Complexity-Theoretic Foundations of Quantum Supremacy Experiments

Scott Aaronson, Lijie Chen

UT Austin, Tsinghua University → MIT

July 7, 2017

Section

Quantum Supremacy

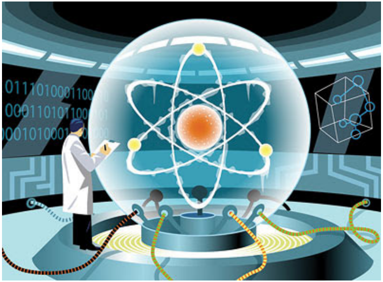

- In this quest, we forget about the applications, only want to find a problem which we can establish a quantum speedup over classical devices as clean as possible.
- The first application of quantum computing:
  - Disprove the QC skeptics!
  - And Extended Church-Turing Thesis.
- An important milestone for QC.

Decision Problem vs. Sampling Problem

- An ideal way for showing quantum supremacy and convincing the skeptics would be:
  - Implement Shor's algorithm [Sho97].
  - Break RSA.
  - Everyone believe your quantum computer works.
- The only problem is that it needs too many qubits.
  - 40 and 4000 are both $O(1)$ in theory, but
  - could require 50 years in the real world.
- Would it be possible to demonstrate quantum supremacy with much less qubits?

Quantum Supremacy via Sampling Problems

- Probably **YES** with a shift to sampling problem.
- **Sampling problem**:
  - Given an input x, you are required to take sample from a certain distribution $\mathcal{D}(x)$ over $\{0,1\}^n$.
- Merits comparing to decision problem:
  - Easier to solve with near-future quantum devices:
    - Do some complicated operations $\Rightarrow$ get a highly entangled quantum state $\Rightarrow$ measure it.
    - Naturally induce a sampling problem.
  - Easier to argue are hard for classical computers:
    - $\mathsf{ExactSampBPP} = \mathsf{ExactSampBQP} \Rightarrow \mathsf{PostBQP} = \mathsf{PostBPP} \Rightarrow \mathsf{PP} \subseteq \mathsf{PH} \Rightarrow$ PH collapses.
- Many works alone this line [TD04, BJS10, AA13, MFF14, JVdN14, FH16, ABKM16].

This talk

- While there are many exciting results, there are still some theoretical challenges for us.

- **Verification for sampling problem**:
  - It is not directly verifiable that our algorithm really takes samples from the predicted distributions $\mathcal{D}(x)$.
  - We have to consider some statistical tests $\mathcal{T}$ on the obtained samples $x_1, x_2, \ldots, x_t$.
  - But then the hardness assumption should imply no classical algorithm can pass $\mathcal{T}$.
  - That is, we ought to talk about relational problems.

This talk

- While there are many exciting results, there are still some theoretical challenges for us.

- **Supremacy Theorem for Approximate Sampling**:
  - PH does not collapse $\Rightarrow$ ExactSampBPP $\neq$ ExactSampBQP.
  - But, real world experiment is **noisy**, hardness for exact version is not convincing enough.
  - Previous results on quantum supremacy for approximate sampling relies on some other unproven conjectures
    - Like in Aaronson and Arkhipov [AA13], they need the hardness of Guassian permanent estimation.
  - Is that necessary? Could there be some simple (relativized) argument for PH does not collapse $\Rightarrow$ SampBPP $\neq$ SampBQP?
    - Or is there an oracle for which the above does not hold?
  - An open question raised in [AA13].

Talk Outline

- Random Quantum Circuit Proposal
  - Heavy Output Generation (HOG)
  - QUAtum THreshold assumption (QUATH)
- Non-Relativizing Techniques Will Be Needed for Strong Quantum Supremacy Theorems.
  - There exists an oracle $\mathcal{O}$, $\mathsf{SampBPP}^{\mathcal{O}} = \mathsf{SampBQP}^{\mathcal{O}}$ and $\mathsf{PH}^{\mathcal{O}}$ is infinite.
  - no relativized way to show quantum supremacy only base on PH doesn't collapse. (unlike the exact version).
- A glimpse on other results.
  - Space-efficient algorithm for simulating quantum algorithm classically.
  - 1 vs. $\Omega(n)$ separation for sampling problems in query complexity.
  - Quantum Supremacy relative to oracles in P/poly.

