

Contribution ID: 96

Type: Talk (15min + 5min)

## lattice QCD software development for heterogeneous supercomputers

Wednesday, March 6, 2024 3:30 PM (20 minutes)

What does it take to develop and maintain a research code for the stochastic simulation of the physics of the strong interaction in Lattice Quantum Chromodynamics (LQCD)? How to make it run on the fastest supercomputers in the world? How many people are involved and what do they contribute when and how? What kind of development and interaction structures are useful? Which kinds of challenges need to be overcome?

In this contribution we present the ongoing and organically cooperative effort between LQCD groups in Bonn and Cyprus, the development team of the QUDA LQCD library and staff at Juelich Supercomputing Center (JSC) in enabling the tmLQCD software framework, the workhorse of the Extended Twisted Mass Collaboration, to successfully run on supercomputers with accelerators by NVIDIA and AMD today as well as machines by these and other vendors in the future.

LQCD has historically been a trailblazer discipline in its adoption of new supercomputing architectures. Some groups have even actively contributed to the development of new machines such as the BlueGene line of supercomputers. At some sites, LQCD practitioners use very large fractions of the total available computing time, such that efficiently implementing the underlying algorithms has a significant impact not only on the possible science output but also on the associated energy consumption.

Many software frameworks for LQCD have been developed by small collaborations of "user-developers" who were able to efficiently target the relatively slowly changing architectures of the time with comparatively simple algorithms, mostly just making use of MPI for inter-process communicaton and some level of hardware specialisation via compiler intrinsincs or even inline assembly to target particular hardware architectures. The rapid proliferation of many-core and accelerated supercomputing systems by multiple vendors and the increased complexity of state-of-the-art algorithms have changed this irrevocably.

In order for LQCD codes to target current and future supercomputers, close interaction between hardware vendors, supercomputing centers as well as library and application developers is mandatory. In addition to addressing performance on individual architectures, questions about correctness testing, provenance tracking, performance-portability, maintainability and programmer productivity arise. Finally, the complexity of these new architectures leads to interesting failure modes which may be difficult or impossible for the developers to diagnose on their own.

We focus on collaborative aspects such as the open development practices of both QUDA and tmLQCD, the interaction between the QUDA development team and the LQCD community, the early access programme organised by JSC for the Juwels Booster supercomputer and the effort required over many months for the diagnosis of a particularly vexing issue with node failures on that machine. In this process we analyse which structural and interactional factors we believe have enabled us to successfully tackle these challenges at different stages and attempt to use our example to characterise what it takes to develop software for LQCD research today.

## **Slot length**

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

**Co-authors:** Dr ALVAREZ, Damian (Juelich Supercomputing Center); Dr BACCHIO, Simone (Cyprus Institute); Dr CLARK, Kate (NIVIDIA); Mr FAHMY, Ahmed (Juelich Supercomputing Center); Dr FINKENRATH, Jacob (University of Wuppertal); Dr GAROFALO, Marco (HISKP, University of Bonn); Dr HERTEN, Andreas (Juelich Supercomputing Center); Dr JOO, Balint (ORNL); Dr PITTLER, Ferenc (Cyprus Institute); Dr ROMITI, Simone (HISKP, University of Bonn); Mr SEN, Aniket (HISKP, University of Bonn); Dr THUST, Kay (Juelich Supercomputing Center); Prof. URBACH, Carsten (HISKP, University of Bonn); Dr WAGNER, Mathias (NVIDIA); Dr WEINBERG, Evan (NVIDIA); Dr HOWARTH, Dean (Berkeley National Laboratory)

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

Session Classification: Parallelization and HPC Infrastructure

Track Classification: HPC and Massive Data: Parallelization and HPC Infrastructure