

———— Master Project ————

Code Optimization for Intel's Xeon Phi Platform

Description

At CERN, particles are collided in order to understand how the universe was created. In the LHCb experiment, proton-proton collisions are used to investigate the matter-antimatter asymmetry of the universe. In 2018, the current readout LHCb system, including the trigger system, will be upgraded. After the upgrade, the LHCb experiment strives for a hardware-trigger-free readout using an event filter farm, which will process an event every 25 ns. A non-optimized implementation of the pattern recognition algorithm for the current CPU hardware exists, written in C++. For this master thesis the algorithm shall be ported to and optimized for the Intel's Xeon Phi Knight's Landing (KNL) processor.

The KNL, an Intel CPU with 64 cores designed for HPC (High Performance Computing), allows the user to run an operating system or any x86 compatible code with up to four hyperthreads per core. Exploiting the features of the Intel Xeon Phi KNL platform, the practical speed-up possible for the RICH pattern algorithm should be analyzed. The RICH algorithm uses an extensive amount of FPU operations, and memory accesses and data representation have to be researched in detail to determine the best optimization techniques to be applied. The hotspots in the code need to be identified and optimized first, and parallelization and vectorization has to be applied in order to make use of all the resources

provided by the Knights Landing processor.

Keywords: SIMD, vectorization, Xeon Phi, Intel, HPC

Required Skills

- knowledge in computer architecture
- programming skills in C/C++

Desired Skills

- understanding of HPC programming and requirements

Contact Person

Angela Pohl
(angela.pohl@tu-berlin.de)