QSS49 - Scott Aaronson - Questions & Answers

Scott Aaronson

As you said, in computer science it's difficult to prove that things are hard. From a computer science perspective, to what extent is "quantum supremacy" something absolute, and to what extent is it a shifting bar that depends on how good our classical hardware and algorithms are?

SCOTT: It's kind of like the question "can computers beat the best humans at chess?" The latter question also involves a non-absolute shifting bar – maybe humans are getting better! Or maybe they simply haven't figured out how to beat the computers yet! So how can we say "computer supremacy" has been decisively achieved? Nevertheless, one can clearly see a phase transition, with Deep Blue in 1996-1997, where the burden of proof shifted massively from one side to the other, and has stayed there ever since. Even if humans **are** getting better, the computers have also gotten better, and at such an insane rate that humans seem to have no chance of ever again catching up. From a theoretical standpoint, that's exactly what we might expect to happen with quantum supremacy experiments – simply because they involve tasks that have polynomial scaling for quantum computers, but exponential scaling for classical computers, where each additional qubit roughly doubles the resources required. In analyzing concrete quantum supremacy experiments, I'd say that the goal is not to pin down the exact factor by which they're beating this or that classical simulation (the answers will change rapidly anyway, and depend on all sorts of low-level details), but simply to figure out whether or not we've entered that phase transition.

What do you generally think about improvements in tensor network methods as a challenge to quantum supremacy and the recent simulation of the supremacy result in Beijing?

SCOTT: I blogged about this here: https://www.scottaaronson.com/blog/?p=5371 It was a clever paper, which showed that **if** you focus narrowly on spoofing the linear cross-entropy benchmark used by Google, then there's a classical algorithm to do that much faster than had been previously pointed out, by generating many samples that all share most of their bits in common (e.g., that all agree on the first 30 of the 53 bits). But many people remain unaware that, if you just changed the benchmark – for example, if you insisted that the returned samples not only pass the linear cross-entropy benchmark, but also be sufficiently different – then this classical spoofing strategy wouldn't work and we'd be back to the previous status quo.

Why do you need a random circuit to test the device, and also why is this more interesting/useful than doing something very specific many times to test the device?

SCOTT: The reasons to use random quantum circuits are simply that (1) they generate complicated entangled states on all the qubits nearly as rapidly as it's possible to do so – indeed, we now have theoretical results that give a detailed understanding of this process, and (2) random circuits seem to have about as little "usable structure" (which a classical simulation might exploit) as it's possible to have. Eventually, of course, we'd like to run actually **useful** circuits (say, those that arise in Shor's factoring algorithm), which will typically have regular patterns and be extremely far from random! But then more qubits will be needed to get an advantage over classical computers. It's not terribly surprising for the **first** advantage over classical to be via random quantum circuits, which in some sense "maximally exploit" the hardware resources available.

Sorry second question (which may yet be answered again): have you heard of/what do you think about the recent NP-verification experiment using a quantum optical setup from Paris? (Nature Comms)

SCOTT: That experiment was actually, based on a protocol that Beigi, Fefferman, Drucker, Shor, and I proposed back in 2008. It's crucial for people to understand that there's no claimed speedup here for **solving** any NP-complete problem. We're talking about a more arcane task: namely, proving to someone that an NP-complete problem has a solution by sending them a small number of qubits. Even there, the protocol crucially depends on the ability to send two states that are guaranteed NOT to be entangled with each other, and also, the communication savings is "only" polynomial rather than exponential (in our original protocol, roughly sqrt(n) qubits where n bits would have been needed). Nevertheless, it's always fun to see a real experiment implementing something that you worked out on paper, even if you already knew it would work!

If the difficulty for classical simulation is related to the Hilbert space dimension, has it been formally proven that an ideal analog classical computer cannot outperform a quantum computer?

