COMPUTERS_CIRCUIT DESIGN

Wired for Speed

As chips shrink, researchers look to optical and radio-frequency interconnects

Every few years scientists and manufacturers from all over the globe draw up what is now called the International Technology Roadmap for Semiconductors, an assessment of semiconductor technology requirements and research goals over the next 15 years. Ironically, one of the biggest challenges the industry faces is traffic congestion on and between the chips themselves.

Thanks to ever shrinking transistors on integrated circuits (ICs), computers have become quicker and more powerful. But as faster and smaller transistors are packed onto a microchip, the layers of wires that connect the transistors must shrink as well. The problem, though, is that the smaller the cross section of a wire, the tougher it is to push an electrical signal through. Capacitance between extremely thin wires can add to the trouble. "The transistors are getting faster, but the wires are getting slower—and that's a prescription for disaster," says Kevin Martin of the Georgia Institute of Technology, who helps to direct the Interconnect Focus Center, an entity created to avert that disaster. Based at Georgia Tech, the center encompasses research at five other universities and is part of the larger Technology Focus Center research program, launched in 1998 with $10 million annual funding per center from the Semiconductor Industry Association's member companies and other groups.

The semiconductor industry has high hopes for the interconnect program. Right now the best commercial interconnects are copper wires, introduced into microchips in late 1998. But although copper is a vast improvement over aluminum interconnects, Martin says, the metal simply won't scale down sufficiently. For example, the intrinsic switching time for transistors having 100-nanometer gate lengths (circuit features referring to distances that electrons must travel) is on the order of 0.1 picosecond, 70 times faster

Q&A_MARK MELLIAR-SMITH

The Route to Finer Lines

Physics may soon put an end to the rapid pace at which manufacturers can double a chip's speed. Today's one-gigahertz microprocessors have up to 20 million transistors and circuit features (specifically, gate lengths) only 140 nanometers long. They are born out of optical lithography using a light source with a wavelength of 248 nanometers. Light shining through a glass mask (essentially a stencil of a chip's features) projects the circuit pattern onto a silicon wafer coated with photoresist, an organic film that hardens when exposed to light. The shorter the wavelength of light projecting through the mask, the smaller the features on the chip.

But etching features much smaller than 100 nanometers by means of optical lithography is a whole new ball game, requiring novel photoresist materials (their sensitivity depends on the wavelength of light). And by the time chips featuring 70 nanometers or smaller are on deck, optical lithography may have to be put out to pasture altogether. Still, Mark Melliar-Smith, president and CEO of the semiconductor research consortium International SEMATECH, expects three to 10 gigahertz logic chips containing five billion transistors to be in production by 2014—one way or another. *—D.P.*

Is the end in sight for chip patterning by optical lithography?

One-hundred-ninety-three-nanometer optical lithography (which can produce transistors with 100-nanometer gates) is about 12 months away from manufacture. Beyond are several alternatives to get down to 50 nanometers, which, according to the International Technology Roadmap for Semiconductors, is about a decade away.

So what's the path from 193- to 50-nanometer lithography?

We have three different choices, each of which has significant technical challenges to be solved. The first, 157-nanometer optical lithography, uses a shorter wavelength of light—essentially more of the same of what we've been doing—but we'll have to find new photoresists and solve some other problems as well. The second one is electron projection lithography (EPL), of which Lucent's SCALPEL is the embodiment in the U.S. That has a different set of problems, including keeping the mask perfectly clean and getting the throughput up to levels comparable to those of optical lithography. The final one is extreme ultraviolet lithography (EUV), using an 11- to 13-nanometer radiation source, which requires special mirrors, a complicated new laser and thin-layer imaging techniques because, at this wavelength, materials are almost all opaque.

Do you fear that you're spreading yourselves too thin?

Having more than one choice gives you the opportunity to hedge your bets. Frankly, though, the two critical reasons we're pursuing all three technologies are that, one, all of them have significant risks, and we don't at this point know which is likely to be the most successful. And second, all three technologies have significant commercial support.

Who's winning so far?

EPL and the 157-nanometer optical lithography should be available in terms of an alpha-chip manufacturing tool within two years. At that point, we can fly them off against each other. EUV may be a little later but has the potential for the highest resolution, and this is important. For instance, 157-nanometer optical lithography only enables us to go 25 percent further down in size than we can today—that's a relatively short life span. But EPL and EUV offer the opportunity to go down several generations.

COURTESY OF SEMATECH

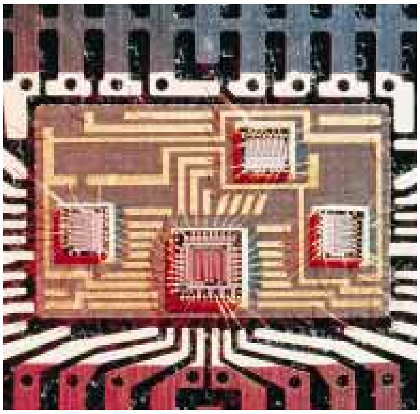

SPEED BUMPS: Interconnects, such as those between integrated-circuit components, could slow future computers.

