QUANTUM LEAPS

Fully fledged quantum computers are still a long way off. But devices that can simulate quantum systems are proving uniquely useful.

BY GEOFF BRUMFIEL

When high-energy physicists announced in July that they had found the long-sought Higgs boson — their biggest find in decades — the thousands of individuals involved rightly held their heads high. But in some sense, they had already been beaten to the prize.

Months earlier, a team of nine physicists had taken a rarefied vapour of rubidium-87 atoms, cooled it down to very near absolute zero and used lasers to arrange the atoms into a tiny grid. The physicists then tweaked the temperature until the atoms neared a critical 'phase transition' — a point between two different behaviours, such as liquid water and solid ice. Monitoring their grid in this in-between region, the researchers saw an unusual wave of energy that appeared momentarily and then died away[1]. Mathematically speaking, this behaviour was the same as the appearance and decay of a Higgs particle inside a particle collider.

"Obviously, it's not at all the Higgs particle," says Immanuel Bloch, the researcher who led the study at the Max Planck Institute for Quantum Optics in Garching, Germany. If nothing else, this particle moved in only two dimensions, whereas the Higgs moves in three. But the experiment is still helpful for particle physicists, says Bloch, because it gives them a new way to explore and test the complex quantum field theories that underlie the Higgs.

This experiment also put Bloch and his team at the vanguard of the rapidly growing field known as quantum simulation. The idea, broadly speaking, is to use orderly systems such as a grid of atoms to model much more complicated things — new particles, for example, or high-temperature superconductors. The behaviour of such systems cannot be derived by hand, and even the world's fastest supercomputers can't model them.

Quantum simulators are the lesser sibling of an idea in physics known as quantum computers, which have been touted for more than three decades as a way to do everything from complex modelling to code-breaking. What the simulators and computers share is an ability to operate by the rules of quantum mechanics. Where they differ is in computational power: quantum computers are general-purpose machines able to carry out any possible algorithm, whereas quantum simulators have to be tailored specifically for the problem at hand. Current-generation simulators are also tough to control, and they may not be able to tackle every problem. Nevertheless, the simulators are much easier to build than quantum computers. And researchers say that the devices will soon be able to solve at least some quantum problems that can't be tackled in any other way.

NUTS AND BOLTS

The world of quantum physics is full of theorems, but one goes unwritten: if you want to get noticed, show that your idea came from Richard Feynman.

Feynman, the mid-twentieth-century's greatest theoretical physicist, came up with the idea of quantum simulation in 1981 when he was asked to deliver a keynote speech at the Massachusetts Institute of Technology (MIT) in Cambridge[2]. He decided to talk about how physics might be simulated with computers and got straight to the core of the problem: computers run on certainty, but at a fundamental level, nature deals in probability. According to the laws of quantum mechanics, he knew, particles very rarely exist in one state or another, but instead live in a 'superposition' of two states at once. When observed, the paradox resolves itself according to the laws of statistics. For example, an electron's 'spin' may orient itself in one direction half the time, and in the other direction for the other half.

It is not hard to program a normal computer to model the probabilistic behaviour of that one electron, said Feynman. But particles do not live in isolation, and in quantum systems their probabilities are linked, or 'correlated'. These correlations mean that every combination of particle states must be computed separately, and this creates an exponential rise in complexity. A system with three electrons has eight possible configurations, with eight probabilities to

QUANTUM BOARD GAMES

The set-ups of quantum simulators are different, but the concept is the same: first take atoms, ions or electrons, cool them to cryogenic temperatures and arrange them in an orderly grid. Then tune the interactions on the grid to mimic a more complex material.

COLD ATOMS

Rubidium atoms are held in place by criss-crossed laser beams, which can also be used to tweak individual particles. A single pair of lasers holds the atoms in a one-dimensional column (top), whereas two pairs hold them in a grid (bottom). Some excitations in the grid system behave like the Higgs particle.

