Yurii Perets

CI/CD improvements for alpaka library, and related projects

Mentors:

Volodymyr Bezguba(KAU)

Jiří Vyskočil(CASUS)

Next generation HEP experiments

- High Energy Physics is facing an enormous challenge in the coming decades as data volumes and processing requirements for precision physics analysis and simulation increase dramatically with experiments such as the DUNE and HL-LHC.
- Heterogeneous computing is one of the key components to meet the computing challenge.
- Programming techniques for GPUs and other massively parallel accelerators are very different from that of traditional CPUs.
- Over the decades HEP experiments have developed millions of lines of code, and in order to use these new computational accelerator facilities would need to rewrite large amounts of their code bases. In order to be able to run on all the different hardware architectures, this task would have to be repeated multiple times in each architecture's preferred language.

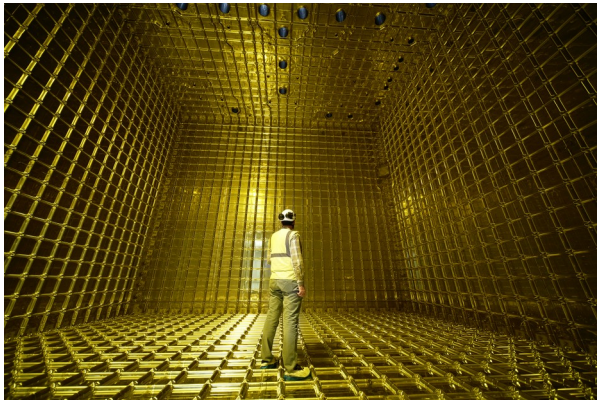

Conceptual overview

Alpaka provides a single abstract C++ interface to describe parallel execution across multiple levels of the parallelization hierarchy on a single compute node. Each parallelization level corresponds to a particular memory level.

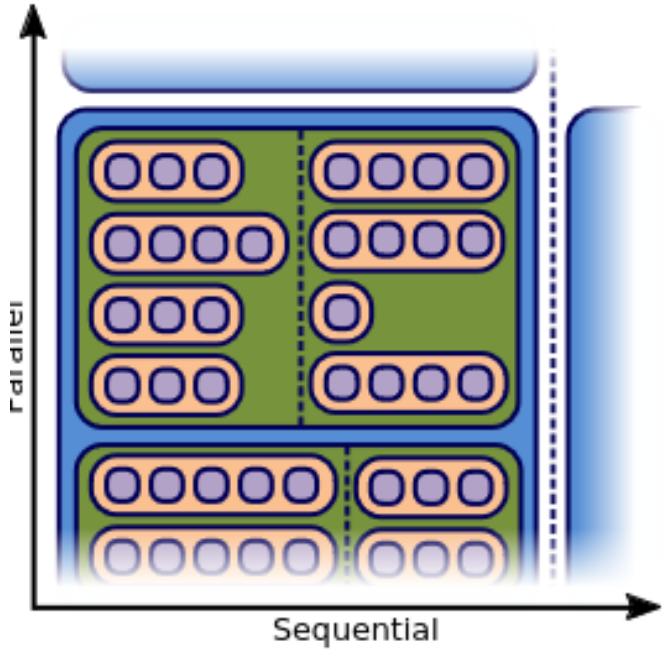

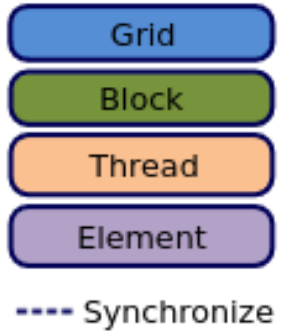

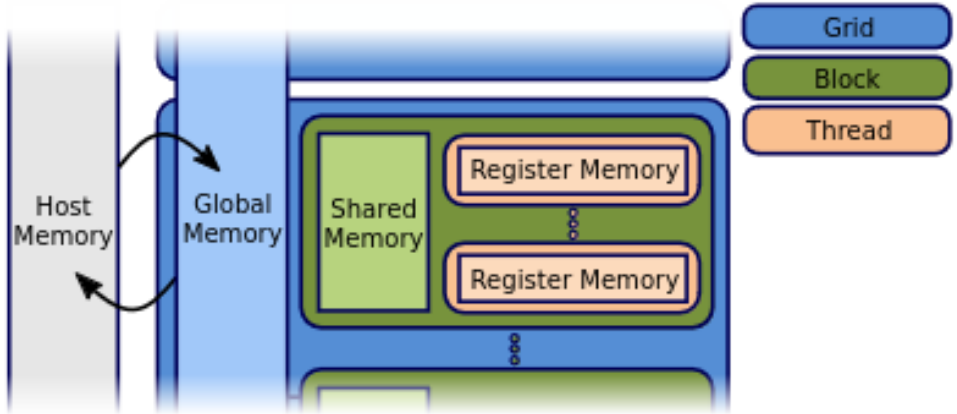

Project aims

This project aims to optimize the test coverage and runtime of alpaka.

Testing alpaka with all possible software dependencies would result in about 2,500,000 different combinations of test jobs.

After implementing an open source CI job generator that works using pair-wise testing, the generator reduced number of jobs to about 170.

As a result of the project, we want to finalize the work on the new open source version of the generator, migrate alpaka to the new generator, and verify that the CI works as expected.

Thank you for your attention!