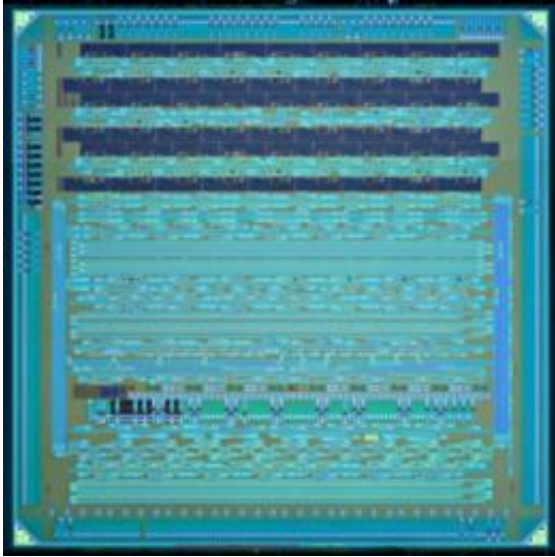


Microchips' optical future

February 15 2012, by Larry Hardesty



A new test chip developed by Vladimir Stojanovic, Rajeev Ram and their colleagues, which monolithically integrates electrical and optical components and was produced on an existing IBM manufacturing line. Image: Vladimir Stojanovic, Rajeeve Ram and Milos Popovic

Computer chips are one area where the United States still enjoys a significant manufacturing lead over the rest of the world. In 2011, five of the top 10 chipmakers by revenue were U.S. companies, and Intel, the largest of them by a wide margin, has seven manufacturing facilities in the United States, versus only three overseas.

The most recent of those to open, however, is in [China](#), and while that may have been a strategic rather than [economic decision](#) — an attempt

to gain leverage in the Chinese computer market — both the Chinese and Indian governments have invested heavily in their countries' chip-making capacities. In order to maintain its manufacturing edge, the [United States](#) will need to continue developing new technologies at a torrid pace. And one of those new technologies will almost certainly be an integrated optoelectronic [chip](#) — a chip that uses light rather than electricity to move data.

As chips' computational power increases, they need higher-bandwidth connections — whether between servers in a server farm, between a chip and main memory, or between the individual cores on a single chip. But with electrical connections, increasing bandwidth means increasing power.

Optoelectronic chips could drastically reduce future computers' power consumption. But to produce the optoelectronic chips used today in telecommunications networks, chipmakers manufacture optical devices — such as lasers, photodetectors and modulators — separately and then attach them to silicon chips. That approach wouldn't work with conventional microprocessors, which require a much denser concentration of higher-performance components.

The most intuitive way to add optics to a microprocessor's electronics would be to build both directly on the same piece of silicon, a technique known as monolithic integration.

In a 2010 paper in the journal *Management Science*, Erica Fuchs, an assistant professor of engineering and public policy at Carnegie Mellon University, who got her PhD in 2006 from MIT's Engineering Systems Division, and MIT's Randolph Kirchain, a principal research scientist at the Materials Systems Laboratory, found that monolithically integrated chips were actually cheaper to produce in the United States than in low-wage countries.

“The designers and the engineers with the capabilities to produce those technologies didn’t want to move to developing East Asia,” Fuchs says. “Those engineers are in the U.S., and that’s where you would need to manufacture.”

During the telecom boom of the late 1990s, Fuchs says, telecommunications companies investigated the possibility of producing monolithically integrated communications chips. But when the bubble burst, they fell back on the less technically demanding process of piecemeal assembly, which was practical overseas. That yielded chips that were cheaper but also much larger.

While large chips are fine in telecommunications systems, they’re not an option in laptops or cellphones. The materials used in today’s optical devices, however, are incompatible with the processes currently used to produce microprocessors, making monolithic integration a stiff challenge.

