

The movement of data in a computer is almost the converse of the movement of traffic in a city. Downtown, in the congested core of the microprocessor, the bits fly at an extraordinary rate. But further out, on the broad avenues of copper that link one processor to another and one circuit board to the next, things slow down to a comparative crawl. A Pentium 4 introduced this spring operates at 2.4 GHz, but the data travels on a bus operating at only 400 MHz. The speed picks up again, though, out on the highways of the world's optical-fiber telecommunications networks. Obviously, the closer engineers can bring the optical superhighway to the microprocessor, the fewer copper bottlenecks can occur, as if you could pull out

"We're already projecting that for certain system requirements data rates are going to be high enough and the link length long enough that we're going to have to use optics," says Modest Oprysko, senior manager of communications technologies at IBM Corp.'s Thomas J. Watson Research Center (Yorktown Heights, N.Y.).

A wire's bit rate is at the mercy of its parasitic resistance, capacitance, and inductance. At low frequencies, the series resistance and shunt capacitance of a circuit board trace dominate its behavior, determining the transition (rise and fall) times and thereby limiting its data rate. At higher frequencies, like those on circuit boards today, the wire's series inductance becomes more important than its resistance as an impeding

# Linking with Light

BY NEIL SAVAGE  
Contributing Editor

Having proven their worth in long-distance communications, photons will soon take over inside the computer

of your driveway straight onto the Autobahn.

So some researchers say that, within just a few years, many of the copper connections in computers will yield to high-speed optical interconnects, in which photons, rather than electrons, will pass signals from board to board, or chip to chip, or even from one part of a chip to another [see illustration, opposite].

The idea is simple in principle, and parallels telecommunications systems. An electrical signal from the processor would modulate a miniature laser beam, which would shine through the air or a waveguide to a photodetector, which would in turn pass the signal on to the electronics. Though at the moment it is more expensive to communicate with light than with electric current, the day is coming when only optical technologies will be able to keep up with the demands of ever-more-powerful microprocessors, just as they are now the only reasonable way to move the world's Internet traffic across the kilometers.

factor, but with the same end result—a limit on the rate at which the trace can transmit pulses.

All these parasitic factors depend heavily on the geometry of the wire, especially its length. Resistance, for instance, is proportional to the wire's length and inversely proportional to its cross-sectional area. Because of this dependence on geometry, a simple wire's ultimate bit rate turns out to be proportional to its cross section, but falls with the square of its length. So thinner—and especially longer—means a lower bit rate. On a chip, the story has some caveats, but the conclusion is the same.

Transition-time limitations can be fought by driving the lines harder, but that's not a very good solution. It adds noise, increases power requirements, and aggravates already serious thermal management problems. Alternatively, the wires can be made fatter, but then you'll run out of space. Photons don't suffer from these limitations; their biggest problems are absorp-

## ● Coming Soon: Optical Interconnects

This approach to signal transfer is moving from longer-distance applications, such as linking separate computers, to joining chips within a computer.

