


CHAMPING AT THE BITS

Despite some remaining hurdles, the mind-bending and frankly weird world of quantum computers is surprisingly close. **Philip Ball** finds out how these unusual machines will earn their keep.



Five years ago, if you'd have asked anyone working in quantum computing how long it would take to make a genuinely useful machine, they'd probably have said it was too far off even to guess. But not any longer.

"A useful computer by 2020 is realistic," says Andrew Steane of the quantum-computing group at the University of Oxford, UK. David Deutsch, the Oxford physicist who more or less came up with the idea of quantum computation, agrees. Given recent theoretical advances, he is optimistic that a practical quantum computer "may well be achieved within the next decade".

This excitement is, however, tempered by the hurdles that have yet to be overcome. Building a quantum computer is still very, very hard to

do. This is partly because it involves making quantum systems do things that don't come naturally to them. "There is progress, but it's still very slow," says physicist Chris Monroe of the University of Michigan in Ann Arbor.

And even if we did have a working quantum computer today, there are hardly any programs that could run on it. In fact, it is likely that even once the machines are available, quantum computers are destined to remain niche products — excellent for certain tasks but not versatile devices like conventional personal computers. "Quantum computers will almost certainly never become general-purpose desktop machines," concedes Isaac Chuang, a quantum physicist at the Massachusetts Institute of Technology (MIT) in Cambridge.


Nevertheless, as a scientific research tool the quantum computer could be revolutionary because of its ability to simulate other quan-

tum systems. In conventional, or classical, computers, information is stored as strings of bits: binary digits each of which can take the value of 0 or 1. The same is true for quantum computers, except that this time the binary digits — 'qubits' — are stored in the quantum states of microscopic systems, such as the electronic state of an atom or ion. So by its very nature, a quantum machine should be much better suited to simulating quantum systems than a classical computer.

A quantum simulator would describe and predict the structure and reactivity of molecules and materials by accurately capturing their fundamental quantum nature. This is the sort of employment the early machines are likely to find: doing calculations of interest to chemists, materials scientists and possibly molecular biologists, says Steane.

"Just a few dozen qubits may shed light on

J. MAGEE



M. BARRETT & J. JOST

other physics problems that are intractable with conventional computers," notes Monroe. "There are models of high-temperature superconductivity and other condensed-matter systems that might be approached in such a quantum simulator."

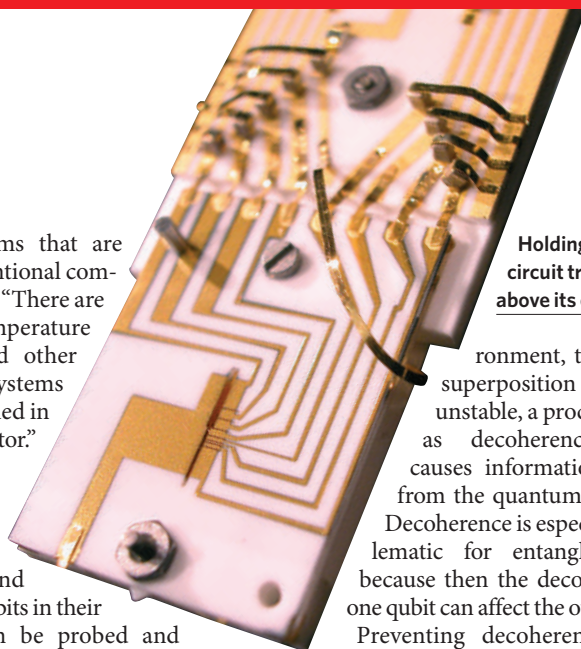
In a spin

In fact, quantum simulations can already be done using atoms and molecules that store qubits in their nuclear spin and can be probed and manipulated using nuclear magnetic resonance (NMR) techniques. In their own terms, these 'computers' "run rings around any classical supercomputer", says Seth Lloyd, a theorist at MIT. He and his MIT colleague David Cory have been using this technique to simulate a variety of quantum systems in crystals of calcium fluoride and other materials. "As the crystal contains a billion billion spins, these simulations remain out of the reach of the most powerful classical computers," says Lloyd. The approach remains limited in terms of the different systems it can simulate, although Lloyd anticipates that fully functioning simulators will be readily available by 2020.

