The future's bright

The processing power of silicon chips and the unmatched speed of photons make an awesome combination. *John McCrone* says computer designers are beginning to see the light

"THE revolution is nigh," proclaimed the computer gurus to anyone who would listen. "The next big thing will be optical computing." According to their predictions, these phenomenally powerful machines would soon replace electricity with light beams and process data at incredible speeds. But the gurus were wrong. Optical processing proved too difficult, and the early enthusiasm has evaporated.

Now optical computers are edging back into the limelight. But they are not the machines that were predicted 10 years ago. As computer chips have become ever faster, it is the moving rather than the processing of data that is becoming the key issue. Wires simply can't satisfy the appetites of these information-hungry number crunchers. So future generations of computers may be hybrids, with their fast silicon chips linked together by thin beams of laser light. There will be no metal wiring, just a maze of lenses, prisms and mirrors to shift data from one chip to another. Hence the name: free-space optics.

Today's machines move millions of bytes every second. Changing to silicon/optical hybrids will allow computer designers to build terabyte and petabyte machines that will move trillions of bytes each second.

Computing with light has three clear advantages over wires. Nothing moves faster than a photon—electrons in wires move at a snail's pace by comparison. More importantly, photons do not interact with each other the way charge-bearing electrons do. Separate beams of light simply pass straight through each other, so incredible amounts of data can be sent through the same narrow corridor of space. Best of all, light does not need a physical connection. If ordinary lenses and lasers can be made small enough to attach to the back of a microchip, then tomorrow's computers could simply route their signals through thin air.

The switch to hybrid computers may occur sooner than people think, says Julian Dines of the optoelectronics team at Heriot-Watt University in Edinburgh: "Most of the fundamental technology problems have actually been solved. The problems now are engineering ones—how to make lasers, lenses and mirrors small enough, reliable enough, and of course cheap enough, to go into commercial systems."

Crammed in

Dines is in no doubt that there is a need for change, and believes that the computer industry recognises this too. "Certainly, chip makers like Intel can see that getting data into and out of their chips is becoming a bottleneck," he says. "Just physically, it's a problem fitting more than a few hundred metal pins around the edge of a microprocessor. But with free-space optics, you could get the equivalent of many thousands of pins going into a chip."

Illustrations: Ian Jackson and Tom Main

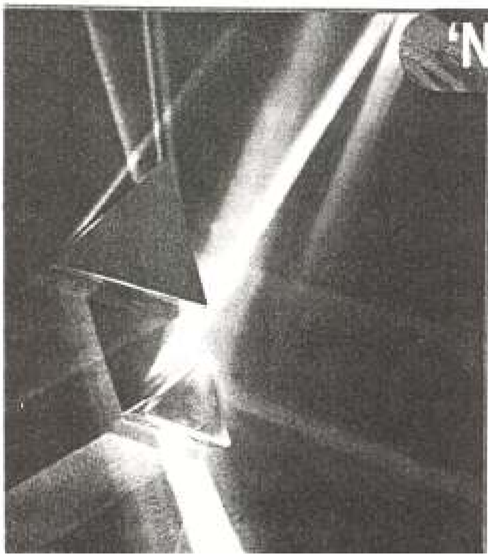

'No metal wiring, just a maze of lenses, prism

Free-space optics can make use of the whole surface of a chip for connections, covering it with an array of thousands of micrometre-sized optical connections, rather than being restricted to its four thin edges. Computers with large numbers of these much faster connections should be able to move more than a terabyte (10^{12} bytes) of data a second, compared with the microprocessors in today's desktop computers, which handle less than a

Light on: Julian Dines reckons hybrid computers are in

and mirrors to shift data around'

gigabyte (10^9 bytes) a second. The difference would be like that between a dripping tap and a firehose, greatly speeding up existing computer designs by keeping the processor fed to capacity with data. It would also allow the development of more communication-intensive designs such as neural networks or parallel-processing machines.

