

Quantum Supremacy

The key to theoretical quantum supremacy arguments is the (apparently) enormous difference between classical and quantum post-selected computation.


Last time we followed the early history of these arguments, where they were used to show that adiabatic computation in the ground state of a stoquastic Hamiltonian cannot be universal unless the PH collapses.

QS also provides evidence for the hardness of sampling output distributions of quantum circuits. There is no classical poly-time algorithm to sample these distributions exactly unless the PH collapses.

But the primary modern application of theoretical QS is to argue for hardness of sampling devices that implement limited computational models that fall short of being universal. One example is sampling bosons in linear optical networks. Another example, which we will focus on, are “commuting” quantum circuits.

Quantum Supremacy

Instantaneous quantum polynomial-time (IQP): a class of circuits which starts in $|+\rangle^n$, applies polynomially many gates which are diagonal in the Z basis, then Hadamard and measure.



$$D = \sum_{k,l} J_{kl} \sigma_z^k \sigma_z^l + \sum_k M_k \sigma_z^k$$

$$H^{\otimes n} U_D H^{\otimes n} |0^n\rangle$$

The name IQP comes from the commuting nature of the intermediate gates. This means that they can be applied in any order, and hence may “all be applied at once” (though this may be experimentally difficult).

Quantum Supremacy

IQP was introduced in “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy”, Bremner, Jozsa, and Shepherd, 2010.

Abstract

We consider quantum computations comprising only commuting gates, known as IQP computations, and provide compelling evidence that the task of sampling their output probability distributions is unlikely to be achievable by any efficient classical means. More specifically we introduce the class post-IQP of languages decided with bounded error by uniform families of IQP circuits with post-selection, and prove first that post-IQP equals the classical class PP. Using this result we show that if the output distributions of uniform IQP circuit families could be classically efficiently sampled, even up to 41% multiplicative error in the probabilities, then the infinite tower of classical complexity classes known as the polynomial hierarchy, would collapse to its third level. We mention some further results on the classical simulation properties of IQP circuit families, in particular showing that if the output distribution results from measurements on only $O(\log n)$ lines then it may in fact, be classically efficiently sampled.

Quantum Supremacy

Therefore the main ingredient needed to establish this result is a proof that $\text{PostIQP} = \text{PP}$.

This is done by showing $\text{PostIQP} = \text{PostBQP}$. In other words, we need to show that adding post-selection to IQP circuits suffices to make them universal.

A convenient universal gate set for BQP is $\{H, Z, CZ, T = \text{diag}(1, e^{i\frac{\pi}{4}})\}$

All of these are diagonal except the Hadamard gate. So what we really need is to use post-selection to intersperse H gates throughout any IQP circuit, to obtain an arbitrary BQP circuit.

Any ideas for using post-selection (together with all the ancilla qubits you want) to insert H gates into an IQP circuit? (we haven't discussed it in class, but some may know the idea from elsewhere...)

Quantum Supremacy

The construction that allows us to insert Hadamard gates into an IQP circuit using ancillas and post-selection is closely related to a general construction called “gate teleportation.”


Gate teleportation is an analog of state teleportation, for unitary operators. It has many uses in quantum computation. It underlies the measurement-based model of quantum computation, and is used in fault-tolerant QC (where “magic states” are consumed to teleport T gates into a Clifford circuit)

As with state teleportation, gate teleportation requires making a measurement with various possible outcomes that may scramble the gate that is being teleported.

In most settings one must use an adaptive correction to descramble the gate after the measurement. But with post-selection, we can just post-select on measurement outcomes that do not require a correction. (as in state teleportation, if we measure $|\Phi^+\rangle$ then there is no correction).

Quantum Supremacy

The most basic form of gate teleportation is the following identity:




Where m is the measurement outcome, 0 or 1. Let $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, so the state after CZ and H is:

$$\alpha|++\rangle + \beta|--\rangle = \frac{1}{\sqrt{2}} (|0\rangle \otimes H|\psi\rangle + |1\rangle \otimes XH|\psi\rangle)$$

Quantum Supremacy

A mild generalization of the previous identity is the following:



This holds because Z_θ commutes with CZ, and so it as if we input the state $Z_\theta |\psi\rangle$ into the circuit identity on the previous slide.

