Also in this section

Electric motors

Reluctant heroes

An electric motor that does not need expensive rare-earth magnets

DYSPROSIUM and neodymium are not exactly the best-known elements in the periodic table, but for makers of high-end electric motors they have become vital. Both are strongly magnetic and thus crucial to the construction of powerful motors of the sort used, for example, in electric cars. Unfortunately, they lurk in the part of the table known as the rare-earth metals and, as that name suggests, workable deposits of them are scarce. At the moment, the main source of supply is in China, whose government has used its near-monopoly to restrict availability and push up the price. So there is a lot of interest in inventing motors that can do without them. And several groups of researchers think they have come up with one.

The device in question is known as a switched reluctance motor. The idea behind it is over 100 years old, but making a practical high-performance version suitable for vehicles has not been possible until recently. A combination of new motor designs and the advent of powerful, fast-switching semiconductor chips, which can be used to build more sophisticated versions of the electronic control systems required to operate a reluctance motor, is giving those motors a new spin.

One of the leading contenders is Inverto, a research and development company based in Ghent, Belgium. Inverto's engineers, led by John De Clercq, the firm's research director, are collaborating with the University of Ghent and the University of Surrey, in Britain, and also with an unnamed carmaker. They already have a motor running in a car. At Newcastle University, also in Britain, researchers are working with several companies to produce reluctance motors for both cars and lorries. And studies are being carried out in America and Japan too. A team led by Nobukazu Hoshi of the Tokyo University of Science, for example, has experimented with a reluctance motor in a Mazda sports car.

Let's twist again

The main practical difference between a reluctance motor and a normal one is that the reluctance motor has no permanent magnets. A normal motor (one that runs on direct current, at least) has two parts. One, the rotor, moves. The other, the stator, does not. The stator usually forms the casing and houses a set of permanent magnets. The rotor, which turns inside that casing, is surrounded by copper windings that act as electromagnets. The motor works by the periodic reversal of the current running through these windings. That reverses the polarity of the electromagnets and causes the rotor to be pulled around by attraction and repulsion between the electromagnets and the permanent magnets.

Traditionally, this current reversal was achieved with a rotary switch known as a commutator, which transmitted electricity to the windings through brushes that made contact with conductive strips which passed under them as the rotor turned. Modern motors use electronic control systems rather than a commutator, but the principle is the same.

A switched reluctance motor is like a normal modern motor in this respect—it is also brushless. But unlike a normal motor it has no need of permanent magnets, rare-earth or otherwise, for it works on a different principle: least magnetic reluctance.

Reluctance in magnetism is analogous to resistance in electricity. Just as a current travels along the path of least resistance, so the flux of a magnetic field (the "lines of force", for those who remember childhood experiments with bar-magnets, sheets of paper and iron filings) takes the path of least reluctance. And iron, a material that is nice and cheap, has very low reluctance.

Inverto's reluctance motor has a rotor made of iron sheets, while the inside of the stator is covered with copper windings. Current is fed to these windings on the say-so of the control system, generating flux. That flux then follows the path of least reluctance—ie, through the sheet-iron rotor. The rotor attempts to align itself to the flux in a way that reduces reluctance to a minimum, and that causes it to turn.

The control system, however, constantly anticipates the rotor's movement and switches the current between windings so as to stop the rotor settling into its preferred alignment. As a result, it keeps on turning.

Reluctance motors still have disadvantages, even now that the control problem has been solved. For example, to deliver a given amount of twisting force—or tor- ▸▸

que—a reluctance motor has to be larger than an equivalent permanent-magnet motor. But the materials needed to build them are significantly cheaper. Moreover, according to Dr De Clercq, the motors' torque characteristics make them particularly suitable for cars.

At high speeds, for example, they do not lose torque as quickly as a permanent-magnet motor would. That helps during overtaking. They are also safer when they fail. If a permanent-magnet motor loses power suddenly it slows down rapidly, creating an unexpected braking effect. That might also damage the motor. A reluctance motor, by contrast, freewheels if the power is cut off. Despite that, it can still act as a generator when slowing down, as permanent-magnet motors do in electric cars. (Generators are, basically, electric motors in reverse. They convert motion into electricity instead of electricity into motion. A slowing motor can thus be used to top up a car's batteries with energy that would otherwise be lost in braking.)

The researchers at Newcastle are working with Tata Steel, part of an Indian conglomerate, on using special steels to exploit the magnetic flux more effectively and hence extract more power. They are also collaborating with a local firm, Sevcon, which makes power electronics, and Cummins Generator Technologies, part of an American group which builds engines, on reluctance motors for hybrid trucks. James Widmer, who leads the Centre for Advanced Electrical Drives, one of the reluctance-engine-research groups at Newcastle, says that by running their new motors at high speeds it should be possible to outperform the best fixed-magnet motors, and to do so at lower cost.

