NEWS FEATURE

THE QUANTUM COMPANY

D-Wave is pioneering a novel way of making quantum computers — but it is also courting controversy.

BY NICOLA JONES

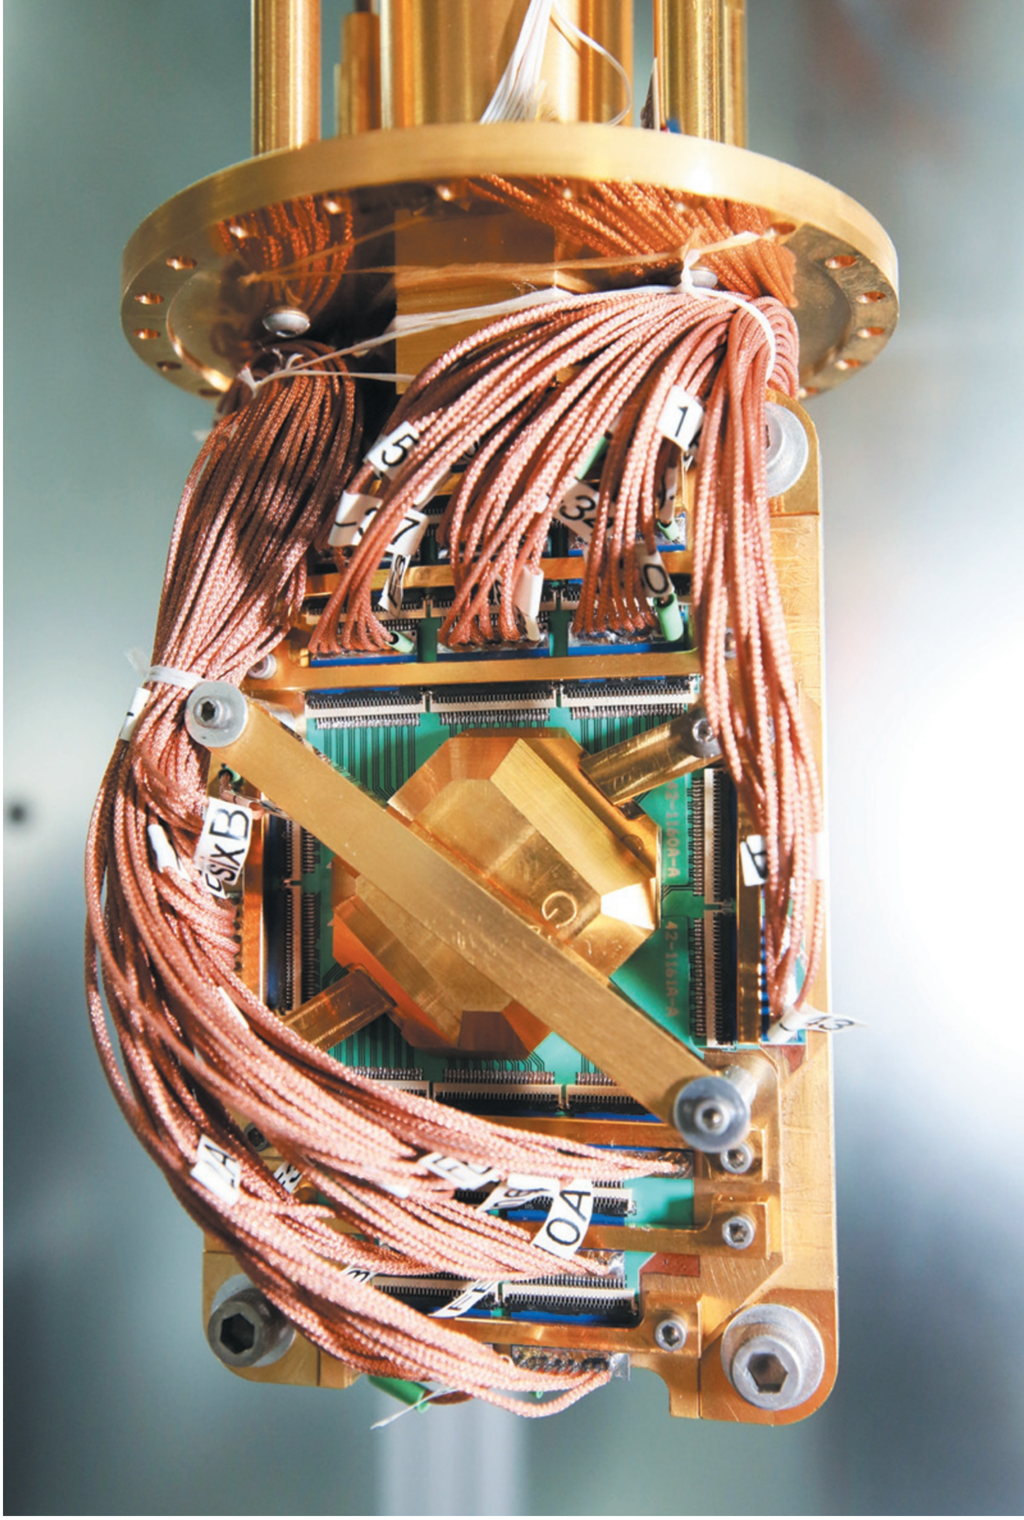

The D-Wave quantum computer processor is 3,600 times faster than classical computers at some tasks.

KIM STALLKNECHT/THE NEW YORK TIMES/REDUX/EYEVINE

“I've been doing combative stuff since I was born,” says Geordie Rose, leaning back in a chair in his small, windowless office in Burnaby, Canada, as he describes how he has spent most of his life making things difficult for himself. Until his early 20s, that meant an obsession with wrestling — the sport that, he claims, provides the least reward for the most work. More recently, says Rose, now 41, “that's been D-Wave in a nutshell: an unbearable amount of pain and very little recognition”.

The problem of lack of recognition is fast disappearing for D-Wave, the world's first and so far only company making quantum computers. After initial disbelief and ridicule from the research community, Rose and his firm are now being taken more seriously — not least by aerospace giant Lockheed Martin, which bought one of D-Wave's computers in 2011 for about US$10 million, and Internet behemoth Google, which acquired one in May.

But the pain has been real — much of it, critics would argue, brought on by Rose himself. In 2007, his company announced its first working computer with a showy public demonstration at the Computer History Museum in Mountain View, California. By the current standards of quantum computing — which in theory offers huge advances in computing power — the device's performance was astonishing. Here was a prototype searching a database for molecules similar to a given drug and solving a sudoku puzzle, while the best machines built using standard quantum approaches could at most break down the number 21 into its factors[1].