According to Andrew Kirk of the photonics group at McGill University in Montreal, the computer industry has suddenly woken up to the potential of free-space optics. He believes that the machines of the future will use some kind of optical/silicon hybrid technology. "Some people are even talking about using optics inside actual chips. The time it takes for an electrical signal to cross a microprocessor is becoming critical, so it'd be useful to be able to hop over areas of silicon with an optical link," he says.

The problem with wiring arises from the way that computers are built. Each contains many subunits, including processor chips which are connected to memory chips, disc drives, the keyboard and other components. These parts chatter constantly to each other along copper wires, exchanging streams of information as electrical pulses that represent binary 1s and 0s. As processors become faster, making a greater number of calculations a second, data must be moved backwards and forwards more quickly than before to keep the processor fully occupied.

Designers know that there is a limit to the rate at which they can move electrical signals along wires. As data bits are transmitted faster and faster, the signals begin to blur. Send them too fast and the difference between 1s and 0s begins to disappear altogether. Wires also start to act like tiny radio antennae, broadcasting their signals to nearby wires. So wires must be shielded if they are to carry data at high speeds without interference, making them bulky and expensive. Problems also occur when designers use narrower wires in an attempt to cram more connections around chips. Thinner wires impede data transmission since their resistance to current flow increases. To break through these physical barriers, Kirk believes that computer engineers must abandon themselves to pure free-space optical connections.

The principle of a free-space connection is simplicity itself. First, the electronic pulses inside a silicon chip are converted into a pencil-thin beam of winking light, "on" for 1, "off" for 0. This stream of bits is projected through a network of mirrors and prisms to where the data are needed. At the receiving end, lenses focus each beam onto a microscopic photocell which converts the flashes back into a series of electronic pulses.

Laser on a wafer

The main requirements are that these optoelectronic converters should use little power and be small, cheap and easy to manufacture. Researchers have tried many options, including light-emitting diodes (LEDs), but the best candidates are multiple quantum well (MQW) devices—a kind of electrically switched shutter—and a microscopic laser known as a vertical cavity surface-emitting laser (VCSEL, pronounced vixel). Both devices are made from semiconductor compounds such as gallium arsenide, giving them the advantage that, like silicon chips, they can be mass-produced on wafers.

MQW devices were pioneered by Bell Laboratories in New Jersey, and were originally intended to be the basic switching component of a true optical computer—one in which much of the processing and data communication were carried out by photons. It now turns out that they may do very nicely as the optoelectronic converters for silicon-based machines.

In essence, an MQW is a stack of semiconductor layers which together form a minute reflective shutter. This can be rapidly switched from shiny to opaque under the control of an electric current (see Diagram). Any light aimed at it will be reflected only when the MQW is in shiny mode, so sending a visible "1" to the photocell on another chip. Each MQW device also has a "window", a small photocell which can be tripped by incoming light pulses, converting them to electrical pulses.

The original idea was to create the optical equivalent of a transistor gate, with light inputs switching straight through to light outputs. But in a hybrid system, an MQW made from gallium