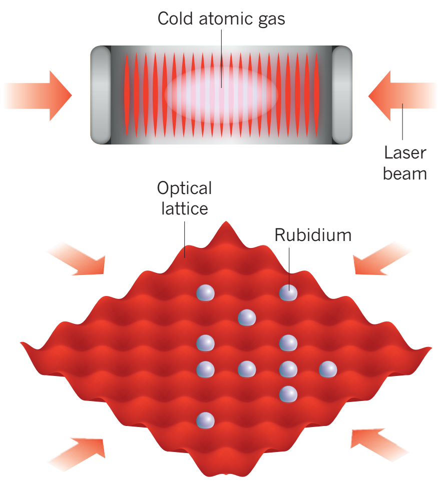

TRAPPED IONS

A combination of electric and magnetic fields trap charged, ionized atoms in an orderly grid. The ions wiggle and rotate in a way that mimics the interactions of quantum magnetism — a phenomenon that can't be simulated in classical systems.

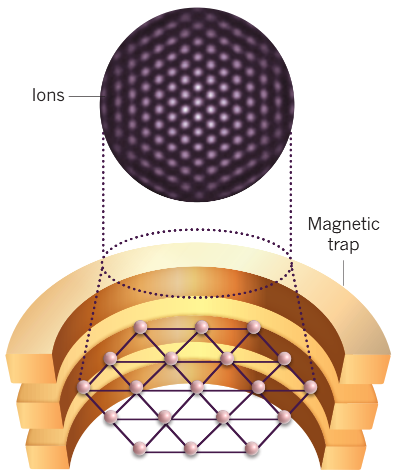

SUPERCONDUCTING LOOPS

A quantized loop of current can flow clockwise, anticlockwise or in a superposition of both in a superconducting circuit (top). An array of such loops (bottom) can be manipulated to simulate various quantum systems — and perhaps even biological processes such as photosynthesis.

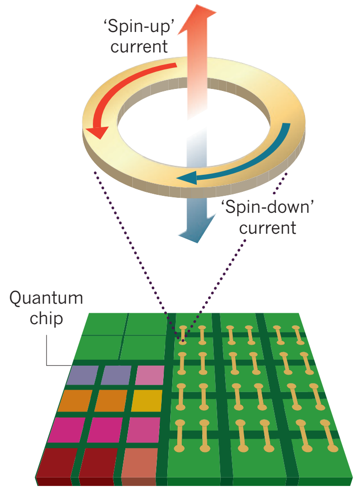

compute; 300 electrons create as many configurations as there are atoms in the known Universe.

Feynman spent most of his lecture trying to find a way out of this conundrum. It is not easy using ordinary computers, he concluded, but there is another possibility: build a computer that thinks in terms of probabilities. This quantum imitator, as he called it, would look a lot like whatever system you were trying to model. It wouldn't need to crunch every outcome, but instead would simply recreate the range of probabilities. Rather than delivering one solution, the imitator would deliver many, and the likelihood of each answer would create a probabilistic picture of how the complex system behaves. Feynman didn't do the maths, but he did conclude that almost any quantum system "can be simulated in every way, apparently, with little latticeworks of spins and other things".

At the time of Feynman's talk, the little lattices of which he spoke didn't exist. Quantum systems are extremely fragile, in the sense that almost any interaction with the outside world will destroy the delicate correlations. It has taken 30 years to develop the technology required to keep the particles isolated enough to finish the simulation unimpeded, yet interactive enough to let physicists extract the answer. But there are now several options. Bloch's group uses neutral atoms, other teams are combining electric and magnetic fields with lasers to trap ions of lighter atoms, such as beryllium. A third technique involves controlling eddies of current inside superconducting microcircuits, and a fourth uses quantum particles of light — photons — moving through microscopic waveguides (see 'Quantum board games').

