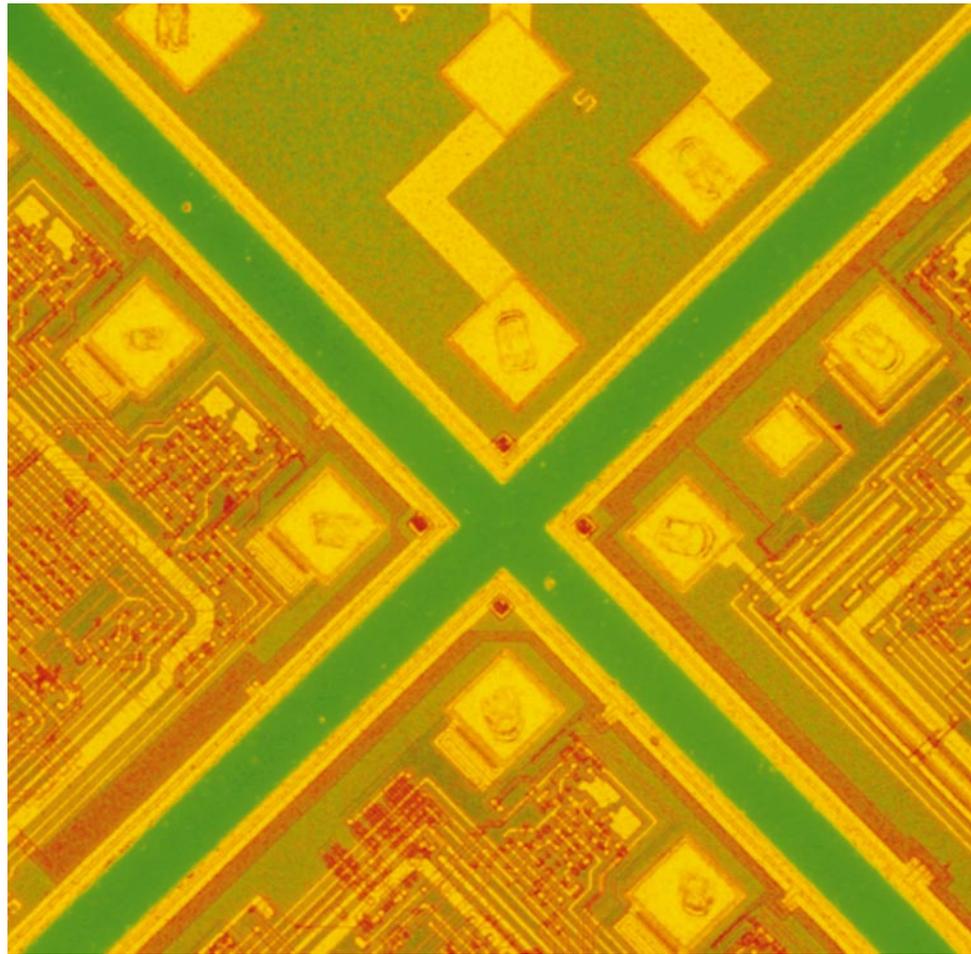


Quantum bits and silicon chips

Quantum computers offer a new kind of processing power. Silicon chips are easy to manufacture. Can the advantages of the two approaches be combined? Jenny Hogan reports.

For all their promise of fast and furious number-crunching, prototype quantum computers are bulky and impractical beasts. Some fill rooms but remain feeble calculators. Others require precision manufacturing that stretches current techniques to the limit. And most must be cooled to within a whisker of absolute zero before they can be used. They may have extraordinary potential, promising to crack complex codes and solve age-old mathematical puzzles, but the devices seem unlikely to slot neatly into desktop machines.

Not unless Marshall Stoneham is right. A materials scientist at University College London, Stoneham is using his experience in semiconductors to marry the worlds of silicon chips and quantum computing. He isn't the first to pursue the silicon approach, but the design he favours can, in theory, be manufactured with the tools available today. It should be powerful enough to do useful calculations and ought to operate at higher temperatures than rival devices — perhaps even at room temperature. Earlier this year, his team won £3.7 million (US\$6 million) from the UK government-funded research councils to put his idea into practice. Confident in his novel design¹, Stoneham taps his desktop computer



Thought processors: could today's silicon chips provide the basis for ultra-fast quantum computers?

and says: "If we succeed, we'll have a quantum processor on people's desks by 2010."

