Quantum Supremacy

Practical QS: perform some computational task on a well-controlled quantum device, which cannot be simulated in a reasonable time by the best-known classical algorithms and hardware.

Theoretical QS: perform a computational task efficiently on a quantum device, and prove that task cannot be efficiently classically simulated.

Since proving $P \neq BQP$ seems to be beyond the capabilities of our current civilization, we lower the standards for theoretical QS. One seeks to provide formal evidence that classical simulation is unlikely.

For example: 3-SAT is NP-complete, so it cannot be efficiently classical solved unless P = NP.

Theoretical QS: perform a computational task efficiently on a quantum device, and prove that task cannot be efficiently classically simulated unless "the polynomial Heierarchy collapses to the 3nd level."

Quantum Supremacy

A common feature of QS arguments is that they consider sampling problems, rather than decision problems. They allow us to characterize the complexity of sampling measurements of quantum states.

Which is more difficult:

Task A: deciding if a circuit outputs 1 with probability at least 2/3s, or at most 1/3s

Task B: sampling from the output of an n-qubit circuit in the computational basis

Sampling from distributions is generically more difficult than approximating observables, since we can use samples to estimate observables, but not the other way around.

One can imagine quantum systems whose local observables are easy to classically compute, but for which sampling the full state is computationally complex.

By moving from decision problems to sampling problems, we make the task of classical simulation much more difficult. But if most useful problems have a decision form, what problem are we solving now?

Quantum Supremacy

All the various examples of theoretical QS are based on the same underlying argument. Outline:

1. We consider the class PostBQP, poly-time quantum circuits with the ability to post-select on measurement probabilities.

2. In 2004, Aaronson showed that PostBQP = PP. By Toda's theorem, $PH \subseteq P^{PP}$.

3. In contrast, we can consider PostBPP, poly-time classical prob computation with the ability to post-select on arbitrary measurement probabilities. But $PostBPP \subseteq BPP^{NP}$, which is in the 3rd level of PH.

4. Now we consider quantum system, which may not be a universal QC. But if we add the fantastical power of post selection, then this device could become universal. Therefore, our device + postselection = PostBQP. Therefore if we could simulate our device classically, then we could simulate our device + postselection in PostBPP. Then PostBPP = PostBQP, and the PH collapses to the 3rd level.

This is the modern version, but the original arguments date back to "Adaptive Quantum Computation, Constant Depth Quantum Circuits and Arthur-Merlin Games." Barbara M. Terhal, David P. DiVincenzo, 2002.

Quantum Supremacy

We begin unpacking this argument by defining classical post-selected computation.

Definition 9 (PostBPP) *A promise problem $L = L_{\text{yes}} \cup L_{\text{no}}$ belongs to the class* PostBPP *iff there exist a polynomial p, predicates $a(x,y)$ and $b(x,y)$ from the class* P *defined for any $y \in \Sigma^{p(|x|)}$, such that*

$$\begin{aligned} x \in L \quad &\Longrightarrow \quad \mathbb{P}[b(x,y)=1] > 0, \\ x \in L_{\text{yes}} \quad &\Longrightarrow \quad \mathbb{P}[a(x,y)=1 \,|\, b(x,y)=1] \geq 2/3, \\ x \in L_{\text{no}} \quad &\Longrightarrow \quad \mathbb{P}[a(x,y)=1 \,|\, b(x,y)=1] \leq 1/3. \end{aligned}$$

where $y \in \Sigma^{p(|x|)}$ is a random uniformly distributed bit string, and $\mathbb{P}[a \,|\, b]$ is the conditional probability.

The point is that the bit b (which is a Boolean proposition computed from the literals x,y) is the bit we post-select on. At the end of the computation, we only compare about events with b = 1.

We also have $\text{PostBPP} = \text{BPP}_{\text{path}}$, where the latter class is define by summing over exponentially many paths, as in our de-randomized definition of BPP. This is used to show that PostBPP is in the 3rd level of PH.

