

Supercomputers take a cue from microwave ovens

December 1 2011

As sophisticated as modern climate models are, one critical component continues to elude their precision—clouds. Simulating these fluffy puffs of water vapor is so computationally complex that even today's most powerful supercomputers, working at quadrillions of calculations per second, cannot accurately model them.

"Clouds modulate the [climate](#). They reflect some sunlight back into space, which cools the Earth; but they can also act as a blanket and trap heat," says Michael Wehner, a climate scientist at the Lawrence Berkeley National Laboratory (Berkeley Lab). "Getting their effect on the climate system correct is critical to increasing confidence in projections of future climate change."

In order to build the break-through supercomputers scientists like Wehner need, researchers are looking to the world of consumer electronics like microwave ovens, cameras and cellphones, where everything from chips to batteries to software is optimized to the device's application. This co-design approach brings scientists and computer engineers into the [supercomputer](#) design process, so that systems are purpose-built for a scientific application, such as climate modeling, from the bottom up.

"Co-design allows us to design computers to answer specific questions, rather than limit our questions by available machines," says Wehner.

Co-design Test Case: Clouds

In a paper entitled "Hardware/Software Co-design of Global Cloud System Resolving Models," recently published in *Advances in Modeling Earth Systems*, Shalf, Wehner and coauthors argue that the scientific supercomputing community should take a cue from consumer electronics like smart phones and microwave ovens: Start with an application—like a climate model—and use that as a metric for successful hardware and software design.

The paper which uses the climate community's global cloud resolving models (GCRMs) as a case-study argues that an aggressive co-design approach to scientific computing could increase code efficiency and enable chip designers to optimize the trade-offs between energy efficiency, cost and application performance.

According to coauthor David Donofrio, a co-designed system for modeling climate would contain about 20 million cores (today's most powerful scientific cluster, Japan's 'K Computer' contains about 705,000 cores) and be capable of modeling climate 1,000 times faster than what is currently possible.

"Most importantly, the system would remain fully programmable so that scientific codes with similar hardware needs to the GCRMs, like seismic exploration, could also benefit from this machine," says Donofrio, a computer scientist at Berkeley Lab.

"Today when we purchase a general purpose supercomputer, it comes with a lot of operating system functions that science applications don't need. When you are worried about power, these codes can be very costly," says Shalf. "Instead of repurposing a chip designed for another market, the scientific HPC (high performance computing) community should specify what they want on a chip—the intellectual property

(IP)—and only buy that."

According to Shalf, a co-designed system for modeling climate would use about one quarter to one tenth the energy required for a conventional supercomputer with the same capabilities.

Consumers Pave the Way for Next Generation Supercomputers

Although innovative for scientific supercomputing, the idea of application-driven design is not new. Electronics like cell-phones and toaster ovens are built of simpler embedded processor cores optimized for one or a few dedicated functions.

"Because the ultimate goal of the embedded market is to maximize battery life, these technologies have always been driven by maximizing performance-per-watt and minimizing cost. Application-driven design is the key to accomplishing this," says Shalf. "Today we look at the motherboard as a canvas for building a supercomputer, but in the embedded market the canvas is the chip."

He notes that the most expensive part of developing a computer chip is designing and validating all of the IP blocks that are placed on the chip. These IP blocks serve different functions, and in the embedded market vendors profit by licensing them out to various product makers. With an application in mind, manufacturers purchase IP block licenses and then work with a system integrator to assemble the different pieces on a chip.

"You can think of these IP blocks as Legos or components of a home entertainment system," says Donofrio. "Each block has a purpose, you can buy them separately, and connect them to achieve a desired result, like surround sound in your living room."

"The expensive part is designing and verifying the IP blocks, and not the cost of the chip. These IP blocks are commodities because the development costs are amortized across the many different licenses for different applications," says Shalf. "Just as the consumer electronics chip designers choose a set of processor characteristics appropriate to the device at hand, HPC designers should also be able to chose processor characteristics appropriate to a specific application or set of applications, like the climate community's global cloud resolving model."

He notes that the resulting machine, while remaining fully programmable, would achieve maximum performance on the targeted set of applications, which were used as the benchmarks in the co-design process. In this sense, Shalf notes that the co-designed machine is less general purpose than the typical supercomputer of today, but much of what is included in modern supercomputers is of little use to scientific computing anyway and so it just wastes power.

"Before this work, if someone asked me when the climate community would be able to compute kilometer scale climate simulations, I would have answered 'not in my lifetime,'" says Wehner. "Now, with this roadmap I think we could be resolving cloud systems within the next decade."

Although climate was the focus of this paper, Shalf notes that future co-design studies will explore whether this will also be cost-effective for other compute intensive sciences such as combustion research.

Provided by DOE/Lawrence Berkeley National Laboratory

Citation: Supercomputers take a cue from microwave ovens (2011, December 1) retrieved 26 April 2024 from <https://phys.org/news/2011-12-supercomputers-cue-microwave-ovens.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.