“Quantum Supremacy” and the Complexity of Random Circuit Sampling

Bill Fefferman (UC Berkeley and U. Maryland)

Joint with Adam Bouland, Chinmay Nirkhe, and Umesh Vazirani

arXiv: 1803.04402

"Quantum Supremacy"

- **Goal**: A practical demonstration of a quantum computation that is prohibitively hard for classical computers
  - Needs to be experimentally feasible
  - Need theoretical evidence for hardness (i.e., problem couldn't be solved efficiently on classical computer)
  - Like early quantum algorithms, no need to be useful!
- Stepping stone to scalable, fault-tolerant, universal quantum computers
- But it's much more than that!

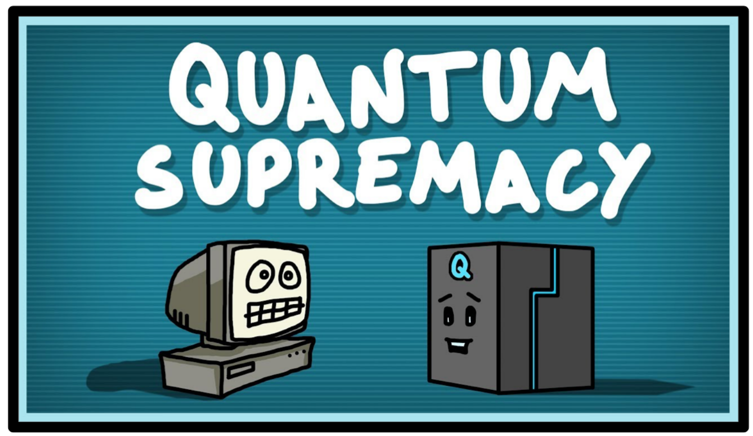

Photo Credit: "Domain of Science"

Importance of quantum supremacy: foundations of computation

- *Experimental* violation of the Extended Church-Turing thesis
  - i.e., If we want to model efficient computation, we must consider quantum mechanics!
- Complements *theoretical* evidence given by earlier speedups (e.g., [Bernstein-Vazirani '93][Simon'94][Shor '94])

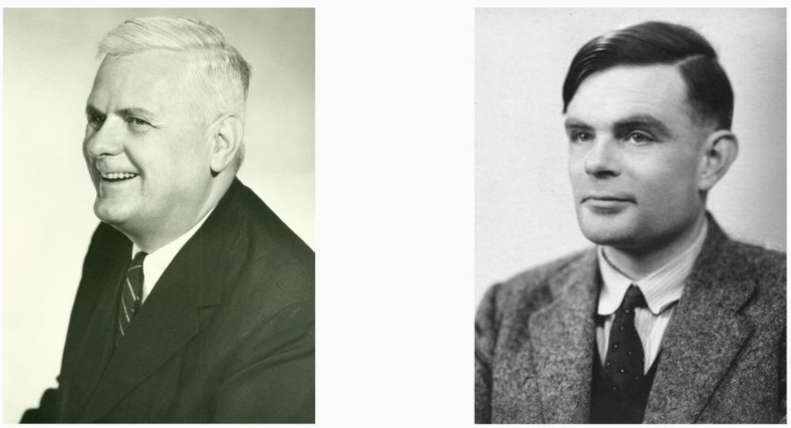

Alonzo Church *Alan Turing*

Importance of quantum supremacy: validating quantum physics

- Exponential growth arguably the most counter-intuitive aspect of quantum mechanics.
  - Is the exponential description of a quantum state really necessary?
- New limit in which to test physics: high complexity.
- *Difficulty*: how to verify something that's exponentially complex?