This definition of PostBPP is taken from the paper which first defined PostBPP, which is “The Complexity of Stoquastic Local Hamiltonian Problems.” Bravyi, DiVincenzo, Oliveira, Terhal, 2006.

Quantum Supremacy

Next we define PostBQP and show that it is equal to PP.

Definition 1 PostBQP *is the class of languages* $L \subseteq \{0,1\}^*$ *for which there exists a uniform*[7] *family of polynomial-size quantum circuits* $\{C_n\}_{n \geq 1}$ *such that for all inputs* x,

(i) After C_n *is applied to the state* $|0 \cdots 0\rangle \otimes |x\rangle$, *the first qubit has a nonzero probability of being measured to be* $|1\rangle$.

(ii) If $x \in L$, *then conditioned on the first qubit being* $|1\rangle$, *the second qubit is* $|1\rangle$ *with probability at least* $2/3$.

(iii) If $x \notin L$, *then conditioned on the first qubit being* $|1\rangle$, *the second qubit is* $|1\rangle$ *with probability at most* $1/3$.

It is immediate that NP $\subseteq$ PostBQP. Also, to show PostBQP $\subseteq$ PP, we can use the same observations used by Adleman, DeMarrais, and Huang [6] to show that BQP $\subseteq$ PP, but sum only over paths where the first qubit is $|1\rangle$ at the end. In more detail:

"Quantum Computing, Postselection, and Probabilistic Polynomial-Time." Aaronson, 2004.

Quantum Supremacy

As Aaronson explains, upper bounding PostBQP by PP is not that different from upper bounding BQP by PP, since PP is already sensitive to low-probability events.

It is immediate that NP $\subseteq$ PostBQP. Also, to show PostBQP $\subseteq$ PP, we can use the same observations used by Adleman, DeMarrais, and Huang [6] to show that BQP $\subseteq$ PP, but sum only over paths where the first qubit is $|1\rangle$ at the end. In more detail:

Proposition 2 PostBQP $\subseteq$ PP.

Proof. By a result of Shi [30], we can assume without loss of generality that our quantum circuit is composed of Hadamard and Toffoli gates.[8] Then the final amplitude α_z of each basis state $|z\rangle$ can be written as a sum of exponentially many contributions, call them $a_{z,1}, \ldots, a_{z,N}$, each of which is a rational real number computable in classical polynomial time. So the final probability of $|z\rangle$ equals

$$\alpha_z^2 = (a_{z,1} + \cdots + a_{z,n})^2 = \sum_{ij} a_{z,i} a_{z,j}.$$

We need to test which is greater: the sum S_0 of α_z^2 over all z beginning with 10, or the sum S_1 of α_z^2 over all z beginning with 11. But we can do this in PP: we simply put the positive contributions $a_{z,i}a_{z,j}$ to S_1 and negative contributions to S_0 on "one side of the ledger," and the negative contributions to S_1 and positive contributions to S_0 on the other side. ■

Quantum Supremacy

The direction $\mathrm{PP} \subseteq \mathrm{PostBQP}$ is more interesting. Recall that PP contains NP (and in fact all of PH), and so we do not expect BQP to contain PP. Therefore post-selection must be used in an essential way to show

$$\mathrm{PP} \subseteq \mathrm{PostBQP}$$

We need to give a reduction, starting from an arbitrary problem in PP and then using quantum computation + post selection to solve that problem.

The definition of PP can be distilled down to the following: let $f : \{0,1\}^n \to \{0,1\}$ be efficiently computable, and let $s = |\{x : f(x) = 1\}|$. Decide whether $s < 2^{n-1}$ or $s > 2^{n-1}$.

