DOT-TO-DOT DESIGN

RESEARCHERS ARE CONNECTING TINY PUDDLES OF ELECTRONS IN A CHIP AND MAKING THEM COMPUTE—THE QUANTUM WAY

BY LIEVEN VANDERSYPEN

THREE AND FIVE! The result was correct. After spending long nights in the lab during the spring of 2001 tweaking and fixing a roomful of equipment, my colleagues and I at Stanford University and the IBM Almaden Research Center had built a computer that could successfully calculate the prime factors of 15. To be sure, you don't need a computer for that—a fifth-grader could give you the answer. What was so remarkable about our machine was that it computed not by toggling a bunch of transistors but by manipulating deep quantum-mechanical properties of individual atomic nuclei. In doing so, this quantum computer prototype factored 15 in a fundamentally different way, and in fewer steps, than any conventional computer was capable of doing.

Six years later, we're still hunkered down in labs—albeit different labs, having dispersed to various research institutions throughout the world—and we're now seeking to build bigger and better quantum computers. We want a computer that can factor not 15 or 21 or 35 but 300-digit-plus numbers. Such a system would in principle be able to break today's most advanced cryptographic codes and could be used to engineer new ways of protecting data. A quantum computer would also easily simulate physical models that today's top supercomputers can't handle—calculating the quantum energy levels of atoms, for example, or simulating the behavior of conventional transistors as they shrink to diminutive dimensions where the laws of quantum mechanics rule. Quantum computers may also speed up key types of search problems in which the correct solution must be found among a vast number of trial solutions.

As we look forward to such possibilities, we often look back to that first Stanford-IBM machine. It taught us a couple of important lessons. The first was that the quantum-mechanical property we used to store the computer's data proved an excellent choice. This property is spin, a kind of intrinsic angular momentum exhibited by atomic nuclei, electrons, and other particles.