The key to the potential success of quantum computers is also the cause of the problems within the field: the quantum nature of data storage and manipulation. In classical computers, bits have clearly defined values of 1 or 0, but the laws of quantum mechanics allow qubits to exist in a 'superposition' of states — a mixture of both 1 and 0 that would be impossible in an everyday computer. This means that a quantum computer has much greater capacity for storing information.

A quantum processor can also compute with more than one qubit at once by exploiting another quantum property called entanglement, which makes qubits interdependent. The weird nature of the entangled state means that a measurement on one qubit instantly affects another, even though their previous individual states were undefined until that moment. Entangled states don't readily exist in nature: quantum engineers have to make them by allowing qubits to interact with one another.

By exploiting superpositions, a single quantum computer in effect mimics a whole suite of classical computers running at once, and by using entanglement these 'parallel computers' can be linked together. Unfortunately, this powerful parallel processor has an Achilles' heel. A quantum superposition has to remain stable for at least as long as it takes to do the computation. But as soon as qubits interact with their envi-



Holding pen: this circuit traps ions above its electrodes.

ronment, the delicate superposition becomes unstable, a process known as decoherence, which causes information to leak from the quantum computer. Decoherence is especially problematic for entangled states, because then the decoherence of one qubit can affect the others too.

Preventing decoherence means reducing uncontrolled interactions with the environment. Cooling the quantum system to very low temperatures helps — but it may also be necessary to shield the qubits from stray electromagnetic fields. In practice, researchers have found it difficult to avoid decoherence of specific qubits for longer than a few seconds. But in principle it should be possible. "For qubits encoded in trapped ions, nobody really believes that we will ever be limited by coherence time," says Monroe.

Despite the fact that qubits need to be isolated from their environment to avoid decoherence, they must interact strongly with one another, to perform computations. And it must be possible for qubits in superposition to interact strongly with the environment when needed, so that the information can be read out. It is an extraordinarily delicate balancing act, which involves rules that defy intuition and aren't even completely understood.

An easy mistake to make


Decoherence also means that, as they process qubits using logic gates, quantum computers will inevitably incur errors at a much higher rate than classical computers. "The modern transistor has an error rate of less than 1 in 10^{14} or more switching events. In comparison, the best quantum gates we currently imagine will optimistically have an error rate of something like 1 in 10^7 ," says Chuang. Some researchers thought at first that this would make quantum computers too error-prone to be useful. But thanks to quantum error-correcting codes devised in the 1990s^{1,2}, it is now possible to correct error rates as high as 1 in 10^5 .

By 2002 the key principles behind a quantum computer had been sketched out by theorists (see 'How to build a quantum computer', overleaf), but how best to implement them in a real device remains a wide-open question. Much of the current effort is focused on making quantum computers using atoms or ions that are held in a trap. In an ion-trap computer, the qubits are encoded in the electronic states

MILESTONES IN SCIENTIFIC COMPUTING

PRE 1960s >>

1946 ENIAC, widely thought of as the first electronic digital computer, is formally unveiled. Designed to compute ballistics during the Second World War, it performs calculations in a variety of scientific fields including random-number studies, wind-tunnel design and weather prediction. Its first 24-hour forecast takes about 24 hours to do.



1951 Marvin Minsky, later of the Massachusetts Institute of Technology (MIT), builds SNARC, the first machine to mimic a network of neurons.

1954 John Backus and his team at IBM begin developing the scientific programming language Fortran.

1956 Building on earlier experiments at the University of Manchester, UK, and elsewhere, MANIAC at the Los Alamos National Laboratory in New Mexico becomes the first computer to play a full game of chess. In 1996, IBM's Deep Blue computer will defeat world chess champion Garry Kasparov.



1959 John Kendrew of the University of Cambridge, UK, uses computers to build an atomic model of myoglobin using crystallography data.



>> PROTOTYPES ...

of ions that are confined by an electromagnetic field. The ions interact with each other through electrostatic repulsion, and can be entangled by using laser beams to make them jiggle in unison. The quantum states of the ions can be read out by using other lasers to excite fluorescence, the wavelength of which depends on the ion's electronic state.

But the more qubits there are, the harder it is to read out their complex, collective vibrational states. One way to get round this is to hold most of the ions in a reservoir, and to perform each computational step using just a few of them, transferred from the reservoir to a processing chamber. This means that the ions have to be shuttled around without their quantum states being affected, so that they don't 'lose their memory' on the journey.

