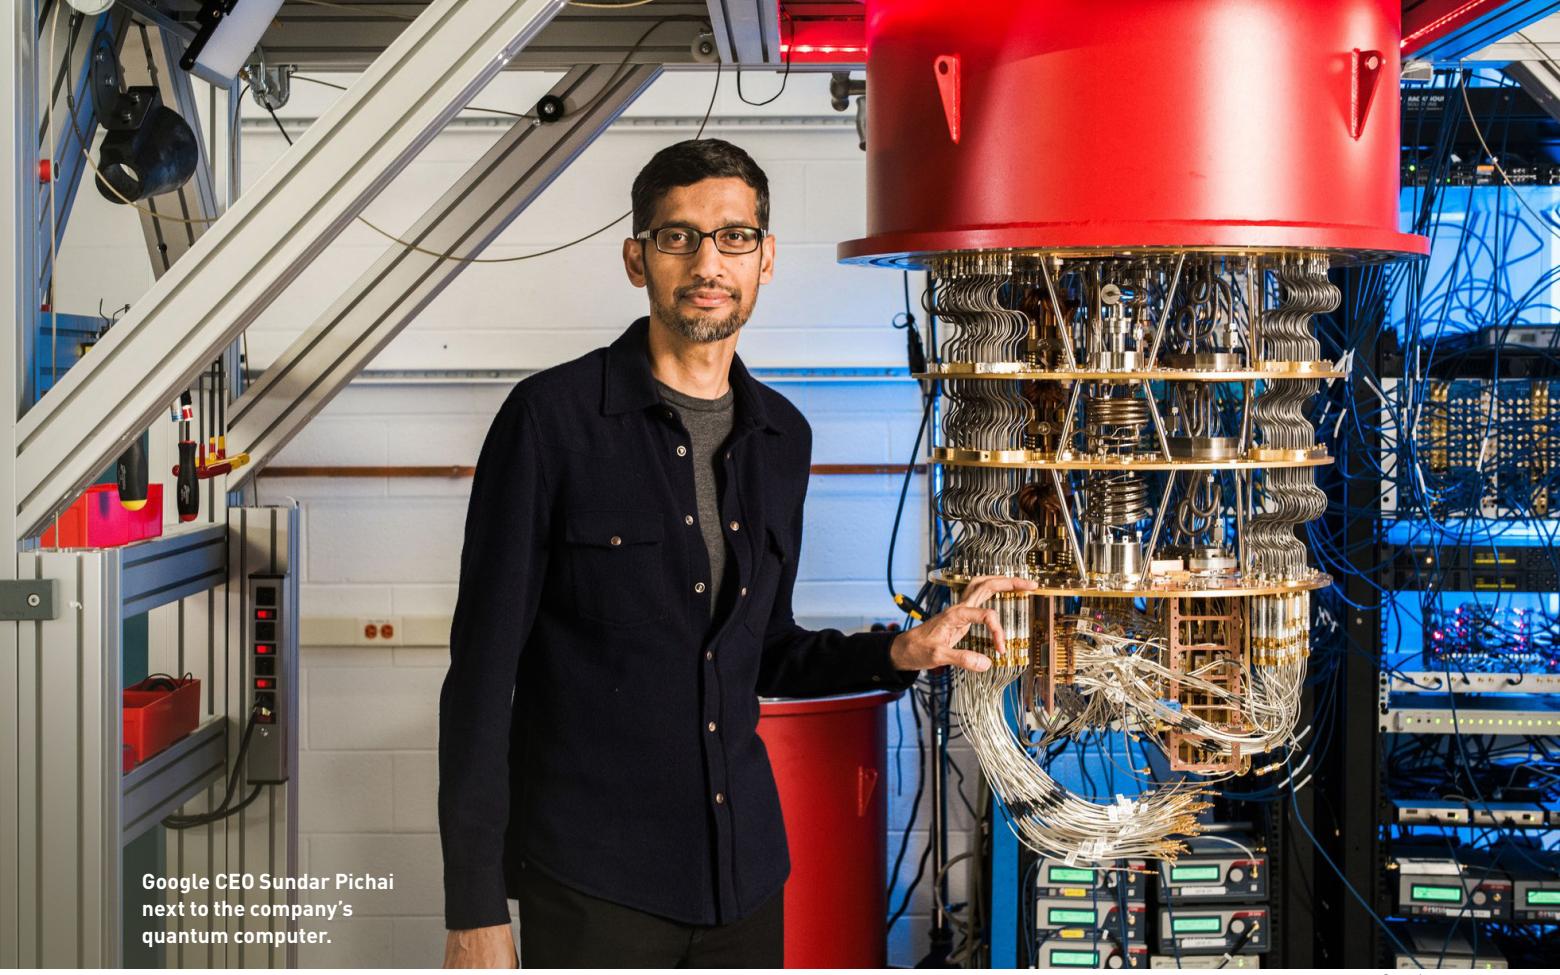




Edwin Cartlidge

Supreme or Unproven?

Despite much recent
fanfare, quantum
computers still need to
show that they can do
something useful.



Google CEO Sundar Pichai
next to the company's
quantum computer.

Google

Judging by the cover of *Nature* that day, 24 October 2019 marked a turning point in the decades-long effort to harness the strange laws of quantum mechanics in the service of computing. The words "quantum supremacy," emblazoned in large capital letters on the front of the prestigious journal, announced to the world that a quantum computer had, for the first time, performed a computation impossible to carry out on a classical supercomputer in any reasonable amount of time—despite having vastly less in the way of processors, memory and software to draw on.

The quantum computer in question, Sycamore, comprised a mere 53 superconducting quantum bits, or qubits. It was built by a group of scientists at Google led by physicist John Martinis, who used it to execute an algorithm that generated a semi-random series of numbers. Those researchers then worked out how long they would have needed to simulate that operation on the IBM-built Summit supercomputer at Oak Ridge National Laboratory in Tennessee, USA, the processors of which include tens of trillions of transistors and which has 250,000 terabytes of storage.

Amazingly, Martinis and colleagues concluded that what Sycamore could do in a little over three minutes, Summit would take 10,000 years to simulate.

Long-sought milestone?

For many scientists, Sycamore's result represents a major milestone on the road to a real-world, general-purpose quantum computer. Having invested millions of dollars in the field over the course of more than 30 years, governments and, increasingly, industry have bet that the exponential speed-up in processing power offered by quantum states in theory can be realized practically.