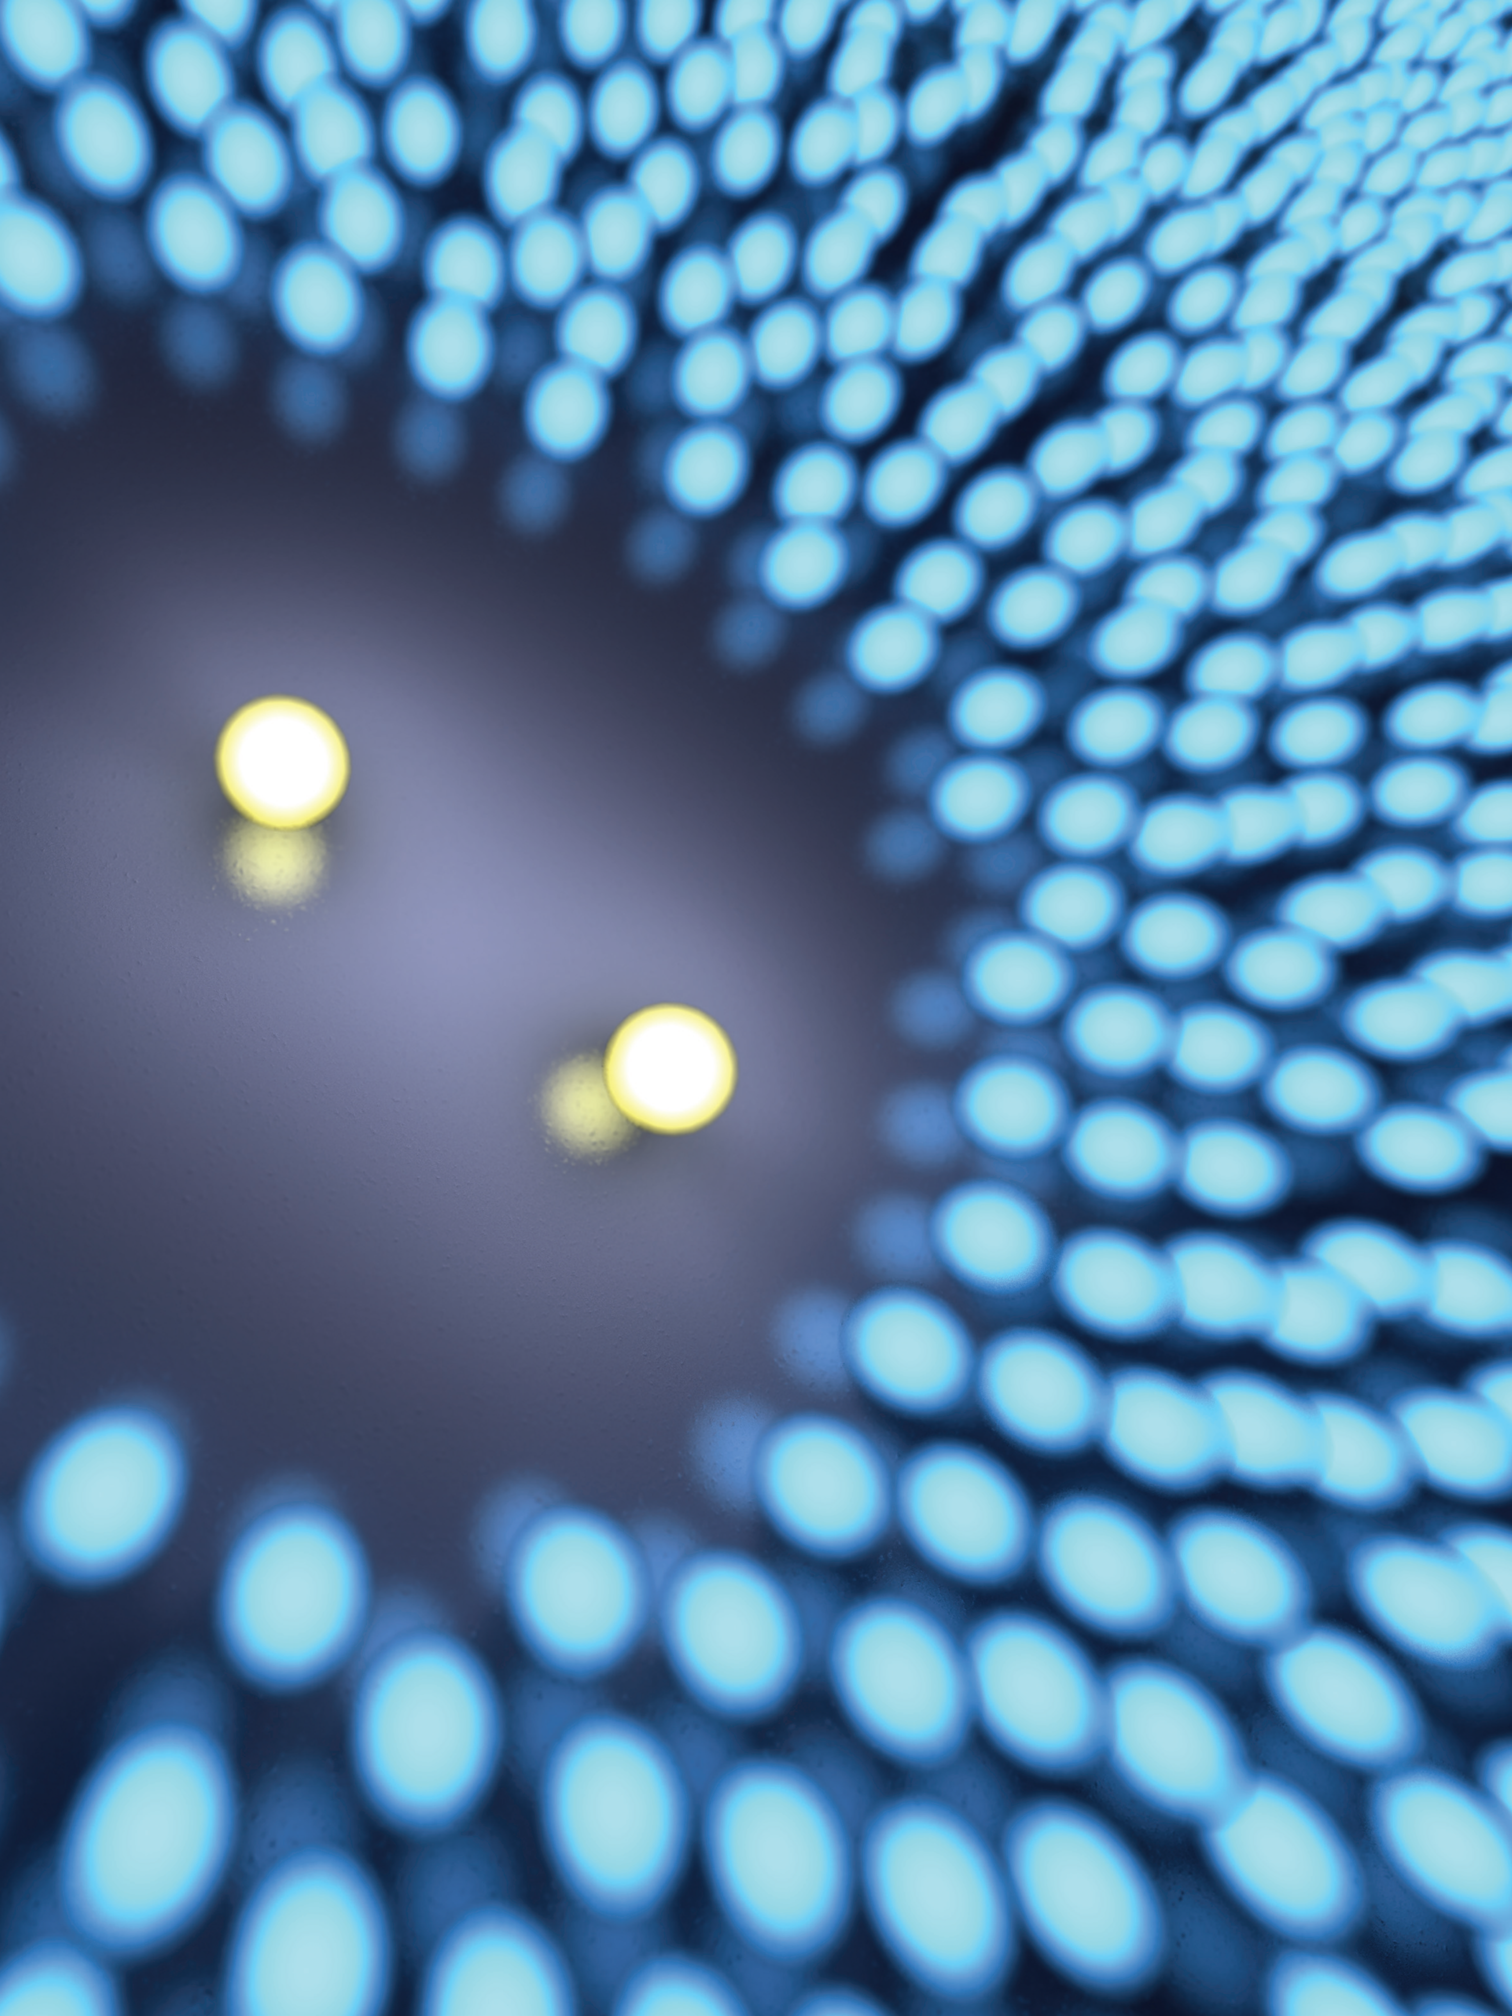

Connecting the (Quantum) Dots

Current quantum computer prototypes based on **quantum dots**—tiny electron puddles in a semiconductor material—are not like your PC. They require cryogenic systems (which can even freeze air) and also superconducting magnets (which can pull a metal pen out of your pocket). Here's a setup used by many quantum computing groups.