Charging ahead

Solving these problems is not easy, but recent progress has been encouraging. "Ion-trap chips look well placed to create useful computers before other methods," says Steane. Last December, for example, Monroe's team reported an ion trap built on a semiconductor chip using standard microfabrication techniques³. The trap held individual cadmium ions for more than an hour, while the researchers were able to move the ions smoothly between trapping sites.

David Wineland's group at the National Institute of Standards and Technology in Boulder, Colorado, is pursuing a similar idea in which the ions are trapped above electrodes etched into a chip's surface⁴. "Both methods have the advantage of using established fabrication techniques," says Wineland. "In the end,

How to build a quantum computer

The current US roadmap for the next decade of quantum computing (<http://qist.lanl.gov>) lists several requirements for a working machine⁸:

1. It must be scalable: it needs a set of qubits that can be added to indefinitely.
2. It must be possible to set all of the qubits to a simple initial state, such as all 0.
3. The interactions between qubits must be controllable enough to make quantum logic gates.
4. To perform operations using these gates, the decoherence times must be much longer than the gate-operation time (typically milliseconds to seconds).
5. There must be some readout capability.
6. To 'wire up' the computer's circuitry, it must be possible to convert memory qubits into processing qubits, and vice versa.
7. It must be possible to move processing qubits accurately between specified locations.

the winner might be determined simply by what is easier to fabricate."

Despite his success so far, Monroe is cautious about the long-term prospects. "Many groups are racing to build complex ion-trap chips," he says. "But it's less clear how trapped ions will ultimately compare with other quantum technologies."

Instead of ions, some researchers are encoding qubits using trapped neutral atoms. Atoms have the advantage that they interact more weakly with their environment than ions, but they also interact more weakly with each other.

They can be trapped by laser beams — and by exploiting the interference pattern generated between crossed laser beams, hundreds of atoms can be held within an 'optical lattice', rather like an egg box. To make the atoms interact, the dimples in the egg box can be shifted closer together by adjusting the trapping beams.

One way to perform quantum computing with atoms is to create discrete clusters of entangled atoms in a larger lattice. This was first suggested in the late 1990s by Hans Briegel, Ignacio Cirac, Peter Zoller and their colleagues at the University of Innsbruck in Austria. It is an approach to quantum computations that Deutsch describes as "far easier to implement physically" than other methods for handling qubits.


Unlike the standard approach, cluster computation does not involve manipulating individual particles. Instead, before the computation is run, several qubits are brought together in a many-particle entangled state. The answer is then read out at the same time as the computation is actually performed, by making a series of measurements on each individual qubit in the cluster. Its 'one-step' nature makes this an appealing approach, but Cirac, who is now based at the Max Planck Institute for Quantum Optics in Garching, Germany, admits that it requires many more qubits than other methods, and that the error-correction procedures are more elaborate.

Join the dots

There is no shortage of other ideas for building a quantum computer. Some are based on superconducting devices, exploiting the fact that superconductivity is itself a quantum phenomenon. Unlike the systems based on individual particles, the qubits here are superconducting circuits, which hold many-particle quantum states of electrical charge or magnetic flux and can interact through classical electromagnetic forces.

Others hope to create optical quantum computers, encoding qubits into the quantum states of photons, or to make qubits from tiny specks of semiconducting material called quantum dots. "I've been particularly impressed by the advances made in quantum-dot systems and by the superconductor-based approaches," says Chuang. Compared with ion-traps, he explains, they may scale up more easily and are perhaps easier to connect to traditional telecommunication systems for readout.


Quantum-dot systems may not produce the first useful computer, says Steane, but they have a naturally faster timescale — largely because the qubits are encoded in electrons, which are much lighter than ions — and so ultimately should outperform other systems such as ion-



Light touch: Chris Monroe aligns laser beams ready to trap ions for a quantum computation.

H.F. MONROE

D. L. STICK



Caught in a trap: a single cadmium ion is held between two electrodes on a semiconductor chip.

trap chips. But, he adds, “the area is still sufficiently open that it is premature to slow efforts on any of the major contenders”.

Because the hurdles in building the hardware are substantial, it is often suggested that the obstacles to making a quantum computer come down to engineering. But there is a bottleneck in theory too. So far, remarkably few specific computational problems have been translated into a form that a quantum machine could run and solve.