All these techniques are rapidly increasing in their capabilities. In April, a group led by John Bollinger at the National Institute of Standards and Technology in Boulder, Colorado, unveiled a two-dimensional system of hundreds of trapped ions that could simulate a form of quantum magnetism[3]. The simulator seems to work well for weak fields of the sort that can already be modelled on classical computers, says Bollinger. Now, with some modifications, he hopes to simulate strong magnetic fields, which are beyond the reach of even the most powerful supercomputers.

Bloch, meanwhile, is considering applications beyond the Higgs for a neutral-atom simulator. For example, the rubidium atoms in his lattice might be used to model a complex class of materials called high-temperature superconductors. These 'high-T_c' materials can conduct electrons with no resistance at temperatures much higher than conventional superconductors can — but for decades nobody has been able to understand why. Theorists have developed a number of competing models to explain the behaviour, but haven't been able to test them: the electrons in the superconductors are just too difficult to isolate and study. So Bloch wants to use atoms as surrogates. By changing the intensity of the criss-crossing laser beams, atoms can be made to tunnel from one point in the lattice to another in a way that mimics the motion of electrons through the atomic lattice of a high-T_c material. At least some theories of high-T_c superconductivity should be checkable with Bloch's set-up.

"IN A QUANTUM COMPUTER YOU HAVE TO MAKE SURE THAT NO PARTICLE MAKES A MISTAKE."

Quantum simulators might even be able to model non-quantum problems, such as protein folding, that still require huge amounts of computing power to decipher. A group at the Canadian company D-Wave Systems in Burnaby and at Harvard University in Cambridge, Massachusetts, recently did just that by mathematically mapping the folding problem onto a quantum system of 128 loops of current spinning on a superconducting chip[4]. Each loop could spin clockwise, anticlockwise or in a superposition of both directions simultaneously. The performance of the system wasn't great; in one of its protein-folding problems, it found the correct, experimentally determined, answer just 13 out of 10,000 times. Still, says Alán Aspuru-Guzik, a theoretical chemist from Harvard and co-author of the paper, "it's remarkable to me that it was possible to do it" at all.

GOAL CHANGE

Despite all the technical progress, however, the existing simulators are at best a limited approximation of Feynman's original vision — a fully fledged quantum computer that is 'universal', or able to execute any quantum algorithm and simulate any conceivable quantum system. Researchers have been exploring the potential applications of such a device ever since Feynman described it. Arguably the most important one came in 1994, when mathematician Peter Shor, now at MIT, laid out an algorithm that would allow a quantum computer to function as a powerful code-breaking machine[5]. Other quantum algorithms have followed, drawing many scientists (and several intelligence services) into the quest for quantum computing and sparking widespread efforts to create such a machine.

"THE MODEL THAT WE'RE ABLE TO IMPLEMENT MIGHT NOT EVEN CORRESPOND TO A REAL MATERIAL."

Yet building a powerful, universal quantum computer has proven to be a tough task. A true Feynman computer would be able to control thousands or millions of atoms at once, but most of the current systems face a trade-off between size and control. Bloch, for example, can hold as many as hundreds of thousands of atoms in his laser lattice, but he can't then set their quantum states individually. Other researchers have more control over individual atoms, but their systems, which use trapped ions of beryllium, can manage only a handful of atoms with exquisite precision. On top of this comes the omnipresent problem of disruptions from the outside world, which ruin delicate quantum states: even the tiniest bump will create a computational error.

With current systems so far from the ideal, quantum simulators have come to be seen as less of a stepping stone, and more of a goal in their own right. Simulators do not need to be as large as computers, and, crucially, because the answer is encoded as an average across all their atoms, they are believed to be tolerant of the outside disruptions. "In a quantum computer you have to make sure that no particle makes a mistake," says Ignacio Cirac, a theorist at the Max Planck Institute for Quantum Optics. "In a quantum simulation, if you have 100 particles and one of them is wrong, then 99 are still right."

