

**ISTeC**



**Colorado  
State  
University**

**The Information Science & Technology Center**

*ISTeC.ColoState.edu*

**Colorado State University's  
Information Science and Technology Center (ISTeC)  
presents two lectures by**



**Dr. Vivek Sarkar**

E.D. Butcher Professor of  
Computer Science  
Rice University

## **ISTeC Distinguished Lecture**

**in conjunction with the  
Electrical and Computer Engineering Department and  
Computer Science Department Seminar Series**

### **“Multicore Programming Models and their Implementation Challenges”**

**Monday, December 8, 2008**

Reception: 10:30 a.m.

Lecture: 11:00 – 12:00 noon

Location: CSU Lory Student Center 203



## **Computer Science Department Lecture**

*sponsored by ISTE C*

### **“A Quick Tour of Modern Multicore Programming Models”**

**Tuesday, December 9, 2008**

Lecture: 9:30 – 10:30 a.m.

Location: Engineering E202

# ABSTRACTS

## “Multicore Programming Models and their Implementation Challenges”

The computer industry is at a major inflection point in its hardware roadmap due to the end of a decades-long trend of exponentially increasing clock frequencies. It is widely agreed that spatial parallelism in the form of multiple power-efficient cores must be exploited to compensate for this lack of frequency scaling. Unlike previous generations of hardware evolution, this shift towards multicore and manycore computing will have a profound impact on software. These software challenges are further compounded by the need to enable parallelism in workloads and application domains that have traditionally not had to worry about multiprocessor parallelism in the past.

In this talk, we will focus on the programming problem for tightly coupled homogeneous and heterogeneous multicore processors. We present early experiences with the new Habanero Multicore Software Research project at Rice University (<http://habanero.rice.edu>) that encompasses work on programming models, compilers, runtimes, and concurrency libraries so as to enable portable software that can run unchanged on a range of homogeneous and heterogeneous multicore systems. The Habanero project takes a two-level approach to programming models, with a high-level model based on Intel Concurrent Collections for parallelism-oblivious domain experts, and a lower-level model based on the high productivity X10 language for parallelism-aware developers. We discuss compiler and runtime implementation challenges that must be overcome to enable mainstream applications to use these models on multicore systems.

## “A Quick Tour of Modern Multicore Programming Models”

This lecture will introduce students to modern multicore programming models by surveying representative examples for homogeneous multicore systems --- Cilk, OpenMP, Java Concurrency Utilities, Intel Thread Building Blocks, .Net Task Parallel Library & PLINQ --- while using the Habanero/X10 execution model as a reference point. It will conclude with a summary of new directions in programming models for heterogeneous multicore systems (CUDA for GPGPU's, and Cn for ClearSpeed), and of research activities under way in the Rice Habanero Multicore Software project ([habanero.rice.edu](http://habanero.rice.edu)).

## SPEAKER BIOGRAPHY

Dr. Vivek Sarkar ([http://www.cs.rice.edu/~vs3/home/Vivek\\_Sarkar.html](http://www.cs.rice.edu/~vs3/home/Vivek_Sarkar.html)) conducts research in programming languages, program analysis, compiler optimizations and virtual machines for parallel and high performance computer systems, and currently leads the Habanero Multicore Software Research project at Rice University ([habanero.rice.edu](http://habanero.rice.edu)). Prior to joining Rice, he was Senior Manager of Programming Technologies at IBM Research. His responsibilities at IBM included leading IBM's research efforts in programming model, tools, and productivity in the PERCS project during 2002 - 2007 as part of the DARPA High Productivity Computing System program. His past projects include the X10 programming language, the Jikes Research Virtual Machine for the Java language, the ASTI optimizer used in IBM's XL Fortran product compilers, the PTRAN automatic parallelization system, and profile-directed partitioning and scheduling of Sisal programs. Vivek became a member of the IBM Academy of Technology in 1995, an ACM Distinguished Scientist in 2006, and the E.D. Butcher Professor of Computer Science at Rice University in 2007. He holds a B.Tech. degree from the Indian Institute of Technology, Kanpur, an M.S. degree from University of Wisconsin-Madison, and a Ph.D. from Stanford University. In 1997, he was on sabbatical as a visiting associate professor at MIT, where he was a founding member of the MIT RAW multicore project.

**To arrange a meeting with the speaker**, please contact Dr. Michelle Strout at (970) 491-7026 [mstrout@cs.colostate.edu](mailto:mstrout@cs.colostate.edu).

**ISTeC (Information Science and Technology Center)** is a university-wide organization for promoting, facilitating, and enhancing CSU's research, education, and outreach activities pertaining to the design and innovative application of computer, communication, and information systems. For more information please see [ISTeC.ColoState.edu](http://ISTeC.ColoState.edu).