FEATURE 1 November 2017

Quantum code: why building the ultimate computer is the easy bit

After decades of hype, quantum computers are poised to prove their superiority over classical machines. Now the race is on to figure out what to do with them

Scott Garrett

By **Michael Brooks**

REMEMBER when Apple launched the iPhone and you first heard people prattling on about things called apps? Back then, there were just a few hundred applications to choose from – a shortage that looked a lot like an opportunity and promptly gave rise to the newfangled job title of "app designer". These days you can select from more than 2 million iPhone apps – yet more proof of the device's runaway success.

Consider now quantum computers, those much-vaunted dream machines that would use the strange laws of quantum physics to make light work of the very hardest problems. They have long been among the most frustrating of all the world-changing technologies we've been promised: perpetually just around the corner. But in the last year or so something has changed. The industry giants racing to build such a machine

quietly started recruiting app designers, which suggests the long quest to make good on the hype about quantum computers might have entered a new stage.

The fact that researchers at Google, IBM, Microsoft and a host of other organisations are even building prototypes shows how far we have come. What's truly exciting, though, is that by challenging a new generation of programmers to go quantum, they are now tackling a question that has largely been brushed under the carpet: when we build the ultimate problem-solving machine, what are we going to do with it?

The quantum leap in computing has been a long time coming. It was first conceived in the 1980s, when theorists predicted that a computer based on quantum effects could vastly outperform classical computers at certain tasks. The trick would be to harness superposition, a quantum property that means particles can exist in several different configurations at once, and entanglement, which lets all the particles work together, to create "massively parallel" processing.

Qubit by qubit

Whereas classical computers encode information as bits that can be in one of two states, 0 or 1, quantum bits, or qubits, can be simultaneously 0 and 1 thanks to superposition. With enough qubits linked together through entanglement, you should be able to do way more calculations at once, resulting in exponentially faster computing.

So far, though, that's just a theory. No one has ever built something capable of properly testing it. "Everyone is assuming that the power of a quantum computer arises from parallel processing," says John Martinis at the University of California, Santa Barbara. "It is crucial to actually check."

And that's exactly what everyone, including Martinis, is now hoping to do. Having worked on quantum computing for decades, he and his team had already made great strides using superconducting qubits – ultra-cold loops of superconducting metal with quantum properties. Then, in 2014, Google bought the lot, and its financial heft has further accelerated progress.

Earlier this year, Martinis announced that his group has been testing a 20-qubit processor. They are now putting the finishing touches to a 49-qubit version with which they plan to do something that no classical computer ever could. This milestone, known as "quantum supremacy", would finally prove the theory and give the field a fillip.

In June, Google said the big moment would arrive by the end of 2017, but now Martinis says it will be "some time into next year".

Whenever it happens, though, supremacy isn't everything. It's one thing to outdo regular computers at predicting the output of chaotic electronic circuits, which is Google's aim; it's quite another to do something practically useful. For that you need to wrangle an increasing number of qubits. This quickly gets devilishly tricky because the more qubits you have, the more difficult it is to stop their fragile quantum states from falling apart.

In practice, qubits are vulnerable to noise and other kinds of disturbance. Vibrations or thermal energy can snap the entanglement or the superposition. This introduces errors in the information they carry, which have to be corrected using networks of more

Ultra-cold loops of superconducting metal are at the heart of most quantum computers
IBM Research

qubits. So an information-carrying "logical qubit" has to be made up from an array of error-correcting "physical qubits", which makes it meaningless to boast that you have a machine with 50 physical qubits if they are so noisy that they can only do computations suited to three logical qubits. "If we want one logical qubit, we need a 2D grid of physical qubits, and have to perform a bunch of measurements on them to keep track of all the errors that are happening," says James Wootton at the University of Basel, Switzerland.

This scaling-up problem looks particularly tricky for Google, according to Winfried Hensinger, who is building his own quantum processor at the University of Sussex, UK. Superconducting qubits have to be cooled to within a whisker of absolute zero, making the computer itself bulky and somewhat impractical in terms of building huge arrays of entangled qubits. "It's a nice approach to get to 100 qubits, but the thought of building a large-scale machine seems very challenging," he says.

Hensinger thinks his approach is more easily scalable. His qubits are charged atoms, or ions, held in magnetic traps, which operate at room temperature. The input and output happens via just a few microwave fields, no matter how many ions are involved in the computation. That means manipulating the qubits should be simple, even for a large number of qubits. His group is two years away from bringing everything together into a single prototype that will have somewhere between 10 and 50 qubits, he says.

Hensinger is not the only one pursuing a different approach. Chris Monroe's group at the University of Maryland is also working with trapped ions, and researchers at the University of South Wales in Sydney, Australia, have qubits made from atoms of phosphorus embedded in a silicon lattice. According to Michelle Simmons, who leads the research, these have the advantage of remaining error-free a million times longer than has been achieved with superconducting qubits. "The noise in our device is

incredibly low," she says.

