



Portable Parallelization Strategies - Overview and Metrics

Martin Kwok(Fermilab)

On behalf of the CCE-PPS team

IRIS-HEP topical meeting

04 August, 2021

What is HEP-CCE?

- HEP-CCE (Center for Computational Excellence)
 - A 3-year (2020-2023) pilot project, funded by US DOE
 - Formed by 4 US labs, 6 experiments, covering cosmic/intensity/energy frontiers
- Develop common strategies to common computational challenges for HEP community
 - Specifically to efficiently run HEP software on HPCs


HEP-CCE



1. PPS: Portable Parallelization Strategy ← Focus of today
2. IOS: HEP I/O and HPC storage issues
3. EG: Event Generation
4. CW: Complex Workflow on HPC


Parallelization: Common challenge to HEP

- Demand for processing power from HEP experiment will outgrow expected resource increase
 - “Buy more CPU” does not work
 - In 2030, LHC experiments will need O(100) PFlops/s
- *Parallelism* is one of the key to meet growing computing power
 - Lesson from the past: single-core CPU frequency has plateaued since ~2005
 - Multi-core CPU flourished



Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten
Dotted line extrapolations by C. Moore



[arxiv:1604.00319](https://arxiv.org/abs/1604.00319)



Parallelization: Common challenge to HEP

- Extend the parallelism:
 - Offload computational intensive tasks to a custom architecture for parallel computation (an accelerator)
 - Could be a GPU or even a FPGA
- Exascale HPC systems could provide the computing power
 - Taking advantage of CPU + GPU
- Mis-match between software and hardware:
 - HEP software predominantly written for x86 platforms
 - Offloading to another architecture is a different programming paradigm
 - Needs more effort to use accelerator efficiently

HEP software



Perlmutter Phase 1 @NERSC
Ranked 5 in the Top-500 list
700k+ cores, ~65 PFlops/s

Portability: HPC landscape

- Accelerator architectures are proliferating
 - Main GPU manufacturers: NVidia, Intel, AMD
 - FPGA is possible but is less mature
- We need portable solutions
 - Cannot afford to re-write the code for each architecture
 - Aim at maintaining a single code-base for all target architectures

| | | Accelerators | | | | |
|-----|---------|--------------|---|------------------------|---------|--------|
| | | Intel | NVIDIA | AMD | FPGA | Other |
| CPU | Intel | Aurora | WLCG (HEP) Cori GPU Piz Daint Tsukuba MareNostrum | | Tsukuba | |
| | AMD | | WLCG (HEP) Perlmutter | Frontier El Capitan | | |
| | IBM | | Summit Sierra MareNostrum | | | |
| | Arm | | Alps | | | Astra |
| | Fujitsu | | | | | Fugaku |

Portability: Software Support

- No universally accepted best way to write performance, portable GPU code
 - Each portability solution has strength and weakness