Sceptics bristled at the ‘science by press conference’ tone of the introduction, and wondered whether the D-Wave device wasn't just a classical computer disguised as a quantum one. “This company from Canada popped out of nowhere and announced it had quantum chips,” says Colin Williams, who published one of the first texts on quantum computing in 1999, and who joined D-Wave last year as business-development director. “The academic world thought they must be crazy.”

Today, those criticisms have been quietened to some degree by the release of more details about D-Wave's technology. But they have been replaced by subtler questions: even if the D-Wave computer is harnessing quantum powers, is it really faster or better than a conventional computer? Will it ultimately crack problems that currently take computers decades or more to solve? Or will its capabilities hit a wall?

UNIVERSAL VISION

When Rose founded D-Wave in 1999, he had an engineering degree, a few years' progress towards a PhD in theoretical physics at the University of British Columbia in Vancouver — and no idea how to build a quantum computer. He did have inspiration, from a class on entrepreneurship that he had taken with Haig Farris, one of Canada's best-known technology venture capitalists. Business, says Rose, "appealed to me as being harder than physics or math. There's no prescription for making people do what you want."

Williams' then-new textbook helped to convince Rose that quantum computing would make a suitable target for a new venture. A cheque for Can$4,059.50 (US$3,991) from Farris let him buy a laptop and printer to produce a business proposal. By the early 2000s, D-Wave had attracted millions of dollars in capital, which Rose invested in 15 different research groups to look for the best technology to pursue. "I was like an evangelist, pitching the vision" of a quantum computer, he says.