Importance of quantum supremacy: validating near-term devices

- Quantum supremacy: necessary to have a large quantity of high quality qubits
  - Achieving both is quite difficult experimentally
- In recent years, tools from quantum supremacy have become more and more central to experimental efforts in *validating* NISQ devices
  - E.g., to "tuning qubits" and "diagnosing errors"

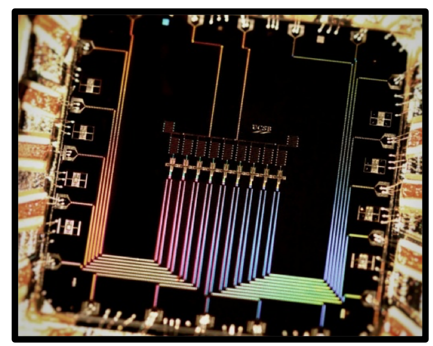

Martinis group: Google/UCSB

Existing quantum supremacy proposals

Broadly speaking, fall into two categories:

1. ***Theoretically driven proposals***
   - Special purpose devices with good evidence for hardness
   - Are not yet experimentally feasible at sufficiently large scale
   - e.g., BosonSampling [Aaronson & Arkhipov '11]
2. ***Experimentally driven proposals***
   - Will be realizable in the near term, on the path to scalable quantum computing
   - But do not yet have as strong theoretical evidence of hardness
   - e.g., Random Circuit Sampling proposal of the Google/UCSB group [Boixo et. al. '16]

Our results

1. We provide ***theoretical backing*** to the leading *experimental* candidate for quantum supremacy: Random Circuit Sampling [Google/UCSB group: Boixo et al '16]

2. We study verification, clarifying when existing proposals work to ***verify these devices***

1. Quantum supremacy from average-case interference patterns

Interference is crucial for quantum algorithms

- Quantum speedups generally come from carefully engineered interference patterns with large amounts of constructive and destructive interference
- **NISQ era:** Random, *average-case* interference patterns
- **Supremacy proposal:** Given random quantum circuit, sample from distribution *close* to its ideal output distribution
- ***Our question:*** How hard is *approximate sampling* for classical computers?

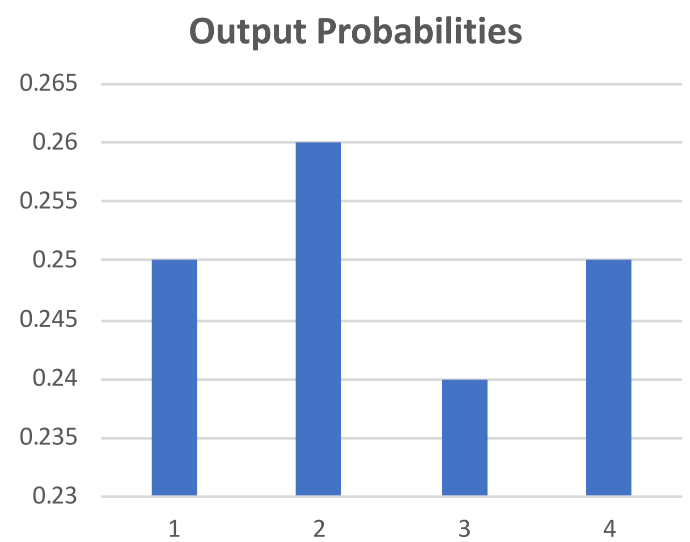

How to prove classical hardness of quantum sampling?

- **Our goal**: to show there's no classical "approximate sampler" algorithm
- **Reduction** [AA'11]: Suffices to prove that ***approximating*** the output probability of ***most*** random quantum circuits is **#P**-hard
- ***Our question*:** Can we give evidence that this true?

BosonSampling [Aaronson & Arkhipov'11]

- ***Task***: sample from the outcome distribution of a random linear optical quantum circuit
- ***Key point*****:** Output probabilities of random linear optical circuits are *permanents* of random matrices
- Permanent has a ***worst-to-average*** case reduction, and so is **#P**-hard to exactly compute permanent of ***most*** random matrices [AA'11, building on Lipton '91]

