



Contribution ID: 39

Type: **not specified**

Data Analysis and Simulations in Exascale Computing: Quō vādis?

Monday, 10 September 2018 14:50 (15 minutes)

We are less than three years apart from the first, double precision Exa-Flop/s supercomputers. Already today, our scientific software stacks are facing the challenge to run efficiently on a potpourri of architectures. But the real troubles might await us at the choke points of extreme data rates, where traditional workflows of data acquisition, filtering, processing and subsequent long-term storage might not be able to be sustained anymore.

How would you like to express your scientific algorithms in a world where Flop/s are increasingly cheap, yet hard to achieve, but data movement and especially data at rest is increasingly in-proportionally expensive? Would you be OK to throw data away and measure twice? Can we in situ compute results with a different prepared question instead of waiting for an always-full and quickly-purged filesystem? How do we ensure reproducibility? How large a mix of programming languages and double-implementations of algorithms can we burden before we are running out of developers (due to lack of maintainability)?

This talk will present our vision for the next years of data-driven scientific computing. Based on our experience with single-source, performance-portable C++ HPC libraries, we will present zero-overhead C++ abstractions that spare code-duplication. Together with light-weight code coupling, possible directions for analyzing resulting data rates are discussed on examples from laser-driven particle accelerator research. With such meta-programming approaches, an underestimated risk lies in cutbacks for both development workflows and user interactivity at runtime, which we want to openly change with interactive Cling-assisted execution in modern environments such as Jupyter, for which we recently enabled CUDA C++ capabilities.

Primary author: HUEBL, Axel (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden)

Co-authors: Mr EHRIG, Simeon (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden); BUSSMANN, Michael (Helmholtz-Zentrum Dresden - Rossendorf)

Presenter: HUEBL, Axel (Helmholtz-Zentrum Dresden-Rossendorf and TU Dresden)

Session Classification: Parallelism, Heterogeneity and Distributed Data Processing

Track Classification: Presentations