At the heart of that vision was quantum computing's promise to solve otherwise-intractable problems by drastically reducing the time required to find an answer. The quintessential example is factorizing: like splitting 21 into 3×7, but with numbers hundreds of digits long. That is the basis of the encryption algorithms widely used to protect digital data. Encryption security rests on the fact that conventional computers have to look at every possible factor in turn — a process that takes exponentially longer as the numbers get bigger.

The bottleneck arises because conventional computers store and process information in an either–or fashion, using 'bits' that can each exist in only one of two states, denoted 1 or 0. In most modern computer chips, each bit is represented by the presence or absence of an electric charge. Quantum computers, by contrast, exploit the fuzzy world of quantum mechanics by using 'qubits' that can exist as both 1 and 0 at the same time. In principle, they can explore different solutions simultaneously — reducing a multi-year calculation to seconds.

By the time Rose began his search for the right technology with which to build a quantum computer, researchers had begun to make qubits from many physical systems, including photons that encode zeroes and ones in the direction of their polarization, and ions that encode them in their electron states.

They were also working on ways to combine and manipulate the quantum information carried by these qubits, in much the same way that transistor logic gates manipulate the flow of bits in a conventional computer. The goal was to produce 'universal' quantum computers that could carry out any conceivable computation, like a modern classical machine.

➲ NATURE.COM
To read another take on quantum computing, see:
go.nature.com/hbsrsz

But this model entailed some huge engineering challenges — starting with the fact that quantum bits are extremely susceptible to outside interference. They are like pencils balanced precariously on their points: the slightest perturbation can knock them off balance, causing an error in the calculation. If each qubit is 99% accurate, an operation involving 10 of them will yield the right answer only 90% of the time, and one with 100 qubits will do so only about 36% of the time. Yet practical applications might require thousands or millions of qubits.

To compensate, developers go to great lengths to shield their qubits from noise, and to devise clever error-correction schemes. But then and now, says Andrew Landahl, who works on quantum computing at Sandia National Laboratories in Albuquerque, New Mexico, "if you look at the redundancy and fidelity you need, it's extremely demanding". Like a rocket that requires tonnes of fuel to hoist a tiny payload, a gate-model quantum computer might need billions of error-correcting qubits just to get 1,000 functional qubits to do something productive.

> "I think it is not too strong to say they were initially ridiculed by the academic community."

By 2003, Rose was convinced that this model was "just a bad, bad, bad idea", he says. So he shifted his focus to what was then a research backwater: adiabatic quantum computing[2]. This technique is best suited to optimization problems — the kind in which the best possible outcome must be found for a number of criteria simultaneously. Examples include trying to arrange the seating for a wedding at which some guests are best friends and others sworn enemies; or finding the most energetically stable way to fold a protein in which the various amino acids attract or repel each other.

All the possible solutions to such problems can be imagined as a mountain range in which the higher elevations correspond to configurations that violate most of the criteria — enemies sitting next to enemies, so to speak — and the lowest points correspond to solutions in which most or all of the criteria are satisfied. The trick is to find those low points. A conventional adiabatic computer can do that through the equivalent of huffing and puffing over the mountain passes, systematically looking for dips. But a quantum adiabatic computer does a rapid global search. It starts with the analogue of tipping water onto a flat landscape — a state in which the qubits are in a perfect quantum superposition of zeroes and ones — then lets the mountains rise slowly, so that the water naturally pools in the best solutions.

The key to such a computer is that its qubits are meant to stay in their lowest energy state at all times — the precariously balanced pencils have already fallen over. This gives it the massive advantage of being relatively resistant to outside interference, so that little or no error correction is needed until the computer has thousands of qubits or more. And although it is not very useful for factorizing large numbers — the thing that spurred research into quantum computers in the first place — its approach could potentially be used on applications ranging from language translation and voice recognition to working out flight plans for spacecraft.