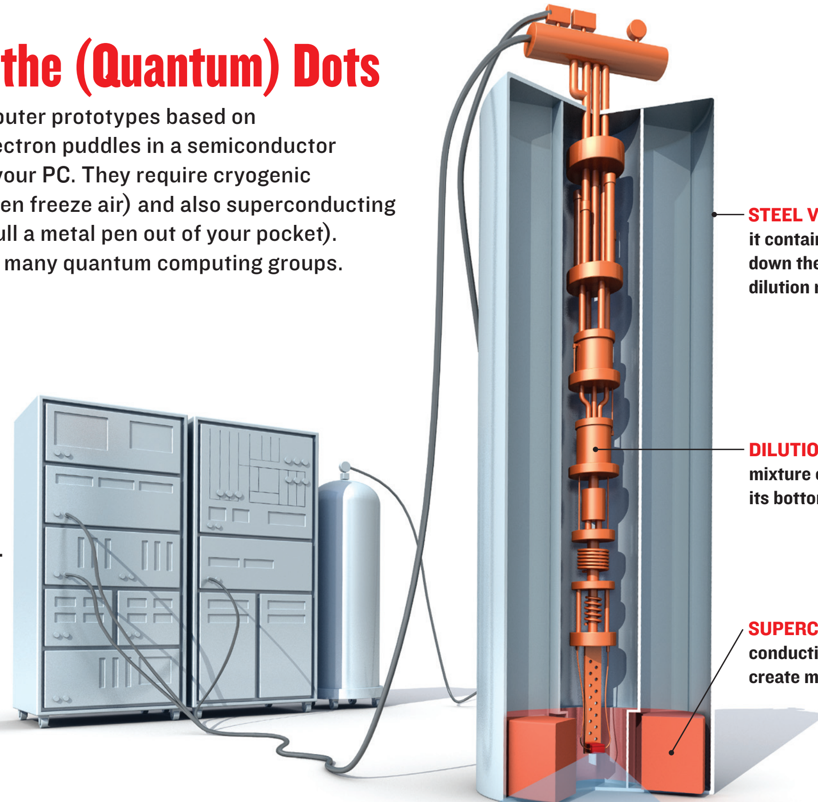

Quantum Chip

A quantum computer would need hundreds of quantum dots to perform useful operations. So far, current prototypes have two or just a handful of dots. But researchers hope that chip-making techniques will allow them to scale up their systems. **Here's how current prototypes look:**

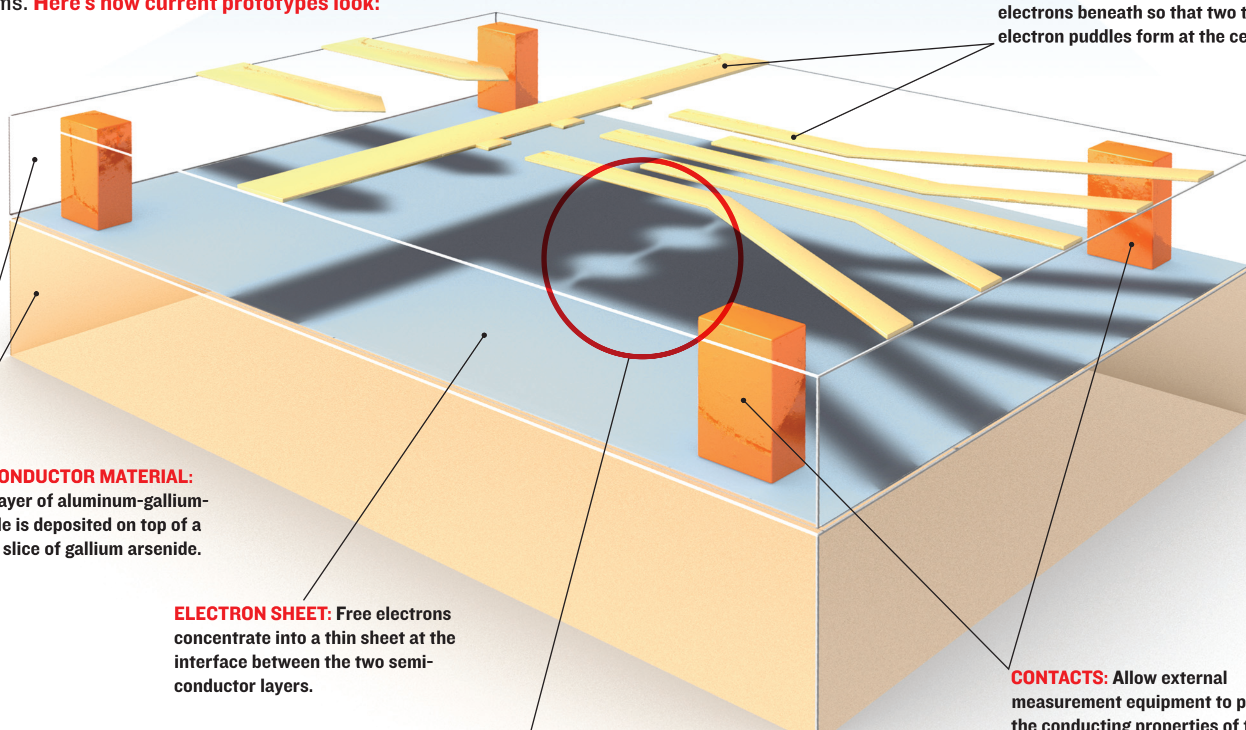

The second lesson was that the way we used spin posed some big challenges. The core of our quantum computer consisted of a custom-synthesized organic molecule in a solution. It had five fluorine and two carbon nuclei whose spins we used to store seven units of information, called quantum bits, or qubits. We blasted the molecule with radio-frequency pulses to alter the spins according to the computational steps of the factoring algorithm. To read out the qubits, we used nuclear magnetic resonance, or NMR, to generate a frequency spectrum of each spin. It worked beautifully for seven qubits, and in fact that system remains the only one to have factored a number to this day. But designing molecules suitable for more complex calculations became just too hard.

