

## Keynote Talk:

### Compiler and Runtime Challenges for Memory Centric Programming

Vivek Sarkar (Georgia Institute of Technology)

#### Abstract:

It is widely recognized that a major disruption is under way in computer hardware as processors strive to extend, and go beyond, the end-game of Moore's Law. This disruption will include new forms of processor and memory hierarchies, including near-memory computation structures. In this talk, we summarize compiler and runtime challenges for memory centric programming, based on past experiences with the X10 project at IBM and the Habanero project at Rice University and Georgia Tech. A key insight in addressing compiler challenges is to expand the state-of-the-art in analyzing and transforming explicitly-parallel programs, so as to encourage programmers to write forward-scalable layout-independent code rather than hardwiring their programs to specific hardware platforms and specific data layouts. A key insight in addressing runtime challenges is to focus on asynchrony in both computation and data movement, while supporting both in a unified and integrated manner. A cross-cutting opportunity across compilers and runtimes is to broaden the class of computation and data mappings that can be considered for future systems. Based on these and other insights, we will discuss recent trends in compilers and runtime systems that point the way towards possible directions for addressing the challenges of memory centric programming.

#### Vivek Sarkar (Georgia Institute of Technology):

Vivek Sarkar is a Professor in the School of Computer Science, and the Stephen Fleming Chair for Telecommunications in the College of Computing at Georgia Institute of Technology, since August 2017. Prior to joining Georgia Tech, Sarkar was a Professor of Computer Science at Rice University, and the E.D. Butcher Chair in Engineering. During 2007 - 2017, Sarkar built Rice's Habanero Extreme Scale Software Research Group with the goal of unifying parallelism and concurrency elements of high-end computing, multicore, and embedded software stacks (<http://habanero.rice.edu>). He also served as Chair of the Department of Computer Science at Rice during 2013 - 2016.

Prior to joining Rice in 2007, Sarkar was Senior Manager of Programming Technologies at IBM Research. His research projects at IBM included the X10 programming language, the Jikes Research Virtual Machine for the Java language, the ASTI optimizer used in IBM's XL Fortran product compilers, and the PTRAN automatic parallelization system. Sarkar became a member of the IBM Academy of Technology in 1995, and was inducted as an ACM Fellow in 2008. He has been serving as a member of the US Department of Energy's Advanced Scientific Computing Advisory Committee (ASCAC) since 2009, and on CRA's Board of Directors since 2015.