

Stanford, tech giants team up to enable software for parallel computers

Stanford and many of the biggest companies in computing will announce Friday, May 2, a joint effort to address a major missed opportunity in information technology: the dearth of software that can harness the parallelism of the multiple processors that are being built into virtually every new computer. The Pervasive Parallelism Lab (PPL) pools the efforts of many leading Stanford computer scientists and electrical engineers with support from Sun Microsystems, Advanced Micro Devices, NVIDIA, IBM, Hewlett Packard and Intel.

Until recently, computers with multiple processors were too expensive for all but specialized uses (e.g. supercomputing) where the high performance of parallel processing was deemed essential. As a consequence, few programmers have learned how to design software that exploits parallelism. The problem has caused serious concern among computer scientists that the progress of computing overall could stall.

"Parallel programming is perhaps the largest problem in computer science today and is the major obstacle to the continued scaling of computing performance that has fueled the computing industry, and several related industries, for the last 40 years," says Bill Dally, chair of the Computer Science Department at Stanford.

Dally will participate in the lab's research, which will be directed by Kunle Olukotun, a professor of electrical engineering and computer science who has worked for more than a decade on multicore computer architecture, in which many processors inhabit the same silicon chip.

Another team member and fellow processor architect, Stanford President John Hennessy, said the lab's work is aimed at restoring the progress that society at large has come to expect from information technology.

"This era is going to be about exploiting some sort of explicit parallelism, and if there's a problem that has confounded computer science for a long time, it is exactly that," Hennessy said. "We need to find ways to overcome this problem so that recent advances in computing hardware can continue benefiting the public and the economy."

Olukotun says he hopes that by working directly with industrial supporters, the work of PPL faculty and students will reach the marketplace where it can have an impact. He emphasized that the lab is open, meaning that other companies can still join the effort and none has exclusive intellectual property rights.

The center, with a budget of \$6 million over three years, will research and develop a top-to-bottom parallel computing system, stretching from fundamental hardware to new user-friendly programming languages that will allow developers to exploit parallelism automatically. In other words, game programmers who already understand artificial intelligence, graphics rendering, and physics would be able to implement their algorithms in accessible "domain-specific" languages. At deeper, more fundamental levels of software—"under the hood," so to speak—the system would do all the work for them to optimize their code for parallel processing.

To enable the research, the team's hardware experts will develop a novel testbed called FARM, for Flexible Architecture Research Machine. The system, which Olukotun said will be ready by the end of the summer, will combine versatility with performance by blending reprogrammable chips with conventional processors.

Olukotun says he hopes the effort will pave the way for programmers to easily create powerful new software for applications such as Artificial Intelligence and robotics, business data analysis, and virtual worlds and gaming. Among the PPL faculty are experts in each of these areas, including Pat Hanrahan, a professor of computer science and electrical engineering whose graphics rendering expertise has earned him two Academy Awards.

"We believe in driving applications," says Hanrahan.. "Among the most interesting are immersive, richly graphical, virtual worlds, both because of the unique experiences for users as well as the challenges in building such demanding parallel applications."

Research in the PPL also will be able to make use of parallelism technologies that Stanford has already developed, as part of years of research on the subject. These include not only Olukotun's work on multicore chips but also his collaboration with computer science and electrical engineering Assistant Professor Christos Kozyrakis to develop a more efficient way for processors to share memory, called "transactional memory." Dally, meanwhile, has developed new ways for the flow, or "streaming," of software instructions from a compiler to parallel processors to work much more efficiently than in conventional supercomputers.

"We have a history here of trying to close this gap between parallel hardware and software," Olukotun says. "It's not enough just to put a bunch of cores on a chip. You also have to make the job of translating software to use that parallelism easier."

Stanford, however, is not the only university trying to solve the problem. The announcement of the PPL comes less than two months after the University of California at Berkeley and the University of Illinois at Urbana-Champaign each received multimillion-dollar grants from Microsoft and Intel to address the issue.

"Clearly this problem is big enough and important enough that we need more people looking at it," Olukotun said. "By having more and different approaches to the problem, we're more likely to find a solution."

Source: Stanford University

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.