


Heterogeneous hardware adoption and portability


Matti Kortelainen (FNAL)
for the HEP-CCE project

Snowmass CompF4 Topical Group Workshop
April 7 2022

The HEP Computing Universe Today Yesterday


HEP-CCE





The HPC Computing Universe Today

Perlmutter NERSC, 2020
AMD CPU, Nvidia Tesla GPU



Frontier ORNL, 2021
AMD CPU, AMD GPU, 1.5 ExaFlop



Leonardo, Cineca, 2021
Intel CPU, NVIDIA GPU, 200+PFlops
Italy



Aurora Argonne, 2022
Intel CPU, Intel Xe GPU, > 1 ExaFlop

USA



El Capitan LLNL, 2023
AMD CPU, AMD GPU, > 1.5 ExaFlop



LUMI, CSC, 2021
AMD CPU, AMD GPU, 550 PFlops
Finland



Alps, CSCS, 2023
NVIDIA Arm CPU+GPU
Switzerland



Challenges:

Hundreds of computing sites (grid clusters + HPCs + clouds)

Hundreds of C++ kernels (several million LOC)

Hundreds of data objects (dynamic, polymorphic)

Hundreds of non-professional developers (domain experts)

Opportunity:

Scale of experiments and community provides significant R&D firepower
scores of active groups, will not attempt to list

Current Focus:

Online event filtering, offline pattern recognition, detector simulation

Wouldn't it be nice to develop and maintain one version of the algorithm code and run that efficiently on all hardware platforms of interest?



Portability technologies

Portable Parallelization Strategies Project

Investigate a range of software portability solutions:

- Kokkos / Raja / Alpaka
- SYCL / dpc++ / hipSYCL
- OpenMP / OpenACC
- `std::parallel::execution` (`std::par`)

Port a small number of HEP testbeds to each portability solution

- Patatrack pixel tracking (CMS) → [arXiv:2008.13461](https://arxiv.org/abs/2008.13461), [arXiv:2104.06573](https://arxiv.org/abs/2104.06573)
- p2r (CMS) → [whitepaper](#)
- [WireCell Toolkit](#) (DUNE) → [arXiv:2104.08265](https://arxiv.org/abs/2104.08265)
- FastCaloSim (ATLAS) → [arXiv:2103.1473](https://arxiv.org/abs/2103.1473)
- [ACTS](#) (ATLAS, sPHOENIX, ...): experiment-independent toolkit for track simulation and reconstruction

Defined [a set of metrics](#) to evaluate portability solutions, as applied to our testbeds

- Productivity, cross-platform performance, broader impact, long-term sustainability, *etc*

Make recommendations to the experiments

- Must address needs of both LHC style workflows (many modules and many developers), and smaller/simpler workflows

Portability Solutions: Software Support Chart

	NVIDIA CUDA	Kokkos	Alpaka	AMD HIP	std::par	dpc++ / SYCL	OpenMP Offload	Python
NVIDIA GPU	Supported	Supported	Supported	Supported	Supported	Supported	Supported	Supported
AMD GPU	Not Supported	<i>select GPUs</i>	Supported	Supported	Not Supported	<i>hipSYCL intel/llvm</i>	Supported	Supported
Intel GPU	Not Supported	Supported	<i>experimental</i>	<i>HIPLZ: early prototype</i>	Supported	Supported	Supported	Not Supported
CPU	Not Supported	Supported	Supported	Not Supported	Supported	Supported	Supported	Supported
FPGA	Not Supported	Supported	<i>experimental (via SYCL)</i>	Not Supported	Not Supported	Supported	Supported	Not Supported

Supported
Under Development
3rd Party
Not Supported

- **products are rapidly evolving**
- some hope of seeing emergence of industry standards at language level

Interim Experiences With Portability Layers

- Deploying portability technologies is currently far from standard
 - But there are encouraging successes in using these solutions
- Scientific success of our experiments and endeavors would be in danger without them
 - Software development could become expert driven and costly to run on available hardware infrastructure
- We think the best solution for the community would be an extension to the C++ standard with low entry bar for users, and supporting all hardware forms and vendors
 - Currently we are very far from that ideal
 - As a community we need to request and work on portability solutions and strive to reach this ideal

Thanks!

<https://www.anl.gov/hep-cce>
hep-cce@anl.gov

Special thanks to **Paolo Calafiura, Taylor Childers, Martin Kwok, Charles Leggett, Meifeng Lin, Vince Pascuzzi** for the material they contributed. All mistakes and misrepresentations are mine only.