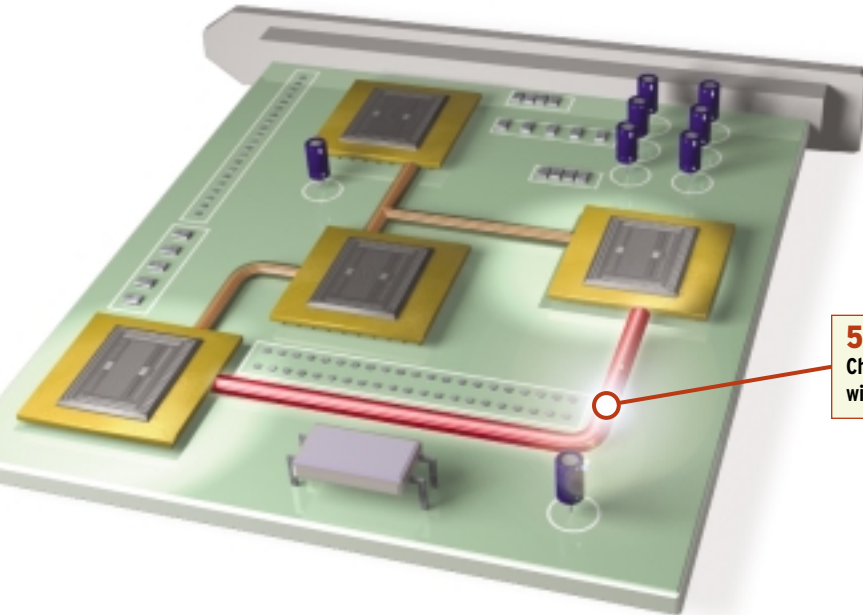
### 2-5 Years

Optical communications will enter the computer, connecting one circuit board to another.



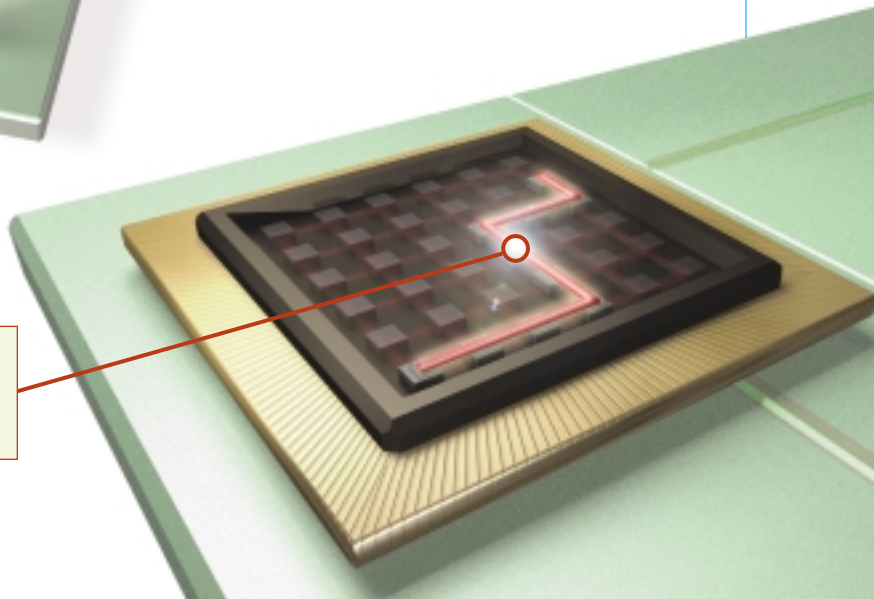
### Today

Optical connections between individual computers are commercially available.



### 5-10 Years

Chip-to-chip communications will enter the market.



### 15+ Years

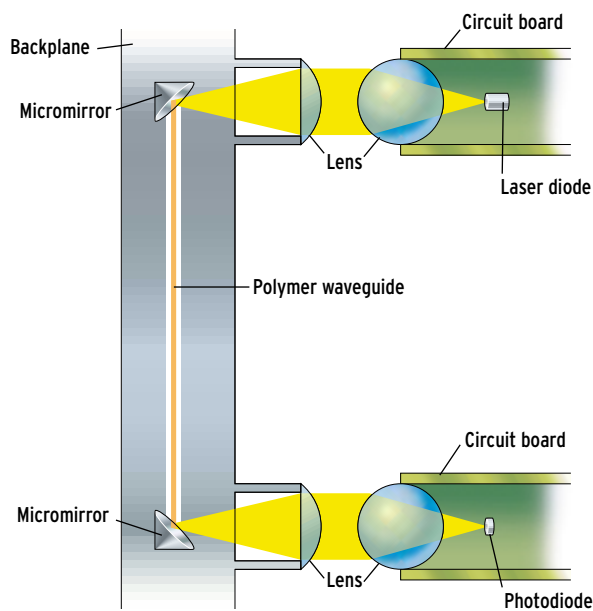
Experts disagree on whether optical interconnects will ever connect the subsystems within a chip.

tion and attenuation, neither of which is an issue over the distances inside a computer, or even across a room.

