Why I dislike what “quantum supremacy” is doing to computing research

Enlarge / This optical setup is not your friend. Many fast switches and delay lines set up to arrange for 20 photons to enter the waveguide circuit simultaneously.

reader comments

60 with 46 posters participating, including story author

Share this story

The big quantum buzzword these days is “quantum supremacy.” (It’s a term I despise, even as I acknowledge that the concept has some utility. I will explain in a moment). Unfortunately, this means that some researchers have focused on quantum supremacy as an end in itself, building useless devices to get there.

Now, optical quantum computers have joined the club with a painstakingly configured device that doesn’t *quite* manage to demonstrate quantum supremacy. But before we get to the news, let’s delve into the world of quantum supremacy.

The quest for quantum supremacy

“Quantum supremacy” boils down to a failure of mathematics, combined with a fear that the well will run dry before we’ve drunk our fill.

In a perfect world, a quantum computer operates perfectly. In this perfect world, you can generate a mathematical proof that

shows a quantum computer will always outperform a classical computer on certain tasks, no matter how fast the classical computer is. Our world, however, is slightly less than perfect, and our quantum computers are not ideal, so these mathematical proofs might not apply.

As a result, to show performance advantages for quantum computers, we have to build an *actual* quantum computer that does something a classical computer can't. Unfortunately, reliable quantum computers were, until recently, limited to just a few quantum bits (qubits). Because of this bit scarcity, any problem solvable on a quantum computer could be solved much faster on a classical computer, simply because the problems were so small.

One solution, of course, is to make quantum computers with a larger number of qubits so that they can handle larger problems. Once that is achieved, quantum computers should be faster than classical computers—provided those tricky mathematical proofs hold for non-ideal quantum computers.

Putting together a general-purpose quantum computer with lots of qubits is easier said than done. Putting together a computer that can solve a single problem, however, is easier than building a general-purpose computer. Such a particular quantum computer allows engineers to show that a quantum computer is faster than a classical computer on *this single problem*, while allowing them to avoid the task of producing a generally useful computer. The ray of hope this success would provide might reassure people who control budgets, providing researchers with the funding necessary to turn a special-purpose (that is, not very useful) quantum computer into a general-purpose (that is, very useful!) quantum computer.

So, I accept the necessity of quantum supremacy—no one wants to pursue a dead end—but I still dislike what it's doing to current research. Now, on to the news…

Fire up the photon mixer

Google recently published a paper claiming to have demonstrated quantum supremacy (though not everyone entirely agreed). Now, another group of scientists has published a second attempt at the problem, this time using an optical quantum computer that, once again, computes nothing useful.

The scientists took a source of single photons and quickly switched the path of the output so that they had 20 different beams, each containing a single photon. The 20 photons simultaneously entered a set of waveguides, where they were repeatedly mixed with each other, creating 60 output beams. The photons repeatedly interfered with each other and influenced each other's path randomly. For any particular pulse of 20 photons, some output beams will be empty and others will have photons.

The researchers looked for simultaneous emission between outputs. Any particular measurement is a sample from a distribution of random numbers (a process called *bose sampling*). By repeating the experiment many times, statistics about the randomness found in photon interference are revealed.

For small experiments, a classical computer can compute the odds of every possible output combination directly. But, for 20 beams with 60 possible outputs, you end up with about 10^{14} output combinations. That number of potential outputs, and, therefore, of the experimental results demonstrated by the researchers, can still be directly calculated using a supercomputer—but would take several hours, which means we're getting close to the edge of what is possible.

Still, it's fast enough to outperform this particular quantum computer. The researchers detected a 14-photon coincidence only once every six hours, which means that accumulating the

statistics from the quantum system sill takes around a day.

If in doubt, add more photons

The future, though, is bigger. The number of inputs can be easily doubled (without doubling the number of beams), and the authors claim that the waveguide circuit can also be extended to accommodate more than 100 outputs. (That’s a key advantage of optical waveguide circuits: ease of scaling.)

Unfortunately, I am not sure that this will gain the researchers as much as they hope. The heart of the system is a source of single photons, which produces millions of photons per second. The speed of calculation, however, is limited by that rate. So, yes, the researchers can extend the number of possible output combinations to well beyond what can be calculated classically. Unfortunately, their photon source may also prevent them from sampling the distribution in a reasonable time.

On a more personal note: sometimes I start reading a paper, and then I stop and give thanks that I am not the grad student who had to *actually do that*. This is one of those papers. The experiment here is one crazy setup!

Physical Review Letters, 2019, DOI: 10.1103/PhysRevLett.123.250503 (About DOIs)

https://arstechnica.com/?p=1636139