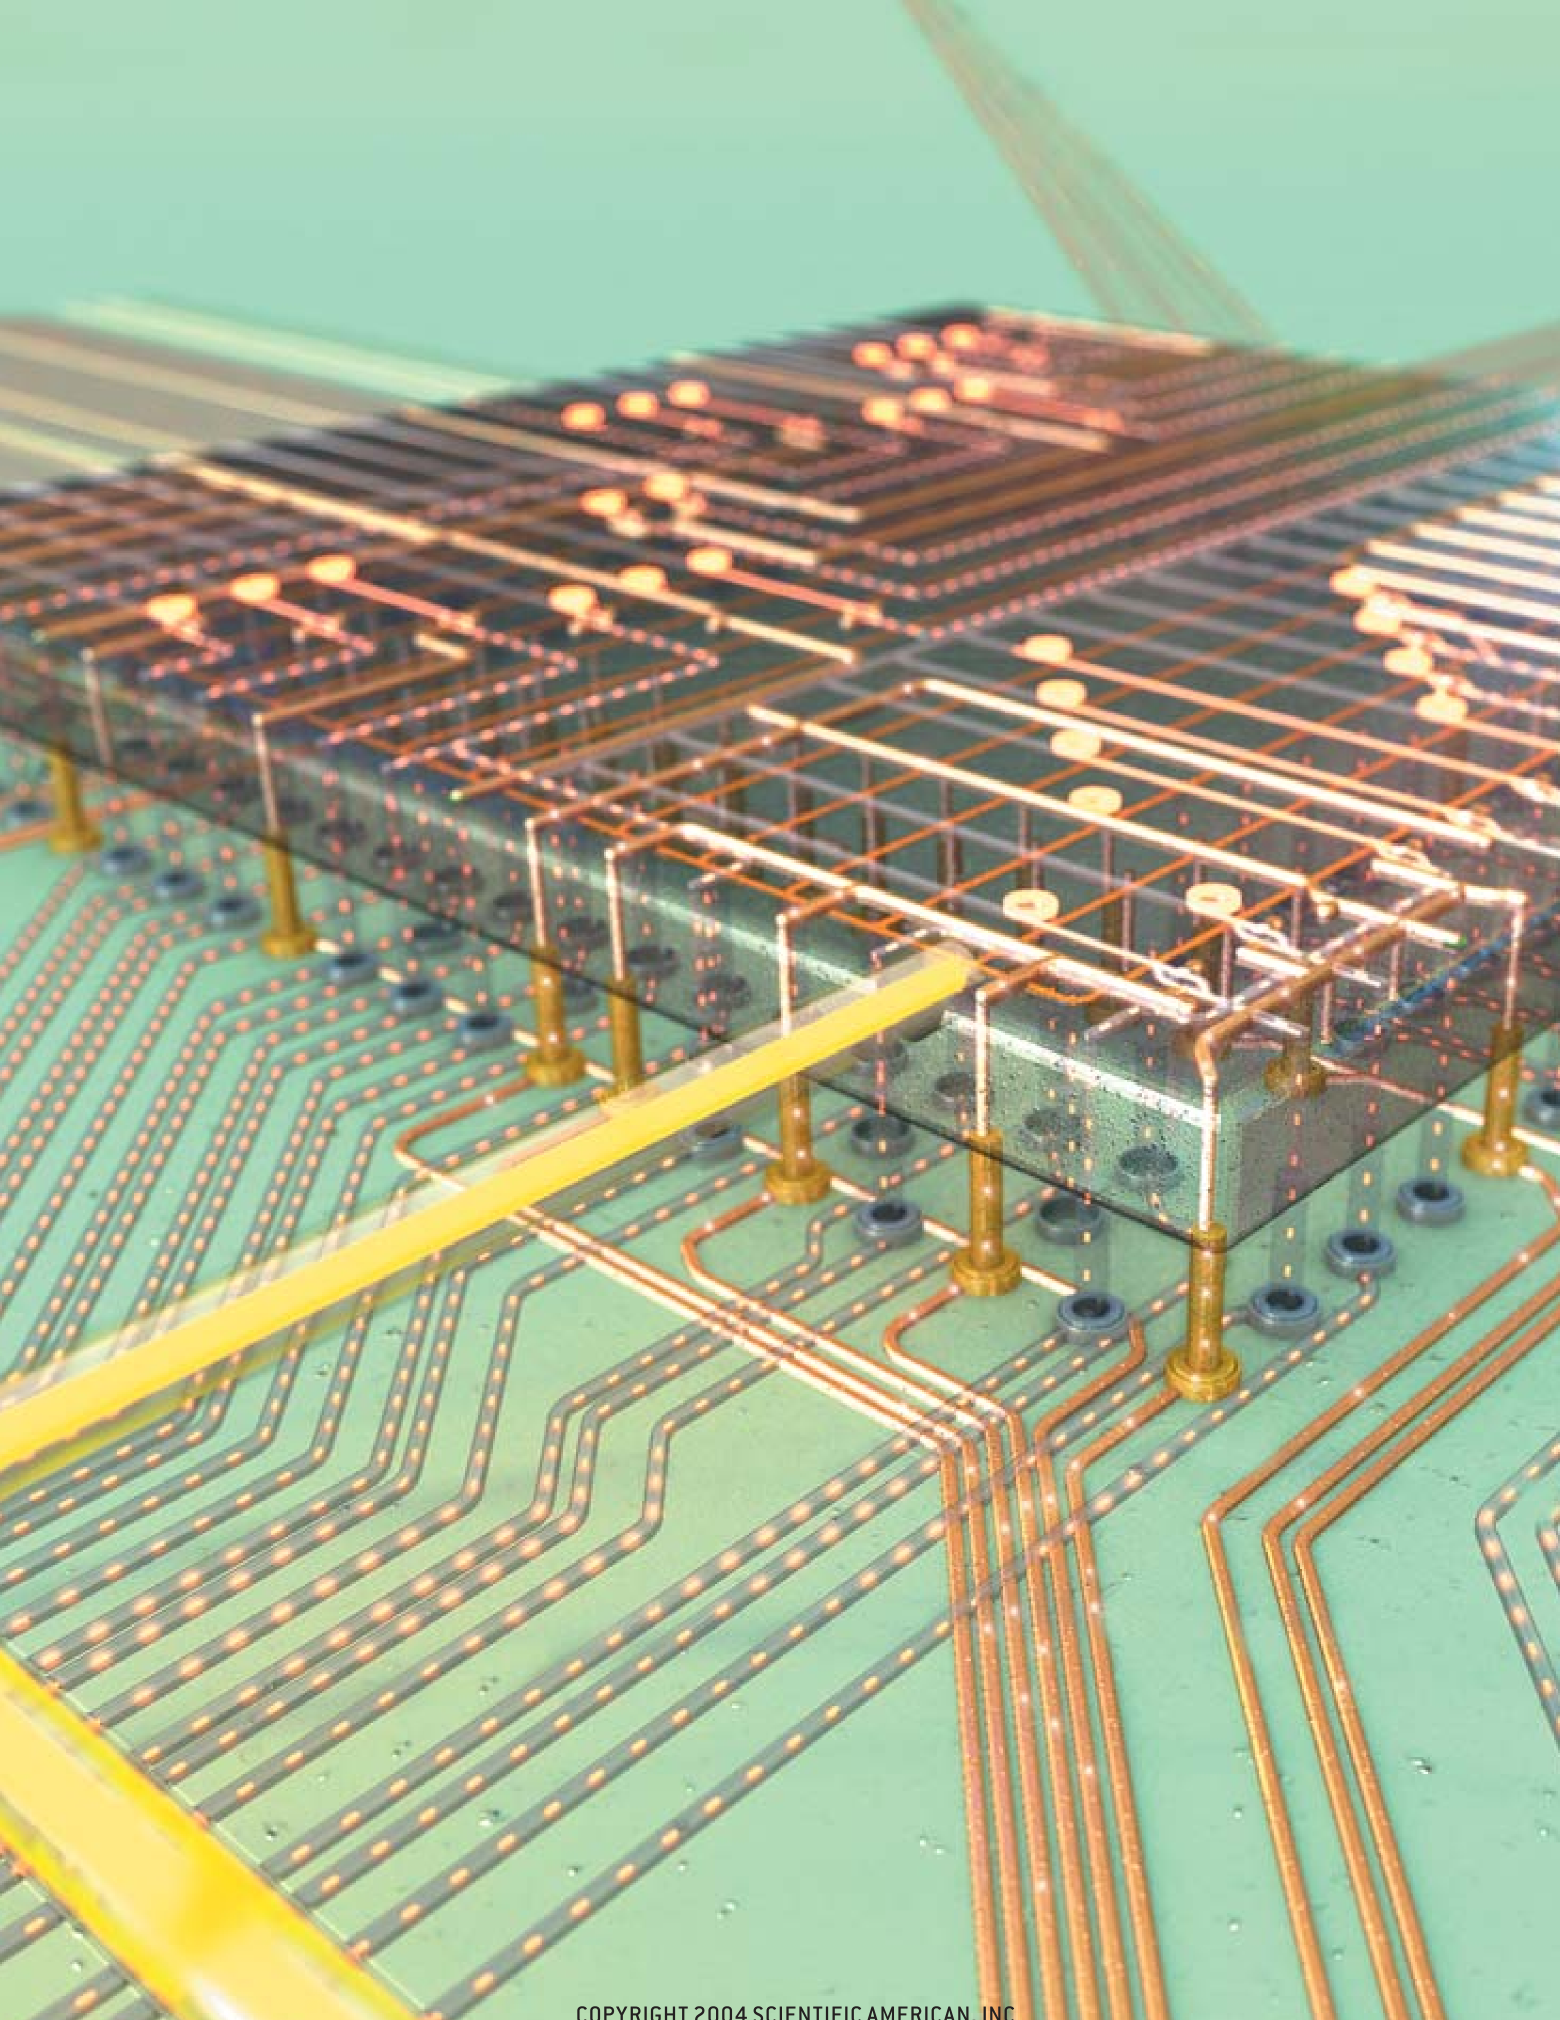

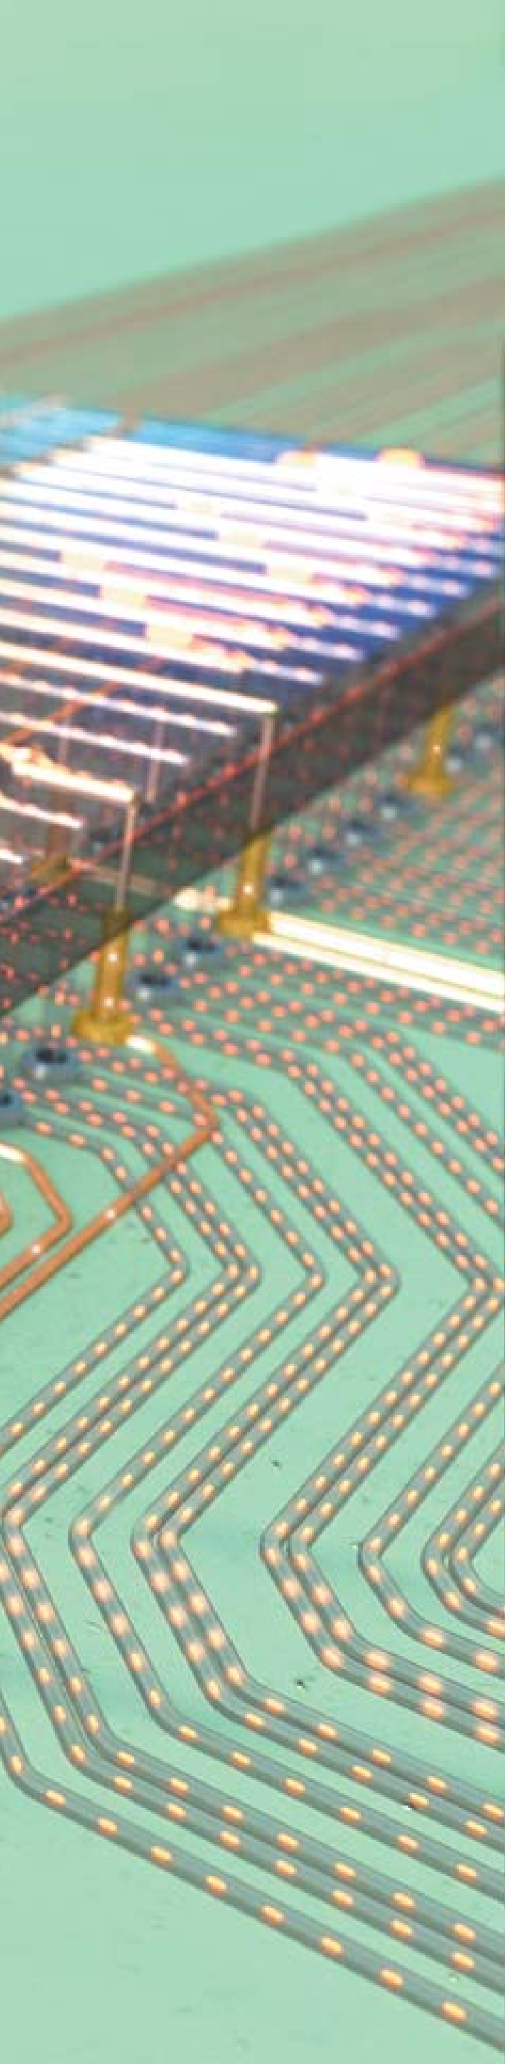

BRYAN CHRISTIE DESIGN

Computing at the Speed of Light

Emerging ways to make photonic connections to electronic microchips may dramatically change the shape of computers in the decade ahead

By W. Wayt Gibbs

FUTURISTIC MICROPROCESSOR might communicate with the rest of the computer via light as well as electricity. Recently invented devices such as microcavity lasers, silicon optical modulators and translucent polymer pillars could be combined to move bits seamlessly from the electronic realm to the photonic, and back again.

Since about 1995, microprocessors have been outrunning the other parts of computer systems by ever increasing margins. The latest processors churn through instructions at up to 3.6 gigahertz (GHz); some operations, such as arithmetic, run at double that rate. But the wiring on the motherboard that connects the processor to its memory chips and other pieces of the system plods along at 1 GHz or less. So the brain of the machine spends as much as 75 percent of its time idle, waiting for instructions and data that are stuck in traffic.

"In the coming years, the imbalance between microprocessor performance and memory access will be driven to a crisis point," physicist Anthony F. J. Levi of the University of Southern California argued in a detailed analysis three years ago. He noted that the plastic material in printed circuit boards squelches high frequencies: for every 2-GHz increase in electrical signal bandwidth, signal strength falls 10-fold. As clock rates rise, so do power consumption, heat production and electromagnetic interference. Those are already three of the biggest headaches for system designers. And International Sematech, an industry consortium, forecasts that processor-to-peripheral links must accelerate by roughly 2 GHz every two years just to keep the bottleneck from tightening further.