In telecommunications, transmission slows as it passes from the all-optical long-haul network to the lower-bandwidth metropolitan-area network, and finally crawls into the home along copper wires. Just as telecom companies want to increase total bandwidth by moving all-optical transmission closer to your house, the designers of optical interconnects want to get as much bandwidth as possible as close as possible to the microprocessor. Already, optics connect computer systems across dis-

## ● Optical Backplanes Speed Data

In DaimlerChrysler's optical backplane, the beam from a laser diode passes through one set of lenses and reflects off a micromirror before reaching a polymer waveguide, then does the converse before arriving at a photodiode and changing back into an electrical signal. A prototype operates at 1 Gb/s.



tances of less than 300 meters, and backplane setups are in the works that will speed up transport of data from one board to another within a computer. Farther down the road are systems for increasing the bandwidth between two microprocessors, or among stacks of chips for massively parallel computing. All the pieces are in place—cheap lasers, sensitive detectors, and the methods needed to transmit from one to the other. Now it's just a question of when optical interconnects will perform well enough, and their production cost fall low enough, for them to replace copper wires.

### Thinking inside the box

Many developers of optical interconnect schemes funnel the signal through some kind of waveguide. For instance, DaimlerChrysler Research (Ulm, Germany) is developing waveguide-based optical backplanes to connect computers in an airplane or to route signals in computers used in telecommunications

[see figure, below left]. As in most optical interconnect schemes, the photons are generated by vertical-cavity surface-emitting lasers (VCSELs) [see "Arrays of Light," p. 36]. Light passes through a lens, which expands the beam to a few millimeters in diameter. The expanded beam hits the waveguide entrance with greater certainty, while at the waveguide exit, it is redirected by a mirror through another lens, which focuses it on a detector. DaimlerChrysler claims the polymer waveguide is easier to integrate into a system than an optical fiber.

Prototypes with data rates of 1 Gb/s per channel across distances less than 100 cm were successfully tested last March. The company is providing development samples with channel rates of 2.5 Gb/s, and has systems with 10-Gb/s channels under development in the laboratory.

Another waveguide-based scheme belongs to two-year-old Primarion Corp. (Tempe, Ariz.). The company is in the early stages of making optical circuits that transmit data at 10 Gb/s across short distances. The goal is to bring fiber all the way to the processor. In its design, an electrical signal travels across a circuit board to a laser driver chip, which processes the signal and prepares to send it out. Connected to the laser driver is another module, containing an array of 12 VCSELs. The laser beams travel across optical fibers to a similar setup on another board, and there a photodetector passes the signal onto a receiver chip for reprocessing and conversion back into electricity.

John Burns, Primarion's vice president of technology development, sees compact optical I/O devices within the computer in just two to three years' time. These would combine the processing circuitry and the laser or receiver into a single package, with electrical contacts on the bottom and fiber-optic sockets on the side. Eventually, the company envisions having a printed optical board, like a printed-circuit board but embedded with polymer waveguides in place of copper wires. Then the signal-processing circuitry could sit right next to the device that needed the data, and the signal would be transmitted through the circuit board as light, minimizing the time data spends as electrons.

Fiber-based interconnects have already made their commercial debut, but not yet within the computer itself. In late 2000, Agilent Technologies (Palo Alto, Calif.) started selling an optical connector to link computers up to 300 meters apart. Developed with Anthony F. J. Levi, head of the Advanced Interconnect and Network Technology program at the University of Southern California (USC, Los Angeles), one version of the plug-in package can transmit data at an aggregate rate of 10 Gb/s. It contains four transmitters and receivers and connects via four optical fibers to a similar package. A 12-channel version operates at up to 40 Gb/s, but transmits in only one direction at a time. Other companies now offer similar products, for use in rack-to-rack connections among servers, in data storage centers, and in network and telecom switching.