Winning that bet, however, depends on being able to protect a quantum computer's delicate superposition states from even the smallest amounts of noise, such as tiny temperature fluctuations or minuscule electric fields. The Google result shows that noise can be controlled sufficiently to enable the execution of a classically difficult algorithm, according to Greg Kuperberg, a mathematician at the University of California, Davis, USA. "This advance is a major blow against arguments that quantum computers are impossible," he says. "It is a tremendous confidence builder for the future."

Not everyone, however, is convinced by the research. A number of experts, including several at IBM, believe that the Google group has seriously underestimated the capacity of traditional digital computers to simulate the kind of algorithms that could be run on Sycamore.

For many scientists, Sycamore's result represents a major milestone on the road to a real-world, general-purpose quantum computer. Not everyone, however, is convinced by the research.

More fundamentally, it still remains to be seen whether scientists can develop a quantum algorithm that is resilient to noise and that does something people are willing to pay for—given how little practical utility the current algorithm is likely to have.

"For me, the biggest value in the Google research is the technical achievement," says Lieven Vandersypen, who works on rival quantum-dot qubits at the Delft University of Technology in the Netherlands. He points out that the previous best superconducting computer featured just 20 quite poorly controlled qubits. "But what we in the field are after is a computer that can solve useful problems, and we are still far from that."

