

A portable LQCD Monte Carlo code using MPI and OpenACC

Thursday, 14 December 2017 17:00 (20 minutes)

HPC systems used by scientific communities often include different kind of accelerators, such as GPUs. Accelerators can deliver an higher level of performance for HPC workloads than ordinary CPUs, but in general they require to rewrite applications using specific programming frameworks.

For example, many HPC clusters nowadays are based on NVIDIA GPUs requiring to use the CUDA framework, which limits to run applications only on a subset of the possible HPC resources.

To avoid this limitation, multiple code versions targeting different architectures could be developed, but this requires a significant programming effort, in particular for scientific applications where code changes are very frequent. To have a single portable code, another approach could be to use the OpenACC parallel programming model. In this case, the original source code is annotated with specific directives to instruct the compiler about how to map instructions on different processors.

In this work, we present the implementation of a state-of-the-art production-level LQCD Monte Carlo application, written using OpenACC and MPI to manage multi-device parallelism. In particular, we give a short introduction to the OpenACC programming model, and we present the design and optimization strategies we have adopted for our code, together with performance results measured on modern HPC clusters based both on CPUs and GPUs.

Primary author: CALORE, Enrico (FE)

Presenter: CALORE, Enrico (FE)

Session Classification: Session 8