If we wanted a quantum computer that we could scale up, we needed a system that would let us precisely manipulate tiny bits of energy, that could be effectively shielded from external interference, and—most important—that could be built by replicating tiny identical building blocks within a small area. We needed something less like a test tube—and more like a microchip.

TO SCALE UP OUR QUANTUM COMPUTER, WE NEEDED SOMETHING LESS LIKE A TEST TUBE AND MORE LIKE A MICROCHIP

A SEMICONDUCTOR quantum computer is now the goal of dozens of research groups worldwide. In the last few years, these groups, including my own at Delft University of Technology, in the Netherlands, have made rapid progress in creating qubits based on materials and processes similar to those used in the microelectronics industry to manufacture standard processors and memory chips. [See "The Trap Technique," *IEEE Spectrum*, August, for the first part of this report.]

The advantage of a solid-state design over the NMR approach is the ability to fabricate large arrays of miniature electronic devices that can be individually addressed and interconnected—just as we do with transistors in an integrated circuit. One promising approach to such a solid-state system was put forward by Daniel Loss of the University of Basel, in Switzerland, and David DiVincenzo of the IBM T.J. Watson Research Center, in Yorktown Heights, N.Y. In their January 1998 paper, "Quantum Computation with Quantum Dots," in *Physical Review A*, they proposed trapping individual electrons in semiconductor structures called quantum dots and then using the electrons' spins as qubits.

With typical dimensions from a few nanometers to a few micrometers—about the size of a virus—a quantum dot is a tiny area in a semiconductor that can hold anything from a single electron to several thousand. To make a quantum dot that's suitable for a quantum computer, you start with a half-millimeter-thick wafer of gallium arsenide and cover it with an even thinner, 100-nm-thick layer of silicon-doped aluminum-gallium-arsenide. Free electrons will concentrate at the interface between the two materials, forming a thin electron sheet. Next, you attach a set of gold electrodes to the top layer and apply negative voltages to them. The electrodes will repel electrons in the sheet underneath and create small islands of electrons isolated from the rest.

Creating such electron puddles is relatively straightforward, but manipulating electron spin is a different matter. Like charge and mass, spin is considered an intrinsic property of electrons, and yet it remains somewhat mysterious. We can measure spin because it interacts with an external magnetic field, much as an ultrasmall magnet rotating about its own axis would. But unlike with a real magnet, when we measure an electron's spin orientation, there will be only two possible outcomes: the spin and the external field are pointing in the same direction, or they are pointing in opposite directions. These two possibilities are also referred to as spin up and spin down, respectively.

More interesting—and bizarre—is that spin can also exist in a combined state of up and down. This superposition state is one of the things that set quantum computers apart from classical ones. A three-bit conventional memory, for example, can hold any combination of three bits at a time: 000, 001, 010, 011, 100, 110, 101, or 111. But using qubits, and representing spin up as 0 and spin down as 1, you can do much better: a three-qubit memory can hold all those eight states simultaneously. As a result, if you perform a calculation using those three qubits, you in effect perform a calculation on all eight states at once. As you add more qubits, this quantum parallel processing increases exponentially.

To perform quantum computations, however, you need to link the qubits somehow. The way researchers do that is by using the quantum phenomenon of entanglement. Two entangled spins can exist in a superposition of, say, up-down and down-up. You don't know which electron has which spin until you measure it. But as soon as you measure one spin, that means the other spin must have the opposite value. How do they "know" which way to point? Scientists devised ingenious experiments to test entanglement and concluded that entangled particles don't carry a "preprogrammed" behavior. Instead, according to quantum mechanics, the pair of electrons forms a single entity. Each electron's spin by itself has no definite orientation until one of them is measured, no matter how far apart they are. Einstein rejected this notion and famously called it "spooky action at a distance."

