

May/June 2008

Riding D-Wave

A pioneer of quantum computing asks: Has a Canadian startup really demonstrated a prototype for a working, commercially viable quantum computer?

By Seth Lloyd

Computers process information by breaking it down into the smallest possible chunks, called "bits." A bit represents the distinction between two possibilities: True and False, Yes and No, or, as they are conventionally represented, 0 and 1.

The end point of Moore's Law (which holds that computers get faster by a factor of two every year and a half or so) is a computer so powerful that it uses individual atoms to store bits of information: one atom, one bit. If we were able to work at subatomic scales and store bits on electrons or quarks, we might go further. But let's stick with what we *know* we can do.

If current rates of miniaturization persist, your PC will store one bit on one atom sometime around 2050. But it's natural to ask whether we can, in fact, achieve a bit-to-atom correspondence. Remarkably, prototype computers that store bits on individual atoms already exist in the laboratory. These computers are called quantum computers, because they store and process information at scales where the laws of quantum mechanics hold sway.

Quantum mechanics is the branch of physics that governs what happens at very small scales. Its principles are famously weird, so it's natural that quantum computers should be odd, too. A conventional electronic computer, in which each bit registers either 0 or 1, is enslaved by binary logic; but a quantum bit, or "qubit," can register 0 and 1 at the same time, a phenomenon known as "superposition." What does it mean for a quantum bit to simultaneously register 0 and 1? The accurate answer is, nobody knows for sure. The counterintuitive nature of quantum mechanics prevents our minds from grasping how quantum bits behave. Nonetheless, because the laws of quantum mechanics are precisely formulated, we can predict what quantum computers will do.

And what they do is remarkable. Since one qubit can simultaneously represent two different values, two qubits can simultaneously represent four (00, 01, 10, and 11, in binary notation); four qubits can represent 16 values; eight qubits 256 values; and so on. Even a relatively small quantum computer, one that had a few tens of thousands of qubits, could consider so many different values at once that it would be able to break all known codes commonly used for secure Internet communication. Quantum computers might also be used for faster database searches, or to tackle hard problems that classical computers couldn't solve with all the time in the universe. My colleagues at MIT and I have been building simple quantum computers and executing quantum algorithms since 1996, as have other scientists around the world. Quantum computers work as promised. If

they can be scaled up, to thousands or tens of thousands of qubits from their current size of a dozen or so, watch out!

Given their power to intercept and disrupt secret communications, it is not surprising that quantum computers have the attention of various U.S. government agencies. The National Security Agency, which supports research in quantum computing, candidly declares that given its interest in keeping U.S. government communications secure, it is loath to see quantum computers built. On the other hand, if they can be built, then it wants to have the first one.

Quantum computation has also attracted commercial interest. At current rates of progress, big, code-breaking quantum computers are at least a decade away, so the private sector is focusing on two types of quantum computation that are easier. The first nontrivial type of quantum computing was proposed by the Nobel laureate Richard Feynman in 1981. Feynman was studying how quantum processes in high-energy physics could be simulated. He noted that classical computers were bad at the job, for the same reason that human beings find quantum mechanics counterintuitive: there is no easy way for either to represent a bit that registers 0 and 1 at the same time. Feynman suggested that if the computer were quantum-mechanical, it might have an easier time dealing with quantum processes. In 1996, I showed that Feynman was correct and created algorithms that would allow a quantum computer to simulate solid-state, chemical, and high-energy systems. Such a simulator would require only a hundred qubits or so to be able to surpass all conventional supercomputers.

A second type of quantum computing, known as adiabatic quantum computing, is not only easier than code breaking but potentially far more powerful. Adiabatic quantum computing is a particularly physical way of trying to solve hard problems.

Like all physical systems, electrons would rather inhabit lower energy states than higher energy states, particularly at low temperatures. The energy of a physical system such as an electron depends on the states of its neighbors. One electron might tell its spinning neighbors, "For a lower energy, spin clockwise." Another electron might say, "For a lower energy, spin counterclockwise." The lowest energy state for the spinning electrons as a community is the one that minimizes the total number of conflicts between neighboring spins. For a group of electrons to find their communal lowest energy state, or "ground state," they must find ways to agree on how to align their spins. In the same way that a complex computational problem can be broken down into flipping bits, it can be posed in terms of finding the ground state of a suitable physical system.

Adiabatic quantum computation attempts to represent problems as the disturbance of a quantum system, so that the answer is represented by the system's assumption of a new ground state. Developed by Eddie Farhi and Jeffrey Goldstone at MIT and Sam Gutmann at Northeastern University, it works by initializing the quantum system to a simple ground state (all spins rotating clockwise, for example) and then gradually, or "adiabatically," turning on the interactions that encode the problem. If this turning-on

process is sufficiently slow, the system will gradually "ooze" from its simple initial state into the complex final state.

The most interesting aspect of adiabatic quantum computation is that no one knows for sure whether it works in practice. It may be that for any meaningful problem, the system would have to ooze so slowly that it would take the age of the universe to return an answer. Conversely, it may be that even the hardest problem will succumb to an adiabatic quantum computer. Despite the concerted attention of a bevy of physicists and mathematicians, the question of whether adiabatic quantum computing works remains open. Most experts suspect that it can't solve the very hardest computational problems. But suspicion is not proof.