Agilent's next aim, according to David Dolfi, manager of its optical interconnects and networks department, is to do something similar for a backplane within a computer. To keep up with faster processors and more complex systems, he says, the bandwidth of backplanes must increase by an order of magnitude, and that will require optical links. Barring some breakthrough in electrical links, Dolfi expects optical interconnects to move inside the box in three to five years.

## Through the air in there

While the nearer-term, less complex interconnect schemes deliver light through a waveguide, engineers interested in high-performance computing are developing chip-to-chip optical communications through the air. So-called free-space interconnects could boost data throughput between chips by a factor of a thousand. This increase stems from the combination of the high data rate of optics and the fact that optical data channels can be much more densely packed than electrical ones. For one, the channels can be formed into two-dimensional arrays, rather than the one-dimensional packing of copper or optical fiber interconnects on a board. Secondly, the channels do not experience crosstalk when close together.

One such system in which three stacks of 16 chips apiece can talk to each other optically has been developed by Sadik Esener, a professor of electrical engineering at the University of California, San Diego, and his colleagues. A 16-by-16 array of VCSELs send a beam through the air, where integrated microlenses focus the beam and diffraction gratings redirect it. The result is a tight, 3-D array of highly interconnected chips for supercomputing and other applications [see figure below].

Most researchers see such free-space interconnects as farther off in the future than those that use fiber to guide the light, in part because they require an accuracy of alignment much greater than that of plug-in cards in the average desktop. And anyone who has ever opened up a home computer and seen the dust built up inside knows another of the obstacles.

Nonetheless, free-space optics intrigues engineers who think about doing computing in revolutionary ways, and it has been an active research area for high-performance computing in North

America, Japan, and especially Europe, which has had a long-standing effort. The appeal is simple. Say you wanted a supercomputer that could be optimized for multiple tasks. You set up the processors to work best with a given algorithm, perhaps for image processing, but when you want to switch to weather modeling, the architecture is no longer ideal. Using free-space optical interconnects, you could reconfigure the way processors talk to one another by redirecting the lasers to different detectors.

John Neff, director of the Optoelectronic Computing Systems Center at the University of Colorado (Boulder), is working on microelectromechanical systems (MEMS) that might be used to steer the beams. In Neff's experimental system, a microscopic mirror is attached to two metals with different thermal properties. Heat them up by running current through them, and one will stretch more than the other, bending it in the desired direction. It's a slow process, however, and the highest switching speed is about a kilohertz. But using electrostatic repulsion instead of heat to bend the mirrors produced speeds of 10 kHz, and Neff thinks he can get that up to 100 kHz.