Some see parallels to the middle of the last century, when scientists such as Vannevar Bush were experimenting with 'analog' computers made from resistors and capacitors. The machines were tailored to specific problems or to a class of problems, and could perform a simple set of operations on an input signal. Some of the devices could even perform mathematical calculations. In retrospect, they seem puny compared with digital computers, which use programmable combinations of transistors to perform practically any program. But they were fast, robust and valuable for applications that matched their architecture, says Seth Lloyd, a theoretical physicist and engineer at MIT. They were particularly good at controlling machinery, for example. "All the control circuits in the Saturn moon rocket were analog," Lloyd says.

Like analog computers, quantum simulators are closely tied to their constituent parts, and are less flexible than a true quantum computer. But Lloyd thinks that they might yet find their 'Moon shot' in problems of quantum complexity. For example, as microprocessors shrink and new materials are engineered at a molecular level, quantum effects become more and more important. That, in turn, will lead to a dramatically growing need for quantum modelling that allows designers to understand and predict the materials' behaviour. At least some of those needs are going to be met by quantum simulators, Lloyd predicts. "What seems to be happening is that quantum simulators work on a variety of special cases," he says, "and the number of cases seems to be growing rather rapidly."

Aspuru-Guzik has one such process in mind: photosynthesis. When light strikes a leaf, it creates a pair of negative and positive charges that travel long distances to reaction centres, where they are used to make energy for the plant. The charge pairs may travel according to the rules of quantum mechanics: some researchers think that the collective wavefunction of the pairs spreads out across the light-absorbing chromophore molecules inside the leaf, allowing the pairs to move more efficiently than they would classically (see *Nature* **474,** 272–274; 2011).

Aspuru-Guzik and others think that a simulator could help them to pin down exactly how this happens. Photosynthesis is what Aspuru-Guzik calls a "dirty quantum system" — that is, it contains both quantum and classical elements. A little matrix of superconducting current loops might be perfect for modelling it, he argues, because the loops, too, are subject to noise from the outside world. It still wouldn't be easy, however: Aspuru-Guzik estimates that something such as photosynthesis would require hundreds of quantum bits to simulate, and those systems, he predicts, are at least a decade away.

The ambitions of the scientists developing quantum simulators are considerably more modest. Most are starting their systems out on models that can be calculated with conventional supercomputers to prove that their simulators produce reliable results. Gradually, they plan to push their atoms, current loops or other little units to the point at which the supercomputers can no longer cope.

At that point, "the model that we're able to implement might not even correspond to a real material, but in a sense, who cares?", says Chris Monroe, a physicist at the University of Maryland in College Park. Even if they don't behave like a superconductor or a Higgs particle, the new systems may still be able to tell researchers a thing or two that their older machines can't. Eventually, Monroe and others believe that simulators will be tailored to model different things. Cold atoms, for example, might work best on superconductors, whereas ions could handle magnetism. Of course, there will still be quantum systems that are too tough for any set-up to tackle.

It may be a vision considerably less flashy than Feynman's universal quantum machine, yet within the physics community, quantum simulators are getting more attention than ever before. "Many physicists who sort of pooh-poohed the idea of quantum computing, especially ten years ago or so, they're now sort of embracing this," says Monroe. The systems may be less ambitious, but that may make them more achievable.

Lloyd puts it another way. "If life doles you quantum lemons, let's make quantum lemonade," he says. Simulators may not be as sweet as quantum computers, but "as long as the lemonade is tasty and refreshing, I think that's fine". ■

Geoff Brumfiel *is a senior reporter at* Nature.

1. Endres, M. *et al. Nature* **487,** 454–458 (2012).
2. Feynman, R. P. *Int. J. Theor. Phys.* **21,** 467–488 (1982).
3. Britton, J. W. *et al. Nature* **484,** 489–492 (2012).
4. Perdomo-Ortiz, A., Dickson, N., Drew-Brook, M., Rose, R. & Aspuru-Guzik, A. *Sci. Rep.* **2,** 571 (2012).
5. Shor, P. W. *SIAM J. Comput.* **26,** 1484–1509 (1997).