In 2003, little was known about how to make or program an adiabatic quantum computer, and no one had put in the money and time to build a prototype. Rose decided that D-Wave should try.

Using qubits made from superconducting loops of niobium, cooled to 20 millikelvin above absolute zero to keep them in their lowest energy states, D-Wave's engineers created a usable computer before even they were sure how it worked. "The name of the game from the outset was to make a functional computer," says Williams. "Then they could probe it to see where it was operating correctly."

From there, D-Wave ramped up quickly. The company's 2007 demonstration used a 16-qubit device. By 2011, the D-Wave One machine purchased by Lockheed Martin had 128 qubits (see *Nature* **474,** 18; 2011). This year's D-Wave Two, the model acquired by Google and collaborators including NASA, has 512 (see *Nature* http://doi.org/mt2; 2013). Their computer looks like the proverbial black box: it is a shiny black cube about the size of a sauna. Most of the space is occupied by a cryogenic cooling system; the quantum chip itself is the size of a fingernail. D-Wave aims to double the number of qubits on that chip every year.

HOSTILE AUDIENCE

From the start, D-Wave generated a lot of bad feeling. "I think it is not too strong to say they were initially ridiculed by the academic community," says Jeremy O'Brien, a physicist at the University of Bristol, UK, who invented the computer that can factorize 21.

The problem was not so much the adiabatic-computing approach — it has a solid, if sparse, academic history — but the company's brash style. Most quantum-computing experts feel that Rose and his colleagues should have started by soberly publishing papers characterizing their qubits, rather than putting out press releases. Scott Aaronson, a computer scientist at the Massachusetts Institute of Technology in Cambridge and a long-time D-Wave sceptic, remains unimpressed by what the company has actually shown that it can do. "They are marketing types who are trying to make the most dramatic claims possible," he bristles.

Rose neither denies nor apologizes for the brashness. He has frequently been quoted as saying, in effect, that his approach is how you build a company. Rose also insists that he has no regrets about the company's 2007 press event — particularly given that it got the attention of Google, which started working informally with D-Wave soon afterwards. "We're not in this business to be popular," he says.

Business style aside, the D-Wave computer is so different from anything else that exists that not even experts know exactly how to judge it. "You do these demonstrations, and how do you know if it's any more significant than factoring 15?" says John Martinis, a physicist at the University of California, Santa Barbara, who heads one of the leading groups working on gate-model quantum computers.

Some of the suspicion is easing as it becomes clearer how the computers operate. In 2011, D-Wave published evidence for quantum behaviour in its 8-qubit chip[3]. Outside the company, the group that has spent the most time on the question is the University of Southern California's Quantum Computing Center in Los Angeles, set up in collaboration with Lockheed Martin when the firm bought its D-Wave computer. In April, a team including the centre's scientific director, Daniel Lidar, circulated results seeming to confirm that the 128-qubit D-Wave One works on a quantum level[4] — although in the fuzzy quantum world nothing is certain, and the results have been challenged[5,6].

Still, D-Wave has chipped away at its credibility problem, concludes O'Brien, "and now they're taken ever more seriously".

ARLEN REDEKOP/VANCOUVER SUN

Geordie Rose expects his company's quantum machine to change the face of computing.

PRACTICAL CONSIDERATIONS

Regardless of how the D-Wave computer works, the practical question is whether it can be used for real-world problems. It can — sort of. In 2009, for example, a Google research team developed a D-Wave algorithm[7] that could learn to judge whether or not a photo showed a car — an example of a 'binary image classifier' that could in principle be used to tell whether a medical image shows a tumour, or a security scan shows a bomb. Finding ever-better ways of doing this sort of task is at the heart of artificial intelligence, and is one area in which an adiabatic quantum computer is expected to excel.

In 2012, researchers at Harvard University in Cambridge, Massachusetts, used a D-Wave machine to find the lowest-energy folding configuration for a protein with six amino acids[8]. They did not have enough qubits to code the problem properly, but even so, on a problem that no other quantum computer could touch, the D-Wave machine found the best solution 13 times out of 10,000 runs. And many of the other answers were good solutions, if not the best.

