

Conjugate Gradient solvers on Intel Xeon Phi and NVIDIA GPUs

Wednesday, 10 September 2014 14:30 (30 minutes)

The runtime of a Lattice QCD simulation is dominated by a small kernel, which calculates the product of a vector by a sparse matrix known as the “Dslash” operator. Therefore, this kernel is frequently optimized for various HPC architectures. In this contribution we want evaluate the performance of the Intel Xeon Phi to current Kepler-based NVIDIA Tesla GPUs running a conjugate gradient solver. By exposing more parallelism to the accelerator by inverting multiple vectors at the same time we obtain a performance >250 GFLOPs/s on both architectures. This more than doubles the performance of the naive separate inversion. A detailed comparison of the performance of the accelerators for different scenarios will be presented in the talk. We also discuss some details of the implementation and the effort required to obtain the achieved performance.

Primary authors: Dr SCHMIDT, Christian (Fakultät für Physik, Universität Bielefeld); Dr WAGNER, Mathias (Department of Physics, Indiana University); SWAGATO, Mukherjee (Brookhaven National Laboratory); Dr KACZMAREK, Olaf (Fakultät für Physik, Universität Bielefeld); Mr STEINBRECHER, Patrick (Fakultät für Physik, Universität Bielefeld)

Presenter: Mr STEINBRECHER, Patrick (Fakultät für Physik, Universität Bielefeld)

Session Classification: GPU in Lattice QCD