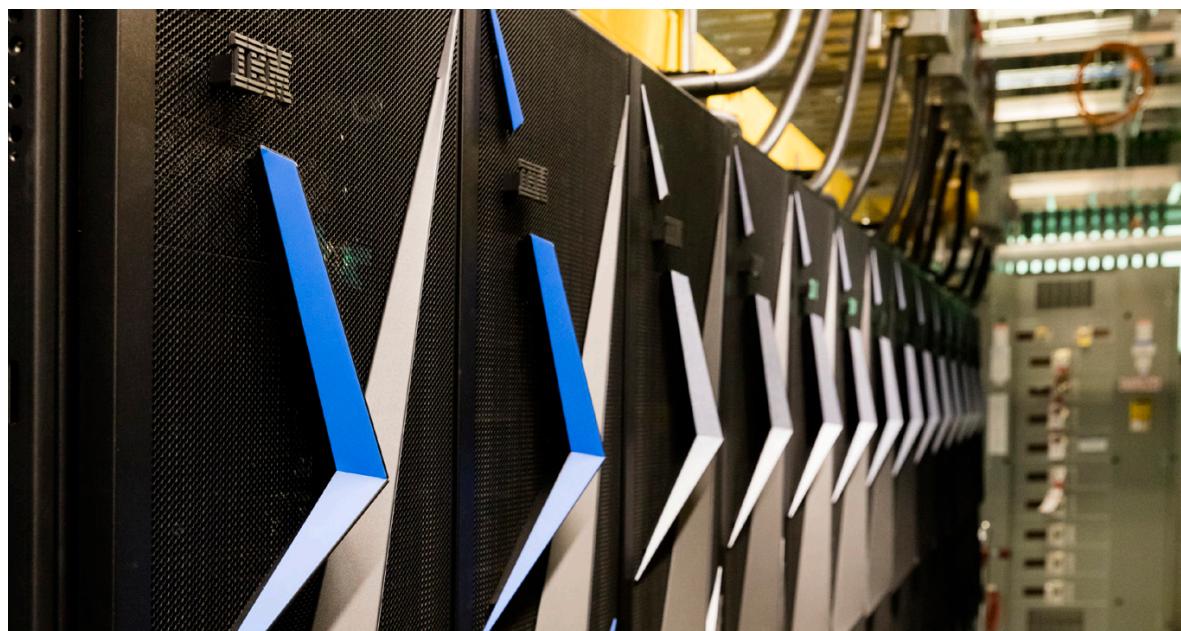
Quantum's power

Quantum computers offer the possibility of carrying out certain tasks far more quickly than is possible with classical devices, owing to a number of bizarre properties of the quantum world. Whereas a classical computer processes data sequentially, a quantum computer

should operate as a massively parallel processor. It does so thanks to the fact that each qubit—encoded in quantum particles such as atoms, electrons or photons—can exist in a superposition of the "0" and "1" states, rather than simply one or the other, and because the qubits are linked together through entanglement.

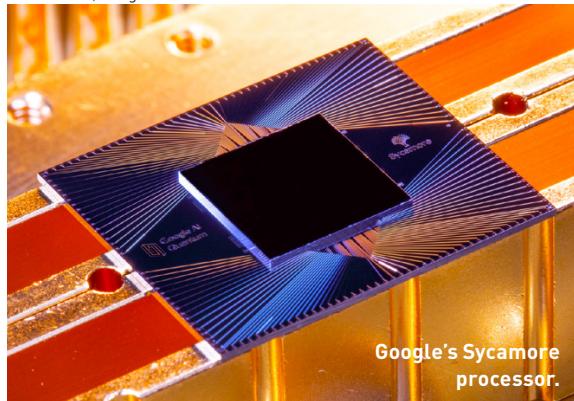
For N qubits, each of the 2^N possible states that can be represented has an associated amplitude. The idea is to carry out a series of operations on the qubits, specified by a quantum algorithm, such that the system's wave function evolves in a predetermined way, causing the amplitudes to change at each step. When the computer's output is then obtained by measuring the value of each qubit, the wave function collapses to yield the result.

The Google experiment, carried out in company labs in Santa Barbara, CA, USA, was designed to execute an algorithm whose answer could only be found classically by simulating the system's wave function. So while running the algorithm on a quantum computer would only take as long as is needed to execute



The IBM-built Summit supercomputer at the Oak Ridge National Laboratory, USA, contains tens of trillions of transistors and can carry out about 200,000 trillion operations a second. ORNL

Erik Lucero, Google



Sycamore—A quantum chip bearing fruit

Google's Sycamore processor consists of a 1-cm² piece of aluminum containing a 2D array of 53 qubits—each acting as a tiny superconducting resonator that encodes the values 0 and 1 in its two lowest energy levels, and coupled to its four nearest neighbors. Cooled to below 20 mK to minimize thermal interference, the qubits are subject to “gate” operations—having their coupling turned on and off, as well as absorbing microwaves and experiencing variations in magnetic flux.

