

Chips as mini Internets

April 10 2012, by Larry Hardesty

Computer chips have stopped getting faster. In order to keep increasing chips' computational power at the rate to which we've grown accustomed, chipmakers are instead giving them additional "cores," or processing units.

Today, a typical chip might have six or eight cores, all communicating with each other over a single bundle of wires, called a bus. With a bus, however, only one pair of cores can talk at a time, which would be a serious limitation in chips with hundreds or even thousands of cores, which many electrical engineers envision as the future of computing.

Li-Shiuan Peh, an associate professor of electrical engineering and computer science at MIT, wants cores to communicate the same way computers hooked to the Internet do: by bundling the information they transmit into "packets." Each core would have its own router, which could send a packet down any of several paths, depending on the condition of the network as a whole.

At the Design Automation Conference in June, Peh and her colleagues will present a paper she describes as "summarizing 10 years of research" on such "networks on chip." Not only do the researchers establish theoretical limits on the efficiency of packet-switched on-chip communication networks, but they also present measurements performed on a test chip in which they came very close to reaching several of those limits.

Last stop for buses

In principle, multicore chips are faster than single-core chips because they can split up computational tasks and run them on several cores at once. Cores working on the same task will occasionally need to share data, but until recently, the core count on commercial chips has been low enough that a single bus has been able to handle the extra communication load. That's already changing, however: "Buses have hit a limit," Peh says. "They typically scale to about eight cores." The 10-core chips found in high-end servers frequently add a second bus, but that approach won't work for chips with hundreds of cores.

For one thing, Peh says, "buses take up a lot of power, because they are trying to drive long wires to eight or 10 cores at the same time." In the type of network Peh is proposing, on the other hand, each core communicates only with the four cores nearest it. "Here, you're driving short segments of wires, so that allows you to go lower in voltage," she explains.

In an on-chip network, however, a packet of data traveling from one core to another has to stop at every router in between. Moreover, if two packets arrive at a router at the same time, one of them has to be stored in memory while the router handles the other. Many engineers, Peh says, worry that these added requirements will introduce enough delays and computational complexity to offset the advantages of packet switching. "The biggest problem, I think, is that in industry right now, people don't know how to build these networks, because it has been buses for decades," Peh says.

Forward thinking

Peh and her colleagues have developed two techniques to address these concerns. One is something they call "virtual bypassing." In the Internet, when a packet arrives at a router, the router inspects its addressing

information before deciding which path to send it down. With virtual bypassing, however, each router sends an advance signal to the next, so that it can preset its switch, speeding the packet on with no additional computation. In her group's test chips, Peh says, virtual bypassing allowed a very close approach to the maximum data-transmission rates predicted by theoretical analysis.

The other technique is something called low-swing signaling. Digital data consists of ones and zeroes, which are transmitted over communications channels as high and low voltages. Sunghyun Park, a PhD student advised by both Peh and Anantha Chandrakasan, the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering, developed a circuit that reduces the swing between the high and low voltages from one volt to 300 millivolts. With its combination of virtual bypassing and low-swing signaling, the researchers' test chip consumed 38 percent less energy than previous packet-switched test chips. The researchers have more work to do, Peh says, before their test chip's power consumption gets as close to the theoretical limit as its data transmission rate does. But, she adds, "if we compare it against a bus, we get orders-of-magnitude savings."

Luca Carloni, an associate professor of computer science at Columbia University who also researches networks on chip, says "the jury is always still out" on the future of chip design, but that "the advantages of packet-switched networks on [chip](#) seem compelling." He emphasizes that those advantages include not only the operational efficiency of the chips themselves, but also "a level of regularity and productivity at design time that is very important." And within the field, he adds, "the contributions of Li-Shiuan are foundational."

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: Chips as mini Internets (2012, April 10) retrieved 4 June 2024 from <https://phys.org/news/2012-04-chips-mini-internets.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.