

Selling chip makers on optical computing

November 24 2009, by Larry Hardesty



In the prototype optical chip shown here, the circles in the top two rows are "ring resonators" that can filter out light of different wavelengths. Image courtesy of Vladimir Stojanovic

(PhysOrg.com) -- Computer chips that transmit data with light instead of electricity consume much less power than conventional chips, but so far, they've remained laboratory curiosities. Professors Vladimir Stojanović and Rameez Ram and their colleagues in MIT's Research Laboratory of Electronics and Microsystems Technology Laboratory hope to change that, by designing optical chips that can be built using ordinary chip-manufacturing processes.

"I don't see anyone else that's doing that," says Michael Watts, a researcher at Sandia National Laboratories who's also working on optical chips. "If they're successful at that, then convincing a major processor or memory manufacturer that this is a viable approach will be much, much

easier.”

Granted access to the same manufacturing facilities that Texas Instruments uses to produce cell phone chips and microprocessors, the MIT researchers have demonstrated that they can put large numbers of working optical components and electronics on the same [chip](#). But so far, the electronics haven’t been able to control the optics directly. That’s something that Stojanović hopes to show with a new batch of chips due back from TI and another major semiconductor manufacturer this winter.

Optical data transmission could solve what will soon be a pressing problem in chip design. As chips’ computational capacity increases, they need higher-bandwidth connections to send data to memory; otherwise, their added processing power is wasted. But sending more data over an electrical connection requires more power.

Smaller transistors are more energy-efficient than larger ones, so over time, chips’ total power consumption has changed little. But “the fraction of power that’s used for communications has grown,” Watts says. “At some point, you have to devote all your power to communications. And that point’s not too far off. And then what’s left for computation? Nothing.” Future chips could simply draw more power, but then they would also be harder to cool, and the battery life of laptops and handheld devices would dramatically shorten.

So chip companies would welcome a more energy-efficient way to move data around -- if they were confident that it was cost-effective. And that’s why demonstrating compatibility with existing manufacturing processes would be so persuasive.

Manufacturers build chips by sequentially depositing layers of different materials — like silicon, silicon dioxide, and copper — on a wafer of

silicon, and then etching the layers away to build three-dimensional structures. The problem with using existing processes to build optical components is that the deposition layers are thinner than would be ideal. “You would want a normal photonic device to be a little bit taller and thinner so that you can minimize the surface-roughness losses,” Stojanović says. “Here you don’t have that choice because the film thicknesses are set by fabrication.”

Optical chips use structures called waveguides to direct light, and researchers trying to add optical components to a silicon chip usually carve the waveguides out of a single crystal of silicon, Stojanović says. But waveguides made from single-crystal silicon require insulating layers above and below them, which standard chip-manufacturing processes like TI’s and Intel’s provide no way to deposit. They do, however, provide a way to deposit insulators above and below layers of polysilicon, which consists of tiny, distinct crystals of silicon clumped together and is typically used in the part of a transistor called the gate. So the MIT researchers built their waveguides from polysilicon instead.

So far, TI has produced two sets of prototypes for the MIT researchers, one using a process that can etch chip features as small as 65 nanometers, the other using a 32-nanometer process. To keep light from leaking out of the polysilicon waveguides, the researchers hollowed out the spaces under them when they got the chips back — the sole manufacturing step that wasn’t possible using TI’s in-house processes. But “that can probably be fixed more elegantly in the fabrication house if they see that by fixing that, we get all these benefits,” Watts says. “That’s a pretty minor modification, I think.”

The MIT researchers’ design uses light provided by an off-chip laser. But in addition to guiding the beam, the chip has to be able to load information onto it and pull information off of it. Both procedures use ring resonators, tiny rings of silicon carved into the chip that pull light of

a particular frequency out of the waveguide. Rapidly activating and deactivating the resonators effectively turns the light signal on and off, and bursts of light and the gaps between them can represent the ones and zeroes of digital information.

To meet the bandwidth demands of next-generation chips, however, the waveguides will have to carry 128 different wavelengths of light, each encoded with its own data. So at the receiving end, the ring resonators provide a bank of filters to disentangle the incoming signals. On the prototype chips, the performance of the filter banks was “the most amazing result to us,” Stojanović says, “which kind of said that, okay, there’s still hope, and we should keep doing this.” The wavelength of light that the resonators filter is determined by the size of their rings, and no one — at either TI or MIT — could be sure that conventional manufacturing processes were precise enough to handle such tiny variations.

Stojanović hopes that the next batch of prototypes, which should give the chips’ electronics control over the optical components, will demonstrate that the resonators perform as well when loading data onto light beams. At the same time, the team is looking to extend its approach to memory chips. “The memory’s a much tougher nut to crack, because it is such a cost-driven business, where every process step matters,” Stojanović says. “Things are a lot harder to change there, and optics really needs to be absolutely compatible with process flow.” But if memory chips as well as processors sent data optically, Stojanović says, then in addition to saving power, they could make computers much faster. “If you just focus on the processor itself, you maybe get a 4x advantage with photonics,” Stojanović says. “But if you focus on the whole connectivity problem, we’re talking 10, 20x improvements in system performance.”

Provided by MIT

Citation: Selling chip makers on optical computing (2009, November 24) retrieved 10 April 2024 from <https://phys.org/news/2009-11-chip-makers-optical.html>

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.