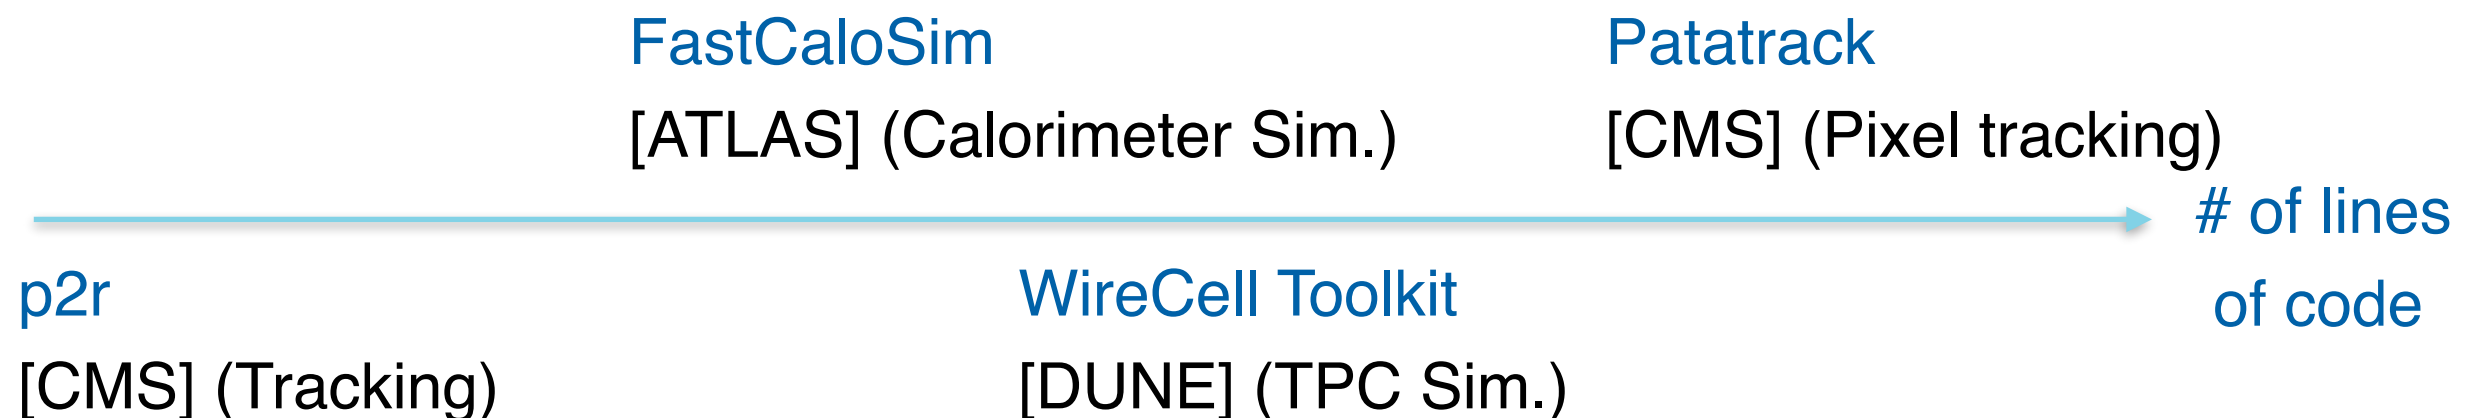
| | OpenMP Offload | Kokkos | dpc++ / SYCL | HIP | CUDA | Alpaka | |
|------------|-------------------|-------------------|----------------|---------------|---------------|------------------------|-------------------|
| NVidia GPU | Supported | Supported | Intel/codeplay | Supported | Supported | Supported | Supported |
| AMD GPU | Supported | prototype | via hipSYCL | Supported | Not Supported | Under Development | Under Development |
| Intel GPU | Supported | Under Development | Supported | Not Supported | Not Supported | very early development | 3rd Party |
| CPU | Supported | Supported | Supported | Not Supported | Not Supported | Supported | Not Supported |
| Fortran | Supported | Not Supported | Not Supported | 3rd Party | 3rd Party | Not Supported | Not Supported |
| FPGA | Under Development | Under Development | Supported | Not Supported | Not Supported | possibly via SYCL | Not Supported |

The technologies are still evolving, platform support is a moving target

Strategy: The plan for PPS

- Investigate a range of software portability solutions
 - Kokkos
 - SYCL/ dpc++
 - Alpaka
 - OpenMP/OpenACC
- Experiment each solution with various HEP applications
 - Covers different HEP experiments, code sizes, physics problems
 - Diversity of the test application tells us different aspects of the technology
- Evaluate each solution with a set of metrics
 - Capture the full porting experience (Porting, building, performance etc.)
- Make recommendations to the experiments
 - Addresses different needs for different workflows

Detailed results
In Haiwang's talk



Kokkos

<https://github.com/kokkos/kokkos>

- Single source C++ template library
- Aims to be descriptive, not prescriptive
 - Kokkos handles the backend (relatively high-level API)
 - Abstraction layer provides handle for efficient data layout for both GPU/CPU + more
- One host(CPU)+parallel (GPU/CPU) backend chosen at compile time
 - Supports NVIDIA/AMD GPU
 - Write once, compile for different architectures
- Active Kokkos development, and user support
 - Open source

Kokkos Abstraction

Pattern: nature of work

Execution Policy:
How computations
are executed

Body: Unit of work

```
Pattern          Policy
for (element = 0; element < numElements; ++element) {
  total = 0;
  Body          for (qp = 0; qp < numQPs; ++qp) {
                 total += dot(left[element][qp], right[element][qp]);
                 }
  elementValues[element] = total;
}
```