Quantum Supremacy

Returning to the setting of IQP, Bremner, Jozsa, and Shepherd use the following identity:

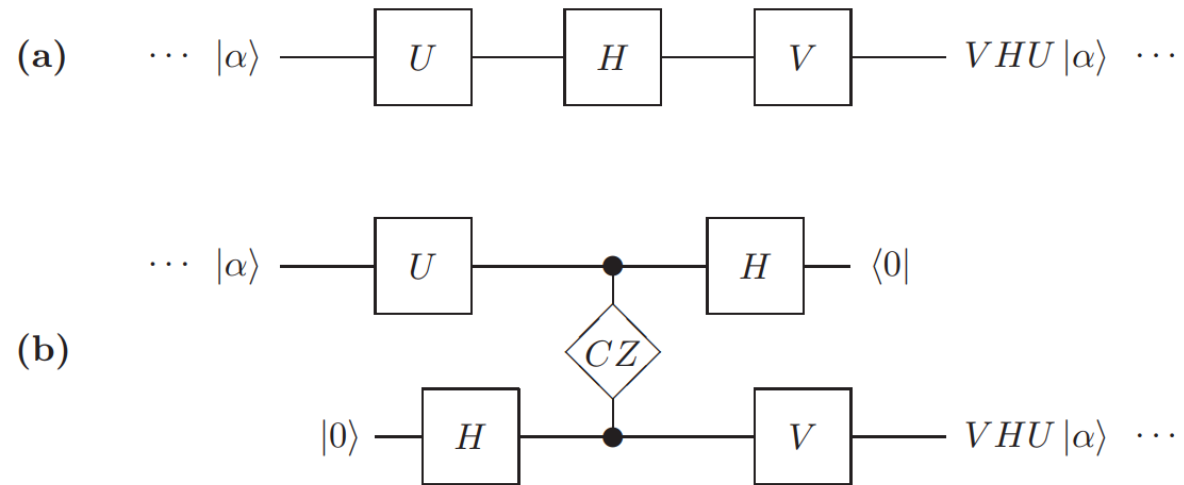


Figure 1: The Hadamard gadget for removal of intermediate H gates. **(a)** $|\alpha\rangle$ represents a general input state to a gate U within the circuit that is followed by an intermediate H gate. **(b)** The lower line is a new ancillary qubit line. The original intermediate H gate may then be replaced by a new CZ gate, a post-selection (denoted by $\langle 0|$) and two H gates that are now both at the ends of lines, as allowed in IQP circuit architecture.

Quantum Supremacy

Once we have $\text{PostIQP} = \text{PostBQP}$, it follows that no classical algorithm can efficiently sample the exact output distribution of an IQP circuit unless the PH collapses.

This “exact” condition can be weakened to “sampling to multiplicative error on each amplitude”,

$$|q_x - p_x| \leq \alpha p_x$$

To get to a more meaningful notion of approximation (additive error), the usual technique is to relate these output distributions to quantities that are believed to be hard to classically compute.

$$Z(\omega) = \sum_{z \in \{\pm 1\}^n} \omega^{\sum_{i < j} w_{ij} z_i z_j + \sum_{k=1}^n v_k z_k}$$

$$Z\left(e^{i\frac{\pi}{8}}\right) = \frac{1}{2^n} \langle 0^n | C_{\text{IQP}} | 0^n \rangle$$

Let Z_R denote partition functions associated with this random choice of weights.

Conjecture 2. *It is #P-hard to approximate $|Z_R|^2$ up to multiplicative error $1/4 + o(1)$ for a $1/24$ fraction of instances over the choice of vertex and edge weights.*

Bremner, Montanaro, Shepherd. 2015

Quantum Supremacy

Note also that the output distribution of nonadaptive MBQC also yields quantum supremacy, for the same reasons as IQP. To my knowledge there is no ideal reference for this statement, because it became common knowledge before anyone wrote about it.

Another point is that if U is a constant-depth quantum circuits that yields quantum supremacy, then

$$H' = UHU^\dagger \quad , \quad H = - \sum_{i=1}^n Z_i$$

Is a k -local Hamiltonian, with $k = O(1)$, which has the ground state $U|0^n\rangle$. Therefore sampling the ground state of this (gapped, frustration-free) Hamiltonian yields QS. Note also that H is stoquastic, so H' samples the ground state of a stoquastic H in a rotated basis.

Quantum Supremacy

Even if we accept the conjectures about complex-temperature partition functions being hard to compute, and so we believe IQP circuits are hard to simulate up to additive error, we can ask what happens to this sampling complexity when we introduce noise into the system.

This is an active area of research. In Feb 2019 there was a notable paper which appears to simulate a large fraction of IQP circuits in classical poly-time, if the input state is slightly mixed

$$(1 - \epsilon)|0\rangle\langle 0| + \epsilon I/2$$

For “a large fraction” of IQP circuits with input states of the above form, the authors give a time

$$n^{\mathcal{O}(\epsilon^{-1} \log \delta^{-1})}$$

To sample IQP output distributions within additive error δ .