Like all quantum computers, Stoneham's design shares a common feature with conventional processors — both carry out calculations by manipulating bits of information, the ones and zeros of binary code. The simplest operation involves flipping the value of a bit, so that 1 becomes 0 or vice versa. More complex operations involve two bits. One bit may be flipped if the value of a second bit is 0, for example, but not if it is 1.

In principle, quantum devices work in the same way. But they also take advantage of the rules of quantum mechanics, which say that particles can exist in two states simultaneously. An electron that moves from one atom to a neighbouring atom, for example,

can actually be in both positions at once. If the electron is used to represent a bit, so that it is given the value of 0 when it is next to one atom, and 1 if it is next to the other, interesting possibilities emerge. Applying an external influence such as an electric field to the electron can flip the value of the quantum bit, or qubit. But because the electron is in two places at once, the operation is performed simultaneously on both states.

When qubits are linked together, this becomes a significant advantage. Two electrons, for example, can exist in four different states — 00, 01, 10 and 11 — depending on the relative positions of the particles. If the electrons interact with each other — a phenomenon called entanglement — then any operation carried out on one electron will

P. SALOUTOS/CORBIS; T. CRADDOCK/SPL

D. PARKER/SPL



simultaneously be carried out on the other. In effect, this means that one operation is carried out on four different states at the same time. Add more qubits and the number soars — a computer with just 30 qubits can perform more than a billion operations at once.

Not all computing tasks would benefit from such 'parallel' processing. A word-processing program, for example, probably wouldn't be enhanced by quantum hardware. But theory suggests that certain problems, such as factorizing large numbers — a crucial part of code-breaking — would yield to quantum techniques.

There are several ways to build a quantum computer, each of which stores and manipulates qubits in a different way. One method stores qubits in the energy states of an ion, and simple programs have already been run in this way. This January, for example, Rainer Blatt and his colleagues at the University of Innsbruck in Austria chilled a calcium ion to within a few millikelvins of absolute zero, and manipulated its energy states using laser pulses². One qubit was stored in the energy levels of the ion's motion, the other in the energy of its orbiting electrons. Blatt's team ran a simple algorithm that calculated whether a particular function always gives the same output, but researchers believe they will soon be able to combine several ions to perform more complex calculations.

In a spin

Others are using the spin of atomic nuclei — a quantum-mechanical property that aligns itself with or against an external magnetic field — to represent a 1 or 0. By adapting the nuclear magnetic resonance (NMR) technology that is used to image molecular structures, Jonathan Jones, a physicist at the University of Oxford, UK, created a liquid-based two-qubit device in 1998. The qubits were stored in the nuclear spins of hydrogen atoms in the biological molecule cytosine. The nuclei were constrained by strong magnetic fields, and radio waves were used to flip their spins³. The result of the calculation can be deduced from the spectrum of the interacting radio waves. NMR computers have since been extended to create the most advanced quantum computer so far — a seven-qubit device⁴.

But despite such progress, these systems have drawbacks. The equipment needed to manipulate ions fills a room, and the magnet-



Marshall Stoneham believes that he can make a viable desktop quantum computer by 2010.

ic fields used in NMR are generated by huge superconducting magnets that require cooling with liquid helium. The nuclear-spins approach also runs into trouble as the number of qubits increases, because the result signal gets lost in the noise from neighbouring molecules. In effect, the computer crashes if it uses eight or more qubits. Even with seven qubits "it begins to fall apart", says Jones.

But would nuclear spins be easier to control if they were embedded in a solid? Silicon is the ideal substance with which to test this idea. Decades of investment in conventional computers have generated a wealth of knowledge about how to fine-tune its properties, and theory suggests that silicon atoms should not interfere with the nuclei of whatever element is chosen to carry the qubits.