In fact, quantum computers are currently little more than two-trick wonders. In 1994, Peter Shor, now based at MIT, devised an algorithm that would allow quantum computers to factor numbers exponentially faster than conventional computers. Factorization is important in cryptography, where it is needed to make and break keys. And in 1996, Lov Grover, who is now at Lucent Technologies in Murray Hill, New Jersey, unveiled a quantum algorithm that can greatly speed up database searches.

Both of these algorithms have already been run using the NMR and optical techniques, but these methods are hard to scale up. At the end of last year, Monroe’s group reported success with a Grover-type search using two cadmium ions in a trap⁵. Admittedly this meant looking through a database of just four entries — hardly a demanding task — but Monroe says the group plans to scale up to dozens of qubits over the next few years.

“It is striking that ten years have passed

since Shor’s invention, and very few new quantum algorithms have been developed,” says Chuang. Among those that have appeared are methods for solving problems in number theory, drawn up by Sean Hallgren at NEC Laboratories in Princeton, New Jersey⁶.

A major stumbling block for those trying to dream up new algorithms is that they first have to identify which problems will benefit most from quantum-computing methods. Theorist Michael Nielsen and his colleagues at the University of Queensland in Australia have recently made progress in this direction by showing that the general problem of finding quantum algorithms can be made easier by borrowing ideas from geometry⁷.

In essence, the number of quantum operations, and thus the length of time, it takes to run an algorithm can be calculated by finding the shortest path between two points in a geometric space defined by all the possible sequences of quantum operations. “It really is a cool idea that has no classical analogue,” says Lloyd. “It opens up a variety of methods for potentially creating new algorithms and for optimizing existing algorithms.”

But not all quantum information processors will need complex algorithms. Many will be purpose-built tools that exploit quantum rules to improve on existing technologies such as atomic clocks and photonic technology. “We’ll probably see rudimentary devices

“Computers for specific applications are likely to come before general-purpose devices. But that doesn’t rule out the possibility that we’ll all be playing quantum Grand Theft Auto in the near future.” — Seth Lloyd

such as a ‘quantum repeater’ that converts photonic qubits to atomic qubits for error correction, and then back to photons to send them on their way down a long length of optical fibre,” Monroe says.

If that seems a far cry from the quantum brains that are sometimes paraded as the next big thing, we may just have to get used to it. But Lloyd remains upbeat about the prospects. “I agree that quantum computers tailored for specific applications are likely to be built before general-purpose devices. But that doesn’t rule out the possibility that we’ll all be playing quantum *Grand Theft Auto* in the near future.” ■

Philip Ball is a consultant editor for Nature.

1. Shor, P. W. *Phys. Rev. A* **52**, R2493–R2496 (1995).
2. Steane, A. M. *Phys. Rev. Lett.* **77**, 793–797 (1996).
3. Stick, D. et al. *Nature Phys.* **2**, 36–39 (2006).
4. Seidelin, S. et al. preprint at <http://www.arxiv.org/quant-ph/0601173> (2006).
5. Brickman, K.-A. et al. *Phys. Rev. A* **72**, 050306(R) (2005).
6. Hallgren, S. in *Proc. 34th Annu. ACM Symp. Theor. Comput.* 653–658 (ACM, New York, 2002).
7. Nielsen, M. A., Dowling, M. R., Gu, M. & Doherty, A. C. *Science* **311**, 1133–1135 (2006).
8. DiVincenzo, D. P. *Fortschr. Phys.* **48**, 771–783 (2000).

1960s >>

1962 Charles Molnar and Wesley Clark at MIT’s Lincoln Laboratory design the Laboratory Instrument Computer (LINC) for researchers at the National Institutes of Health. It is the first lab-based computer to process data in real time.



1963 In California, the Rancho Arm becomes the first artificial robot arm to be controlled by a computer.


1966 Cyrus Levinthal at MIT designs the first program to represent and interpret protein structures.

1967 ARPANET — the predecessor of the Internet — is proposed by the US Department of Defense for research networking.

1969 Results of the first coupled ocean-atmosphere general circulation model are published by Syukuro Manabe and Kirk Bryan, paving the way for later climate simulations that become a powerful tool in research on global warming.

1970s >>

1971 Computing power shows its potential in medical imagery with a prototype of the first computerized tomography (CT) scanner.



1971 The Protein Data Bank is established at Brookhaven National Laboratory in Upton, New York.

1972 Hewlett Packard releases the HP-35, the first hand-held scientific calculator, rendering the engineer’s slide rule obsolete.



>> MAIN FRAMES . . .

>> WORKSTATIONS . . .