What's more, the semiconductor industry is so used to working with silicon that scaling up production will be much easier than with other technologies, says Simmons. "We believe that if you can build a quantum computer, this is the best way to do it."

Then there's the wild card, an exotic approach to quantum computing that would pretty much sidestep the scaling and error correction obstacles – although some say it's wishful thinking. Researchers at Microsoft's Station Q laboratory in California are building topological qubits based on the properties of particles called non-abelian anyons, which can encode quantum information in the intricate way they move past one another. Microsoft has recently employed Leo Kouwenhoven from the University of Delft in the Netherlands, the first person to claim to have created them. However, not everyone is convinced that Kouwenhoven has really made topological qubits – and some doubt they can ever actually exist.

"A wild card approach sidesteps the problem, but might be wishful thinking"

Even so, the topological quantum computer is worth pursuing, says Station Q's Todd Holmdahl. The anyons' quantum states are fixed by their relative trajectories in the past, he says, and therefore not held in properties like spin or charge, which can easily be disturbed. This means the architecture is relatively error-resistant and resilient to external influence. "You're doing all your error correction down at the hardware," he says. That also means fewer qubits are necessary to achieve useful quantum information processing.

Although Holmdahl is happy to admit that his group is bringing up the rear for now, he expects a surge towards the finish line just as his rivals are flagging. "There are people running the marathon already, but they're all wearing army boots or hip waders," he says. "We're still sitting there putting on our running shoes, but once we've got them on, we're going to go much faster."

So who should you put your money on? "The truth is, we don't really know yet," says Scott Aaronson of The University of Texas, Austin. "One approach could pull ahead of everyone else, or multiple approaches could succeed, or it could even require a hybrid of multiple approaches to get all the way there."

However it pans out, the front runners are increasingly confident. "Given the momentum we have, and how many smart people want to get involved, we are poised to do something amazing," says Jerry Chow, who leads IBM's quantum computing venture.

What, exactly, is not at all clear. And herein lies the problem with quantum computers: we don't know what to do with them.

For starters, killer apps are not as abundant as you might assume. Physicists have known for decades that quantum computers could solve particular kinds of problems: optimisation, which involves finding the lowest or highest point in a landscape of possibilities, is one. But it's not revolutionary; just a little bit better than what we can do with classical machines. Another is "backwards search", where a huge, unsorted database can be searched faster than is possible with a classical computer. Then there is the factorisation of large numbers (see "Quantum codebreakers"), but even this

requires a huge number of qubits to do significantly better than a classical computer. In effect, these are solutions looking for an application. To fulfil the immense potential of quantum computers, what we need now is applications requiring a quantum solution.

We do have ideas. Simulating other quantum objects, such as a chemical molecules, was one of the first suggestions. The hope is that drug discovery will eventually be radically accelerated by quantum processors. IBM's seven qubit quantum processor – which is also based on superconducting loops – was able to calculate the lowest energy state of a three-atom molecule, beryllium hydride (BeH_2) – a key step towards understanding the full range of reactions the molecule will undergo. This is not a breakthrough in the quantum supremacy vein; the same calculation can be done on a classical machine. But it is a step towards much bigger calculations when more qubits become available.

IBM's quantum computing lab has made its 16-qubit chip (below) available to programmers
IBM Q

IBM has also made progress towards quantum versions of machine learning. This field, which lies at the heart of artificial intelligence, relies on identifying patterns in vast piles of data. The team has shown that quantum processors are potentially much better than classical ones when searching for patterns in noisy data.

But the truth is that many uses lie beyond our puny imaginations. "I think we are only beginning to scratch the surface of possible applications," says Joe Fitzsimmons, a quantum programmer at Singapore University of Technology and Design.

And even with a series of problems ripe for the quantum-treatment in hand, there is something missing: operating instructions. Algorithms are sets of operations performed to solve a problem. We have all manner of them for ordinary, classical computers. But quantum computers work in fundamentally different ways, so we need new algorithms if we want to take advantage of the massively parallel processing they make possible.

