



Contribution ID: 414

Type: Poster

Performance portability with alpaka

Thursday 27 October 2022 16:10 (30 minutes)

The alpaka library is a header-only C++17 abstraction library for development across hardware accelerators (CPUs, GPUs, FPGAs). Its aim is to provide performance portability across accelerators through the abstraction (not hiding!) of the underlying levels of parallelism. In this talk we will show the concepts behind alpaka, how it is mapped to the various underlying hardware models, and show the features introduced over the last year. In addition, we will also (shortly) present the software ecosystem surrounding alpaka.

Experiment context, if any

References

https://doi.org/10.1007/978-3-319-67630-2_36 - latest peer-reviewed paper about alpaka https://github.com/alpaka-group/alpaka - Main software repository https://www.hzdr.de/publications/Publ-33634 - Introductory lecture given at the ESC21 school in Bertinoro, Italy

Significance

The alpaka library has been adopted by the CMS experiment at CERN to be integrated in CMSSW as a performance portability library. It is also used in other projects, for example the PIConGPU project for particle-incell simulations.

Authors: GRUBER, Bernhard Manfred (Technische Universitaet Dresden (DE)); Mr STEPHAN, Jan; VYSKOČIL, Jiří (CASUS - Center for Advanced Systems Understanding); Dr BUSSMANN, Michael; WIDERA, René (Helmholtz-Zentrum Dresden-Rossendorf); BASTRAKOV, Sergei (Helmholtz-Zentrum Dresden-Rossendorf); EHRIG, Simeon; DI PILATO, Tony (CASUS - Center for Advanced Systems Understanding (DE))

Presenter: Mr STEPHAN, Jan

Session Classification: Poster session with coffee break