

**Events**[Current Events](#)[Lectures](#)[Events Archive](#)

Other

Speeding Up Nek5000 with Autotuning and Specialization**Jaewook Shin, Hewlett Packard**

Core Expertise Faculty Candidate

Johnston Hall 338
April 20, 2012 - 10:30 am**Abstract:**

Autotuning technology has emerged recently as a systematic process for evaluating alternative implementations of a computation, in order to select the best-performing solution for a particular architecture. Specialization optimizes code customized to a particular class of input data set. In this talk, I will talk about how compiler-based autotuning that incorporates specialization for expected data set sizes of key computations can be used to speed up Nek5000, a spectral-element code. Nek5000 makes heavy use of what are effectively Basic Linear Algebra Subroutine (BLAS) calls, but for very small matrices. Through autotuning and specialization, we can achieve significant performance gains over hand-tuned libraries (e.g., Goto, ATLAS, and ACML BLAS). We demonstrate more than 2.2X performance gains on an Opteron over the original manually tuned implementation, and speedups of up to 1.26X on the entire application running on 256 nodes of the Cray XT5 Jaguar system at Oak Ridge.

Speaker's Bio:

Dr. Jaewook Shin is a performance engineer at Hewlett Packard. He earned his Ph.D. degree from the University of Southern California in 2005 with his work on compiler optimizations for multimedia extension architectures. He had been a postdoctoral researcher and Enrico Fermi scholar in the Mathematics and Computer Science division at Argonne National Laboratory until he joined HP in 2010.