"Our engineers think they will eventually be able to squeeze 20 GHz out of wires as long as 20 inches," says Michael Morse, a photonics researcher at Intel. According to the Sematech roadmap, 20 GHz would be just sufficient for the 32-nanometer generation of microchips, three steps down the road from the 90-nanometer chips that arrived earlier this year. Mark T. Bohr, director of process architecture at Intel, reports that his company is on track to bring that generation to market by 2010.

The stage thus appears set for photonic connections, which exchange data via laser light, to take over for copper wiring in the next decade. "I'm a big fan of optical connections at the system level," says Patrick P. Gelsinger, Intel's chief technical officer, although he remains unconvinced that they will take over the very high speed but short-distance hop from the processor to the memory bank. Exactly when the transition will occur, in which connections, and at what price, depends in large measure on how the photonic devices are made.

Data already often move between electronic and photonic forms at the periphery of a computer system, on their way to or from a CD or DVD, display monitor, mouse, camera, stereo amplifier or fiber-optic network. But the core of most computers—the processor, the main memory, and the motherboard that connects those to the various peripheral devices—remains an all-electron show.

The reason for this is simple: optical interconnections, though often many times faster than copper wires and traces, tend to be 10 to 100 times more expensive. For some applications, such as switching thousands of telephone calls or shuttling billions of Internet packets, capacity trumps cost. That is why long-distance communications in rich countries now travel primarily over optical fiber. And it is why Cisco spent half a billion dollars over the past four years to create an optical router, debuted this past May, whose 30 fiber-optic lines run at 40 gigabits per second (Gbps)—in