Spooky indeed. But those are the rules of quantum mechanics, and we might as well use them to our advantage. Quantum researchers not only accept spin's weirdness, but they also embrace it. They think of spin as a vector in a mathematical domain called a Hilbert space. Basically, this vector describes the probabilities of obtaining spin up or down when a particle's spin is measured. The researchers perform a host of mathematical transformations to those vectors to concoct quantum computing algorithms. But as physicist Asher Peres has put it, "Quantum phenomena do not occur in a Hilbert space, they occur in a laboratory." And it's in the lab that our group and many others set out to build a practical quantum computer.

OUR STARTING POINT was Loss and DiVincenzo's proposal and related concepts. Clearly, it would be too difficult to build a whole computer at once. So the idea has been to develop a set of basic functions that any working system would need. These are an initialization mechanism to set all of the qubits to a known state before computations begin, a readout scheme to measure the individual spins, and a set of spin-manipulation techniques capable of carrying out any possible quantum computations.

Here's the basic design: the core of the machine will consist of a single chip, which will sit inside an ultracold receptacle called a dilution refrigerator, which in turn will be encircled by a powerful superconducting magnet. (As DiVincenzo once said, "This is not going to be a laptop computer!")

Whereas a conventional microchip is packed with transistors, the quantum-computing chip will be packed with quantum dots. The dots—dozens, hundreds, or perhaps thousands—will

A New Spin on Computing

To perform quantum computations you need quantum bits, or qubits. One possibility is to use a quantum-mechanical property of electrons called spin. Whereas a conventional bit can be either 0 or 1, a qubit can be 0, 1, or both at the same time, a state called superposition.

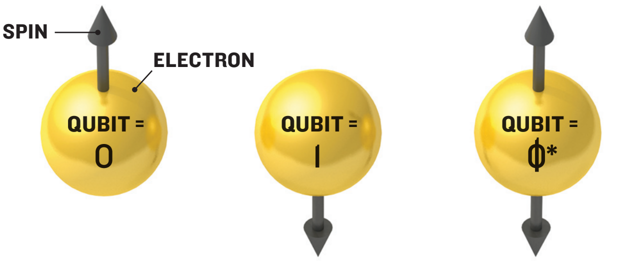

*Superposition of 1 and 0

Spinning It

To compute with qubits, your computer needs to be able to perform four critical tasks.

INITIALIZATION: Before computing, you need to initialize the qubits. This is necessary because spins initially point in random directions. A system based on ultracold temperatures and powerful magnetic fields orients all spins in the same direction: up.

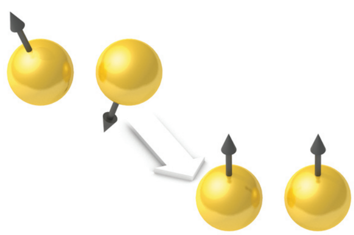

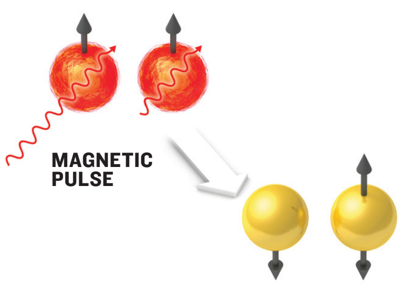

ROTATION: An oscillating magnetic pulse applied to an individual electron can reorient the direction of its spin. The duration of the pulse determines the final direction. The left spin was switched from up to down, and the right spin was put in a superposition state. Rotation is required to load the initial data and perform single-spin operations in quantum algorithms.

SWAP: Voltage changes to electrodes push the electrons closer together, and their quantum properties become coupled. By being separated after a precise period of time, the two electrons will have swapped their spins. Rotation and swap are basic operations that let you carry out any quantum-computing algorithm.

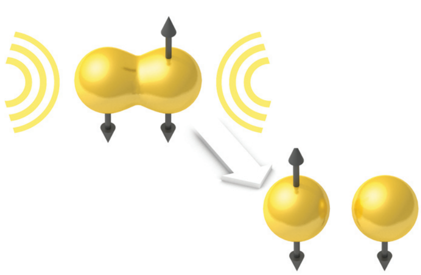

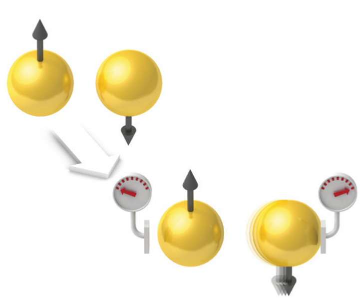

READOUT: After carrying out a quantum algorithm, you need to read out the qubits. Voltage changes to electrodes can push electrons out of their dots. Only the electron with "spin down" makes the jump, because it has a more energetic state. A highly sensitive electrometer detects a current increase, indicating a spin-down electron. No change in current means "spin up."

each hold one electron. Electrodes near each dot will control how its electron interacts with its neighbors.