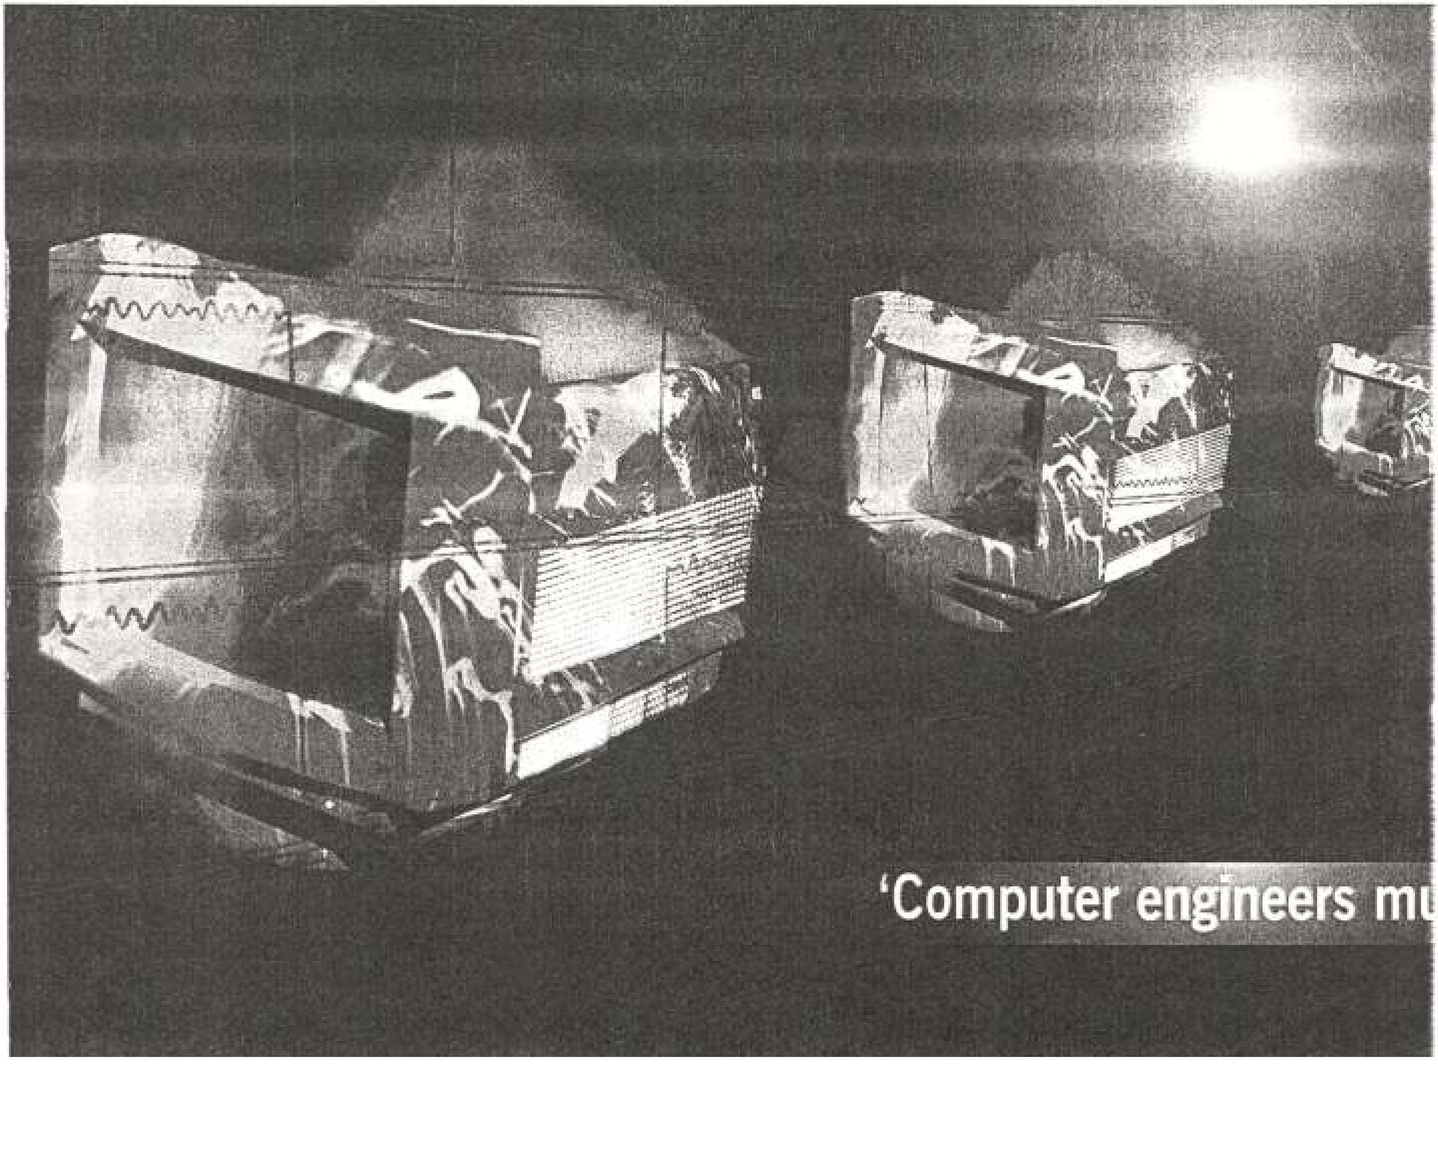

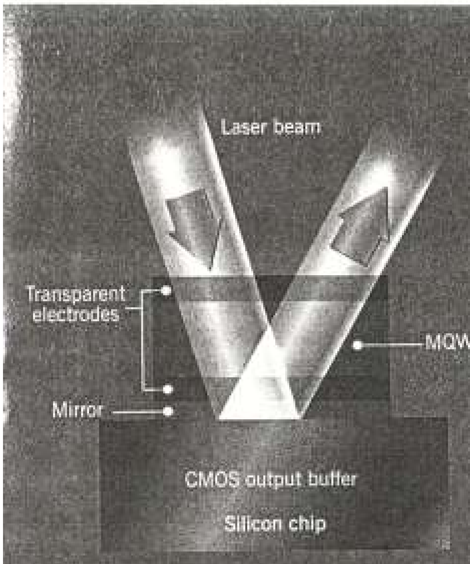

Light signals: multiple quantum well devices, which are made from alternating layers of aluminium gallium arsenide (AlGaAs) and indium gallium arsenide (InGaAs) semiconductors, are normally transparent to infrared light. In this state, light passes straight through and is reflected off the mirror. But applying a voltage across these devices changes their electronic structure, making them absorb the light, and preventing reflection. Connect the output of the silicon chip to the device and you can convert an electrical signal into light pulses. You can also use the same device to build a photodiode. Measure the tiny voltages that are generated in the semiconductor when it absorbs light, and you can convert light pulses back into electrical signals

arsenide is soldered to the back of a silicon processor chip so that the output channels of the chip drive the shutter directly. The photocell window is the optical input for the chip, converting light pulses to electrical pulses. More than a thousand of these MQW devices, each about 15 micrometres across, have already been attached to the back of a single silicon chip by researchers at Bell Labs.

The big drawback of MQW devices is that they need an external light source. In Heriot-Watt's demonstration system, a large laser sits outside the computer. Its beam is split into thousands of smaller beams and a series of lenses and mirrors then aim each beam onto a MQW window. With a 32-by-32 array of MQWs, their latest machine has a throughput of trillions of bits a second, eclipsing the latest Cray supercomputer, which can only handle billions of bits a second. Dines believes there is no reason why such a setup could not work commercially.