The Google team executed a series of cycles, each involving a random selection of one-qubit gates and a specific two-qubit gate. After completing the last cycle, they then read out the value of each qubit to yield a 53-bit-long string of 0s and 1s. That sequence appears random, but quantum entanglement and interference dictate that some of the 2^{53} permutations are much more likely to occur than others. Repeating the process a million times builds up a statistically significant number of bit strings that can be compared with the theoretical distribution calculated using a classical computer.

Measuring Sycamore’s “fidelity” to the theoretical distribution over 14 cycles, Martinis and coworkers found that the figure, 0.8%, agreed with calculations based on the fidelities of individual gates—and used that fact to estimate that after 20 cycles, the fidelity would have been about 0.1% [as the fidelity is gradually eroded by gate errors]. At this level of complexity and fidelity, the team calculated, the classical Summit supercomputer would require a whopping 10,000 years to simulate the quantum wave function—whereas Sycamore needed a mere 200 seconds to take its 1 million samples.

its limited number of steps, simulating that algorithm classically would involve tracking the 2^N probability amplitudes. Even with just 53 qubits that is an enormous number— 9×10^{15} , or 9,000 trillion.

Sycamore is not the first processor to have harnessed quantum interference to perform a calculation considered very difficult, if not impossible, to do using a classical computer. In 2017, two groups in the U.S. each used about 50 interacting, individually controllable qubits to simulate collections of quantum spins. Christopher Monroe and colleagues at the University of Maryland, College Park, manipulated electrically trapped ions using laser pulses, while OSA Fellow Mikhail Lukin of Harvard University and coworkers used a laser to excite neutral atoms. Both groups used their devices to determine the critical point at which a magnetic-phase transition occurs.

However, these systems were designed to carry out very specific tasks, somewhat akin to early classical analog computers. Google's processor, in contrast, is a programmable digital machine. By employing a handful of different logic gates—specific operations applied either to one or two qubits—it in principle can execute many types of quantum algorithms.

Martinis and colleagues showed that they could use these gates to reliably generate a sample of numbers from the semi-random algorithm. Crucially, they found that they could prevent errors in the gates from building up and generating garbage at the output—leading them to declare that they had achieved quantum supremacy.

“We are thrilled,” says Martinis, who is also a professor at the University of California, Santa Barbara. “We have been trying to do this for quite a few years and have been talking about it, but of course there is a bit of pressure on you to make good on your claims.”

Classical shortcuts

When the Google team published its results—a preliminary version of which had been accidentally posted online at NASA a month earlier—rivals lost little time in criticizing them. In particular, researchers at IBM, which itself works on superconducting qubits, posted a paper on the arXiv server arguing that Summit could in fact simulate Sycamore's operations in just 2.5 days (and at higher fidelity). Google's oversight, they said, was to not have considered how much more efficiently the supercomputer could track the system's wave function if it fully exploited all of its hard disk space.

Turning NISQ computers into practical devices will need more than improvements in hardware. Developing new algorithms that can exploit these devices for commercial ends will be essential.

Kuperberg argues that Sycamore's performance still merits the label "supremacy" given the disparity in resources available to the two computers. (In fact, the IBM researchers didn't actually carry out the simulation, possibly because it would have been too expensive.) Kuperberg adds that with just a dozen or so more qubits, the simulation time would climb from days to centuries. "If this is what passes as refutation, then this is still a quantum David versus a classical Goliath," he says. "This is supremacy enough as far as I am concerned."

Indeed, in their paper Martinis and colleagues write that while they expect classical simulation techniques to improve, they also expect that "they will be consistently outpaced by hardware improvements on larger quantum processors." Others, however, suggest that quantum computers might struggle to deliver any meaningful speed-up over classical devices. In particular, argue critics, it remains to be seen just how "quantum mechanical" future quantum computers will be—and therefore how easy it might be to imitate them.

