

# More chip cores can mean slower supercomputing, simulation shows

January 14 2009

---



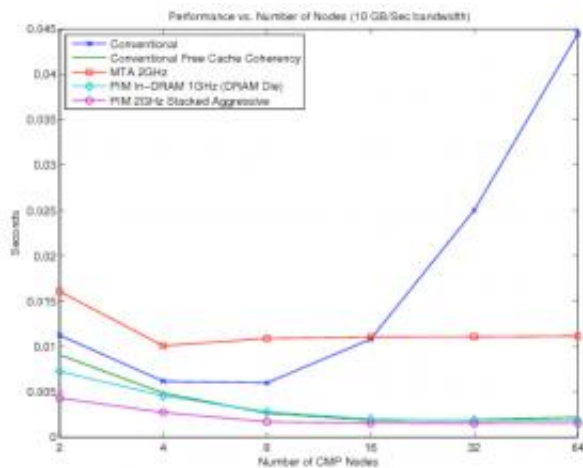
**THE MULTICORE DILEMMA:** more cores on a single chip don't necessarily mean faster clock speeds, a Sandia simulation has determined. (Photo by Randy Montoya)

(PhysOrg.com) -- The worldwide attempt to increase the speed of supercomputers merely by increasing the number of processor cores on individual chips unexpectedly worsens performance for many complex applications, Sandia simulations have found.

A Sandia team simulated key algorithms for deriving knowledge from large data sets. The simulations show a significant increase in speed going from two to four multicores, but an insignificant increase from four to eight multicores. Exceeding eight multicores causes a decrease in speed. Sixteen multicores perform barely as well as two, and after that, a

steep decline is registered as more cores are added.

The problem is the lack of memory bandwidth as well as contention between processors over the memory bus available to each processor. (The memory bus is the set of wires used to carry memory addresses and data to and from the system RAM.)



The graph depicts simulations of four potential multicore computers: the “conventional,” which adds more standard cores to a single processor socket; an MTA, which looks like the processor used the exotic Cray XMT supercomputer; and a PIM, which is based on Sandia's X-caliber processor design and includes memory tightly integrated with the processor. The fourth line simulates a conventional processor that represents a theoretical ideal. (Graphic by Richard Murphy et al, Sandia National Laboratories)

## A supermarket analogy

To use a supermarket analogy, if two clerks at the same checkout counter are processing your food instead of one, the checkout process should go faster. Or, you could be served by four clerks.

Or eight clerks. Or sixteen. And so on.

The problem is, if each clerk doesn't have access to the groceries, he or she doesn't necessarily help the process. Worse, the clerks may get in each other's way.

Similarly, it seems a no-brainer that if one core is fast, two would be faster, four still faster, and so on.

But the lack of immediate access to individualized memory caches — the "food" of each processor — slows the process down instead of speeding it up once the number of cores exceeds eight, according to a simulation of high-performance computers by Sandia's Richard Murphy, Arun Rodrigues and former student Megan Vance.

"To some extent, it is pointing out the obvious — many of our applications have been memory-bandwidth-limited even on a single core," says Rodrigues. "However, it is not an issue to which industry has a known solution, and the problem is often ignored."

"The difficulty is contention among modules," says James Peery, director of Sandia's Computations, Computers, Information and Mathematics Center. "The cores are all asking for memory through the same pipe. It's like having one, two, four, or eight people all talking to you at the same time, saying, 'I want this information.' Then they have to wait until the answer to their request comes back. This causes delays."

"The original AMD processors in Red Storm were chosen because they had better memory performance than other processors, including other Opteron processors," says Ron Brightwell. "One of the main reasons that AMD processors are popular in high-performance computing is that they have an integrated memory controller that, until very recently, Intel processors didn't have."

Multicore technologies are considered a possible savior of Moore's Law,

the prediction that the number of transistors that can be placed inexpensively on an integrated circuit will double approximately every two years.

"Multicore gives chip manufacturers something to do with the extra transistors successfully predicted by Moore's Law," Rodrigues says. "The bottleneck now is getting the data off the chip to or from memory or the network."

A more natural goal of researchers would be to increase the clock speed of single cores, since the vast majority of applications are designed for single-core performance on word processors, music, and video applications. But power consumption, increased heat, and basic laws of physics involving parasitic currents meant that designers were reaching their limit in improving chip speed for common silicon processes.

"The [chip design] community didn't go with multicores because they were without flaw," says Mike Heroux. "The community couldn't see a better approach. It was desperate. Presently we are seeing memory system designs that provide a dramatic improvement over what was available 12 months ago, but the fundamental problem still exists."

In the early days of supercomputing, Seymour Cray produced a superchip that processed information faster than any other chip. Then a movement — led in part by Sandia — proved that ordinary chips, programmed to work different parts of a problem at the same time, could solve complex problems faster than the most powerful superchip. Sandia's Paragon supercomputer, in fact, was the world's first parallel processing supercomputer.

Today, Sandia has a large investment in message-passing programs. Its Institute for Advanced Architectures, operated jointly with Oak Ridge National Laboratory (ORNL) and intended to prepare the way for

exaflop computing, may help solve the multichip dilemma.

ORNL's Jaguar supercomputer, currently the world's fastest for scientific computing, is a Cray XT model based on technology developed by Sandia and Cray for Sandia's Red Storm supercomputer. Red Storm's original and unique design is the most copied of all supercomputer architectures.

Provided by Sandia National Laboratories

Citation: More chip cores can mean slower supercomputing, simulation shows (2009, January 14) retrieved 19 May 2024 from <https://phys.org/news/2009-01-chip-cores-slower-supercomputing-simulation.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.