



Contribution ID: 172

Type: **Poster**

QUDA 1.0

Tuesday 18 June 2019 18:30 (20 minutes)

Eleven years after its inception, the QUDA library for Lattice QCD on NVIDIA GPUs has achieved a historic 1.0 release. In this span QUDA has evolved into an open-source framework for full QCD simulations. The library has been fully re-written in a new highly optimized C++11 framework, superseding python-generated routines and ushering in a new age of rapid algorithm prototyping and development. QUDA supports nearly all fermion discretizations, features cutting-edge algorithms such as adaptive multigrid, deflation, and block Krylov-space methods, and contains native support for mixed precision and symmetry-inspired data compression. In preparation for the exascale era, there is full-featured support for communication-mitigating methods, intra-node peer-to-peer support, and inter-node GPUDirect MPI. On this poster we will explicate these features, reinforced by performance results on state-of-the-art hardware. We strongly encourage discussions about user requests and contributions.

Primary authors: CLARK, Kate (NVIDIA); WAGNER, Mathias (NVIDIA); WEINBERG, Evan (NVIDIA Corporation)

Presenter: CLARK, Kate (NVIDIA)

Session Classification: Poster

Track Classification: Algorithms and Machines