Section

Random Quantum Circuit Proposal

High level picture:

- Generate a random quantum circuit C on $\sqrt{n} \times \sqrt{n}$ grid.
- Apply C to $|0\rangle^{\otimes n}$ for t times to obtain t samples $x_1, x_2, \ldots, x_t$.
- Apply a statistical test on $x_1, \ldots, x_t$.
  - This step may takes exponential classical time, but would be OK for $n \approx 40$.
- Publish C, to challenge skeptics to pass the same test classically with reasonable amount of time.

The Heavy Output Generation Problem

More specifically:

Problem (HOG, or Heavy Output Generation)

Given as input a random quantum circuit C (will be specified later), generate output strings $x_1, \ldots, x_k$*, at least a* $2/3$ *fraction of which have greater than the median probability in C's output distribution.*

- The verification can be done in exponential time classically.
- We want to find a clean assumption that implies HOG is hard.

The Random Circuit Distribution

We use $\mu_{\text{grid}}^{n,m}$ to denote the following distribution of random circuit on $\sqrt{n} \times \sqrt{n}$ with m gates. (Assuming $m \gg n$).

- A gate can only act on two adjacent qubits.
- For each $t \leq n$, we pick the t-th qubit and a random neighbor of it. (The purpose here is to make sure that there is a gate on every qubit.)
- For each $t > n$, we pick a uniform random pair of adjacent qubits in the grid.
- In either case, we set the t-th gate to be a uniform random 2-qubit gate.

Some notations: Heavy Output, and $\mathrm{adv}(|u\rangle)$

- For a pure state $|u\rangle$ on n qubits, we define $\mathrm{probList}(|u\rangle)$ to be the list consisting of 2^n numbers, $|\langle u|x\rangle|^2$ for each $x \in \{0,1\}^n$.
- Given N real numbers $a_1, a_2, \ldots, a_N$, we use $\mathrm{uphalf}(a_1, a_2, \ldots, a_N)$ to denote the sum of the largest $N/2$ numbers among them, and we let

$$\mathrm{adv}(|u\rangle) = \mathrm{uphalf}(\mathrm{probList}(|u\rangle)).$$

- We say that an output $z \in \{0,1\}^n$ is *heavy* for a quantum circuit C, if it is greater than the median of $\mathrm{probList}(C|0^n\rangle)$.
- We abbreviate $\mathrm{adv}(C|0^n\rangle)$ as $\mathrm{adv}(C)$.
- The simple quantum algorithm's output is heavy w.p. $\mathrm{adv}(C)$.

Lower bound on adv(C)

- What we can prove, is that the expectation of adv(C) is high.

Lemma

For $n \geq 2$ and $m \geq n$:

$$\mathop{\mathbb{E}}_{C \leftarrow \mu_{\text{grid}}^{n,m}}[\text{adv}(C)] \geq \frac{5}{8}.$$

- But we conjecture that adv(C) is large with an *overwhelming* probability.

Conjecture

For $n \geq 2$ and $m \geq n^2$, and for all constants $\varepsilon > 0$,

$$\Pr_{C \leftarrow \mu_{\text{grid}}^{n,m}}\left[\text{adv}(C) < \frac{1 + \ln 2}{2} - \varepsilon\right] < \exp\{-\Omega(n)\}.$$

Lower bound on adv(C)

- But we conjecture that adv(C) is large with an *overwhelming* probability.

Conjecture

For $n \geq 2$ and $m \geq n^2$, and for all constants $\varepsilon > 0$,

$$\Pr_{C \leftarrow \mu_{\text{grid}}^{n,m}} \left[\text{adv}(C) < \frac{1 + \ln 2}{2} - \varepsilon \right] < \exp\{-\Omega(n)\}.$$

- Basically, the above inequality holds when C is replaced by a uniform random unitary on n qubits.
- So what we conjecture is that a random quantum circuit is pseudo-random in a certain sense.
- We provide some evidence by numeric simulation in the Appendix.
- In the following we will assume this conjecture.

Easiness for Quantum Algorithm

We are going to argue that HOG problem is a good quantum supremacy experiment.

Proposition

There is a quantum algorithm that succeeds at HOG with probability $1 - \exp\{-\Omega(\min(n, k))\}$.

- From the conjecture, w.h.p., $\mathrm{adv}(C) > 0.7$.
- In that case, A random sample from C is heavy w.p. 0.7.
- Then a Chernoff bound suffices.