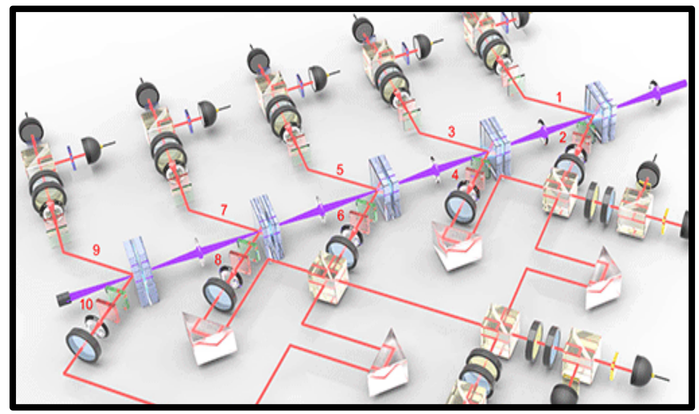

Photo credit: X.-L. Wang *et al*. (2016)

But BosonSampling seems hard to experimentally implement…

- *We've yet to see sufficiently large experiments to test extended Church Turing thesis*
- Further, it's a special purpose device – not necessarily on path to universal scalable quantum computation

Random Circuit Sampling [e.g., Boixo et. al., ‘16]

- ***Task*****:** sample from the output distribution of a random quantum circuit
  - Generate a quantum circuit C on n qubits on a 2D lattice, with $d=n^{1/2}$ layers of Haar random nearest-neighbor gates
  - Start with $|0^n\rangle$ input state and measure in computational basis
- E*xperimentally compelling:* large systems of superconducting qubits coming soon [e.g., Google/UCSB]
- **RCS Conjecture**: **#P**-hard to ***estimate*** output probability of ***most*** random quantum circuits
  - But unlike BosonSampling, no connection to permanents
  - ***Missing*****:** average-case hardness!

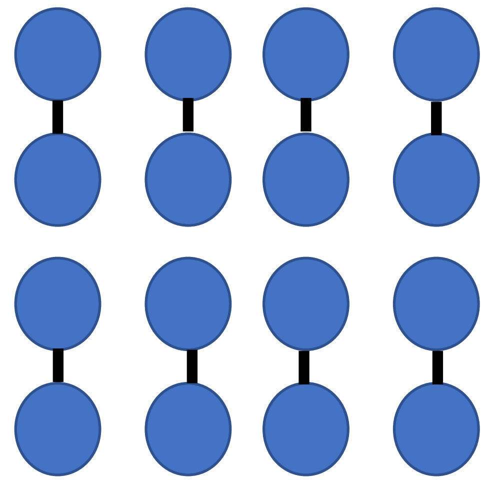

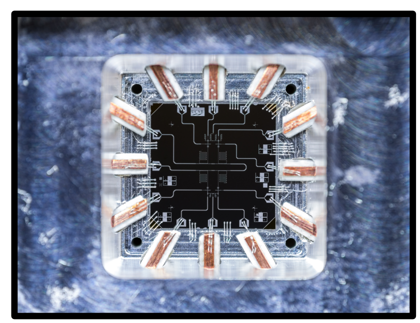

Photo Credit: Michael Fang

Main result: Average-case exact hardness for RCS

- ***We prove:*** “Worst-case to average-case reduction” for exactly computing quantum output probabilities
- Provides a **rigorous** foundation for the classical hardness of RCS!
  - Raises RCS to level of BosonSampling and has a property called anti-concentration [e.g., BHH’12, HBVSE’17, HM’18]
  - *Remaining hurdle*: Extend **exact** to **approximate** average-case hardness.

Major ideas used in proof

- ***Goal***: Use the ability to compute output probabilities of *typical* quantum circuits
  - To compute output probability for worst-case C, $|\langle 0^n|C|0^n\rangle|^2$
- Want to "scramble" worst-case C so that it looks *typical*!
- First attempt:
  - Choose $\{H_i\}_{i\in[m]}$ Haar random gates
  - Now take each gate in C and set $C_i' = C_i H_i$
  - But now C' is completely uncorrelated with C !!