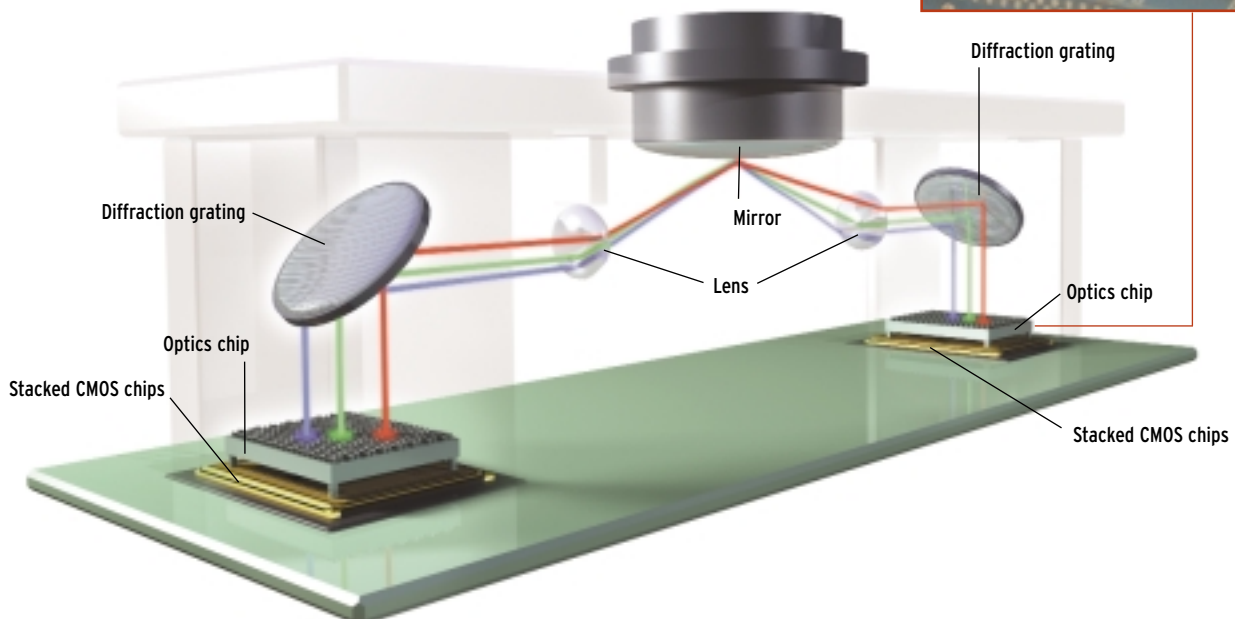
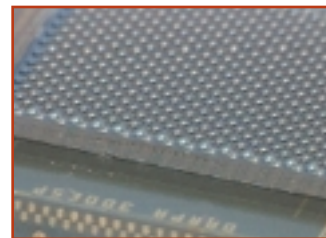
## Up for debate

The exact point at which the economics of moving data dictate optical interconnects is debatable. According to USC's Levi, when microprocessors reach speeds of 10 GHz, their inability to communicate quickly with computer memory or with other chips in a multiprocessor system will start to stifle their usefulness. "At that point, [it will be] increasingly difficult...for Intel to make useful systems out of those processors," Levi said.

Naturally, Intel Corp. is loath to sound the death knell for electrical interconnects quite yet. Says Wilfred Pinfeld, technology

## ● Free-Space Interconnects Pack in Data Channels

An experimental module from the University of California, San Diego, just 2 cm high, connects stacks of CMOS chips. Each stack is topped with an optics chip [right] consisting of 256 lasers (VCSELs) and photodiodes. Light from the VCSELs makes a vertical exit from one stack [below, left] and a vertical entry into the other. In between it is redirected via a diffraction grating, lenses, an alignment mirror [center], and another grating. Each of the device's 256 channels operates at 1 Gb/s.



# Arrays Of Light

**M**ost optical interconnect schemes envision using vertical-cavity, surface-emitting lasers (VCSELs) as their light source along with one of two types of semiconductor detectors. All three technologies use III-V materials like gallium arsenide rather than silicon, and so cannot be directly integrated with the silicon CMOS devices with which they communicate. Nonetheless, developing silicon CMOS-based lasers and detectors is an active research area that could smooth the introduction of optical interconnects.

Many detector designs use p-i-n photodiodes, a sandwich of p-type, intrinsic, and n-type semiconductor material that generates current when struck by a photon. A sandwich of metal-semiconductor-metal creates MSM photodetectors that work much the same way, at a higher speed than p-i-n photodiodes but with less sensitivity. Researchers use both in their designs, depending on which attributes are most important to them.

GaAs VCSELs have been commercially available for the last few years. The devices are diode lasers that, unlike most used in long-haul telecommunications, emit from the surface of the laser rather than the edge, so they can be placed on top of a chip. It's relatively easy to build an array of hundreds or thousands together in a plane and to make them close to the size of transistors, a few hundred nanometers across. A small percentage of the individual VCSELs

in an array, however, may not work properly. Bond a bum array to an expensive microprocessor and you've wasted a lot of money, so if optical interconnects are to take off, VCSEL manufacturers must reduce the failure rate.