Meanwhile, Lockheed Martin and University of Southern California researchers have developed an algorithm that allows D-Wave machines to tell whether a piece of software code is bug-free[9] — something that, Lockheed Martin notes, is impossible with classical computers. "You would never know" for sure if a piece of classical-computer code was clean, says Ray Johnson, chief technology officer for Lockheed Martin in Bethesda, Maryland. All anyone could say was that no fault had been found after years of testing. "But now you can say with certainty," says Johnson. "We have great hope, and confidence, in the ability of the computer to scale to real-world complex problems."

D-Wave also competes well against conventional computers in terms of speed, although direct comparisons are difficult. Earlier this year, D-Wave asked Catherine McGeoch, a computer scientist at Amherst College in Massachusetts, to put the D-Wave Two through its paces to satisfy Google before the Internet giant confirmed its deal. McGeoch found that in the optimization-type problems that the D-Wave was designed to solve, it came up with the right answers in half a second, compared with 30 minutes for a top-level IBM machine[10]. "That's one of the most exciting things to happen in quantum computing," says O'Brien.

It is far from clear how long that advantage will last, however, if only because there is no good theory to describe how quantum adiabatic computers will behave on a larger scale. "We are absolutely certain we can build the next generation of this device, but we have absolutely no idea how well it will work," laughs Rose. And since McGeoch presented her results[10] at a meeting in May, other computer scientists have been trying to write yet-faster codes for classical computers. Aaronson says that speed should not be taken as proof of how the device is working. "Even if the machine does get to a solution faster than an ordinary laptop," he says, "then you still face the question of whether that's because of quantum effects, or because a team of people spent $100 million designing a special machine optimized to these types of problems."

In the meantime, work continues to make qubits for universal gate-model quantum computers more reliable, or easier to mass-produce. O'Brien, who admits that his 4-year-old daughter can factorize 21 faster than his computer, is optimistic about the future. "In 10 years' time, I'd be hugely disappointed if we didn't have a machine capable of factoring a 1,000-bit number, involving millions of qubits," he says.

But Rose remains a devotee of the adiabatic church — and is convinced that D-Wave's next generation will prove that it can solve exponentially more difficult problems without taking exponentially more time. "There's going to be absolutely no hope for classical computers if this thing next year behaves as we expect," he says. Rose goes so far as to consider the hardware problem solved: the real challenge, he says, will be the software. "Programming this thing is ridiculously hard," he admits; it can take months to work out how to phrase a problem so that the computer can understand it. But D-Wave has teams working on that — including Rose.

Rose expects tough competition. But with his instinct for fighting, he seems ready for it. ■

Nicola Jones *is a freelance writer near Vancouver, Canada.*

1. Martín-López, E. *et al. Nature Photon.* **6,** 773–776 (2012).
2. Farhi, E. *et al. Science* **292,** 472–475 (2001).
3. Johnson, M. W. *et al. Nature* **473,** 194–198 (2011).
4. Boixo, S. *et al.* Preprint at http://arxiv.org/abs/1304.4595 (2013).
5. Smolin, J. A. & Smith, G. Preprint at http://arxiv.org/abs/1305.4904 (2013).
6. Wang, L. *et al.* Preprint at http://arxiv.org/abs/1305.5837 (2013).
7. Neven, H., Denchev, V. S., Rose, G. & Macready, W. G. Preprint at http://arxiv.org/abs.0912.0779 (2009).
8. Perdomo-Ortiz, A., Dickson, N., Drew-Brook, M., Rose, G. & Aspuru-Guzik, A. *Sci. Rep.* **2,** 571 (2012).
9. Pudenz, K. L. & Lidar, D. A. *Quantum Inform. Process.* **12,** 2027–2070 (2013).
10. McGeoch, C. G. & Wang, C. *Proc. ACM Intl Conf. Comput. Front.* http://dx.doi.org/10.1145/2482767.2482797 (2013).