Quanta magazine

Why I Called It 'Quantum Supremacy'

By **John Preskill**

October 2, 2019

Researchers finally seem to have a quantum computer that can outperform a classical computer. But what does that really mean?

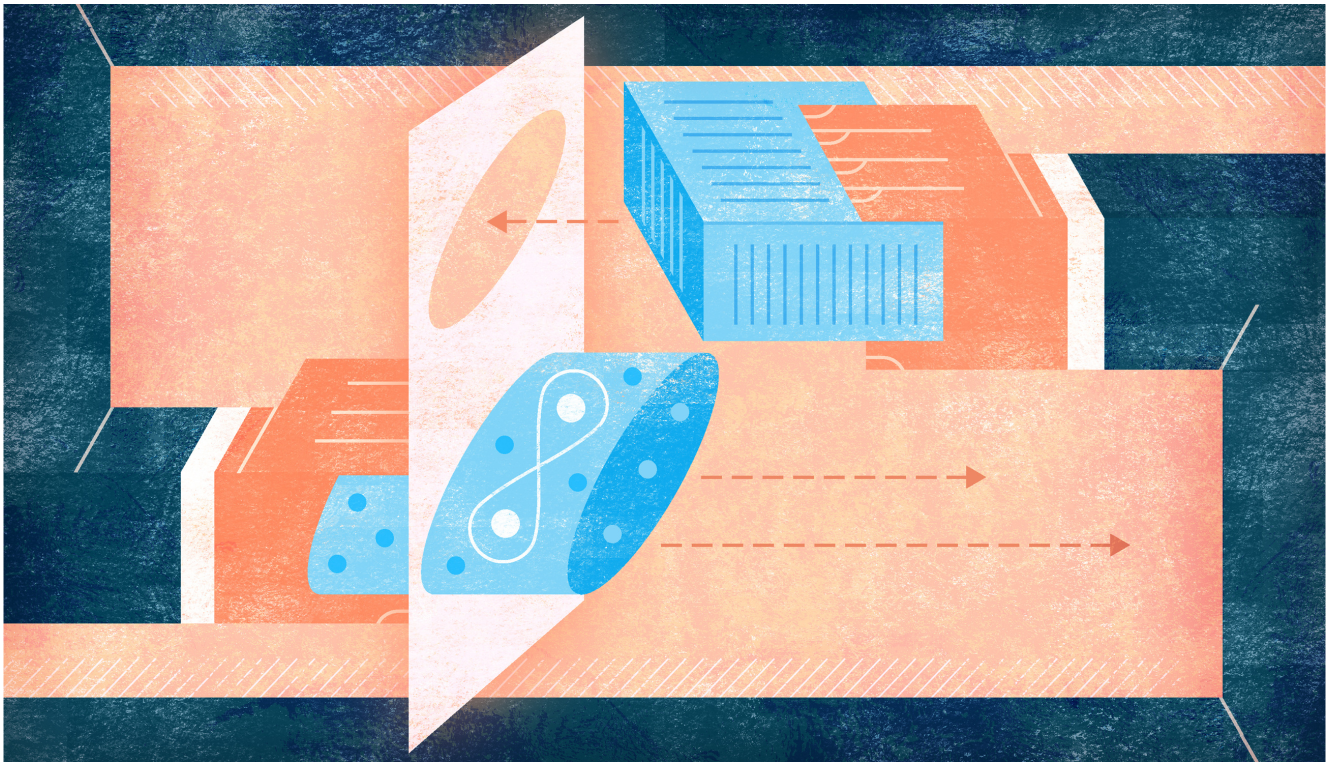

James O'Brien for Quanta Magazine

A recent paper from Google's quantum computing lab announced that the company had achieved quantum supremacy. Everyone has been talking about it, but what does it all mean?

In 2012, I proposed the term "quantum supremacy" to describe the point where quantum computers can do things that classical computers can't, regardless of whether those tasks are useful. With that new term, I wanted to emphasize that this is a privileged time in the history of our planet, when information technologies based on principles of quantum physics are ascendant.

The words "quantum supremacy" — if not the concept — proved to be controversial for two reasons. One is that supremacy, through its association with white supremacy, evokes a repugnant political stance. The other reason is that the word exacerbates the already overhyped reporting on the status of quantum technology. I anticipated the second objection, but failed to foresee the first. In any case, the term caught on, and it has been embraced with particular zeal by the Google AI Quantum team.