principle, enough aggregate bandwidth to handle the Internet traffic of 1.6 million DSL-equipped households. For distances greater than 100 meters, nothing beats the switching speed of light. But over short links, such as those in office networks and inside computer cases, copper still reigns king.

A change of regime now seems more likely, however, because scientists have at last succeeded in making a wide range of photonic devices that could be manufactured by existing microchip factories and thus be cost-competitive. "We want to drive optics all the way down to chip-to-chip communications," says Mario Paniccia, head of Intel's silicon photonics research group.

If that happens, computers may look and operate very differently a decade from now. Some changes will be of the "faster, smaller" variety. Video cameras and portable video players might plug their fiber-optic cables into the photonic successors to USB ports. Some machines may have holographic disk drives that can archive hundreds of gigabytes on one removable CD-size platter. For those people fortunate enough to have a direct connection to the international fiber-optic telecommunications grid, an optical network card may provide Internet access at more than a gigabit per second (Gbps)—about 1,000 times the speed of today's DSL and cable modems.

Other changes could be more dramatic. The maximum practical speed of electronic connections falls off quickly as cable lengths increase. So memory chips and graphics cards have to be close to the processor that shovels information to them. "But once you have data in the optical realm, distance doesn't matter," Paniccia observes. "A low-cost photonics technology is low cost both at one foot and at 1,000 miles." Many of the components of a computer that are now crammed into a two-foot-high rectangular box could in principle be spread across a car, throughout a building or all over a city, with data flowing among them on pulses of light.

