Plots from [NJF ’20] (2)

- To see this saturation behavior more directly we plot Number of qubits vs *Max MPO Entanglement Entropy*
- Each curve represents a different noise rate at optimal depth (from prior plot)
- Again, we see there’s a maximum system size, determined by the noise rate, after which we don’t gain in complexity using this measure

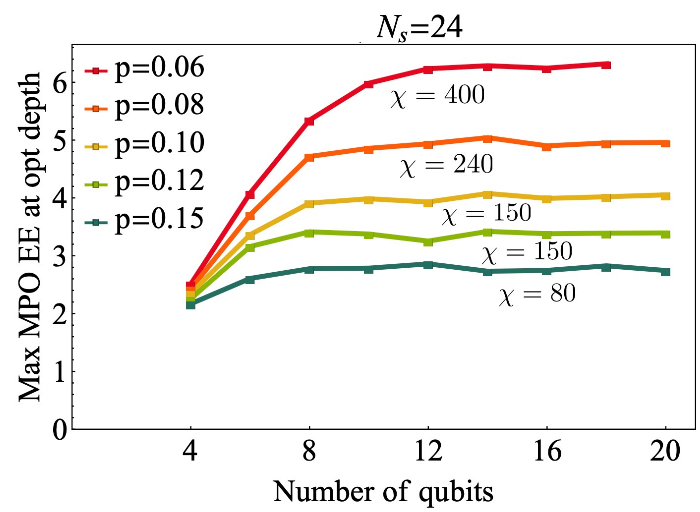

FIG. 8: Maximum achievable MPO entanglement entropy at the optimal circuit depth $\mathcal{S}^{\star}_{\max}$ for various two-qubit gate error rates $0.06 \leq p \leq 0.15$ and number of qubits $4 \leq n \leq 18$. The bond dimension χ used in each case is specified next to each curve.

Conclusions

- Numerically, we observe that for noisy 1D random circuits there is a sense in which quantum correlations peak at a particular system size
- We can make use of this observation to compute noisy output probabilities using Matrix Product Operator (MPO) methods
- On the other hand, we can prove that computing noisy output probabilities (to precision $2^{-O(m \log m)}$) is hard for 2D random circuits below a noise threshold
  - We think this precision is (essentially) tight.

Thanks!