nia is to use a chemical called cyanuric acid, which is commonly used in swimming pools. The Sandia group, led by Dr Robert Perry, recently reported that it could remove virtually all the NOx from the exhaust of a small diesel engine at temperatures above 450°C. America's Energy Department has granted Dr Perry all rights to the idea. Once he has raised the capital, he plans to start his own company, Technor, to market it. He is aiming at diesel engines first and coal-fired boilers later. He thinks vehicles could simply carry refillable tanks of cyanuric acid as they do petrol tanks.

Cyanuric acid has its drawbacks. It is still fairly costly, at nearly $1 a lb or about $2,000 a tonne of NOx removed. This is cheaper than catalytic reduction but not by much. It is not water-soluble, which makes it hard to inject into a boiler, and its efficiency varies with temperature. These are all points on which another technology scores.

The rival process is touted by Fuel Tech, a company based in Stamford, Connecticut, that is financed by about $50m, most of it from British investors. America's Electric Power Research Institute developed a technology to reduce NOx using urea instead of ammonia, but this was limited in temperature range and efficiency. So Fuel Tech tested a variety of chemicals to make urea more effective, to broaden the temperature range and to cut the production of ammonia as an irksome side-effect. It has devised a cocktail of molecules that it calls NOxOUT. Though it will not say exactly what they are, all are cheap (less than one-tenth the cost of cyanuric acid), readily available from the chemical industry, harmless and water-soluble.

Fuel Tech, which has disclosed its plans for the first time to *The Economist*, is much closer to market than the Sandia group. It has tested NOxOUT in several types of boilers, including a 110-megawatt boiler fired by gas and oil in California. It is now conducting a full-scale test on a coal-fired boiler in Europe. Its aim is to reduce NOx by 75% with no significant ammonia "slippage". In ideal conditions, it can achieve that goal over a wide range of temperatures (500-1,000°C), but the company's chairman, Mr William Haney, stresses that each application must be tailor-made for a particular power plant or industrial boiler, so it is hard to draw general conclusions.

Nonetheless, rumours about Fuel Tech's cocktail are making European utilities and equipment manufacturers prick up their ears. It promises efficiency similar to catalysts at one-tenth of the cost (see chart on previous page), one-hundredth as much equipment and one-third of the installation time.

Data protection

Enigma variations

Here is a brutal way to show that your product performs better than its competitors: make those competitors look pretty useless. A young British cryptographer and computer programmer, Mr Martin Kochanski, has done approximately that. He has taken the five best-selling data-encryption programs for IBM personal computers and cracked them. To the annoyance of his rivals, *Cryptologia*, an American codebreakers' monthly, has published the method he used. Unravelling these codes is a new game for computer buffs.

Not so cryptic after all

There are half-a-dozen coding programs on the market. The programs take computer files and jumble up the binary digits that are used to represent letters and numbers inside a computer. The result is an apparently unintelligible—and apparently secure—splurge of characters on the screen. Mr Kochanski wrote a program that works out the key used to do the coding from a piece of jumbled text. He claims it took less than two hours to do this for each rival system.

Mr Kochanski is clearly better than average at code-cracking, but many keen amateurs can rival him. Manufacturers are wrong when they claim it takes long hours on a $15m Cray supercomputer to crack their codes. A spare Sunday afternoon and a $500 personal machine are all that are needed.

The first trick is to create a file with a monotonous stream of characters—say 30,000 zeroes. The result is likely to betray some sort of repeating pattern, because of the way most ciphers are designed. Each character in the text to be encoded (called the plaintext) gets paired off with a character from an endlessly repeating sequence called a key. Many of the available packages use keys that are partly inspired by the program's serial number. Another part of the key may be generated automatically from the date, or the name of the file to be coded. A third part is usually thought up by the user.

After matching a character of plaintext with a character of key, the coding algorithm combines them to create a single character of ciphertext. As the key repeats, so will the ciphertext—so long as the plaintext is a repeating string of uniform characters. That reveals the first clue: the length of the key. By substituting new characters in the plaintext one by one, and then re-encoding the string, it is possible to see just what the algorithm does with the key. Mr Kochanski says he was amazed at the simplicity of some of the algorithms used.