I considered but rejected several other possibilities, deciding that quantum supremacy best captured the point I wanted to convey. One alternative is "quantum advantage," which is also now widely used. But to me, "advantage" lacks the punch of "supremacy." In a race, a horse has an advantage if it wins by a nose. In contrast, the speed of a quantum computer vastly exceeds that of classical computers, for certain tasks. At least, that's true in principle.

The recent Google paper illustrates the point. They used a device with 53 qubits (the quantum analogues of a classical computer's bits), and they report that it took just minutes to perform quantum computations that would take today's most powerful supercomputers thousands of years. Assuming it's true, this is a remarkable achievement in experimental physics and a testament to the brisk pace of progress in quantum computing hardware; I offer my hearty congratulations to everyone involved.

Quantized Columns

A regular column in which top researchers explore the process of discovery. This month's columnist, John Preskill, is a professor of theoretical physics at the California Institute of Technology, and his main research area is quantum information science. You can follow him on Twitter @preskill.

See all Quantized Columns

The catch, as the Google team acknowledges, is that the problem their machine solved with astounding speed was carefully chosen just for the purpose of demonstrating the quantum computer's superiority. It is not otherwise a problem of much practical interest. In brief, the quantum computer executed a randomly chosen sequence of instructions, and then all the qubits were measured to produce an output bit string. This quantum computation has very little structure, which makes it harder for the classical computer to keep up, but also means that the answer is not very informative.

However, the demonstration is still significant. By checking that the output of their quantum computer agrees with the output of a classical supercomputer (in cases where it doesn't take thousands of years), the team has verified that they understand their device and that it performs as it should. Now that we know the hardware is working, we can begin the search for more useful applications.

Why is it so important to verify the performance of the hardware? It's because precisely controlling a quantum computer is notoriously difficult. In a sense, merely looking at a quantum system unavoidably disturbs it, a manifestation of Heisenberg's famous uncertainty principle. So if we want to use such a system to store and reliably process information, we need to keep that system nearly perfectly isolated from the outside world. At the same time, though, we want the qubits to interact with one another so we can process the information; we also need to control the system from the outside and eventually measure the qubits to learn the results of our computations. It is quite challenging to build a quantum system that satisfies all of these desiderata, and it has taken many years of progress in materials, fabrication, design and control to get where we are now.

The quantum supremacy milestone allegedly achieved by Google is a pivotal step in the quest for practical quantum computers. I thought it would be useful to have a word for the era that is now dawning, so I recently made one up: NISQ. (It rhymes with risk.) This stands for "noisy intermediate-scale quantum." Here "intermediate-scale" refers to the size of quantum computers that are now becoming available: potentially large enough to perform certain highly specialized tasks beyond the reach of today's supercomputers. "Noisy" emphasizes that we have imperfect control over the qubits, resulting in small errors that accumulate over time; if we attempt too long a computation, we're not likely to get the right answer.

The Google team has apparently demonstrated that it's now possible to build a quantum machine that's large enough and accurate enough to solve a problem we could not solve before, heralding the onset of the NISQ era.

Where do we go from here? Naturally, Google and other hardware builders hope to find practical applications for their quantum devices. A much larger quantum computer might help researchers design new materials and chemical compounds or build better tools for machine learning, but a noisy quantum computer with a few hundred qubits might not be capable of anything useful. Still, we have ideas about how to use NISQ computers that we're eager to try, which might yield better methods for optimization or more accurate physical simulations, but we're not sure if any of these will pan out. It will be fun to play with NISQ technology to learn more about what it can do. I expect that quantum computers will have a transformative effect on society, but this may still be decades away.

In the 2012 paper that introduced the term "quantum supremacy," I wondered: "Is controlling large-scale quantum systems merely *really, really hard*, or is it *ridiculously hard*? In the former case we might succeed in building large-scale quantum computers after a few decades of very hard work. In the latter case we might not succeed for centuries, if ever." The recent achievement by the Google team bolsters our confidence that quantum computing is merely really, really hard. If that's true, a plethora of quantum technologies are likely to blossom in the decades ahead.

This article was reprinted on Wired.com.