

Mastering multicore

April 26 2010, by Larry Hardesty



Graphic: Christine Daniloff

(PhysOrg.com) -- MIT researchers have developed software that makes computer simulations of physical systems run much more efficiently on so-called multicore chips. In experiments involving chips with 24 separate cores -- or processors -- simulations of fluid flows were at least 50 percent more efficient with the new software than they were with conventional software. And that figure should only increase with the number of cores.

Complex computer models — such as atom-by-atom simulations of physical materials, or high-resolution models of weather systems — typically run on multiple computers working in parallel. A software management system splits the model into separate computational tasks

and distributes them among the computers. In the last five years or so, as multicore chips have become more common, researchers have simply transferred the old management systems over to them. But John Williams, professor of information engineering in the Department of Civil and Environmental Engineering (CEE), CEE postdoc David Holmes, and Peter Tilke, a visiting scientist in the Department of Earth, Atmospheric and Planetary Sciences, have developed a new management system that exploits the idiosyncrasies of multicore chips to improve performance.

To get a sense of what it might mean to split a model into separate tasks, consider a two-dimensional simulation of a weather system over some geographical area — like the animated weather maps on the nightly news. The simulation considers factors like temperature, humidity and wind speed, as measured at different weather stations, and tries to calculate how they will have changed a few minutes later. Then it takes the updated factors and performs the same set of calculations again, gradually projecting its model out across hours and days.

Changes to the factors in a given area depend on the factors measured nearby, but not on the factors measured far away. So the computational problem can, in fact, be split up according to geographic proximity, with the weather in different areas being assigned to different computers — or cores. The same holds true for simulations of many other physical phenomena.

Smaller is better

When such simulations run on a cluster of computers, the cluster's management system tries to minimize the communication between computers, which is much slower than communication within a given computer. To do this, it splits the model into the largest chunks it can — in the case of the weather simulation, the largest geographical regions —

so that it has to send them to the individual computers only once. That, however, requires it to guess in advance how long each chunk will take to execute. If it guesses wrong, the entire cluster has to wait for the slowest machine to finish its computation before moving on to the next part of the simulation.

In a multicore chip, however, communication between cores, and between cores and memory, is much more efficient. So the MIT researchers' system can break a simulation into much smaller chunks, which it loads into a queue. When a core finishes a calculation, it simply receives the next chunk in the queue. That also saves the system from having to estimate how long each chunk will take to execute. If one chunk takes an unexpectedly long time, it doesn't matter: The other cores can keep working their way through the queue.

Perhaps more important, smaller chunks means that the system is better able to handle the problem of boundaries. To return to the example of the weather simulation, factors measured along the edges of a chunk will affect factors in the adjacent chunks. In a cluster of computers, that means that computers working on adjacent chunks still have to use their low-bandwidth connections to communicate with each other about what's happening at the boundaries.

Multicore chips, however, have a memory bank called a cache, which is relatively small but can be accessed very efficiently. The MIT researchers' management system can split a simulation into chunks that are so small that not only do they themselves fit in the cache, but so does information about the adjacent chunks. So a core working on one chunk can rapidly update factors along the boundaries of adjacent chunks.

E pluribus unum

In theory, a single machine with 24 separate cores should be able to

perform a simulation 24 times as rapidly as a machine with only one core. In the February issue of [Computer Physics Communications](#), the MIT researchers report that, in their experiments, a 24-core machine using the existing management system was 14 times as fast as a single-core machine; but with their new [management system](#), the same machine was about 22 times as fast. And, Williams says, the new system's performance advantage compounds with the number of cores, "like compound interest over time."

Geoffrey Fox, professor of informatics at Indiana University, says that the MIT researchers' system is "clever and elegant," but he has doubts about its broad usefulness. The problems of greatest interest to many scientists and engineers, he says, are so large that they will still require clusters of computers, where the MIT researchers' system offers scant advantages. "State-of-the-art problems will not run on single machines," Fox says.

But Holmes points out that the model that he and his colleagues used in their experiments was a simulation of fluid flow through an oilfield, which is of immediate interest to the oilfield services company Schlumberger, which helped fund the research and employs Tilke when he's not on loan to MIT. "We're running problems with 50, 60 million particles," Holmes says, "which is on the order of 20, 30 gigabytes." Holmes also points out that 24-core computers "will not remain the state of the art for long." Manufacturers have already announced lines of 128-core computers, and that could just be the tip of the iceberg.

Williams adds that, even for problems that still require clusters of computers, the new system would allow the individual machines within the clusters to operate more efficiently. "Cross-machine communication is one or two orders of magnitude slower than on-machine communication," Williams says, "so it makes sense to keep cross-machine communication to a minimum, which is what our solution

allows.”

More information: Project website: geonumerics.mit.edu/

Provided by Massachusetts Institute of Technology

Citation: Mastering multicore (2010, April 26) retrieved 23 April 2024 from <https://phys.org/news/2010-04-mastering-multicore.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.