



Contribution ID: 397

Type: **Talk**

Adaptive Multigrid Solvers for LQCD on GPUs

Monday, 23 June 2014 16:50 (20 minutes)

There has been much progress to date in developing scalable sparse linear solver algorithms on GPUs, utilizing well-known mathematical methods such as mixed precision, domain decomposition and pipelining to improve performance, allowing efficient use of large GPU installations such as Blue Waters and Titan. However, there has been less focus on deploying ‘mathematically optimal’ linear solvers, that have optimal $O(N)$ complexity. In this work we utilize the QUDA framework to deploy adaptive multigrid solvers on GPUs, in particular we describe the architecture abstractions that allow for deployment on heterogeneous systems, utilizing both GPUs and CPUs. We discuss in general the suitability of heterogeneous architectures for hierarchical algorithms, and compare performance against a highly optimized CPU implementation.

Primary author: Dr CLARK, M (NVIDIA)

Co-authors: Dr CHENG, Michael (Boston University); Prof. BROWER, Richard (Boston University)

Presenter: Dr CLARK, M (NVIDIA)

Session Classification: Algorithms and Machines

Track Classification: Algorithms and Machines