Parallel
execution
backends

CUDA
NVIDIA

OpenMP

HIP
AMD

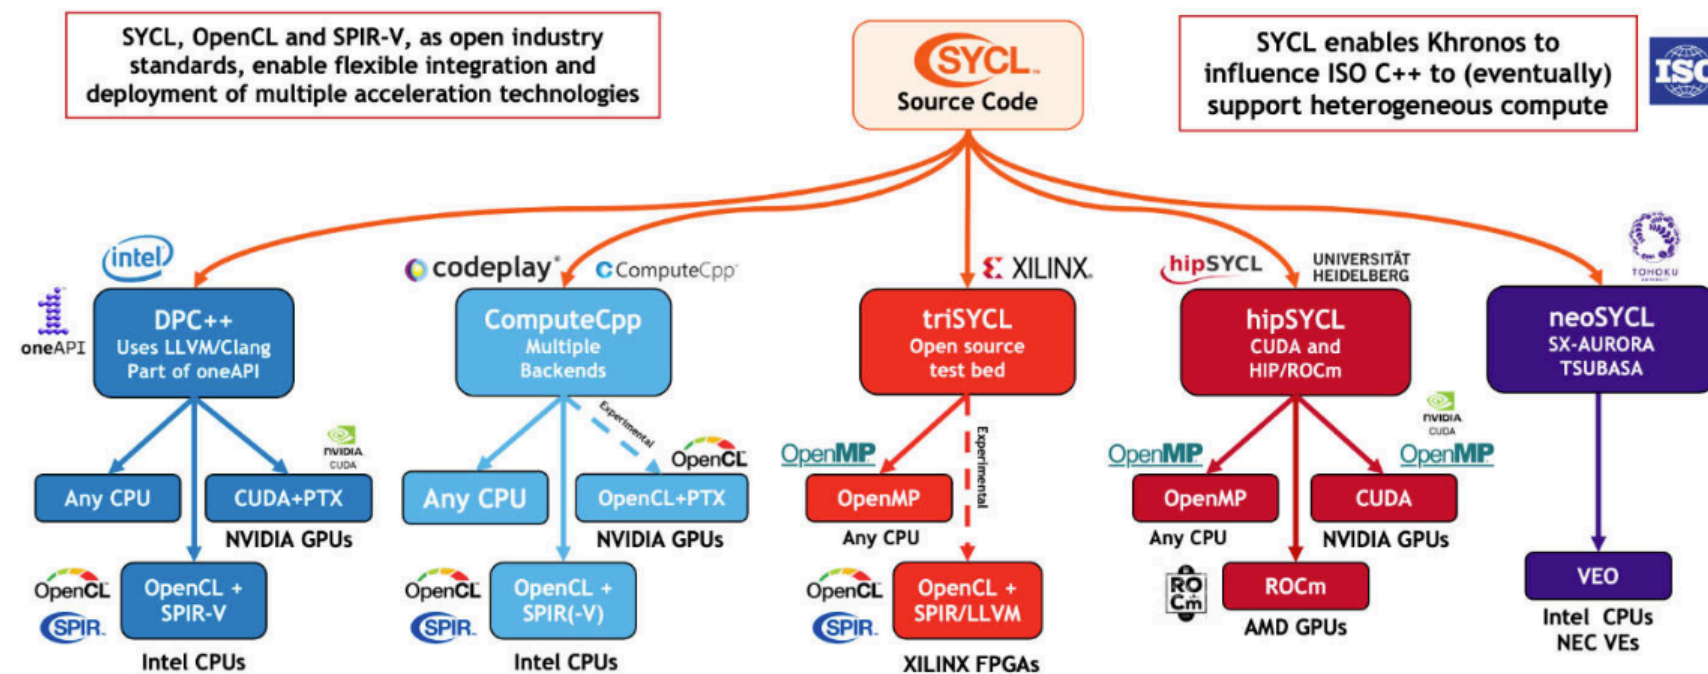
CPU
pThread



SYCL / Data Parallel C++ (DPC++)