Overview/*Optical Computing*

- Computer engineers expect that within the next decade, the copper wiring that now connects components inside machines will reach the practical limits of its bandwidth.
- Until recently, connecting microchips with light meant using lasers and detectors made from exotic semiconductors. Such devices are affordable only for niche applications, such as high-speed telecommunications hubs. But this year engineers unveiled new classes of photonic devices that could be made in the same factories used to make low-cost microchips.
- Researchers have also begun demonstrating novel schemes for guiding laser pulses to and from microprocessors and circuit boards.
- Because optical connections can run at very high bandwidth over both long and short distances, the addition of photonics could fundamentally change the shape of computers over the long term.

Opening the Bottleneck

CURRENT OPTICAL CHIPS, of the kind used as lasers in CD players and as photodetectors in telecommunications switches, are manufactured from III-V semiconductors. These compounds pair one or more elements from the third column of the periodic table (such as aluminum, gallium or indium) with an element from the fifth column (typically phosphate, arsenic or antimony).

At first glance, III-V chips might seem ideal for photonics. Electrons move faster in them than in silicon, so III-V processors can operate at much higher frequencies. And they not only emit laser light from cavities in their surface but also con-

THE PHOTONIC PC: WHAT'S HERE AND WHAT'S AHEAD

Computers today already use optical devices at a few spots around the periphery of the machine. But photonic components now in labs or near release will probably make their way to the core of machines within the next decade.

NOW: Microprocessors operate at three gigahertz (GHz) or more but often idle while waiting for data to arrive from RAM running at just 0.4 GHz. Newer processors can retrieve data from memory over multiple wires, at aggregate rates up to 51 gigabits per second (Gbps)

IN PROTOTYPE: Photonic connections between memory and processor operating at 1.25 GHz per waveguide

IN LABS: New methods for fabricating thousands of polymer pillars onto processors, allowing both high-speed electrical and optical connections to a motherboard

NOW: Fiber-optic networks run at speeds up to 10 Gbps, but most machines use lower-cost Ethernet connections of just 0.1 Gbps

IN PROTOTYPE: Silicon-optical modulators built at Intel could lead to affordable optical networks that run at 2.5 to 10 Gbps