The alternative to MQW shutters would be to have a tiny solid-state laser mounted on each output channel. Until recently, the smallest lasers were simply too big. But advances in chip-making techniques mean that much smaller lasers can now be formed as finger-like projections, each a stacked-up tower of gallium and aluminium arsenide poking up vertically from the semiconductor surface like a microscopic Manhattan skyline ("How chips build better lasers", *New Scientist*, 11 January 1992, p 25).

VCSEL lasers are still fairly large—about 250 micrometres across compared with between 10 and 20 micrometres for an MQW. But scientists at Bell Labs and other research institutions have already proved that such lasers can be made just as small as an MQW. And like MQWs, VCSELs could be etched as a single block of circuitry and soldered straight onto the back of a chip, making them a cheap way to add an optical output to silicon chips.

The optics used to guide light beams from one chip to another are also borrowing the techniques of chip manufacturing. At the integrated photonics laboratory of the University of California at Los Angeles, researchers are working with lenses only a few hundred micrometres across and other Lilliputian optical devices. The shape of each component is first etched onto the surface of a block of silicon, just like a printed circuit. But each part has a hinge on one edge so that, once acid has eaten it free of its silicon base, it can be flipped upright. Supporting struts, also etched onto the silicon, are then clicked up to lock the lens or mirror into place. In this way, all of the optical components can be made cheaply, accurately and extremely small, ready to be mounted between the lasers and photocells.