To make classical simulation more competitive, the IBM researchers, as well as counterparts at the Chinese tech company Alibaba, are looking to make better use of supercomputer hardware. But Graeme Smith, a theoretical physicist at the University of Colorado and the JILA research institute in Boulder, USA, thinks that more radical improvement might be possible. He argues that the noise in Google's gates, low as it is, could still swamp much of the system's quantum information after multiple cycles. As such, he reckons it may be possible to develop a classical algorithm that sidesteps the need to calculate the 53-qubit wave function. "There is

nothing to suggest that you have to do that to sample from [Google's] circuit," he says.

Indeed, Itay Hen, a numerical physicist at the University of Southern California in Los Angeles, USA, is trying to devise a classical algorithm that directly samples from the distribution output by Google's circuit. Although too early to know whether the scheme will work, he says it would involve calculating easy bits of the wave function and interfering them to generate a succession of individual data strings very quickly. "I am guessing that lots of other people are doing a similar thing," he adds.

As Hen explains, Martinis and colleagues had to make a compromise when designing their quantum-supremacy experiment—making the circuit complex enough to be classically hard, but not so complex that its output ended up being pure noise. And he says that the same compromise faces all developers of what



Collaborating scientists from Intel and QuTech at the Delft University of Technology with Intel's 17-qubit superconducting test chip. Courtesy of Intel Corp.



A look at optics

As qubits, photons have several virtues. Because they usually don't interact with one another they are immune to stray electromagnetic fields, while their high energies at visible wavelengths make them robust against thermal fluctuations—removing the need for refrigeration. But their isolation makes them tricky to manipulate and process.

Two startups are working to get around this problem—and raising tens of millions of dollars in the process. PsiQuantum in Palo Alto, CA, USA, aims to make a chip with around 1 million qubits. Because photons are bosons and tend to stick together, their paths combine after entering 50-50 beam splitters from opposite sides, effectively interacting. Xanadu in Toronto, Canada, instead relies on the uncertainty principle, generating beams of “squeezed light” that have lower uncertainty in one quantum property at the expense of greater uncertainty in another. In theory, interfering these beams and counting photons at the output might enable quantum computation.

Both Xanadu and PsiQuantum have major, if different, technical hurdles to overcome before their computers become reality, according to OSA Fellow Michael Raymer, an optical physicist at the University of Oregon, USA, and a driving force behind the U.S. National Quantum Initiative.

Raymer adds that photons might also interact not directly, but via matter intermediaries, potentially enabling quantum-logic operations between single photons. Or they might be used to link superconducting processors to slower but longer-lived trapped-ion qubits (acting as memory). Alternatively, photon-matter interactions could be exploited in the quantum repeaters needed to ensure entanglement between distant particles—potentially a boost for both communication and sensing.

“Whether or not optics will be used to create free-standing quantum computers,” says Raymer, “I will defer prediction on that.”

is hoped will become the first generation of useful quantum computers—a technology known as “noisy intermediate-scale quantum,” or NISQ.

Such devices might consist of several hundred qubits, perhaps allowing them to simulate molecules and other small quantum systems. This is how Richard Feynman, back in the early 1980s, originally envisaged quantum computers being used—conceivably allowing scientists to design new materials or develop new drugs. But as their name suggests, these devices, too, would be limited by noise. The question, says Hen, is whether they can be built with enough qubits and processor cycles to do something that a classical computer can't.

Dots, ions and photons

To try and meet the challenge, physicists are working on a number of competing technologies—superconducting circuits, qubits encoded in nuclear or electronic spins, trapped atoms or ions—each of which has its strengths and weaknesses (see OPN, October 2016, p. 30). Vandersypen, for instance, is hopeful that spin qubits made from quantum dots—essentially artificial atoms—can be scaled up. He points out that such qubits have been fabricated in an industrial clean room at the U.S. chip giant Intel, which has teamed up with him and his colleagues at the Delft University of Technology to develop the technology. “We have done measurements [on the qubits],” he adds, “but not yet gotten to the point of qubit manipulation.”

Trapped-ion qubits, meanwhile, are relatively slow, but have higher fidelities and can operate more cycles than their superconducting rivals. Monroe is confident that by linking up multiple ion traps, perhaps optically, it should be possible to make NISQ devices with hundreds of qubits. Indeed, he cofounded the company IonQ with OSA Fellow Jungsang Kim from Duke University, USA, to commercialize the technology.