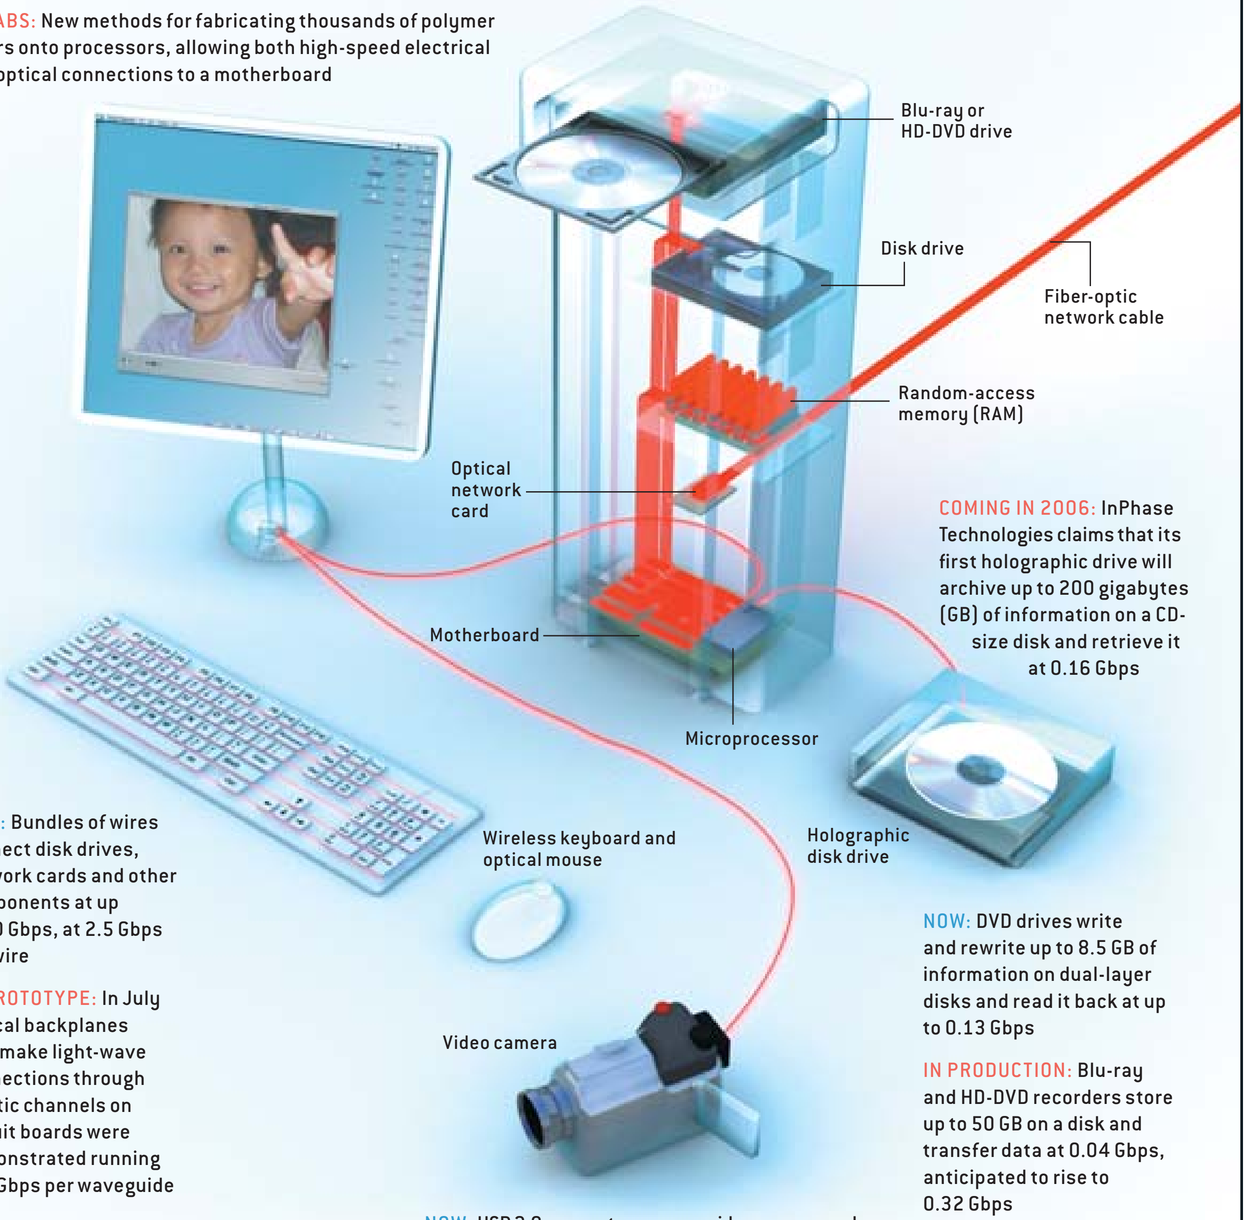

COMING IN 2006: InPhase Technologies claims that its first holographic drive will archive up to 200 gigabytes (GB) of information on a CD-size disk and retrieve it at 0.16 Gbps

NOW: Bundles of wires connect disk drives, network cards and other components at up to 40 Gbps, at 2.5 Gbps per wire

IN PROTOTYPE: In July optical backplanes that make light-wave connections through plastic channels on circuit boards were demonstrated running at 8 Gbps per waveguide

NOW: DVD drives write and rewrite up to 8.5 GB of information on dual-layer disks and read it back at up to 0.13 Gbps

IN PRODUCTION: Blu-ray and HD-DVD recorders store up to 50 GB on a disk and transfer data at 0.04 Gbps, anticipated to rise to 0.32 Gbps

NOW: USB 2.0 connects a mouse, video camera and other gadgets to the computer at about 0.48 Gbps, over distances as long as five meters

IN PRODUCTION: High-priced fiber-optic connectors manufactured by Xanoptix transfer data at up to 245 Gbps, over distances up to 500 meters

BRYAN CHRISTIE DESIGN

vert incoming flashes to electrical signals at blistering speeds. For that reason, photonics researchers have turned first to III-Vs to build optical integrated circuits.

Using indium phosphide, for example, a group led by Daniel Blumenthal and Larry Coldren at the University of California at Santa Barbara last year constructed a "photon copier." The device accepts photonic bits at one wavelength, regenerates them if they have faded, and uses a tunable laser to translate them to a different wavelength without ever converting the information into electronic form. Such a device would be very handy in a future photonic computer.

But in comparison with silicon, III-V semiconductors are finicky and recalcitrant materials to manufacture, and that makes them expensive. A microchip that costs $5 to make from silicon, using the standard complementary metal oxide semiconductor (CMOS) process, would cost about $500 to fabricate from indium phosphide. And with the performance of silicon continually improving, "competing against mainstream CMOS is like lying on a railroad track," bemoans Ravindra A. Athale, who manages photonics programs at the Defense Advanced Research Projects Agency. "Sooner or later the train runs you over."

If photonics is ever to find its way onto $100 motherboards, it must hop aboard that train. So in recent years, much of the research in optical computing has focused on finding CMOS-compatible ways to integrate electronic and photonic devices. This strategy has begun to pay off.

"We are at a stage in the field now that was unthinkable two years ago," says Salvatore Coffa, who directs the silicon photonics laboratory at STMicroelectronics in Catania, Sicily. "We are talking about going to market soon with the first silicon-based device with optical functions."

All Aboard the CMOS Express

AT LEAST THREE WAYS exist to make photonic components passengers on the CMOS train, and each has been making impressive progress. The most conservative approach, called hybrid integration,

FROM WIRES TO WAVEGUIDES

Photonic microchips will likely find their first uses in special-purpose computers that must quickly process huge amounts of data, such as those used in medical imaging. An MRI scanner, for example, might one day use new kinds of microscopic lasers (*below left*) and silicon modulators (*below center*) to send its images over an optical fiber to a computer. New kinds of connectors are being developed to bring such vast amounts of data directly to the central processor (*far right*). Such super-speedy interconnections should make it easier for doctors to consult with distant colleagues.