Quantum Supremacy

“Efficient simulation of Clifford circuits with nonstabilizer input states.” Bu, Koh. 2019

Example 2—IQP circuits have a simple structure with input states $|0\rangle^{\otimes n}$ and gates of the form $H^{\otimes n}DH^{\otimes n}$, where the diagonal gates in D are chosen from the gate set $\{Z, S, T, CZ\}$. It has been shown that $\text{postIQP} = \text{postBQP}$ [9] and thus, the output probabilities are #P-hard to approximate up to some constant relative error [24–26]. Also, if there is some depolarizing noise acting on each input state $|0\rangle$, i.e., each input state is a mixed state $(1 - \epsilon)|0\rangle\langle 0| + \epsilon\frac{I}{2}$, then Theorem 1 implies that there exists a classical algorithm to approximate the output probability up to l_1 norm δ in time $n^{O(\log(1/\delta)/\epsilon)}$ for a large fraction of such IQP circuits. (The proof is presented in Appendix B in detail, which depends on the output distribution of IQP circuits in Appendix C.)

Quantum Supremacy

The term “Quantum Supremacy” is credited to a lecture by John Preskill, which is on the arxiv titled “Quantum computation and the entanglement frontier”, 2012.

Quantum information science explores the frontier of highly complex quantum states, the "entanglement frontier." This study is motivated by the observation (widely believed but unproven) that classical systems cannot simulate highly entangled quantum systems efficiently, and we hope to hasten the day when well controlled quantum systems can perform tasks surpassing what can be done in the classical world. **One way to achieve such "quantum supremacy" would be to run an algorithm on a quantum computer which solves a problem with a super-polynomial speedup relative to classical computers, but there may be other ways that can be achieved sooner, such as simulating exotic quantum states of strongly correlated matter.** To operate a large scale quantum computer reliably we will need to overcome the debilitating effects of decoherence, which might be done using "standard" quantum hardware protected by quantum error-correcting codes, or by exploiting the nonabelian quantum statistics of anyons realized in solid state systems, or by combining both methods. Only by challenging the entanglement frontier will we learn whether Nature provides extravagant resources far beyond what the classical world would allow.



John Preskill

Quantum Supremacy

Preskill's words inspired Sergio Boixo and other researchers at Google to propose a scheme for achieving near-term practical quantum supremacy. "Characterizing Quantum Supremacy in Near-Term Devices", 2016.

A major downside of theoretical QS, besides all of the assumptions that are involved, is that it only holds rigorously in the model of asymptotically large computations.



Sergio Boixo, UNM Physics PhD 2008

Instead, the Google team set out to show that a superconducting chip with 50 qubits, built by John Martinis (whose UCSB lab was acquired by Google around this time), could perform some well-defined task much more quickly than the worlds largest supercomputers.

The task they chose to focus on is the one that is native to their device: random circuit sampling. Several dozen random 1 and 2 qubit gates (taken from the Clifford + T architecture).

Quantum Supremacy

Given a random quantum circuit U_1, \dots, U_t , how can we verify the QC implements it accurately?

This is a hard problem, and Google's solution is to use an exponential-time brute force classical simulation to verify the quantum computation.

This may seem at first like circular logic: we need a classical computer to verify the output of the quantum computer, to make sure it is a distribution that no classical computer could sample. (?)

Fortunately the practical setting allows us to break out of the cycle: the quantum computer will run for a few seconds, and the best classical supercomputer running the best classical algorithms will take several weeks.

We can approach the limits of the classical supercomputer along various curves (e.g. # of qubits vs depth), and then extrapolate beyond the classical limit of feasibility to declare quantum supremacy.

Quantum Supremacy

In more detail, we want to devise a metric or a YES/NO test that decides whether or not a classical computer could simulate the given quantum circuit.

The output of a random circuit is not uniformly random over bit strings. Some bit strings are more likely than others. The histogram of these output probabilities is known to be a Porter-Thomas distribution, $N e^{-Np}$.

To perform the experiment, we sample bit strings from the device x_1, \dots, x_m , and we use an exponential-time classical computation to compute the probabilities of these output strings,

$$p(x_1), \dots, p(x_m)$$

Because the PT distribution is highly nonuniform, we expect to see several strings x_i (e.g. a 2/3s fraction) with output probabilities $p(x_i)$ that are above the median. This metric is called “heavy-output generation” (HOG) and it comes from a simplification of Google’s original proposal by Aaronson and Chen in 2018.