We need to show that this problem is in PostBQP for any efficiently computable f. The first step prepares:

$$2^{-n/2} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$$

Quantum Supremacy

Before we can responsibly prepare the following state and proceed with the proof:

$$2^{-n/2} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$$

We need to take the safety training course on applying functions in superposition. As long as f is an efficiently computable function, we can construct the following unitary as an efficient quantum circuit:

$$U|x, x\rangle \rightarrow |x, f(x)\rangle$$

Keeping around a copy of x in the extra register makes the function reversible, and hence unitary. Note that classical computers also let you apply functions in parallel:

$$2^{-n} \sum_{x \in \{0,1\}^n} |x, x\rangle\langle x, x| \rightarrow 2^{-n} \sum_{x \in \{0,1\}^n} |x, f(x)\rangle\langle x, f(x)|$$

Quantum Supremacy

Now that we all have our training certificates, we can resume the proof:

$$2^{-n/2} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$$

Apply a Hadamard to n qubits, and postselect on them being in the state 00…0. This results in $|0^n\rangle|\psi\rangle$,

$$|\psi\rangle = \frac{(2^n - s)\,|0\rangle + s\,|1\rangle}{\sqrt{(2^n - s)^2 + s^2}}.$$

At this point we have a single qubit state, whose amplitudes encode the accept/reject probabilities of the original circuit. We are done if we learn these amplitudes by a measurement. But because the interesting case is very close uniform, we use postselection to help the final readout.

Quantum Supremacy

This produces the state $|0\rangle^{\otimes n}|\psi\rangle$ where

$$|\psi\rangle = \frac{(2^n - s)|0\rangle + s|1\rangle}{\sqrt{(2^n - s)^2 + s^2}}.$$

Next, for some positive real numbers α, β to be specified later, prepare $\alpha|0\rangle|\psi\rangle + \beta|1\rangle H|\psi\rangle$ where

$$H|\psi\rangle = \frac{\sqrt{1/2}(2^n)|0\rangle + \sqrt{1/2}(2^n - 2s)|1\rangle}{\sqrt{(2^n - s)^2 + s^2}}$$

is the result of applying a Hadamard gate to $|\psi\rangle$. Then postselect on the second qubit being $|1\rangle$. This yields the reduced state

$$|\varphi_{\beta/\alpha}\rangle = \frac{\alpha s|0\rangle + \beta\sqrt{1/2}(2^n - 2s)|1\rangle}{\sqrt{\alpha^2 s^2 + (\beta^2/2)(2^n - 2s)^2}}$$

in the first qubit.

Quantum Supremacy

Suppose $s < 2^{n-1}$, so that s and $\sqrt{1/2}\,(2^n - 2s)$ are both at least 1. Then we claim there exists an integer $i \in [-n, n]$ such that, if we set $\beta/\alpha = 2^i$, then $|\varphi_{2^i}\rangle$ is close to the state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$:

$$|\langle +|\varphi_{2^i}\rangle| \geq \frac{1+\sqrt{2}}{\sqrt{6}} > 0.985.$$

For since $\sqrt{1/2}\,(2^n - 2s)/s$ lies between 2^{-n} and 2^n, there must be an integer $i \in [-n, n-1]$ such that $|\varphi_{2^i}\rangle$ and $|\varphi_{2^{i+1}}\rangle$ fall on opposite sides of $|+\rangle$ in the first quadrant (see Figure 1). So the worst case is that $\langle +|\varphi_{2^i}\rangle = \langle +|\varphi_{2^{i+1}}\rangle$, which occurs when $|\varphi_{2^i}\rangle = \sqrt{2/3}\,|0\rangle + \sqrt{1/3}\,|0\rangle$ and $|\varphi_{2^{i+1}}\rangle = \sqrt{1/3}\,|0\rangle + \sqrt{2/3}\,|0\rangle$. On the other hand, suppose $s \geq 2^{n-1}$, so that $\sqrt{1/2}\,(2^n - 2s) \leq 0$. Then $|\varphi_{2^i}\rangle$ never lies in the first or third quadrants, and therefore $|\langle +|\varphi_{2^i}\rangle| \leq 1/\sqrt{2} < 0.985$.

