Contribution ID: **38** Type: **not specified**

Adding CUDA Support to Cling: JIT Compile to GPUs

Monday, 10 September 2018 12:20 (15 minutes)

We present the results of a diploma thesis adding CUDA (runtime) C++ support to cling. Today's HPC systems are heterogeneous and get most of their computing power from so-called accelerator hardware, such as GPUs. Programming GPUs with modern C++ is a perfect match, allowing perfectly tailored and zero-overhead abstractions for performance-critical "kernels".

Nevertheless, tool complexity in development and debugging can be discouraging for new users. We are addressing this by not only adding low-level support for accelerators but also by going up the open source software-stack enabling interactive, CUDA C++ Jupyter notebooks, e.g. through xeus-cling.

Primary author: Mr EHRIG, Simeon (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden)

Co-authors: NAUMANN, Axel (CERN); Mr HUEBL, Axel (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden)

Presenter: Mr EHRIG, Simeon (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden)

Session Classification: Parallelism, Heterogeneity and Distributed Data Processing

Track Classification: Presentations