A completely different approach is to encode quantum information in light rather than matter. Photonic qubits are naturally resistant to certain types of noise, but being harder to manipulate they may ultimately be more suited to communication and sensing rather than computing (see “A look at optics,” left).

Yet turning NISQ computers into practical devices will need more than just improvements in hardware, according to William Oliver, an electrical engineer and physicist at the Massachusetts Institute of Technology,

The only realistic aim, Hen argues—and one that all experts see as the ultimate goal of quantum computing—is to build large, fault-tolerant machines.

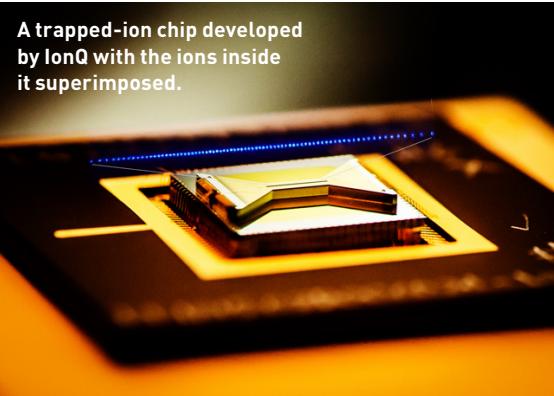
USA. Also essential, he says, will be developing new algorithms that can exploit these devices for commercial ends—be those ends optimizing investment portfolios or simulating new materials. “The most important thing,” Oliver says, “is to find commercial applications that gain advantage from the qubits we have today.”

According to Hen, though, it remains to be seen whether any suitable algorithms can be found. For simulation of chemical systems, he says, it is not clear if even hundreds of qubits would be enough to reproduce the interactions of just 40 electrons—the current classical limit—given the inaccuracies introduced by noise. Indeed, Smith is pessimistic about NISQ computers being able to do anything useful. “There is a lot of hope,” he says, “but not a lot of good science to substantiate that hope.”

Erring on the side of caution

The only realistic aim, Hen argues—and one that all experts see as the ultimate goal of quantum computing—is to build large, fault-tolerant machines. These would rely on error correction, which involves spreading the value of a single “logical qubit” over multiple physical qubits to make computations robust against errors on any specific bit (since quantum information cannot simply be copied). But implementing error correction will require that the error rate on individual qubits and logic gates is low enough that adding the error-correcting qubits doesn’t introduce more noise into the system than it removes.

Vandersypen reckons that this milestone could be achieved in as little as a year or two. The real challenge, he argues, will be scaling up—given how many qubits are likely to be needed for full-scale fault-tolerant computers. Particularly challenging will be making a machine that can find the prime factors of huge numbers, an application put forward by mathematician Peter Shor in 1994 that could famously threaten internet encryption. Martinis himself estimates that a device capable of finding the prime factors of a 2000-bit number in a day would need about 20 million physical qubits, given a two-qubit error probability of about 0.1%.



Kai Hudek, IonQ, Inc., and Emily Edwards, JQI and the University of Maryland

Despite the huge challenges that lie ahead, Martinis is optimistic about future progress. He says that he and his colleagues at Google are aiming to get two-qubit error rates down to 0.1% by increasing the coherence time of their qubits—doubling their current value of 10–20 microseconds within six months, and then quadrupling it in two years. They then hope to build a computer with 1,000 logical qubits within 10 years—a device that he says wouldn’t be big enough to threaten internet security but could solve problems in quantum chemistry. “We are putting together a plan and a timeline and we are going to try to stick to that,” he says.

However, Oliver is skeptical that such an ambitious timeframe can be met, estimating that a full-scale fault-tolerant computer is likely to take “a couple of decades” to build. Indeed, he urges his fellow scientists not to overstate quantum computers’ near-term potential. Otherwise, he fears, the field could enter a “quantum winter” in which enthusiasm gives way to pessimism and the withdrawal of funding. “A better approach,” according to Oliver, “is to be realistic about the promise and the challenges of quantum computing so that progress remains steady.” **OPN**

Edwin Cartlidge (edwin.cartlidge@yahoo.com) is a freelance science journalist based in Rome.

For references and resources, go online:
www.osa-opn.org/link/quantum-supremacy.