When the theorists can't agree, experimentalists forge ahead. Because the whole point of adiabatic quantum computation is to go slow rather than fast, adiabatic quantum computers are in principle significantly easier to build than general-purpose code-breaking quantum computers. Realizing this key point, in 2002 my graduate student Bill Kaminsky and I created a design for an adiabatic quantum computer based on superconducting technology. Last year, D-Wave Systems, a quantum-computing startup in Burnaby, British Columbia, announced that it had constructed an adiabatic quantum computer based on our design. At that point, things got interesting.

D-Wave was founded a little less than a decade ago, with the express purpose of building a commercial quantum computer. After toying with the idea of building a quantum computer to factor large numbers, its researchers sensibly settled on the more straightforward and still potentially profitable tasks of quantum simulation and adiabatic quantum computing. In February 2007, at Silicon Valley's Computer History Museum, the company demonstrated a 16-qubit device that it claimed could solve reasonably complex optimization problems. It could even do Sudoku puzzles!

D-Wave has raised about \$60 million in funding from venture capitalists such as Draper Fisher Jurvetson. As a private company, it is responsible primarily to its investors rather than to the scientific community. So it was no surprise that in announcing its success in building an adiabatic quantum computer, D-Wave focused on commercial applications rather than scientific details. While venture capitalists were impressed by the announcement, treating the company to another round of funding, scientists were less excited. The press release provided no device specifications that would allow the scientific accuracy of its claims to be assessed. It seemed possible that the computer was simply finding solutions by being cooled down to its ground state, a fairly dull and not-so-quantum-mechanical process, rather than performing the more subtle adiabatic procedure described above. When D-Wave neglected to supply any concrete evidence that the device was actually performing a quantum computation, even the most charitable scientific observers simply assumed that its scientists didn't know whether it was or not. (See "[Desultory D-Wave](#)") Less charitable observers uttered words I cannot report in this publication. For my part, I was conflicted. I would really like to know whether adiabatic quantum computation works. Even if this approach can't solve the very hardest problems, if D-Wave's system could perform a well-defined demonstration of adiabatic quantum

computation in some simple instances, that would be a validation of Kaminsky's and my design. As matters stood, however, D-Wave seemed to be muddying the quantum well for money.

Last fall, the waters became clearer. D-Wave's chief theoretician, Mohammad Amin, and its chief experimentalist, Andrew Berkley, visited the quantum-computing community at MIT. They discussed the scientific issues frankly. No, they admitted, they couldn't prove that what they were doing was true adiabatic quantum computation--but it looked as if it probably was. How could they answer the question conclusively?

The pioneers of superconducting quantum computation had been able to demonstrate the quantum nature of their devices by zapping them with fast microwave pulses and looking at their responses. But those devices weren't adiabatic; they operated at speeds comparable to those of a conventional computer. The D-Wave device, by contrast, is purposefully slow: therefore, no zapping is possible. As a result, there are a limited number of experiments that can indicate whether the device is really doing quantum computation. One, however, is to vary the slowness with which the device oozes from its initial state to its final state. Halfway through the oozing process, the computer arrives at a point where it must start making the hard choices that lead to the problem's solution. Here the computer is in a weird quantum state, in which every bit registers 0 and 1 at the same time. I urged the D-Wave researchers to explore this critical point and search for the telltale signs.

More recently, I spoke with Herb Martin, the CEO of D-Wave, and Geordie Rose, the company's chief technology officer and cofounder, and emphasized the need for them to pursue these experiments if they are truly interested in explaining how their devices work. One experiment that I recommended to Rose is a specific protocol for creating and verifying the presence of a so-called Schrödinger's-cat state, a specific instance of the state in which all the qubits register both 0 and 1 simultaneously. (The name comes from a thought experiment proposed by one of the founders of quantum mechanics, Erwin Schrödinger, who imagined a quantum cat that could be both dead and alive at the same time.) Both Martin and Rose seem enthusiastic: they are well aware that if they can't prove that their device is really doing something quantum-mechanical, then their name within the scientific community will remain mud.

In November of last year, D-Wave demonstrated what it claimed was a 28-qubit adiabatic quantum computer. Now, the company's scientists are attempting to demonstrate the fundamentally quantum-mechanical nature of their device. There is a strong motivation for doing the science and getting it right. Engineering is science so well established that even engineers like me can do it. If you can't get the science of a 16-qubit quantum computer right, then your chances of building 512-qubit and 1,024-qubit devices (D-Wave's next planned steps) are nil. On the other hand, if D-Wave can confirm that its current system enters the state where all its qubits are 0 and 1 at the same time, then it has a good shot at building quantum devices that are more complex.

And a 16-qubit superconducting Schrödinger's cat would be pretty cool.

Seth Lloyd is a professor of mechanical engineering and the director of the Center for Extreme Quantum Information Theory at MIT.

Copyright Technology Review 2008.