The would-be codebreaker is still only half way there. Once the coding algorithm is cracked, he still has to work out the key. The relative frequencies of different characters provide a clue. For instance, the space character is by far the commonest in files with a lot of text. Next (in English) come E, T, A, O, and so on. Zero is the most common character in files devoted to data. Knowing such things helps cut down the number of possibilities a code-breaking computer has to work through in order to find the key. Mr Kochanski found that once he had deduced the programs' algorithms, he could write programs that would find the keys used for particular files in less than a minute.

Where does that leave his own product, Ultralock, which encodes files for the IBM PC? Mr Kochanski does not claim that his codes are unbreakable. He notes that even the Data Encryption Standard used by America's federal government is breakable—so long as the cracker has a Cray supercomputer and a few months. He does say his algorithm is many times more complicated than his rivals', and that he could not crack it himself. Perhaps his rivals can: they can buy his program for £190 ($290) and get cracking.

Surface science

The secret seven

At the moment, electronics researchers prefer not to think about what happens on the surfaces of silicon crystals. They will soon have to. The atoms inside a millimetre-sized crystal outnumber those on the surface by about a billion to one. But as circuits get smaller, surfaces become more important—and chips' properties will eventually be dominated by their surfaces.

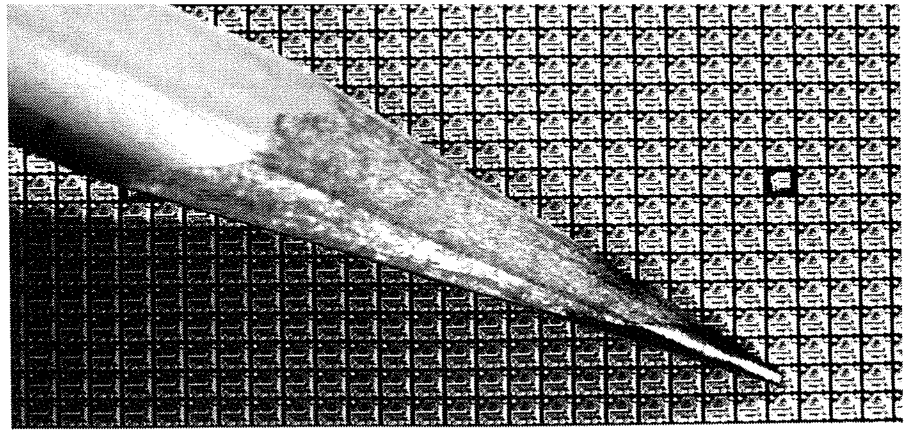

Surfaces merit a closer look

To see why, consider a deflating balloon. As it gets smaller, it has a larger surface-to-volume ratio: the volume decreases with the cube of the radius, but the surface area falls off more slowly with the square of the radius. The same is true of the circuits on a chip as they shrink out of sight. Despite the R&D money lavished on silicon chips, seemingly easy questions about the surfaces of silicon crystals stumped researchers.

A crystal is made of an orderly stack of building blocks called unit cells. The simplest silicon unit cell contains only two atoms. Of the many ways to cut a silicon crystal, the one that most interests surface scientists is done by cleaving it apart. The resulting surface is cleaned by heating it to about 1,000°C in a vacuum chamber (in air, it would oxidise). This rearranges the atoms on the surface to form a giant unit cell containing not two but about 100 atoms—so the new unit cell is dubbed the silicon 7x7 reconstruction, because it is 49 times larger than the usual unit cell.

That much was clear more than 25 years ago. Yet the detailed atomic structure of the 7x7 unit cell has eluded scientists until now. The search for an answer culminated in the award of the 1986 Nobel prize in physics to Dr Gerd Binnig and Dr Heinrich Rohrer of the IBM Research Laboratory in Zurich, inventors of the scanning tunnelling electron microscope. Much simpler in principle than the traditional electron microscope, the scanning tunnelling microscope is a fine tungsten needle that is drawn across the surface, displaying the atomic topology in much the same way as a record-player stylus picks up the modulation of the grooves of a record.