Up to five years ago such a system was mere speculation, but in 1998 a paper appeared that described a way in which the device could be made⁵. The author was Bruce Kane, then at the University of New South Wales (UNSW) in Sydney. Kane proposed burying phosphorus atoms 20 nanometres apart in a silicon grid. Each atom would have an electrode positioned above it. The qubits would be stored in the nuclear spins of the phosphorus atoms, and the voltage on the electrodes used to tune each nucleus to respond to a particular frequency of radio wave. In this way, the states of individual nuclei could be flipped using a radio signal sent to the chip. A second elec-

trode between adjacent phosphorus atoms would mediate the interaction between two neighbours, allowing operations to be performed using linked qubits.

The proposal has some distinct advantages. It would, for example, be possible in principle to build a structure containing thousands of qubits. And the qubits can be controlled by conventional electronics, which could, in theory, be easily integrated into the chip. "Even if there are some disadvantages to using silicon, the approach would probably win out over better technology because it would take advantage of existing infrastructure," says Stan Williams, director of quantum-science research at Hewlett-Packard Laboratories in Palo Alto, California.

Solid progress

With so much going for it, Kane's proposal was guaranteed attention. "There was a phase transition when the paper came out," says Henry Everitt, associate director of the physics division of the US Army Research Office in Durham, North Carolina, and chair of the US government's Quantum Information Science Coordination Group. "People who thought quantum computing was just a pipedream started thinking it could happen."

Robert Clark, a colleague of Kane's at UNSW, was one of the researchers to follow up the new idea. He is now director of Australia's Centre for Quantum Computer Technology, a multi-university initiative set up in 2000 that embraces about 150 researchers and that has made good progress towards making the electrodes and phosphorus grids that Kane's design requires.

There are two basic approaches to the construction of the device⁶. Bottom-up techniques use tools, such as a scanning tunnelling microscope (STM), that can manipulate individual atoms. STMs create atomic-resolution images by scanning a microscopic tip over a surface, but they can also manipulate individual atoms. The researchers use the STM to remove an atom from the layer of hydrogen that caps the surface of silicon wafers. When phosphine gas (PH₃) is blown over the wafer, the molecules drop into the gaps. The hydrogen can then be removed by heating the wafer. The team is currently trying to work out how to add a top layer of crystal-perfect silicon, which will be used to protect the phosphorus atoms.

Such manipulations represent the more



precise approach, but prototype devices are easier to make top-down. In this method, the Australian team covers the surface of a silicon wafer with a thin layer of a polymer containing an array of nanometre-sized holes. Next, the wafer is bombarded with phosphorus atoms. One atom shoots down each hole, punching its way about 15 nm into the silicon. The polymer around the hole is replaced and then carved away to make a stencil through which the circuits that control and read the qubit can be painted. “Many people looked at the proposal and said that the technology was decades away,” recalls Gerard Milburn, a theoretical physicist at the University of Queensland and deputy director of the Centre for Quantum Computer Technology. “Well, I think we’ve proved them wrong.”

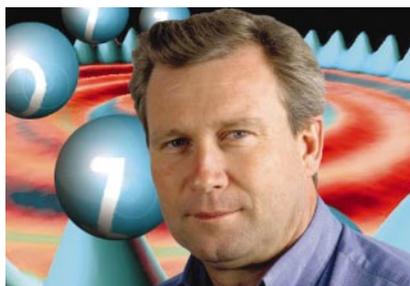
Rather than working with nuclear spins, which are relatively difficult to manipulate, the Australian team is focusing on the outermost electrons of the phosphorus atoms. The ones and zeros are stored in the position of an electron that can hop between neighbouring phosphorus atoms. Clark’s team is currently attempting to track the motion of the electron using a charge detector, and says that initial results are promising.

Under control

Researchers are impressed with the Australian team’s progress. Adding phosphorus to silicon is common in the computer-chip industry, but no one knew how to do it one atom at a time when Kane made his proposal. “What they’ve done is spectacular,” says Thomas Schenkel, a physicist at Lawrence Berkeley National Laboratory in California, who is working on a different implantation technique. “For the first time we can attempt to do electronics on a single-atom level.”

Impressive, yes, but not yet quantum — the Australians have yet to obtain a working qubit. Other groups are also working on ideas stimulated by Kane’s paper. Crispin Barnes, a semiconductor physicist at the University of Cambridge, UK, is working with colleagues to implant sodium atoms in silicon and store qubits in electron spin. And Kane, now at the University of Maryland in College Park, is working on techniques for measuring the spin of a single electron⁷.