As chips continue to shrink, transistors are squeezed closer together, as would be the VCSELs communicating with them. Such spacing raises the specter of a problem shared with electrical wires: as they get closer and operate at higher rates, the threat of electrical interference between adjacent channels increases. Though the light from the lasers itself isn't subject to crosstalk, the electric current that drives it is.

The power driving the VCSELs also creates heat, which is worse for the lasers than for electronic circuitry. As the GaAs heats up, its material properties change, and the lasers can detune, wandering to a different wavelength. Beyond about 80 °C they stop working altogether. So the challenge is to increase their efficiency so they can be driven at even lower powers.

Perhaps the biggest difficulty with VCSELs is that the GaAs, and in fact any laser semiconductor, has a lattice that doesn't match with silicon's. That means they have to be manufactured in a separate process, adding enormously to production expense. They can be bonded to CMOS chips, but they will undergo thermal expansion at a different rate from the silicon, stretching out of alignment or even breaking the device. The photodetectors are also based on GaAs, so they suffer from the same drawbacks, although some researchers are looking into silicon-based detectors.

Naturally, optoelectronics engineers would love to make VCSELs from silicon, but the material doesn't emit light easily. Unlike such semiconductors as GaAs, silicon has an indirect band gap, making a free electron more likely to give up its energy as heat than to combine with a hole and emit a photon.

## Glimmer of hope

But Wilfred Pinfold, technology director of Intel Corp.'s microprocessor research laboratory, sees some hope in work last year from Britain and Australia. In March 2001, Kevin Homewood of the University of Surrey (Guildford, England) reported success in making a silicon light-emitting diode (LED). He achieved this by bombarding the silicon with boron ions. These created barriers in the silicon lattice that trapped electrons in a small area, increasing their likelihood of finding holes to combine with.

Last August, Martin Green of the University of New South Wales (Sydney, Australia) reported a new silicon LED. The device has inverted pyramids etched into it that reflect some emitted photons back into the silicon, where they can stimulate further emission. It converted electrons to photons with 1 percent efficiency—very low, but 100 times that of previous silicon LEDs.

LEDs, unfortunately, are not as useful as more powerful, more focused lasers. And, after decades of struggle, some researchers are skeptical of silicon's chances. "If God wanted ordinary silicon to efficiently emit light, he would not have given us gallium arsenide," says Elias Towe of Carnegie Mellon University (Pittsburgh). —N.S.

director in its microprocessor research laboratory, "We still see a lot of opportunity to continue to move copper signaling forward."

For a start, engineers could rework the architecture of chips and boards to shorten the distances over which signals have to travel. Designers could add more pins to increase the number of I/O links, though not indefinitely. Intel can move data at 10 GHz over a copper wire on the board, and should be able to push that to at least 20 GHz. Pinfold says: "Where that peters out, we don't know." Still, his lab is striving for advances in both electrical and optical interconnects, as are IBM, Advanced Micro Devices, and other competitors.

Much of the research into optical interconnects in the U.S. has been funded by a US \$70 million program run by the U.S. Defense Advanced Research Projects Agency (DARPA, Arlington, Va.). The massively parallel computing systems it has in view

could find use in modeling nuclear reactions. But Ravindra Athale, who oversees the research, expects it to filter down to the consumer as well. The progression should be from expensive, intensive computer systems to simpler, cheaper ones, as well as from longer to shorter distances. Rack-to-rack connections are already here, board-to-board links seem likely in a few years, and chip-to-chip interconnects may happen in a decade or so.

As for optical connections between one system and another on a single chip, the view of most experts is dim. "It think that's beyond what's viable given today's technology," says USC's Levi. But finding out whether across-chip optical interconnects will ever be practical is one of the reasons for advanced research, says DARPA's Athale. "The answer is maybe, and that's final." ●

Samuel K. Moore, *Editor*



TO PROBE FURTHER, SEE PAGE 63