The Quantum Threshold Assumption

Assumption (QUATH, or the QUAntum THreshold assumption)

There is no polynomial-time classical algorithm that takes as input a description of a random quantum circuit C, and that guesses whether $|\langle 0^n|C|0^n\rangle|^2$ is greater or less than the median of all 2^n of the $|\langle 0^n|C|x\rangle|^2$ values, with success probability at least $\frac{1}{2} + \Omega\left(\frac{1}{2^n}\right)$ over the choice of C.

Hardness for Classical Algorithm : Proof Sketch

Theorem

Assuming QUATH, no polynomial-time classical algorithm can solve HOG with probability at least 0.99*.*

- Suppose for contradiction that there exists such an algorithm A, we construct an algorithm to violate QUATH.
- Given a circuit C.
- Apply a random "xor"-mask z on C to get a circuit C' such that $\langle 0|C'|z\rangle = \langle 0|C|0\rangle$.
  - i.e. Hide the amplitude we care about.
- Run A on C', to get a list of outputs $x_1, x_2, \ldots, x_t$, pick one of them x_i at uniformly random.
  - We guess it's greater than median, if $z = x_i$.
  - Take a uniform random guess otherwise.
- Violates QUATH.

Section

SampBPP and SampBQP

Definition (Sampling Problems, SampBPP, and SampBQP)

- *A sampling problem S is a collection of probability distributions $(\mathcal{D}_x)_{x \in \{0,1\}^*}$, one for each input string $x \in \{0,1\}^n$, where $\mathcal{D}_x$ is a distribution over $\{0,1\}^{p(n)}$, for some fixed polynomial p.*
- *Then SampBPP is the class of sampling problems $S = (\mathcal{D}_x)_{x \in \{0,1\}^*}$ for which there exists a probabilistic polynomial-time algorithm B that, given $\left\langle x, 0^{1/\varepsilon} \right\rangle$ as input, samples from a probability distribution $\mathcal{C}_x$ such that $\|\mathcal{C}_x - \mathcal{D}_x\| \leq \varepsilon$.*
- *SampBQP is defined the same way, except that B is quantum now.*

Our goal and what we have

- Our goal is to construct an oracle $\mathcal{O}$ such that:
  - $\mathsf{PH}^{\mathcal{O}}$ is infinite.
  - $\mathsf{SampBPP}^{\mathcal{O}} = \mathsf{SampBQP}^{\mathcal{O}}$.

- What we know is:
  - For a random oracle $\mathcal{O}$, $\mathsf{PH}^{\mathcal{O}}$ is infinite by Rossman, Servedio and Tan [RST15].
  - For a PSPACE-complete language L, $\mathsf{SampBPP}^{L} = \mathsf{SampBQP}^{L}$.

Intuition

- Naive idea:
  - Simply let our oracle be a combination of both a PSPACE-complete language and a random oracle.
  - Problem: SampBPP and SampBQP now get access to a random oracle, it can be proved they are not equal in this case.
- Trying to fix it, can we somehow hide the random oracle so that:
  - An algorithm in PH has access to it, so PH is still infinite.
  - SampBQP algorithm cannot access it (or with very small probability), so SampBQP and SampBPP are not re-separated.

Construction

- Given a string $w \in \{0,1\}^N$, we hide it in a random matrix $\mathcal{M}_w$ of $\{0,1\}^{N \times N}$ as follows:
  - If $w_i = 1$, a uniform random position of i-th row is 1, other positions are 0.
  - If $w_i = 0$, the entire i-th row is 0.
- A random oracle $\mathcal{O}$ can be viewed as a list of functions

$$\{f_n : \{0,1\}^n \to \{0,1\}\}_{n=1}^{\infty}$$

- Or a list of strings

$$\{w_n : \{0,1\}^{2^n} \to \{0,1\}\}_{n=1}^{\infty}$$

- By hiding each w_n into a random matrix of $\{0,1\}^{2^n \times 2^n}$, we can obtain another oracle $\mathcal{M}_{\mathcal{O}}$ (actually a distribution on oracles).

Construction

- $\mathcal{M}_\mathcal{O}$ is just what we want:
  - An algorithm in PH can recover w from $\mathcal{M}_w$ (simply by a OR layer), hence PH is still infinite.
  - Meanwhile, since OR is hard for quantum algorithms [BBBV97], use a BBBV-type argument, one can show that essentially a quantum algorithm with oracle accesses to $\mathcal{M}_\mathcal{O}$ can be simulated efficiently by a classical randomized algorithm.
- Need to work out many technical details, but the idea is very clean.