Need to correlate many "random looking" circuits with worst-case circuit C

- Uses polynomial structure coming from the ***Feynman path integral*** and also the ***uniquely quantum*** ability to implement "small fraction of quantum gate"
  - Pick many small angles θ
  - For each θ consider the circuit C' in which $C_i' = C_i H_i e^{-ih_i\theta}$
- ***Observation***: Each circuit C' individually looks random, but they are all correlated (can express each $|\langle 0^n|C'|0^n\rangle|^2$ as a fixed function of θ)
- Use ideas from polynomial interpolation to recover the output probability of worst-case C

2. Using statistical tests to verify RCS

Verifying RCS in the NISQ era

- ***Challenge:*** Need to develop a statistical measure to verify the RCS output distribution from samples of device, but…
  - ***Constraint 1***: don't know the output distribution (only given a description of circuit)
  - ***Constraint 2***: can only take a small (poly(n)) number of samples from the quantum device
- ***Compromise***: OK to use exponential postprocessing time on supercomputer to compute "a few" ideal output probabilities (doable for n=49 qubits)
- Complexity arguments require closeness in total variation distance. **But we can't hope to unconditionally verify this with few samples from the device.**

A candidate test for verifying RCS: cross-entropy [Boixo et. al., 16]

- Proposal is to compute:

$$CE(p_{dev}, p_{id}) = \sum_x p_{dev}(x) \log \frac{1}{p_{id}(x)} = \mathbb{E}_{p_{dev}} \log \left(\frac{1}{p_{id}} \right)$$

- This can be well-approximated in few samples using concentration of measure arguments!
- Then accept if score is sufficiently close to the expected ideal cross-entropy, which can be calculated analytically

Why Cross-Entropy?

- This is a “one-dimensional projection” of observed data
- Does not verify closeness in total variation distance directly
- (Theorem: exist distributions score well on CE but are far in total variation)
- [Boixo et al. ’16]: Assume that

$$\rho_{dev} = \alpha\rho_{id} + (1-\alpha)\, Id$$

In this case, achieving near-perfect cross-entropy certifies closeness in total variation distance

Deeper reasons to believe in Cross-Entropy?

Claim: If Cross-Entropy is close to ideal **and $H(p_{dev}) \geq H(p_{id})$,** then the output distribution is close to ideal in total variation distance

This assumption would follow from certain noise models (e.g., local depolarizing noise) but not from others (e.g., correlated noise, erasure channel etc…)

Proof:

- Pinsker's inequality: $|p_{dev} - p_{id}|_{TV} \leq \sqrt{\frac{1}{2}|p_{dev} - p_{id}|_{KL}}$
  - Where $|p_{dev}-p_{id}|_{KL}=CE(p_{dev},p_{id})-H(p_{dev})$
- So if we find cross-entropy ε-close to ideal, we've certified closeness in total variation distance to error $O(\varepsilon^{1/2})$

More on verification

- The output distributions of RCS are "Porter-Thomas"
  - $\Pr[p_x = q/N] = e^{-q}$
- This is ***not*** a signature of quantum effects
  - Can be reproduced classically (e.g., by Poisson processes)
- However, pairs of distributions scoring highly on CE test share similar "heavy" outcomes
  - This intuition was sharpened by a recent proposal of Aaronson & Chen called "HOG"

$$\mathbb{E}_{p_{dev}} \delta(p_{id} \text{ is ``heavier than median''})$$

- Scoring above some threshold conjectured to be intrinsically hard
  - Can be connected to nonstandard complexity assumptions [AC'16]
  - But don't know how to connect to well-studied complexity assumptions
- Can create a generalized measure which simultaneously verifies both Cross-Entropy and HOG which we call "Binned Output Generation" and is, in some sense "optimal"

Conclusions

- Average case hardness gives evidence that circuit sampling hard even for random circuits which exhibit generic interference patterns.

- For sufficiently small supremacy experiments we can verify supremacy if we make strong enough assumptions about the device output distribution: e.g., experiment only increases entropy

Open Questions

Missing piece: extend hardness of *exactly* computing typical quantum output probability to *approximate* case (this is open for *all* supremacy proposals!)

At what system size should we conclude “quantum supremacy”? What is the importance of implementing a particular system size, like 49 qubits?

Can recent classical heuristics for RCS simulation, such as those of the Alibaba group [Chen et. al., ‘18] be used to verify RCS experiments for larger system sizes?

Thanks!