SCOTT: This is not a question of "formal proof" but of physics. In my view, we already know deep reasons why, in our universe, analog classical computers are unlikely to be able to do anything that can't be efficiently simulated using a standard digital computer. (In other words, why analog classical computers don't violate the "Extended Church-Turing Thesis.") Those reasons have to do with nonlinear dynamics chaotically amplifying even the tiniest errors in an analog device. Or, if you really want to push this discussion to the bitter end, they have to do with the breakdown in our picture of a smooth spacetime that's expected to occur at the Planck scale, of ~10^{-33} centimeters and ~10^{-43} seconds, for reasons of black hole thermodynamics and quantum gravity. Crucially, neither of these issues apply to quantum computation. The former doesn't apply because of quantum error-correction and fault-tolerance, which have the effect of "discretizing" continuous errors; while the latter doesn't apply because as far as anyone knows today, quantum mechanics (unlike theories that assume a smooth spacetime) is exactly true. In a sense, the IBM thought-experiment simulation connects to this question we received: [In the comparison between quantum and classical computation], what do we mean by "classical resources" here? Do we mean something parochial like "resources obeying non-quantum laws that are available to us humans on Earth?" Or is any classical resource fair game, no matter its size and classical equations of motion? In that case, doesn't demonstrating quantum supremacy require demonstrating that the quantum computer exceeds the capabilities of, say, a classical computer the size of the solar system that exploits some CTC? If CTCs were possible, that would be a revolution in physics even greater than quantum mechanics! Leaving CTCs aside, though, cosmology and quantum gravity seem to impose a limit of roughly 10^{122} on the number of bits (and the number of operations on the bits) that any computer that fit inside the observable universe could possibly have. And if you envision that cosmological computer as a **classical** one, then it shouldn't take impossibly long for us to build quantum computers that can outperform it on some tasks: indeed, a device with ~400 qubits should already be enough! But of course we're not there yet: with 50-60 qubits, QCs right now are "merely" challenging the largest classical computers that are currently available on earth.

Why is a quantum state (superposition or entangled state) inherently more fragile than a classical state?

SCOTT: Because when information from the state (say, whether a qubit is 0 or 1, or which path a photon takes through a beamsplitter network) "leaks out" into the environment, the information effectively becomes **entangled** with the environment, which damages the state in a measurable way. Indeed, it now appears to an observer as a mixed state rather than a pure state, so that interference between the different components can no longer happen. This is a fundamental, justly-famous feature of quantum information that's not shared by classical information.

Given that we have noisy-intermediate scale quantum devices, what do you see as the fundamental role of noise in the discussion on quantum advantage or quantum supremacy. Is it simply a question of less noise is better, or are there things that cannot be done if there is noise?

SCOTT: There will always be **some** noise. The big question, about any given platform, is whether the noise is low enough that you can start usefully error-correcting it away, or whether the noise is so high that there's no point (i.e., whether you're above or below the "fault-tolerance threshold"). Until you start doing error-correction, most experts believe there's a severe limit to how far you can scale: probably to quantum computations involving a few hundred qubits at most. Whereas once you have error-correction, at least in principle the sky's the limit.

You used the Wright brothers as an example where an airplane that was not practically useful itself pioneered the path to useful flight. In what sense to the sampling experiments of Google and USTC also pioneer that path for what we need for the future of quantum computing, or to what extent do you see them as niche examples of quantum supremacy?

SCOTT: I think the majority of what Google did, in integrating 53 superconducting qubits, making them programmable, etc. – and especially in characterizing the noise in a system at that scale – will be directly useful going forward. Indeed, that's a large part of why they did! Likewise, a lot of what USTC did, in integrating hundreds of beamsplitters, photon sources, and photodetectors, could be directly relevant to building a universal optical quantum computer. On the other hand, it's true that both groups took shortcuts with the immediate goal of quantum supremacy in mind. As an example, Google's chip uses a particular 2-qubit gate that was chosen, not because it shows up naturally in any application, but simply because it's extra-hard to simulate using tensor network contraction algorithms, so it let them get to quantum supremacy faster than if they'd used a more conventional 2-qubit gate like the CNOT.

To what extent does the measured circuit fidelity of 0.2% in the Google experiment limit the usability of this system for other computations?

SCOTT: Oh, we don't know of **anything** particularly useful to do with Google's Sycamore chip – that is, anything that you couldn't do much more easily without it – other than (1) quantum supremacy demonstrations, (2) **possibly** the generation of cryptographically certified random bits, and (3) of course, calibration experiments that tell you about the behavior of integrated superconducting qubits and thereby help you iterate to the next device. But things are developing rapidly – the best circuit fidelity that was achievable in 2019 is not necessarily the best now, or the best that will be achieved in another year or two.