ALFRED PASIEKA *Science Photo Library*

than the response time of a typical one-millimeter-long copper interconnect wire. And the pressure is on—the International Technology Roadmap calls for chips with 100-nanometer gate lengths next year.

Leading the race for new interconnects are optical ones—replacing wires with fiber-optic cables that are resistance-free. Optics are ideal for high-bandwidth applications and are not constrained by long distances, unlike wire interconnects. Research at the Massachusetts Institute of Technology is focusing on sending signals between transistors on the chip itself, whereas David A. B. Miller, an electrical engineer at Stanford University, has directed his efforts at enabling separate chips to talk to one another at the necessary speed without having to be crammed closely together. "Using optics instead of wires is like being able to put in a 1,000-lane highway where you previously had a one-lane freeway," Miller remarks.

There are two main approaches to optical interconnects, albeit with myriad variations. One is transmitting light beams, generated by five- to 20-micron-high vertical cavity-surface-emitting lasers, or VCSELs, down waveguides built onto the chips. The other paradigm is based on freespace optics. Light from an external source can be reflected by tiny structures called quantum-well light modulators, which rapidly switch on and off in response to small voltages. Alternatively, patterns of light generated on one chip by VCSELs can be imaged on the other chip by a lens. "The second chip behaves like your retina," Miller explains. Though not yet ready for prime time, optical interconnects have been successfully demonstrated at several universities.

Just out of the gate, so to speak, is wireless-interconnect technology using radio-frequency (RF) signals. Various groups are working on this concept, including M. C. Frank Chang of the University of California at Los Angeles under the auspices of the Interconnect Focus Center. One example of how RF interconnects would work was presented in March by Kenneth K. O of the University of Florida and graduate students Brian A. Floyd and Kihong Kim at the International Solid-State Circuits Conference. They delivered a paper on the use of RF signals in massively parallel computers. With such large computer systems, maintaining a constant clock signal throughout numerous microprocessors becomes difficult. O's group hopes to get around that by broadcasting a clock signal from one IC to others using microwaves. One design integrates millimeter-size receivers and antennae on each IC in a multichip module. "By propagating the signal at the speed of light, we're trying to reduce the clock skew," O says. "You could send a wave down to a multichip module and provide equal clock phase to a very large area."

The group recently demonstrated on-chip wireless transmission and reception of a 7.4-gigahertz clock signal. O believes the same technology could be modified for data transfer between chips as well. Not surprisingly, the biggest antagonist to wireless interconnects is noise. Both the chip's silicon substrate and the switching of the transistors themselves degrade and taint the radio signal. The materials in chips "are just not very friendly to radio reception," O says. Whether optics or RF, researchers will undoubtedly find ways to keep traffic moving on tomorrow's computer systems. *—David Pescovitz*

DAVID PESCOVITZ is a contributing editor to Wired *and writes frequently for* Scientific American. *He is based in Oakland, Calif.*

Chilly Crystals

Thermoelectrics could double computer speeds

Another trick that could speed up computers is the use of semiconducting, thermoelectric materials to cool microprocessor chips. Unlike most metals, which become hot when an electric current passes through them, these substances have the ability to carry away heat while conducting electricity. Since the 1950s thermoelectric materials, fashioned into miniature heat pumps, have chilled solid-state lasers, infrared detectors and other electronic devices, which tend to run best cold. Unfortunately, the lowest temperature achieved by existing materials hovers around –50 degrees Celsius, a drop not large enough to justify routine use of these expensive minirefrigerators in today's computers.

Now a team led by Mercouri G. Kanatzidis, a chemist at Michigan State University, has concocted a new compound that can beat out the existing competition by cooling to a record –100 degrees C and make faster chips a reality. "This new technology has the potential to increase computer speeds by 100 percent simply by cooling the chip," Kanatzidis notes.

The new crystal—a mixture of bismuth, tellurium and cesium—is a breakthrough because it enhances the thermoelectric effect by being both a good conductor and thermal insulator. But a thermoelectric cooler cannot be made yet, Kanatzidis reveals, because his team has only developed one of the necessary conductors. For heat pumps to work, two different material types (technically called *n*- and *p*-types) are needed to create a temperature difference.

Kanatzidis believes that a cooling device made from his new material could be sandwiched between a microprocessor and a heat exchanger, such as a fan. Heat, generated from the superfast chip sitting on the surface of the semiconductor, would travel from top to bottom and be dispersed by the fan. Direct cooling of the chip would translate into higher speed because the mobility of the electrons would increase in a chillier environment.

Speed freaks, though, will have to keep it in the slow lane for a while longer. A workable prototype for the general market will take several more years to develop. Still, since the story broke out, "people have been calling me and asking when I can ship them 2,000 of these things," Kanatzidis says with a laugh. "We have a material, not a device." *—Diane Martindale*