

Contribution ID: 3

Type: Talk

Autotuning multigrid parameters in the HMC on different architectures

Tuesday, 30 July 2024 16:15 (20 minutes)

Multigrid-preconditioned solvers have proven crucial for the efficient generation of ensembles of gauge configurations at physical quark mass parameters. A highly efficient implementation of such a solver for GPUs by different vendors and for different types of Wilson fermions is provided in the QUDA library. It includes functionality for updating and evolving the multigrid setup in the Hybrid Monte Carlo algorithm together with the gauge field. In the force calculation for the most poorly conditioned systems in simulations with Wilsonclover twisted mass fermions the solver outperforms mixed-precision CG by up to two orders of magnitude at the physical light quark mass, leading also to a large overall speedup of the HMC as a whole.

QUDA provides an autotuner which selects optimal launch parameters and communication policies for each kernel, problem size and domain decomposition, ensuring optimal performance of the underlying kernels. The multigrid solver, however, depends on a large number of choices such as block sizes, numbers of vectors, maximum iterations as well as thresholds and, in the case of twisted mass fermions, a scaling of the twisted quark mass on the coarse grids. As these parameters are generally defined on a per-level basis the search space is large, making exhaustive scans expensive. In addition, even if a good parameter set for a particular situation is found, in general it will fail to be optimal on a different machine or for a different domain decomposition.

We present an autotuner for these solver parameters implemented in tmLQCD which finds good parameter sets relatively quickly, requiring only some intution on the order in which parameters are to be tuned and on the step sizes to be used in the tuning procedure. By comparing the performance of the resulting setups on machines based on NVIDIA and AMD GPUs we further demonstrate its practical applicability.

Primary author: KOSTRZEWA, Bartosz (High Performance Computing & Analytics Lab, University of Bonn)

Co-authors: SEN, Aniket (HISKP, University of Bonn); GAROFALO, Marco (University of Bonn); ROMITI, Simone (University of Bern)

Presenter: KOSTRZEWA, Bartosz (High Performance Computing & Analytics Lab, University of Bonn)

Session Classification: Software development and machines

Track Classification: Software Development and Machines