

Toward Improved Performance Solutions (TIPS): Productive Programming on Manycores Architectures

Instructors: Bilel Hadri, Saber Feki, KAUST Supercomputing Laboratory

Prerequisites: Programming knowledge in C and/or Fortran, Basics in parallel computing

Content level of the tutorial: 65% beginner, 35% intermediate

1. Abstract:

The HPC community is now using powerful supercomputing systems composed of heterogeneous nodes built from multi-core processors and accelerator like NVIDIA GPUs and Intel Xeon Phi, providing an aggregate node performance of more than one TeraFlop/s. This tremendous computational power can only be efficiently utilized with the appropriate programming models and software infrastructure. Indeed, it requires end-users to rethink and redesign their algorithms to take advantage of the underlying architecture while keeping productivity in mind.

This proposed tutorial is designed for the HPC community interested on the best practices to speed up efficiently their parallel codes. A comprehensive overview on GPU architecture and the Intel Xeon Phi coprocessor, also known as Intel Many Integrated Core (MIC) will be presented. Technical sessions will include interactive examples using high-level programming models such as OpenACC and OpenMP. With a focus on high productivity in programming manycores architectures, examples of current supported software packages such as numerical libraries (e.g. cuBLAS, MAGMA), and a broad range of applications (e.g., computational electromagnetics, computational chemistry, CFD, seismic imaging) will be also highlighted. Participants are encouraged to bring laptop computers and follow live demonstrations with detailed examples.

2. Justification

Most of the materials of this tutorial have been presented in the past during various venues, such as workshops held at the Oak Ridge National Lab (March and October 2011), University of Tennessee (March 2012) and KAUST (September and October 2013); and conferences like XSEDE in July 2012 and 2013. By presenting the best practices on using the accelerators, this tutorial will bring together researchers from architecture, compilers, applications and languages at PACT 2014. We expect to gather at least 20 participants at this tutorial.

3. Detailed outline of the half-day tutorial

08:00 General Introduction and Motivation on Accelerators

08:15 Overview on GPU architecture and the OpenACC programming model

09:00 Overview on the Intel Many Integrated Core architecture and OpenMP 4.0 programming model

09:45 Break

10:00 Numerical Libraries on accelerators

10:45 Examples of porting scientific applications on GPUs and Intel Xeon Phi

12:00 Adjourn