The first step consists of initializing all of the electron spins to a known state, say, up. The computer will then "load" the initial data: it will leave some spins pointing up, make some point down, and put others in superposition up-down states. The computation comes next. By varying the electrodes' voltages and applying magnetic fields, it will manipulate individual spins or pairs of spins. These manipulations correspond to sums, multiplications, and other operations specified by a quantum algorithm. The final step is reading the output qubits.

We're far from being able to build such a system. The chips developed to date are still too rudimentary, having at most a few quantum dots. Also, even though the basic functions have been successfully demonstrated, integrating them into a single system will take lots of time in the lab. To understand why, it's worth looking at each in more detail.

Consider the task of placing a single electron in each dot. Today that's a routine operation, but it took Leo Kouwenhoven at Delft and Seigo Tarucha at the University of Tokyo a lot of effort to accomplish that. In 1996, these two researchers demonstrated how to apply negative voltages to metal electrodes near a quantum dot to expel the electrons from the dot one by one until just a single electron remained. By comparison, a conventional memory chip uses 10 000 to 100 000 electrons to store one bit of data. So manipulating a single electron was no small achievement.

Having isolated one electron per dot, it's then necessary to set their spins to the same initial state. The superconducting magnet generates a static magnetic field of several teslas that acts as a frame of reference for the spins in the dots. You also need to cool the whole thing to keep the electrons from wiggling around too much. Hence the dilution refrigerator, which circulates a mixture of helium isotopes to remove heat from its surroundings, cooling the chip to about 30 millikelvins. Under these conditions, the spins will assume the lowest energy state, which by convention is spin up.

Measuring the spin of each electron was long considered difficult, because its interaction with external fields is so tiny. In 2004, my colleagues and I found a way around this problem. The trick is to measure the spin indirectly. To do that, we pulsed the electrodes near a quantum dot with 0.5-microsecond, 5-millivolt signals. These pulses give the electron just enough energy to make it slop out of the dot if its spin is down, but not if its spin is up. That's because in a magnetic field, a spin-down electron has a higher energy than a spin-up one. The presence or absence of a single electron in a dot in turn changes the current flowing through a nanoscale electron channel next to the dot by about 300 picoamperes. We measured this tiny current using highly sensitive electronics, which told us the spin state. We got the correct answer about 82 percent of the time, and that performance could be boosted to around 99 percent using even faster and lower-noise electronics.

By 2004 researchers attempting to build a quantum computer out of quantum dots had accomplished two main tasks: initialization and readout of spins. These were important steps, but the essence was still missing: the quantum dots didn't compute.

IN CONVENTIONAL COMPUTERS, any operation you would wish to perform to a group of bits—the AND, OR, and NOT operations of Boolean logic—can be implemented using a universal digital logic gate. One example is the NAND gate, and you can make one with a couple of transistors and resistors. But NAND gates, or any other conventional gate for that matter, can't be used for quantum computations, because they can't handle bits in superposition states. A new type of gate—a quantum gate, capable of operating on superposition qubits—is needed.

Theorists have already identified several universal quantum gates, as well as sets of gates. One such set consists of a spin-rotation gate and a spin-swap gate. The first lets you rotate a spin by a controlled amount. It could be a full flip from up to down, for example, or half a flip to a superposition state. The other gate lets you couple the spins of two electrons, making them swap their states. In the past two years, experimentalists have successfully demonstrated these two gates using quantum dots.

First came the spin swap. In 2005, a team led by Charles Marcus at Harvard coupled two neighboring spins using a phenomenon called the exchange interaction. To understand this, you have to think of electrons not as particles with well-defined locations but rather as waves with fuzzier positions and energy levels. Electrons as waves are described mathematically by wave functions, equations that define the behavior of electrons in terms of probabilities. When two electrons are very close to each other, their wave functions partially overlap, and they can exchange their spins.