"The people leading the race are wearing army boots or hip

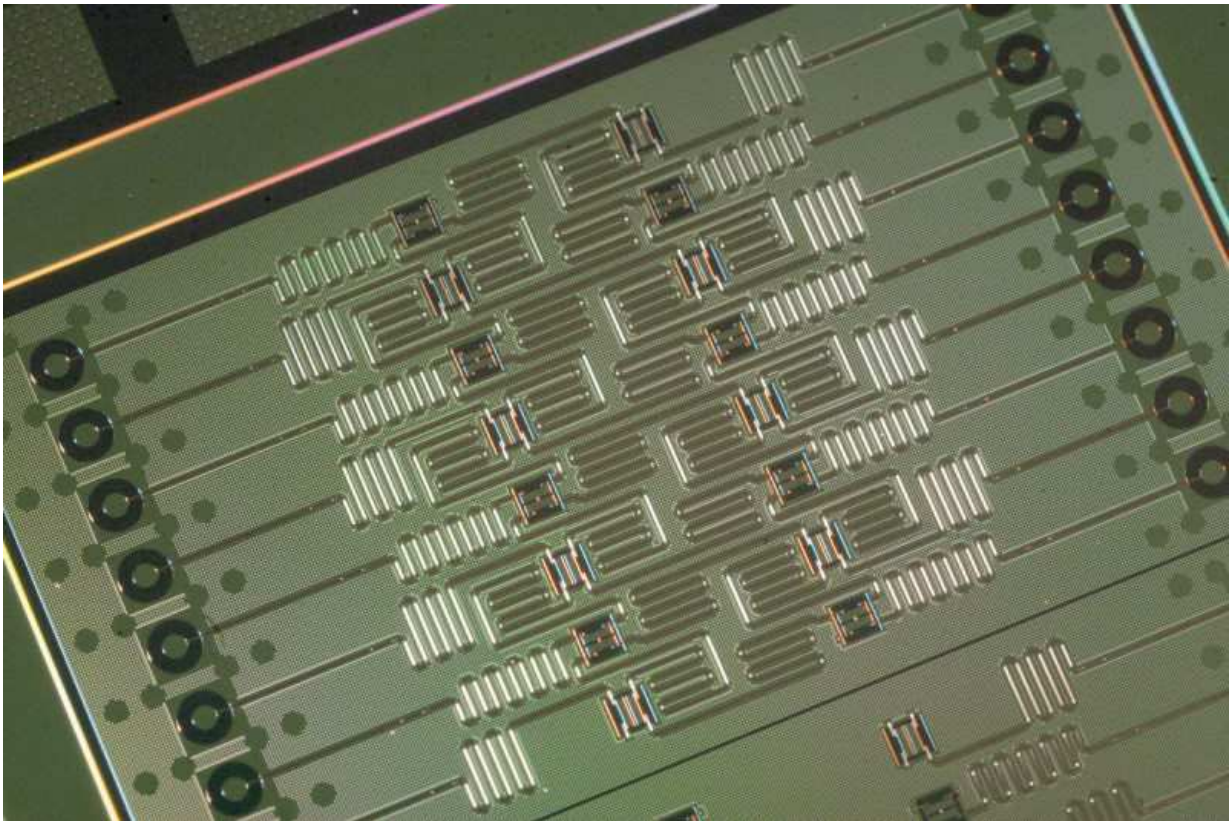

IBM Research

waders”

As things stand, there are precious few useful ones. In fact, we are roughly where programming for classical computers was in the 1950s: you have to understand what the processor hardware is actually doing – physically – to manipulate the qubits, and so you have to talk to them in the quantum version of the 1s and 0s that we feed to regular computers. “It’s extremely difficult at the moment,” Fitzsimmons says.

The big guns are aware of the problem. But real progress will depend on making these new devices accessible and appealing to people with no quantum experience. It’s something that Wootton is pursuing, especially now IBM has made its five-qubit computer accessible. “You can play with real quantum computers now,” he says.

In 2011, Wootton designed a quantum algorithm to simulate the behaviour of the exotic anyons that Microsoft hopes to put to work. He has since implemented it with real qubits, and now he is exploring algorithms that can be implemented with IBM’s latest 16-qubit device.

The experience has made him an evangelist for the joy of quantum programming, and he is actively recruiting people to the cause because he believes the killer ideas will not come from insiders. “We’re used to thinking in a certain way, and people now accessing the field are not thinking that same way,” says Wootton. “When the right people get interested we’ll see some great things happening on these devices.”

What are quantum computers good for? That might be up to you.

Quantum Codebreakers

In 2016, after decades of watching and waiting, the US National Security Agency finally decided that quantum computing is a serious threat. That's because quantum processors have the potential to render all our tricks for protecting online transactions, securing financial systems and email encryption as useful as a chocolate strongbox.

The cryptographic codes behind these systems are based on a mathematical oddity: that there is no known algorithm for efficiently finding the prime factors of a large number. Factors are smaller numbers that multiply together to make a larger one. All you can do is try various combinations, one by one.

But there is an algorithm for a quantum computer that could. Peter Shor concocted one in 1994 that could efficiently find factors of large numbers. That's not an immediate problem, because the Shor algorithm requires hundreds, if not thousands, of qubits to be any use – and current machines only have a handful at best (see main story). But there is no room for complacency.

"NSA does not know if or when a quantum computer of sufficient size to exploit public key cryptography will exist," said the recent NSA document, issued to encourage businesses to consider quantum cryptography.

This article appeared in print under the headline "The problem with the ultimate problem-solving machine"

Michael Brooks is a consultant for *New Scientist*