Section

Space-efficient algorithm for simulating quantum algorithm classically

- Given a n qubit and m gates circuit, how to simulate it classically and efficiently?
- "Schrodinger way":
  - Store the whole wave-function.
  - $O(m2^n)$ time and $O(2^n)$ space.
- "Feynman way":
  - Sum over paths.
  - $O(4^m 1)$ time and $O(m+n)$ space.
- We show:
  - "Savitch way": $O((2d)^n)$ time and poly space, (d is the depth).
  - Can be further improved on circuit on grids.
  - Trade-off between space and time:
    - A d factor in time $\Leftrightarrow$ a 2 factor in space.

1 vs $\Omega(n)$ Separation in query complexity

- Here we consider sampling problems in query complexity.
- The Fourier Sampling problem introduced by Aaronson and Ambainis [AA14], requires only 1 query for a quantum algorithm.
- It is also shown in [AA14] that it requires $\Omega(N/\log N)$ queries for classical randomized algorithms.
- We improve it by showing that Fourier Sampling requires $\Omega(N)$ queries in fact.
- Hence, in the world of query complexity, classical and quantum sampling algorithm has the maximum possible separation.

Quantum Supremacy with respect to oracles in P/poly

- We ask: is there an oracle $\mathcal{O}$ in P/poly, such that $\mathrm{BQP}^{\mathcal{O}} \neq \mathrm{BPP}^{\mathcal{O}}$?
- An intermediate case between black-box (oracle separation) and non-black-box arguments (real world, no oracle) by requiring the oracle to "exist in real world".
- Previous works [Zha12, SG04] imply that the answer is YES when one-way function exist.
- We show that at least some computational assumptions are needed by proving that the answer is NO if $\mathrm{SampBPP} = \mathrm{SampBQP}$ and $\mathrm{NP} \subseteq \mathrm{BPP}$.

Any Questions?

Thank you

S. Aaronson and A. Arkhipov.
The computational complexity of linear optics.
Theory of Computing, 9(4):143–252, 2013.
Earlier version in Proc. ACM STOC'2011. ECCC TR10-170, arXiv:1011.3245.

S. Aaronson and A. Ambainis.
Forrelation: a problem that optimally separates quantum from classical computing.
arXiv:1411.5729, 2014.

Scott Aaronson, Adam Bouland, Greg Kuperberg, and Saeed Mehraban.
The computational complexity of ball permutations.
arXiv preprint arXiv:1610.06646, 2016.

C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani.
Strengths and weaknesses of quantum computing.
SIAM J. Comput., 26(5):1510–1523, 1997.
quant-ph/9701001.

M. Bremner, R. Jozsa, and D. Shepherd.
Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy.
Proc. Roy. Soc. London, A467(2126):459–472, 2010.

arXiv:1005.1407.

Edward Farhi and Aram W Harrow.
Quantum supremacy through the quantum approximate optimization algorithm.
arXiv preprint arXiv:1602.07674, 2016.

Richard Jozsa and Marrten Van den Nest.
Classical simulation complexity of extended clifford circuits.
Quantum Information & Computation, 14(7&8):633–648, 2014.

Tomoyuki Morimae, Keisuke Fujii, and Joseph F Fitzsimons.
Hardness of classically simulating the one-clean-qubit model.
Physical review letters, 112(13):130502, 2014.

Benjamin Rossman, Rocco A Servedio, and Li-Yang Tan.
An average-case depth hierarchy theorem for boolean circuits.
In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1030–1048. IEEE, 2015.

Rocco A Servedio and Steven J Gortler.
Equivalences and separations between quantum and classical learnability.
SIAM Journal on Computing, 33(5):1067–1092, 2004.

P. W. Shor.

Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.
SIAM J. Comput., 26(5):1484–1509, 1997.
Earlier version in Proc. IEEE FOCS'1994. quant-ph/9508027.

B. M. Terhal and D. P. DiVincenzo.
Adaptive quantum computation, constant-depth circuits and Arthur-Merlin games.
Quantum Information and Computation, 4(2):134–145, 2004.
quant-ph/0205133.

Mark Zhandry.
How to construct quantum random functions.
In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 679–687. IEEE, 2012.