

Hydra: A framework for data analysis in massively parallel platforms

Tuesday 10 July 2018 11:15 (15 minutes)

Hydra is a templated header-only, C++11-compliant library for data analysis on massively parallel platforms targeting, but not limited to, the field High Energy Physics research.

Hydra supports the description of particle decays via the generation of phase-space Monte Carlo, generic function evaluation, data fitting, multidimensional adaptive numerical integration and histogramming.

Hydra is open source and the code is hosted in GitHub.

The library deploys a series of techniques in order to achieve optimal performance in both computing and management of memory resources.

The overall design exploits heavily C++ variadic templates in order to implement static polymorphism, kernel fusion and coalesced memory access patterns, avoiding completely the usage of function pointers, virtual methods calls and other known potential performance degrading constructs.

Hydra is developed on top of the Thrust library and runs on Linux systems and can deploy transparently NVidia CUDA-enabled GPUs as well as multicore CPUs and accelerators.

This contribution summarizes the main features of Hydra. A basic description of the user interface and some examples of applications are provided, along with measurements of performance in a variety of environments.

Primary author: ALVES JUNIOR, Antonio Augusto (University of Cincinnati (US))

Presenter: ALVES JUNIOR, Antonio Augusto (University of Cincinnati (US))

Session Classification: T5 - Software development

Track Classification: Track 5 –Software development