The first pictures of the silicon 7x7 surface dramatically confirmed a model proposed by Dr Kunio Takayanagi of the Tokyo Institute of Technology. The model requires extra silicon atoms, called adatoms, to stick to the surface at specific sites, making the surface chemically more stable. In pictures through the scanning tunnelling microscope, these adatoms showed up clearly. They surprised researchers, who have always had to look at surface structure in indirect, error-prone ways.

Once the adatoms were spotted, the last pieces of the silicon 7x7 puzzle fell into place. More adatoms can be fitted on to the surface by introducing stacking faults—regions where the stacking of the surface-layer atoms is out of step with the underlying crystal structure. But too many stacking faults make the surface unstable. To the theorists' delight, it turns out that the crystal is most stable electronically when the stacking-fault regions criss-cross the surface in a 7x7 pattern.

Dr Ian Robinson, a British researcher working at AT&T's Bell Laboratories in New Jersey, came up with results that support the same model, using a technique called surface X-ray diffraction. Since X-rays penetrate deep into the crystal, they might not seem the ideal tool for studying surfaces. But this handicap can be used to advantage to study "interfaces"—the structures at the joins between materials.

Metal-oxide-silicon (MOS) chips use an interface between a silicon crystal and silicon oxide. Dr Robinson's team has confirmed that even if a thick oxide layer is deposited on silicon the 7x7 structure is preserved. This contrasts with many other surfaces, which are violently disrupted by an oxide layer, with damaging consequences for the electronic properties of the interface. All handy stuff to know, in a chip generation or two.

Dancing on a pin-head

Common sense reveals that a man cannot easily walk through a wall. Early this century, quantum mechanics revealed that electrons are more agile. Minute particles such as electrons can spontaneously—though only occasionally—tunnel through barriers that an earlier generation of physicists thought impenetrable. The scanning tunnelling microscope (STM) exploits this oddity in increasingly useful ways.

The STM "looks" through a thin needle of tungsten placed just above the surface to be analysed. Thanks to a difference in electrical potential between the surface and the needle, electrons tunnel out of the lattice of atoms that forms the surface, across the gap and up into the needle itself. The needle registers this trickle of charged particles as a current: the smaller the gap, the bigger the current. By scanning the needle across the surface and keeping the current constant, the bumps and dips caused by individual atoms can be mapped.

The needle must be manipulated to within billionths of a metre. This is achieved by means of piezoelectric materials that can be made to change their size—and hence the position of the needle—by tiny amounts when a voltage is passed through them. Because the tunnelling effect works only if the gap between needle and surface is exactly the right size, a tiny protuberance of a few tungsten atoms can affect the current. To provide precise information about the surface, the needle has to be as thin as possible so that the width of the stream of tunnelling electrons is kept narrow. Researchers at IBM in Zurich (where the STM was invented) are developing needle points that taper down to a single atom.

Three scientists at AT&T's Bell Laboratories in Murray Hill, New Jersey, have used an STM to demonstrate a virtuoso piece of micro-engineering. In the magazine *Nature*, Dr Russell Becker, Dr Jene Golovchenko and Dr Brian Swartzentruber report that they placed one or two germanium atoms (the "pile" is so small that not even the STM can yet tell precisely how many) on to a surface of the same material, which is used in semiconductors. They found that by increasing the gap voltage they could evaporate the germanium atoms off the needle, across the gap and on to the surface.

Attempts to carry out the same process with silicon have so far failed—the scientists speculate that this is because atoms of silicon bond more tightly to tungsten. The precise mechanism which allows atoms to be deposited in this fashion remains mysterious. Nevertheless, some potential applications are already clear. Physicists and materials scientists can look forward to examining the effect of a single extra atom on the well-ordered atomic arrays in surfaces. And, more practically, as Dr Becker and his colleagues speculate, new devices dependent on the electronic properties peculiar to single-atom impurities become feasible—as do high-density computer memories, in which a single atomic rearrangement represents a piece of information—the smallest "bit" possible.