Institute for Advanced Simulation
Jülich Supercomputing Centre

IAS Seminar

Topic: **Perplexed by today's programming models for the Xeon Phi? Beware, it may get worse soon!**

Speaker: Dr. Lars Koesterke, Texas Advanced Computing Center (TACC), Austin, TX, USA

Contents: I will start my presentation with a very brief overview of the Stampede project (9.5 PFlops, 6400+ Xeon Phi KNC cards) at the Texas Advanced Computing Center (TACC) and will share our experiences with Xeon Phi and GPU accelerators. I will then report on the status of heterogeneous computing today and in the near future. Many in the community, experts and average users alike, have concerns regarding the current direction, the speed by which features are implemented, and generally the software environment for Xeon Phi accelerators.

In particular, I will discuss an ongoing research project in which I try to ascertain how self-hosted MICs may be programmed within the 'offload' software paradigm. In about a year, we will have self-hosted KNL nodes available, but to this day it remains uncertain if and how Intel will support 'offloading'. In my project, I use plain/standard MPI calls to enable 'offloading' across an interconnect. My approach is in many ways similar to the approach taken by the DEEP project, but everything is programmed by the user without any additional support from an enhanced compiler and/or software stack. My main interest is in gauging the burden that is put on the programmer when using today's standard for distributed computing (MPI) for the imminent next generation of Xeon Phi's. After all, software support for today's hardware is already sketchy which means that we may have to make due with homemade workarounds for a good while.

Time: Friday, 13 June 2014, 11:00

Venue: Jülich Supercomputing Centre, Hörsaal, building 16.3, room 006

Anyone interested is cordially invited to participate in this seminar.

sgd Prof. Dr. Dr. Thomas Lippert