- SYCL is a specification of single-source C++ programming model for heterogeneous computing
 - Different compilers implements the specification
 - Intel's oneAPI and Data Parallel C++ (DPC++) supports Intel CPU, GPU, FPGA
 - Parallelism based on C++17 standard
- Execute on most architectures
 - Including CPU, GPU, FPGA
- Exact behavior of program depends on compiler
 - Compilers are evolving rapidly
- C++ standard for heterogeneous computing is still evolving
 - Features are growing

Complex ecosystem



First year of CCE-PPS

- Projects typically start with parallel CPU/ CUDA implementation
 - Target a portable solution
 - Could take a long time, depending the size of project
 - Benchmarking/Optimize implementation
 - See results from Haiwang's slides
- Designed metric for evaluating different solutions
 - Lots of useful lessons learned from hands-on experience
 - Collecting results among ourselves from different projects

| | | <i>Implementation</i> | | | |
|---|-----------------|-----------------------|-------------|-------------|---------------|
| | <i>Projects</i> | <i>CPU</i> | <i>CUDA</i> | <i>SYCL</i> | <i>Kokkos</i> |
| 1 | FastCaloSim | Completed | Completed | Completed | Completed |
| 2 | Patatrack | Completed | Completed | Planned | Completed |
| 3 | Wirecell | Completed | Completed | Planned | In progress |



Planned



In progress



Completed

Metrics: Evaluation

1. Ease of learning and extend of code modification
2. Code conversion
3. Impact on existing code
 - Control of main(),
 - Threading/execution model
4. Impact on EDM/ data
5. Impact on existing toolchain and build infrastructure
6. Hardware mapping
 - Current and future backend support
7. Feature availability
 - Reduction, concurrent kernels etc.
 - Interface to commonly used math libraries
8. Ease of debugging
9. Address needs of all workflows
 - Scaling with # kernels / applications
10. Long-term sustainability and code stability
11. Compilation time
12. Performance/Run time
 - Compares with native implementation (e.g. CUDA for NVidia GPUs)
13. Interoperability
 - Run different technologies in same application?
14. Aesthetics

Metric progress



Needs input



Some input there



Good amount of content

| | <i>Topics</i> | <i>Metric formulation</i> | <i>SYCL</i> | <i>Kokkos</i> |
|----|---|---------------------------|------------------------|------------------------|
| 1 | Ease of Learning Language | Good amount of content | Good amount of content | Good amount of content |
| 2 | Code conversion | Good amount of content | Good amount of content | Good amount of content |
| 3 | Extent of modification to existing code | Good amount of content | Good amount of content | Good amount of content |
| 4 | Extent of modifications to EDM / Data | Some input there | Good amount of content | Some input there |
| 5 | Extent of modifications to build rules | Good amount of content | Good amount of content | Good amount of content |
| 6 | Hardware Mapping | Good amount of content | Good amount of content | Good amount of content |
| 7 | Feature Availability | Some input there | Some input there | Some input there |
| 8 | Address needs of all workflows | Needs input | Needs input | Needs input |
| 9 | Long term sustainability and code | Good amount of content | Some input there | Some input there |
| 10 | Compilation time | Good amount of content | Good amount of content | Good amount of content |
| 11 | Run time | Some input there | Some input there | Some input there |
| 12 | Ease of debugging | Good amount of content | Some input there | Some input there |
| 13 | Aesthetics | Good amount of content | Some input there | Some input there |
| 14 | Flexibility | Good amount of content | Needs input | Needs input |

Summary and moving forward

- PPS stands for Portable Parallelization Strategy
 - Develop a common strategy for parallelization of HEP software
 - One of the key ingredients to meet the demand for HEP computing in the next 10 years
 - Fruitful results from 1st year: See Haiwang's talk up next
- What's next for PPS?
 - New use-case: A Common Tracking Software (ACTS)
 - HEP common software, instead of experiment-specific application
 - Propagate-to-R (P2R)
 - Very light-weight tracking application [$O(1000)$ lines] to speed-up exploration of technologies
 - Explore new technologies
 - Alpaka: Europe-backed C++ parallel programming library, similar to Kokkos
 - `std::par`: Parallel execution policy in the new C++ standard
 - Summarize experience from difference projects into metrics/recommendations
 - Documentation of specific details could be useful for future decision making