Making the case

According to Vladimir Stojanovic, an associate professor of electrical engineering, microprocessor manufacturers are all the more reluctant to pursue monolithic integration because they’ve pushed up against the physical limits of the transistor design that has remained more or less consistent for more than 50 years. “It never was the case that from one generation to another, you’d be completely redesigning the device,” Stojanovic says. U.S. chip manufacturers are so concerned with keeping up with Moore’s Law — the doubling of the number of transistors on a chip roughly every 18 months — that integrating optics is on the back burner. “You’re trying to push really hard on the transistor,” Stojanovic says, “and then somebody else is telling you, ‘Oh, but you need to worry about all these extra constraints on photonics if you integrate in the same front end.’”

To try to get U.S. chip manufacturers to pay more attention to optics, Stojanovic and professor of electrical engineering Rajeev Ram have been leading an effort to develop techniques for monolithically integrating optical components into [computer chips](#) without disrupting existing manufacturing processes. They've gotten very close: Using IBM's chip-fabrication facilities, they've produced chips with photodetectors, ring resonators (which filter out particular wavelengths of light) and waveguides (which conduct light across the chip), all of which are controlled by on-chip circuitry. The one production step that can't be performed in the fabrication facility is etching a channel under the waveguides, to prevent light from leaking out of them.

But Stojanovic acknowledges that optimizing the performance of these optical components would probably require some modification to existing processes. In that respect, the uncertainty of the future of transistor design may actually offer an opportunity. It could be easier to add optical components to a chip being designed from the ground up than to one whose design is fixed. "That's the moment it has to come in," Stojanovic says, "at the moment where everything's in flux, and soft."

Loyal opposition

Another of Stojanovic and Ram's collaborators on the monolithic-integration project is Michael Watts, who received his PhD from MIT in 2005 and returned in 2010 as an associate professor of electrical engineering after a stint at Sandia National Labs. Stojanovic and Watts are also collaborating with researchers at MIT's Lincoln Laboratory on a different project — with Watts as the primary investigator — in which optical and electrical components are built on different wafers of silicon, which are then fused together to produce a hybrid wafer.

This approach falls somewhere between monolithic integration and the piecemeal-assembly technique used today. Because it involves several

additional processing steps, it could prove more expensive than fully realized monolithic integration — but in the near term, it could also prove more practical, because it allows the performance of the optics and electronics to be optimized separately. As for why the researchers would collaborate on two projects that in some sense compete with each other, Watts says, “Sometimes the best policy is: When you come to a fork in the road, take it.”

And indeed, Ram thinks that the two approaches may not compete with each other at all. Even if optical and electrical components were built on separate chips and fused together, Ram explains, the optical chip would still require some electronics. “You will likely have the electronics that do the detection on the same chip as the photodetector,” Ram says. “Same for the electronics that drive the modulator. Either way, you have to figure out how to integrate photonics in the plane with the transistor.”

Getting on board

With both approaches, however — monolithic integration or chip stacking — the laser that provides the data-carrying beam of light would be off-chip. Off-chip lasers may well be a feature of the first optoelectronic chips to appear in computers, which might be used chiefly to relay data between servers, or between a processor and memory. But chips that use optics to communicate between cores will probably require on-chip lasers.

The lasers used in telecommunications networks, however, are made from exotic semiconductors that make them even more difficult to integrate into existing manufacturing processes than photodetectors or waveguides. In 2010, Lionel Kimerling, the Thomas Lord Professor of Materials Science and Engineering, and his group demonstrated the first laser built from germanium that can produce wavelengths of light useful for optical communication.

Since many chip manufacturers already use germanium to increase their chips' speed, this could make on-chip lasers much easier to build. And that, in turn would reduce chips' power consumption even further. "The on-chip laser lets you do all the energy-efficiency tricks that you can do with electronics," Kimerling says. "You can turn it off when you're not using it. You can reduce its power when it doesn't have to go that far."

But the real advantage of on-chip lasers, Kimerling says, would be realized in microprocessors with hundreds of cores, such as the one currently being designed in a major project led by Anant Agarwal, director of MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL). "Right now, we have a vision that you need one laser for one core, and it can communicate with all the other cores," Kimerling says. "That allows you to trade data among the cores without going to [memory], and that's a major energy saving."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Microchips' optical future (2012, February 15) retrieved 3 May 2024 from <https://phys.org/news/2012-02-microchips-optical-future.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--