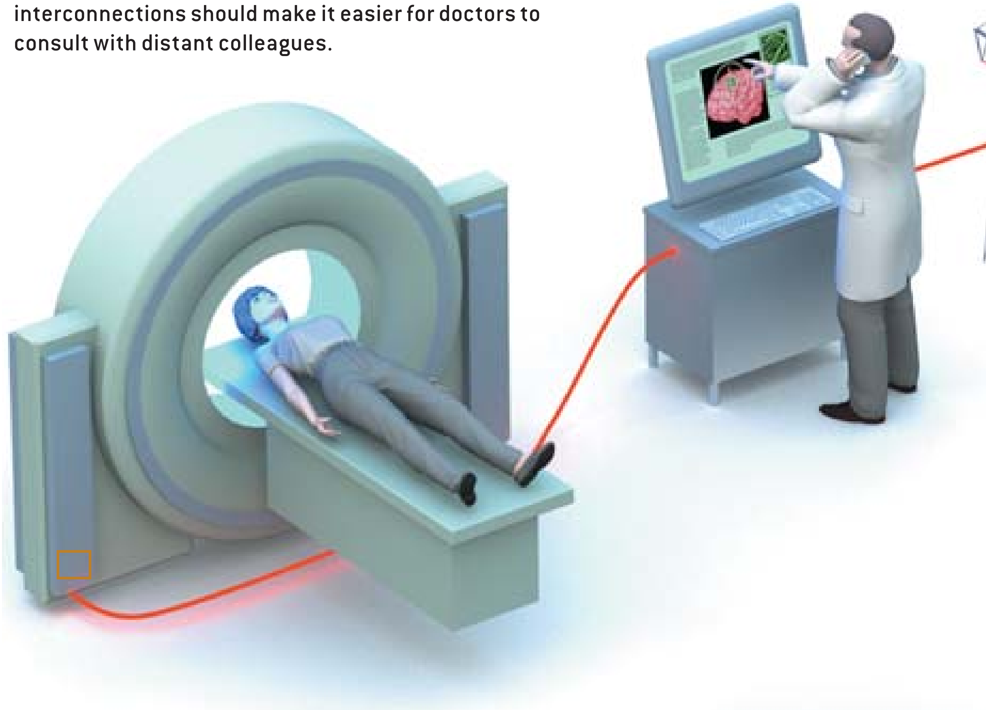

PRODUCING LIGHT

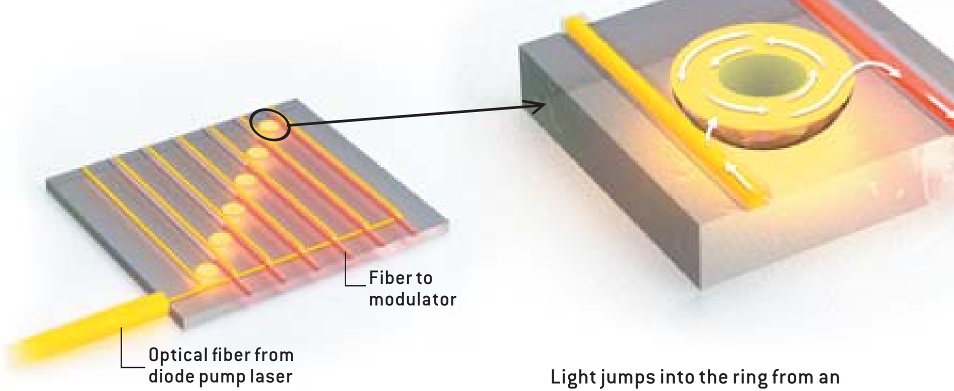

Microlasers built by Kerry J. Vahala of the California Institute of Technology can be constructed by the thousands on standard silicon microchips. The tiny rings can purify light pumped in from an inexpensive diode laser and change its color to match the standard wavelength used by other photonic components.

Light jumps into the ring from an ultrathin optical fiber (*yellow*) that passes nearby. Light of a particular frequency resonates within the ring (*arrows*) and stimulates it to emit a laser beam on a different fiber (*red*). In a working microchip, the ring lasers would probably be made on the rims of holes, with connecting fibers embedded in the chip surface.

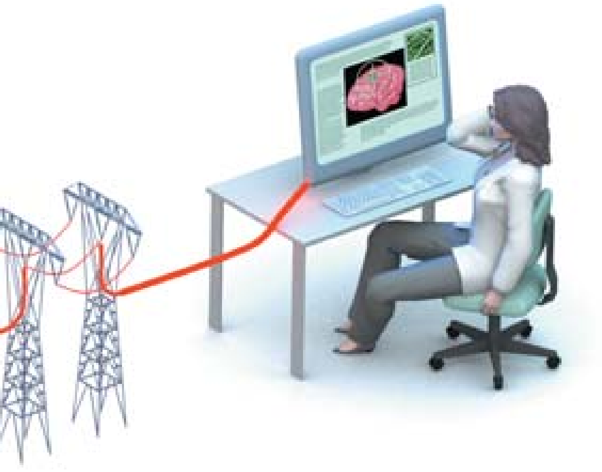

MODULATING AN OPTICAL SIGNAL

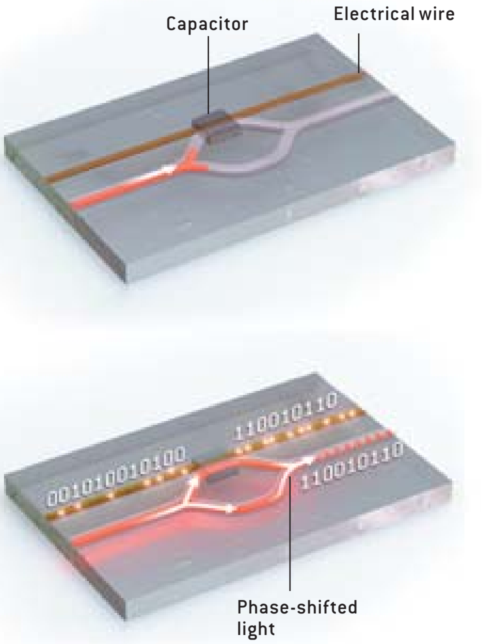

An optical modulator, such as one made last year by Mario Paniccia and his co-workers at Intel, takes digital bits in electrical form and encodes them onto a light beam. First the beam is split into two arms (*top*). Digital signals arrive at a capacitor on one or both of the arms. The signal alters the phase of light passing near each capacitor. When the arms recombine, the phase-shifted light beams interfere, creating pulses in the outgoing beam (*bottom*).

PHOTONIC MICROPROCESSING

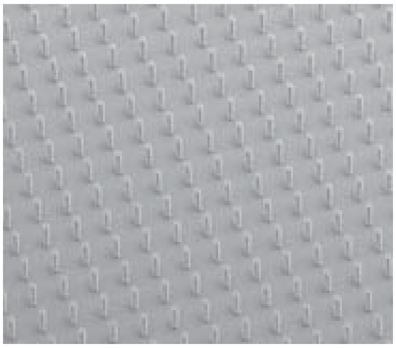

Minuscule pillars made of translucent plastic could connect microprocessors to the circuit boards on which they rest. Thousands of such pillars, developed by James D. Meindl's group at the Georgia Institute of Technology, would be attached to the bottom of the processor.

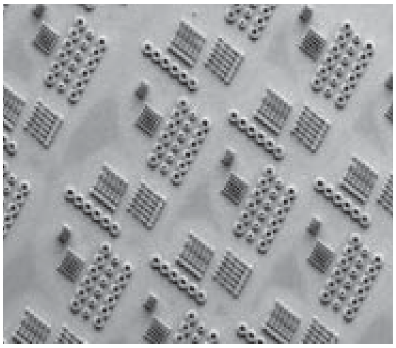

The pillars would fit into corresponding sockets on the surface of the circuit board.

