

Berkeley Lab researchers propose a new breed of supercomputers

Three researchers from the U.S. Department of Energy's Lawrence Berkeley National Laboratory (Berkeley Lab) have proposed an innovative way to improve global climate change predictions by using a supercomputer with low-power embedded microprocessors, an approach that would overcome limitations posed by today's conventional supercomputers.

In a paper published in the May issue of the *International Journal of High Performance Computing Applications*, Michael Wehner and Lenny Oliker of Berkeley Lab's Computational Research Division, and John Shalf of the National Energy Research Scientific Computing Center (NERSC) lay out the benefit of a new class of supercomputers for modeling climate conditions and understanding climate change. Using the embedded microprocessor technology used in cell phones, iPods, toaster ovens and most other modern day electronic conveniences, they propose designing a cost-effective machine for running these models and improving climate predictions.

In April, Berkeley Lab signed a collaboration agreement with Tensilica®, Inc. to explore such new design concepts for energy-efficient high-performance scientific computer systems. The joint effort is focused on novel processor and systems architectures using large numbers of small processor cores, connected together with optimized links, and tuned to the requirements of highly-parallel applications such as climate modeling.

Understanding how human activity is changing global climate is one of the great scientific challenges of our time. Scientists have tackled this issue by developing climate models that use the historical data of factors that shape the earth's climate, such as rainfall, hurricanes, sea surface temperatures and carbon dioxide in the atmosphere. One of the greatest challenges in creating these models, however, is to develop accurate cloud simulations.

Although cloud systems have been included in climate models in the past, they lack the details that could improve the accuracy of climate predictions. Wehner, Oliker and Shalf set out to establish a practical estimate for building a supercomputer capable of creating climate models at 1-kilometer (km) scale. A cloud system model at the 1-km scale would provide rich details that are not available from existing models.

To develop a 1-km cloud model, scientists would need a supercomputer that is 1,000 times more powerful than what is available today, the researchers say. But building a supercomputer powerful enough to tackle this problem is a huge challenge.

Historically, supercomputer makers build larger and more powerful systems by increasing the number of conventional microprocessors — usually the same kinds of microprocessors used to build personal computers. Although feasible for building computers large enough to solve many scientific problems, using this approach to build a system capable of modeling clouds at a 1-km scale would cost about \$1 billion. The system also would require 200 megawatts of electricity to operate, enough energy to power a small city of 100,000 residents.

In their paper, "Towards Ultra-High Resolution models of Climate and Weather," the researchers present a radical alternative that would cost less to build and require less electricity to operate. They conclude that a supercomputer using about 20 million embedded microprocessors would deliver the results and cost \$75

million to construct. This “climate computer” would consume less than 4 megawatts of power and achieve a peak performance of 200 petaflops.

“Without such a paradigm shift, power will ultimately limit the scale and performance of future supercomputing systems, and therefore fail to meet the demanding computational needs of important scientific challenges like the climate modeling,” Shalf said.

The researchers arrive at their findings by extrapolating performance data from the Community Atmospheric Model (CAM). CAM, developed at the National Center for Atmospheric Research in Boulder, Colorado, is a series of global atmosphere models commonly used by weather and climate researchers.

The “climate computer” is not merely a concept. Wehner, Olikier and Shalf, along with researchers from UC Berkeley, are working with scientists from Colorado State University to build a prototype system in order to run a new global atmospheric model developed at Colorado State.

“What we have demonstrated is that in the exascale computing regime, it makes more sense to target machine design for specific applications,” Wehner said. “It will be impractical from a cost and power perspective to build general-purpose machines like today’s supercomputers.”

Under the agreement with Tensilica, the team will use Tensilica’s Xtensa LX extensible processor cores as the basic building blocks in a massively parallel system design. Each processor will dissipate a few hundred milliwatts of power, yet deliver billions of floating point operations per second and be programmable using standard programming languages and tools. This equates to an order-of-magnitude improvement in floating point operations per watt, compared to conventional desktop and server processor chips. The small size and low power of these processors allows tight integration at the chip, board and rack level and scaling to millions of processors within a power budget of a few megawatts.

Source: Lawrence Berkeley National Laboratory

This document is subject to copyright. Apart from any fair dealing for the purpose of private study, research, no part may be reproduced without the written permission. The content is provided for information purposes only.