High-Performance Parallel Computing

.....

Overview

With the saturation of core clock frequency, it is imperative to go parallel. However, parallelizing a program and optimizing it for concurrent processing is challenging. One needs to worry about locality, memory latency, access patterns, and data representation , all at the same time. One needs methodical ways to extract high performance.

The course is structured to provide a participant with in-depth knowledge and techniques to optimize parallel programs. In particular, the course uses OpenMP for shared memory systems and CUDA / OpenCL for graphics processing units. It would emphasize the and enable participants to appreciate the importance of data locality .

Course participants will learn these topics through lectures and hands-on experiments.

More details about the course are available at http://www.cse.iitm.ac.in/~rupesh/teaching/hpc/jun16

Dates for the Course	**20th June to 1st July 2016**
Host Institute	**IIT Madras**
No. of Credits	**2**
Maximum No. of Participants	**100**
You Should Attend If…	▪ you are an engineer working in high-performance computing. ▪ you are a student or a faculty member from academic institution interested in learning how to optimize programs for performance.
Course Registration Fees	The participation fees for taking the course is as follows: **Student Participants:** Rs.2000 **Faculty Participants:** Rs.5000 **Government Research Organization Participants**: Rs.5000 **Industry Participants:** Rs.10000 The above fee is towards participation in the course, the course material, computer use for tutorials and assignments, and laboratory equipment usage charges. The participants may be provided with hostel accommodation, depending on the availability, on a separate payment basis.

Course Faculty

Prof. Sadayappan is a Professor in the Department of Computer Science and Engineering at The Ohio State University. His primary research interests center around performance optimization and compiler/runtime systems for high-performance computing, with special emphasis on high-performance frameworks that enable high productivity for application developers in scientific computing. Two recent projects include a polyhedral framework for automatic parallelization and data locality optimization, and the Tensor Contraction Engine – a domain-specific compiler/runtime system to automatically transform high-level specifications into efficient parallel programs, for a class of high accuracy ab initio models in *quantum chemistry*.

Rupesh Nasre is a faculty member at the Department of Computer Science and Engineering at IIT Madras. His primary research interests are in compilers and parallelization. In particular, he has been exploring parallelization of graph algorithms on multi-core as well as GPUs. Towards improving productivity of end-users, he has developed a domain-specific language for graph algorithms , in collaboration with other researchers. He is a recipient of NVIDIA Special Prize, has won Yahoo! HackU Award, and holds five US patents.