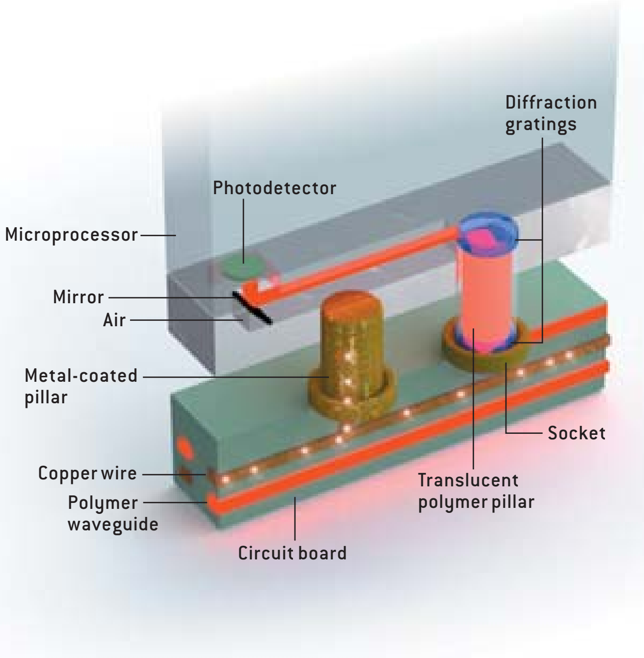

Both electronic and photonic signals could pass through pillars, carrying data to and from the microprocessor. Standard electrical connections would be made from wires to metal-coated pillars. Light pulses would flow through polymer waveguides, be turned by plastic diffraction gratings or metallic mirrors, and be detected by silicon photodiodes.

BRYAN CHRISTIE DESIGN; MUHANNAD S. BAKIR AND JAMES D. MEINDL (*polymer pillars*)

is closest to commercial success, as it has already produced chips for the telecommunications industry.

Hybrid devices cram logic-bearing silicon microchips into a small package that also houses III-V chips, which perform all of the optical duties. A CMOS plant would have to be significantly modified before managers could let gallium arsenide or indium phosphide anywhere near their multibillion-dollar equipment, because those compounds can contaminate a silicon production line. But the two halves of a hybrid device can be manufactured in separate factories and then assembled later.

Xanoptix, a small start-up in Merrimack, N.H., has used this technique to mate gallium arsenide lasers to silicon control chips. The result is a thumb-size optical connector that looks similar to a USB plug. But whereas USB cables top out at less than 0.5 Gbps, Xanoptix says its photonic jack can pump up to 245 Gbps through a pencil-width bundle of 72 optical fibers.

The hybrid approach faces a long-term problem, however—the faster microprocessors become, the hotter they run. The hottest spots on some chips already rise above 80 Celsius, the temperature at which III-V lasers start to burn out. So hybrid optoelectronic chips may find a niche in slower external connectors and peripheral devices, rather than at the center of the computer.

Intel has ordered its scientists to stick to CMOS, in the hope that it might one day build entire photonic systems right into microprocessors or motherboard chips using its existing factories. To make this so-called monolithic integration work, engineers have been tricking silicon and the few other elements that are CMOS-friendly into emitting, manipulating and detecting light.

That first step is a doozy: "We think we can do everything in silicon—except the laser," Paniccia says. Silicon by itself lacks the quantum-mechanical wherewithal to make light. Coffa's group at STMicroelectronics has discovered a way around part of that problem, however. By infusing small amounts of cerium or erbium into a layer of silicon dioxide laced with silicon nanocrystals, the researchers constructed silicon chips that glow green or blue in response to a small voltage.

Because the luminescence is incoherent, these are light-emitting diodes (LEDs), not lasers. "But they are as efficient as gallium arsenide LEDs," Coffa reports. And because they are CMOS-compatible, he adds, "we can incorporate them directly into our existing electronic parts." By next year, STMicroelectronics plans to introduce silicon-based optocouplers that allow computers to control high-voltage machinery.

The silicon LEDs might also serve as a light source for a CMOS-compatible laser demonstrated earlier this year by Kerry J. Vahala and his co-workers at the California Institute of Technology. Vahala and others have been experimenting with microscopic disks of silicon dioxide perched on silicon pillars. By smoothing the edges of each disk and carefully controlling its diameter, Vahala turned it into the optical equivalent of a whispering gallery. Light passing through a nearby optical fiber leaks into the disk and circles its edge over and over, building in intensity by a factor of a million or more [*see illustration on page 84*] before it is emitted as laser light.

Incoherent sources such as LEDs could feed light to the disks, or they could be used to purify laser light coming in from outside the chip, regenerate it as it fades and tune it to a new wavelength. "Instead of etching disks, we could create holes. The laser would then form on the interior rim of the void," Vahala says. That would make it easier to connect the devices to waveguides and other photonic components on the surface. "These 'microcavity' lasers could be sources of carrier signals for information launched from the chip," he suggests.

For that to work, engineers will need a way to transfer information from electronic to optical form. Until this year, that was hard to do in silicon, which offered only a slow and unsteady lever for manipulating light. But in February, Paniccia's team unveiled a way to use silicon to modulate a laser beam—to flicker it in step with a digital signal—50 times faster.

CHANGING THE SHAPE

Over the long term, high-speed photonic connections may obviate the need to pack all the components of a computing system into a single box. "With an electrical signal, it is a totally different physical phenomenon to go four meters than it is to go four inches," notes Ravindra A. Athale of the Defense Advanced Research Projects Agency. "But once you launch into the optical domain, there is no significant difference between four inches and four meters"—or even 400 meters. So optically connected pieces could be widely scattered and still function as a seamless machine.

In a car, for example, multiple processors, memory banks and disk drives could be embedded within the body and linked by fiber optics. "I understand that Daimler Chrysler is already planning to have laser-based, gigabit optical links inside their 2005-model S-class Mercedes cars," Athale reports.

In principle, each worker in a future office could use a computer assembled by making photonic connections among parts spread

"There are about 100 modulators on this," Paniccia says, as he picks up a silicon chip the size of a postage stamp. We are in Optics Lab 1A at Intel's research center in Santa Clara. "And here's one in action," he continues, motioning to a workstation at the far end of the room. The computer is playing a high-definition DVD of *Terminator 3*. As it processes the video stream, the machine sends a copy of each bit down an Ethernet cable and into a tiny circuit board containing a single modulator.