However, the big issue is whether such delicate technology can stand up to the bumps and insults of everyday use. With thousands of crisscrossing shafts of light, each just a few micrometres wide, aimed at photocells which are themselves only micrometres across, the rumble of a passing truck or hot sun swelling its casing could crash the computer.

Self-aligning beams

Prototypes now being built at Heriot-Watt, McGill and other laboratories around the world are showing what a hybrid computer might look like. The parts have to be laboriously aligned by hand, with only strong magnets to hold them in place. "The system can be happy for weeks or sometimes months," says Dines. "But obviously that's not much good for mass-production." Ideally, he says, the optics would be clicked into place at the factory, sealed to keep out dust, and never touched again. But, Dines believes, in practice some sort of active alignment system will have be built into every free-space optical computer. Sensors would monitor each channel to detect when a light beam has wandered from its target, then tiny motors would tilt mirrors to bring the connections back into alignment. "The system would have to have the intelligence to check itself constantly," says Dines. "That would add to the cost and complexity, but it would also give you advantages. For example, you would be able to plug new cards into the system just as you do with PCs at the moment. The system would adjust itself."

Thanks to the camcorder industry's attempts to deal with the problem of camera shake, says Dines, the technology for a self-aligning system is far more advanced than many realise. Some of the latest video cameras can cancel out even the vibrations of filming from a moving car. They contain a small liquid-filled prism that reacts to camera movement by changing shape, rapidly expanding and contracting like an accordion. The path of light beams passing through the prism is bent accordingly, keeping the

abandon themselves to pure free-space optical connections'

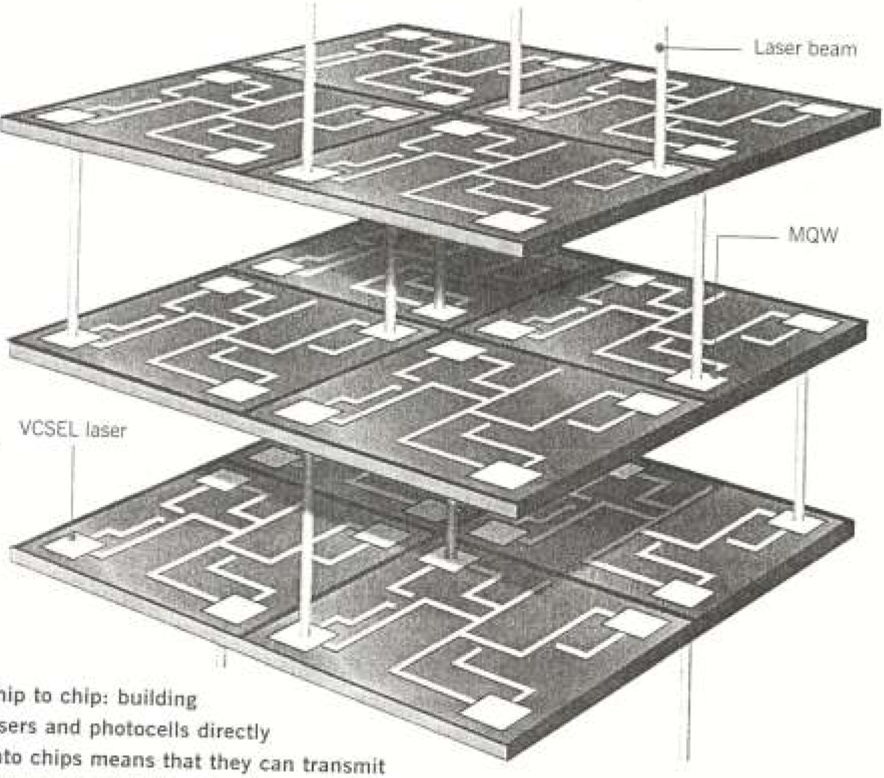

Chip to chip: building lasers and photocells directly onto chips means that they can transmit data back and forth using light rather than wires. Communication is faster and processors work more efficiently

image steady on the video sensor.

