

1 million trillion 'flops' per second targeted by new Institute for Advanced Architectures

Preparing groundwork for an exascale computer is the mission of the new Institute for Advanced Architectures, launched jointly at Sandia and Oak Ridge national laboratories.

An exaflop is a thousand times faster than a petaflop, itself a thousand times faster than a teraflop. Teraflop computers —the first was developed 10 years ago at Sandia — currently are the state of the art. They do trillions of calculations a second. Exaflop computers would perform a million trillion calculations per second.

The idea behind the institute —under consideration for a year and a half prior to its opening — is “to close critical gaps between theoretical peak performance and actual performance on current supercomputers,” says Sandia project lead Sudip Dosanjh. “We believe this can be done by developing novel and innovative computer architectures.”

Ultrafast supercomputers improve detection of real-world conditions by helping researchers more closely examine the interactions of larger numbers of particles over time periods divided into smaller segments.

“An exascale computer is essential to perform more accurate simulations that, in turn, support solutions for emerging science and engineering challenges in national defense, energy assurance, advanced materials, climate, and medicine,” says James Peery, director of computation, computers and math.

The institute is funded in FY08 by congressional mandate at \$7.4 million. It is supported by the National Nuclear Security Administration and the Department of Energy’s Office of Science. Sandia is an NNSA laboratory.

One aim, Dosanjh says, is to reduce or eliminate the growing mismatch between data movement and processing speeds.

Processing speed refers to the rapidity with which a processor can manipulate data to solve its part of a larger problem. Data movement refers to the act of getting data from a computer’s memory to its processing chip and then back again. The larger the machine, the farther away from a processor the data may be stored and the slower the movement of data.

“In an exascale computer, data might be tens of thousands of processors away from the processor that wants it,” says Sandia computer architect Doug Doerfler. “But until that processor gets its data, it has nothing useful to do. One key to scalability is to make sure all processors have something to work on at all times.”

Compounding the problem is new technology that has enabled designers to split a processor into first two, then four, and now eight cores on a single die. Some special-purpose processors have 24 or more cores on a die. Dosanjh suggests there might eventually be hundreds operating in parallel on a single chip.

“In order to continue to make progress in running scientific applications at these [very large] scales,” says Jeff Nichols, who heads the Oak Ridge branch of the institute, “we need to address our ability to maintain the balance between the hardware and the software. There are huge software and programming challenges and our goal is to do the critical R&D to close some of the gaps.”

Operating in parallel means that each core can work its part of the puzzle simultaneously with other cores on a chip, greatly increasing the speed a processor operates on data. The method does not require faster clock speeds, measured in faster gigahertz, which would generate unmanageable amounts of heat to

dissipate as well as current leakage.

The new method bolsters the continued relevance of Moore's Law, the 1965 observation of Intel cofounder Gordon Moore that the number of transistors placed on a single computer chip will double approximately every two years.

Another problem for the institute is to reduce the amount of power needed to run a future exascale computer.

"The electrical power needed with today's technologies would be many tens of megawatts — a significant fraction of a power plant. A megawatt can cost as much as a million dollars a year," says Dosanjh. "We want to bring that down."

Sandia and Oak Ridge will work together on these and other problems, he says. "Although all of our efforts will be collaborative, in some areas Sandia will take the lead and Oak Ridge may lead in others, depending on who has the most expertise in a given discipline." In addition, a key component of the institute will be the involvement of industry and universities.

A spontaneous demonstration of wide interest in faster computing was evidenced in the response to an invitation-only workshop, "Memory Opportunities for High-Performing Computing," sponsored in January by the institute.

Workshop organizers planned for 25 participants but nearly 50 attended. Attendees represented the national labs, DOE, National Science Foundation, National Security Agency, Defense Advanced Research Projects Agency, and leading manufacturers of processors and supercomputing systems.

Ten years ago, people worldwide were astounded at the emergence of a teraflop supercomputer — that would be Sandia's ASCI Red — able in one second to perform a trillion mathematical operations.

More recently, bloggers seem stunned that a machine capable of petaflop computing — a thousand times faster than a teraflop — could soon break the next barrier of a thousand trillion mathematical operations a second.

Source: Sandia National Laboratories

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.