The breakthrough, reported in the 30 September 2005 issue of *Science*, was that the researchers controlled the duration of the overlap by tweaking the voltage on an electrode separating the dots. When the exchange interaction was switched on for a precisely timed interval—just a few hundred picoseconds—the two spin states were swapped. And when the exchange interaction was switched on for half as long, the spins were "half swapped," assuming an entangled state.

The other quantum gate is the spin-rotation gate. My colleagues and I had struggled since 2003 to create such a device, and last year we finally succeeded. The method we devised relies on magnetic resonance, the same technique we used to manipulate the atomic nuclei of organic molecules in the 2001 Stanford-IBM prototype. Magnetic resonance, which is also used in hospitals for medical imaging, is based on the fact that when spin is in a static magnetic field, it wobbles about the axis of the field with a certain speed—picture a top wobbling about its axis. Now, if you apply an oscillating magnetic field with the same frequency as the wobbling, the spin can be gradually rotated. Again, this rotation can be a full flip from up to down or down to up, or a partial rotation, say, from up to a superposition, depending on how long the magnetic field was applied.

One challenge we had to overcome was to generate on-chip an oscillating magnetic field of around 1 millitesla. Even that small a field heated up our device. Another problem was reading the spins; the oscillating magnetic fields unavoidably generate stray electric fields, which kick the electron out of the dot. After trying many chip designs, we found a configuration that minimized those problems, and we're now able to rotate the spin in every possible direction. This achievement, which we reported in *Nature* (17 August 2006), was the first time the spin of a single electron was controlled using a semiconductor nanostructure.

So researchers are able to control single and coupled spins and put them in specific superposition states, but how long would the spins remain that way? Not long at all, as it turns out. Electron spins in quantum dots are strongly disturbed by the spins of the atomic nuclei in the host semiconductor. The result: computation errors. We know that such errors begin to crop up after just tens of nanoseconds from the time you put your spin in a desired state.

Some clever tricks exist to lengthen this so-called coherence time. One involves a technique called spin echo. After creating a spin superposition, you wait a short period of time and then apply a control pulse that rotates the spin 180 degrees. You then let the same amount of time elapse, and the errors accrued during the two intervals will cancel out each other. By correcting the spin this way, the coherence time can be extended to a few microseconds. This is still rather short, but for now it is sufficient to proceed.

IN THE FOUR YEARS of research on electron spins in quantum dots, all of the essential ingredients for a quantum computer have been realized. The next steps are clear. First, we need to integrate all of the basic functions into a single system. Then we need to expand the system from two quantum dots to a large array of dots.

And we need to find better ways of overcoming the environment's effects on the fragile spin states—the most fundamental challenge we researchers face now. One possibility is to construct the quantum-computing chip out of materials that have no nuclear spin, such as isotopically pure silicon-28 or carbon-12. Eventually we'll need to reduce the number of errors to at most one in every 10 000 elementary operations. At that point, we could use a technique called quantum error correction to guarantee reliable calculations.

ALL OF THE ESSENTIAL INGREDIENTS FOR A QUANTUM COMPUTER HAVE BEEN REALIZED. NOW WE NEED TO INTEGRATE THEM

As we progress, a particularly profound theme that we'll explore in depth is entanglement. We will need to come up with better ways of detecting that this connection is actually there, that two spins are indeed entangled. Entanglement is routinely studied in laboratories using photons and optical equipment, but controlling this ghostly linkage between scores of electrons in a fingernail-size chip still remains uncharted terrain, ready to be explored in the next few years.

Given the many uncertainties ahead, we still don't know exactly where quantum dots and electron spins will take us. Based on the progress of the last several years, though, we venture to say that the creation of a practical quantum computer may be possible within the next few decades. To get there, we'll keep on connecting the (quantum) dots. ■

ABOUT THE AUTHOR

LIEVEN VANDERSYPEN is an associate professor at Delft University of Technology's Kavli Institute of NanoScience, in the Netherlands.

TO PROBE FURTHER

Visit "The Delft Spin Qubit Project" Web page for more technical details: http://qt.tn.tudelft.nl/research/spinqubits.

For an overview of the field, see "Prospects for Quantum Computing" by David P. DiVincenzo (2000 IEEE International Electron Devices Meeting) and "Challenges for Quantum Computing with Solid-State Devices" by Robert W. Keyes (*IEEE Computer*, January 2005).

An extensive review of spin in few-electron quantum dots is available at http://arxiv.org/abs/cond-mat/0610433.