Although it is powered by a laser, the device works on the same principle as AM radio [*see illustration on preceding page*]. It splits the microscopic laser beam into two arms. A CMOS capacitor under each arm—electrically connected to the Ethernet cable—stores and releases static electricity. "When those regions are highly charged, electrons interact with the light," Paniccia explains, shifting the relative position of the light waves. As the two halves rejoin, the peaks and troughs in their waves interfere, causing the out-

BRYAN CHRISTIE DESIGN

OF COMPUTERS

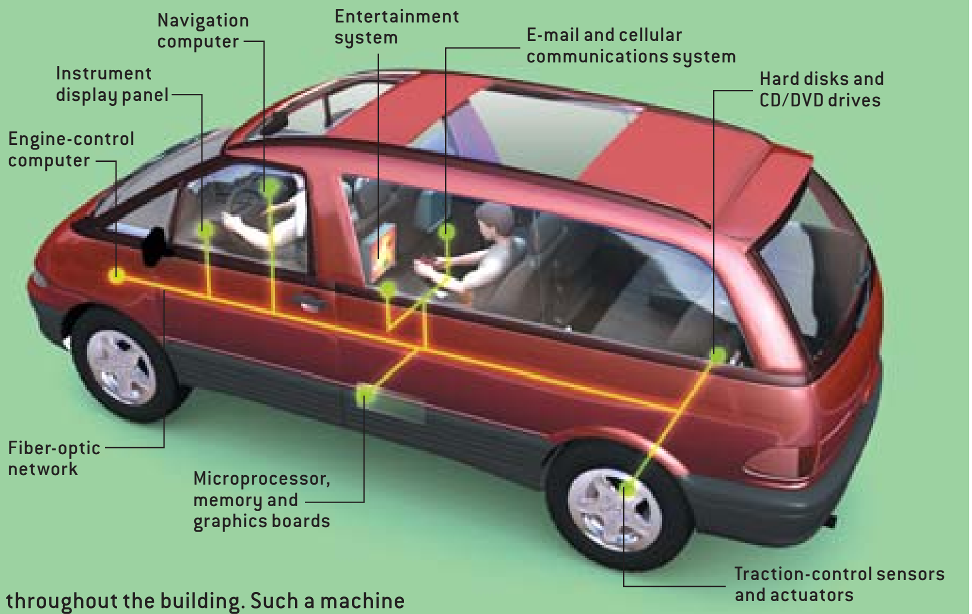

throughout the building. Such a machine could be temporarily "upgraded" at a keystroke to a faster processor or a larger memory bank, if a particular task demanded extra resources.

The ramifications of photonic links for computer architecture are still largely speculative. "System architects tend not to think along these lines because they know that such a technology does not yet exist," Athale says. But it is on its way.

put beam to pulse in the same pattern as the bits in the video stream.

Those pulses carry the data over a single fiber, thin as spider silk, that leaves the chip and connects to a photodetector attached to a second computer a few feet away. The two computers show Arnold Schwarzenegger leaping from his car in perfect synchrony.

To date, the modulator has run at rates up to 2.5 GHz. "But we can make it much smaller, and we are confident that we can scale it up to 10 GHz," Paniccia asserts. "By combining all these elements in a single chip, we'll be able to make this," he says, raising an Ethernet plug up to his eye: "little silicon-optical devices you can plug in anywhere and a $250 network interface card that replaces a $25,000 router."

"Of course, if we're going to transmit at 10 GHz, we need to be able to receive at that speed, too," Morse points out. Silicon is as transparent as glass at the infrared wavelengths typically used in photonic devices. But with the addition of germanium to the mix—which chipmakers have begun doing anyway to help speed up their processors—CMOS-compatible photodetectors have been built to convert the light pulses back into electronic bits.

The Best of Both Worlds

AS A PHOTONIC MATERIAL, silicon has come a long way in two years. But it has much farther yet to go if it is to handle optical data at more than 20 GHz. So it may be that a relatively new method of pulling photonics into the electronics, known as polylithic integration, will prove the most economical.

The general idea is to attach a CMOS processor to the motherboard with a dense array of both optical and electronic connections. Light could then be pumped into the processor from small (and thus relatively affordable) III-V chips, mounted a safe distance away so they do not overheat.

James D. Meindl and Muhannad S. Bakir of the Georgia Institute of Technology, working with Anthony V. Mulé of Intel, have demonstrated several polylithic schemes. One is called a sea of leads: thousands of microscopic S-shaped springs of metal are etched onto the processor as a final step in its manufacture. Electrical signals pass through the metal springs; light signals shoot through the holes in their centers and hit diffracting gratings that deflect the pulses into waveguides buried within the chip or motherboard.

In a second scheme, the processor rests on thousands of transparent plastic pillars, which fit into circular plastic sockets on the circuit board [*see illustration on page 85*]. Meindl's group has fabricated regular arrays of pillars that are just five microns wide and 12 microns apart. The researchers have also demonstrated how some of the cylinders and sockets could be coated with metal to make electrical, rather than optical, connections.

Bristling with tens of thousands of such minuscule pillars, a microprocessor 10 or 15 years from now may throb with infrared flashes even as it hums with high-frequency electrons. Microchip factories may etch transistors and wires in the spaces between lasers and waveguides. The long separation between laboratory photonics and consumer electronics seems to be closing at last, and our machines will be better for it. SA

W. Wayt Gibbs is senior writer.

MORE TO EXPLORE

A High-Speed Silicon Optical Modulator Based on a Metal-Oxide-Semiconductor Capacitor. Mario Paniccia et al. in *Nature*, Vol. 427, pages 615–618; February 12, 2004.

Ultralow-Threshold Microcavity Raman Laser on a Microelectronic Chip. Kerry J. Vahala et al. in *Optics Letters*, Vol. 29, No. 11, pages 1224–1226; June 1, 2004.

Sea of Polymer Pillars Electrical and Optical Chip I/O Interconnections for Gigascale Integration. M. S. Bakir and J. D. Meindl in *IEEE Transactions On Electron Devices*, Vol. 51, No. 7, pages 1069–1077; July 2004.

Silicon Photonics. Edited by Lorenzo Pavesi and David J. Lockwood. Springer, 2004.