The chances are good, therefore, that powerful switched reluctance motors will emerge. That does not automatically mean they will be commercialised, but if the price of rare earths remains high, there is a strong chance that will happen. And if it does, the need for dysprosium and neodymium will diminish quite a lot. China may thus find that, like the farmer in the fable who killed the goose that laid golden eggs, its desire for short-term advantage brings long-term regret. ■

Supercomputing

Deeper thought

The world has a new fastest computer, thanks to video games

SPEED fanatics that they are, computer nerds like to check the website of Top500, a collaboration between German and American computer scientists that keeps tabs on which of the world's supercomputers is the fastest. On November 12th the website released its latest list, and unveiled a new champion.

The computer in question is called "Titan", and it lives at Oak Ridge National Laboratory, in Tennessee. It took first place from another American machine, IBM's "Sequoia", which is housed at Lawrence Livermore National Laboratory, in California. These two machines have helped reassert America's dominance of a list that had, in the past few years, been headed by computers from China and Japan.

Titan is different from the previous champion in several ways. For one thing, it is an open system, meaning that scientific researchers with sufficiently thorny problems will be able to bid for time on it, in much the same way that astronomers can bid for time on telescopes. Sequoia, by contrast, spends most of its time running hush-hush simulations of exploding nuclear weapons, and is therefore rarely available for public use.

Titan has an unusual design, too. All supercomputers are composed of thousands of processor chips harnessed together. Often, these are derivatives of the central processing units, or CPUs, that sit at the heart of modern, desktop machines. But Titan derives the majority of its oomph—more than 90%—from technology originally developed for the video-game industry. Half of its 37,376 processors are ordinary CPUs. But the other half are graphics processing units, or GPUs. These are specialised devices designed to cope with modern video games, which are some of the most demanding applications any home machine is ever likely to run. China's "TianHe-1" machine, a previous Top500 champion, was built in the same way, as are 60 other machines in the Top500 list.

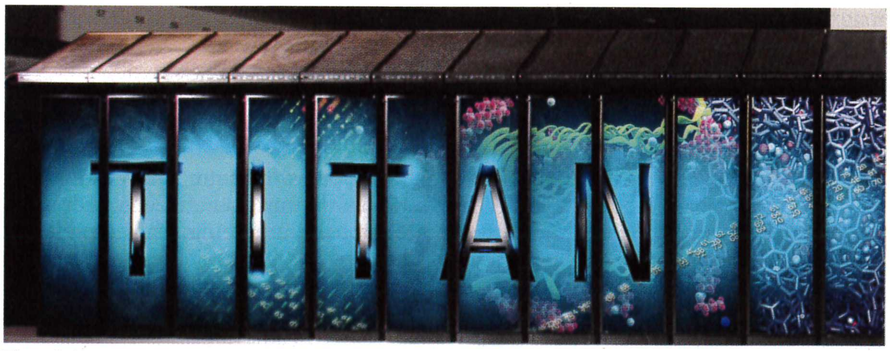

The ultimate games machine

Parallel worlds

Broadly speaking, a CPU—which will be expected to run everything from spreadsheets to voice-recognition software to encoded video—has to be a generalist, competent at every sort of mathematical task but excelling at nothing. A GPU, by contrast, is designed to excel at one thing only: manipulating huge numbers of the triangles out of which all modern computer graphics are made.

Several years ago researchers at Nvidia and AMD (the two companies that produce most of the world's high-performance GPUs) realised that many scientific problems which demand huge amounts of computing power—everything from climate simulations and modelling combustion in an engine to seismic analysis for the oil-and-gas industry—could be translated into a form that was digestible by their GPUs. Soon after, supercomputer builders such as Cray (which put Titan together using Nvidia's GPUs) began to take notice.

Borrowing from the games industry in this way brings several benefits. One big one is efficiency. Titan is an upgrade of Oak Ridge's existing "Jaguar" machine. Upgrading Jaguar with ordinary CPUs would have meant building a computer that sucked around 30MW of electricity when running flat out—enough juice to power a small town. Because GPUs are so good at their specialised tasks, Titan can achieve its blistering performance while sipping a (relatively) modest 8.2MW.

It makes sense financially, too, says Sumit Gupta, head of supercomputing at Nvidia. The chips that the firm sells to supercomputer-makers are almost identical to those it sells to gamers. As Dr Gupta observes, "The history of high-performance computing is littered with the bodies of firms that tried to build products just for the supercomputing market. By itself, it's just too small a niche."

It is not all upsides, though. Machines like Titan achieve their speed by breaking a problem into thousands of tiny pieces and farming each out to a single processor. A ▸▸