It follows that, by repeating the whole algorithm $n\,(2n+1)$ times (as in Proposition 3), with n invocations for each integer $i \in [-n, n]$, we can learn whether $s < 2^{n-1}$ or $s \geq 2^{n-1}$ with exponentially small probability of error. ■

Quantum Supremacy

Once we have PostBQP = PP, we know that quantum computing with postselection can solve problems that are believed to be outside of the polynomial hierarchy.

This gives us formal evidence that PostBQP, $\mathrm{PostBPP} \neq \mathrm{PostBQP}$. We would like "un-post-select both sides of this equation" to obtain $\mathrm{BPP} \neq \mathrm{BQP}$.

The argument is: if a BPP machine can simulate a BQP machine, then that same BPP with postselection can simulate that same BQP machine with postselection.

This argument works formally, but comes with a strong caveat: the BPP simulation must be to within multiplicative error on the amplitudes. If the BQP distribution is p, we demand there is some such that for all events x the BPP simulation samples from q with

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

Why is this level of precision needed? (we can see the reason from the definitions we have so far).

Quantum Supremacy

Therefore these QS arguments allow us to rule out super-precise classical simulations of various quantum systems. Either exact simulation, or "simulation up to multiplicative error", are likely impossible.

But two quantum states that are close in trace distance are effectively indistinguishable, so what we would really like to do is rule out classically sampling distributions that are close in trace distance.

For this we need additional arguments and assumptions. Typically we want to go from a worst-case hardness (or approximating an amplitude) to an average-case hardness (the amplitudes are hard to compute on average).

Therefore one way to have all quantum supremacy arguments collapse is to show that, sure exact classical simulation is inefficient, but approximate classical simulation could be. Despite dating back to 2002, QS arguments only became prominent in 2011 and later.

The gold standard problems are the ones like factoring: in $BQP \cap NP$ and apparently not in P. The shift to considering exact or super-precise sampling problems is in some sense a sign of desperation.

Quantum Supremacy

As mentioned, the arguments about PostBPP and PostBQP originated from investigations of stoquastic Hamiltonians by the group at IBM. It's worth understand their motivation.

In 2000, Farhi, Goldstone, Gutmann, and Sipser proposed a quantum algorithm for solving optimization problems. They called it the quantum adiabatic algorithm because it operates by the adiabatic theorem.

The goal is to solve discrete optimization problems: minimize $f : \{0,1\}^n \to \mathbb{R}$. To do this we define

$$H_p = \sum_{x \in \{0,1\}^n} f(x)|x\rangle\langle x|$$

So that the ground state of H_p corresponds to the bit string that minimizes f. Now introduce a simple transverse field, begin its ground state, and slowly interpolate to the final Hamiltonian:

$$H_B = -\sum_{i=1}^{n} X_i \quad , \quad H(s) = (1-s)H_B + sH_p$$

Quantum Supremacy

Begin in the ground state of the beginning Hamiltonian at s = 0, then increase s sufficiently slowly so as to remain in the ground state of H(s) until s = 1, at which point the ground state solves the problem.

$$H_p = \sum_{x \in \{0,1\}^n} f(x)|x\rangle\langle x| \qquad H_B = -\sum_{i=1}^{n} X_i \quad , \quad H(s) = (1-s)H_B + sH_p$$

To characterize "sufficiently slowly" set s = t/T for a time scale T > 0, and evolve with H(t). The adiabatic theorem implies this algorithm works if

$$T = \mathrm{poly}(n, \Delta^{-1}) \quad , \quad \Delta = \min_{0 \le s \le 1} E_1(s) = E_0(s)$$

We could either imagine building this Hamiltonian in a hardware device ("analog"), or using Hamiltonian simulation to evolve with H(t) on a digital quantum computer. Note that the latter can handle sparse matrices and so writing H_p in terms of Pauli Z operators is not necessarily required.