Not surprisingly, those involved in research into free-space optics believe their technology will spark a new computer revolution. Having got over the psychological barrier of using light beams to maintain connections, designers will be free to exploit the special properties of photons more creatively. Kirk says the ability to split a beam and broadcast a single bit of data to any number of chip inputs, or to have many channels reading the same memory location simultaneously, opens up the way to all kinds of novel parallel processing and neural network systems.

For example, free-space optics have a bright future in mainframes and supercomputers. PetaFLOPS is an ambitious project backed by NASA to build a massively parallel computer by 2010 that will be capable of a million billion calculations a second. NASA engineers have concluded that optical connections are a must. Such a machine would have to shift about 10 petabytes (10^{16} bytes) of data a second. A petabyte is the equivalent of a billion books or 2300 years of video footage.

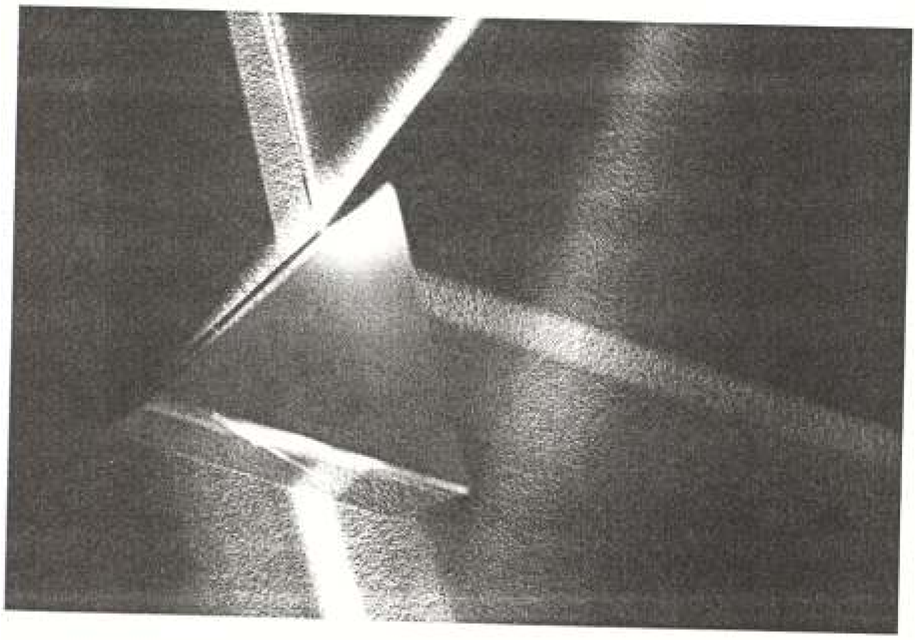

All change

But PC makers could opt for a different route to higher performance. Mark Bohr, a research fellow at Intel, says one solution to the communications bottleneck is simply to cram more of the total system on a single chip. This has already happened to some extent with today's microprocessors, which have an onboard cache memory to store frequently used information. This lessens the need for slow, wire connections between components. If PC makers wanted to move to free-space optics, warns Bohr, there would be huge difficulties. "There are a lot of companies making chips, and a lot of board manufacturers and power supply manufacturers. We would have to get the whole industry shifting to a new kind of interconnect technology together, which would be a big problem."

But attitudes could change rapidly if the trend towards more communications-intensive computing continues. Today's PCs already need to be fast enough to handle huge, high-resolution images and Internet access. If they are to become intelligent office assistants, the desktop machines of the next few years will have to be capable of video conferencing and even vision, taking their cues from what is happening in the room where they are sitting. Kirk believes that within a few years of free-space optics proving its worth inside expensive special-purpose computers, it will make its way down the scale to PCs.

As a sign of how fast things change, Kirk notes that a few years ago nobody in the industry was showing the slightest interest in free-space technology. But at a research conference last spring, he found engineers from IBM, Cray and Digital all sidling up to quiz him. Suddenly it is no longer a question of whether the free-space optics revolution will happen, but when.

John McCrone is a freelance writer based in London