The researchers pursuing these and other approaches believe that they will succeed, but there may be bigger problems ahead. Some believe that the electric fields assoc-



Gerard Milburn: rapid progress is being made towards silicon-based quantum computers.

iated with the control apparatus could interfere with the qubits, disrupting the calculation. Others warn that slight variations between supposedly identical qubits, introduced by defects in the silicon or errors in aligning the electrodes, could render silicon quantum computers unworkable⁸.

Follow the light

This is where Stoneham’s plan may have the edge. He proposes to manipulate the spins of electrons using laser light. In his design, the embedded atoms — he is currently considering which element to use — are arranged randomly in silicon, with no electrodes on top of them. Unlike the other silicon-based systems, which rely on ordered arrays, explains his colleague Andrew Fisher, Stoneham’s design benefits from this disorder. The qubits are stored in the outermost electron on each atom. Because every atom has a unique set of neighbours, the qubit electrons each experience different forces and so can be manipulated individually using a specific frequency of laser light.

Linking the qubits together, however, is difficult. Stoneham plans to use ‘control atoms’ that are implanted across the wafer. As with the other embedded atoms, control atoms should be sensitive to a particular frequency of laser light. By tuning the pulse, the control atoms can be excited so that their electrons interact with the electrons that store qubits in neighbouring embedded atoms. This links nearby qubits, creating a two-qubit system. The states of the system, and hence the result of a calculation, are read by looking at the way in which photons scatter off the linked electrons.

Stoneham thinks that his processor should run at much higher temperatures than other silicon quantum computers. In

principle, other designs could work at up to 4 K. Above this, the electrons that carry the qubits will be disrupted by vibrations of the silicon lattice. But different atoms are sensitive to vibrations of different energies. Stoneham believes that some elements — perhaps bismuth — will bind their electrons in such a way that they are immune to the interfering vibrations. He is confident that his processor will work above 4 K, and maybe all the way up to room temperature. With the tools needed to build his processor already available, he is aiming to have a three-qubit system running in four years’ time.

But that is a long way from a quantum processor with tens of bits, which Stoneham suggests he will have by 2010. He is a “brave man” to make such a claim, says Milburn. “I know I’m sticking my neck out,” acknowledges Stoneham. But not everyone thinks he is being unrealistic. “The physics has not got far enough for us to say to Intel or Motorola: ‘this is the chip we want you to build,’” says Everitt, “but that day may come in the next decade.”

So what effect will Stoneham’s ideas, and the other silicon devices, have on the race to build a practical quantum computer? The issue was addressed in a report published last year by the Advanced Research and Development Activity, a US government body that funds high-risk information technology research. It concluded that it is too soon to pick a winner, and that the ultimate technology may not have been invented yet.

Even Kane is hedging his bets. This June, the US Defense Advanced Research Projects Agency organized a meeting on quantum computing in Los Angeles, California. At it, Kane was asked whether the agency should initiate a quantum-computing version of the Manhattan project — the Second World War scheme that brought together high-level scientists and provided them with the investment needed to develop the first nuclear weapons. How did Kane respond? “I said no. We can’t do that because we haven’t even discovered fission yet,” he says. ■

Jenny Hogan is a reporter for *New Scientist*.

1. Stoneham A. M., Fisher, A. J. & Greenland, P. T. *J. Phys. Condens. Matter* **15**, L447–L451 (2003).
2. Gulde, S. *et al. Nature* **421**, 48–50 (2003).
3. Jones, J. A. & Mosca, M. J. *Chem. Phys.* **109**, 1648–1653 (1998).
4. Vandersypen, L. M. K. *et al. Nature* **414**, 883–887 (2001).
5. Kane, B. E. *Nature* **393**, 133–137 (1998).
6. Clark, R. G. *et al. Phil. Trans. R. Soc. Lond. A* **361**, 1451–1471 (2003).
7. Kane, B. E. *et al. Phys. Rev. B* **61**, 2961–2972 (2000).
8. Keyes, R. W. *Appl. Phys. A* **76**, 737–741 (2003).