Quantum Supremacy

This algorithm was immediately controversial, because it proposed to solve NP-complete problems with a quantum computer. The computer scientists immediately set out to prove the algorithm fails in general (which happens if the spectral gap is exponentially small).

But this quantum adiabatic algorithm (which later became known as quantum adiabatic optimization, or quantum annealing) is just too compelling. No matter how many problems or failures people find with the algorithm, it continues to hold interest.

Shortly after the work of Farhi et al. in the MIT physics and CS departments , Seth Lloyd (who is also at MIT, in the EE department) proposed an implementation of Farhi's algorithm using superconducting flux qubits.

Lloyd plugged in units to describe energy gaps in terms of Ghz, and concluded there was no way for the algorithm to work: even if the gap closes as 1/n, it will be lost to noise above 10s of qubits.

Quantum Supremacy

Lloyd did not patent his proposed architecture of Farhi's adiabatic algorithm.

Perhaps it would have been a good idea to do so, because around 2004 a quantum computing startup was founded in Canada: D-Wave Systems Inc. From the beginning their vision was to accumulate hardware patents on route to building a quantum annealer with flux qubits.

From 2004-2012 D-Wave attracted 10s of millions in venture capital, but they did so at the expense of their reputation: they made outlandish claims about what quantum annealing could do.

This greatly irked the quantum computer scientists, who set out to prove that D-Wave's model of quantum computation is classically simulable.

Quantum Supremacy

By 2004, IBM had built up an expert team of quantum researchers. About corporate politics, who can say, but this directly preceeded the burst of discovery by the IBM group about stoquastic Hamiltonians…

The D-Wave quantum annealing Hamiltonian is stoquastic. In fact Bravyi et al. make a point of saying that every effective spin interaction that can be engineered from superconduting flux qubits is stoquastic.

We've already seen the fact that the stoquastic LH problem is in AM, which provides formal evidence for the reduced complexity of ground states used in stoquastic quantum annealing.

But AM is still too powerful to place limits on quantum annealing, since AM contains NP. Therefore we would like to limit the complexity of stoquastic adiabatic computation, without relying on a wizard.

Quantum Supremacy

This was the reason to consider PostBPP in "The Complexity of Stoquastic Local Hamiltonian Problems." There it is shown that sampling the outputs of stoquastic adiabatic computation is in PostBPP, so StoqAQC cannot be universal for quantum computation unless the PH collapses.

The proof is based on path integral Monte Carlo and rejection sampling:

$$\pi(z_1, ..., z_L) = \frac{1}{Z} \prod_{i=1}^{L} \langle z_i | e^{-\beta H/L} | z_{i+1} \rangle$$

$$\pi(z_1) = \sum_{z_2,...,z_L} \pi(z_1, ..., z_L) = \langle z_1 | e^{-\beta H/L} | z_1 \rangle / Z$$

We already said we can sample from the ground state (or thermal state) of stoquastic H if we can sample from π. We can write down a random walk to do this, but we can't determine the number of steps the walk needs to converge. But with postselection we can replace the Markov chain by rejection sampling.

Quantum Supremacy

One of the greatest open problems in Hamiltonian complexity is to show that stoquastic adiabatic computation can be simulated in polynomial-time. I have worked on this problem for 10 years!

In 2012, D-Wave opened their devices with over 100 qubits to academic researchers. Now they have over 4000 qubits. The device is interesting but has not yet found any quantum speedup.

The reason they have not found a speedup is that QMC works very well in practice. Therefore since 2012 we have accumulated a lot of empirical evidence for this "stoquastic simulation conjecture."

However, due to Hastings, we have (pathological) counterexamples to solving this conjecture by any existing QMC method. Therefore we need new ideas for QMC methods, or for ruling out counterexamples.

In APS March Meeting 2019, D-Wave announced new QA Hamiltonians with nonstoquastic interactions. But the stoquastic simulation conjecture continues to hold great interests amongst the QI complexity theorists.