



Contribution ID: 302

Type: **Talk**

Achieving strong scaling in many-GPU calculations in lattice QCD

Monday, 23 June 2014 17:30 (20 minutes)

We describe recent additions to the QUDA software library that are aimed at extending strong scaling in multi-GPU lattice calculations. These include the addition of CPU-thread support in order to increase concurrency and improve the overlap of computation and communication in Krylov solver routines, as well as the modifications needed to enable the GPUDirect RDMA feature recently introduced by NVIDIA and Mellanox. However, we focus in particular on the implementation and performance of so-called S-step variants of common Krylov solvers on current NVIDIA hardware. The S-step formulations are designed to reduce the number of global synchronizations associated with the calculation of vector inner products. These formulations may, when combined with communication-reducing methods such as additive Schwarz preconditioning, form the basis for a set of optimal Krylov solvers for many-GPU calculations.

Primary author: Dr FOLEY, Justin (Microway and Nvidia)

Co-author: Dr CLARK, Mike (Nvidia)

Presenter: Dr FOLEY, Justin (Microway and Nvidia)

Session Classification: Algorithms and Machines

Track Classification: Algorithms and Machines