Summary of Compf01, ExpAlgos parallel session

Katrin Heitmann (ANL, LSST DESC)
Walter Hopkins (ANL, ATLAS),
Giuseppe Cerati (FNAL, SBN/DUNE, CMS)

August 11, 2020

CompF01 workshop agenda

The compf01 sessions

Input from frontiers on main needs and challenges for the experiments (list of frontiers covered is not exclusive!)

Discussion of structured R&D efforts, plans and early results

Open discussion about challenges and plans for LOIs

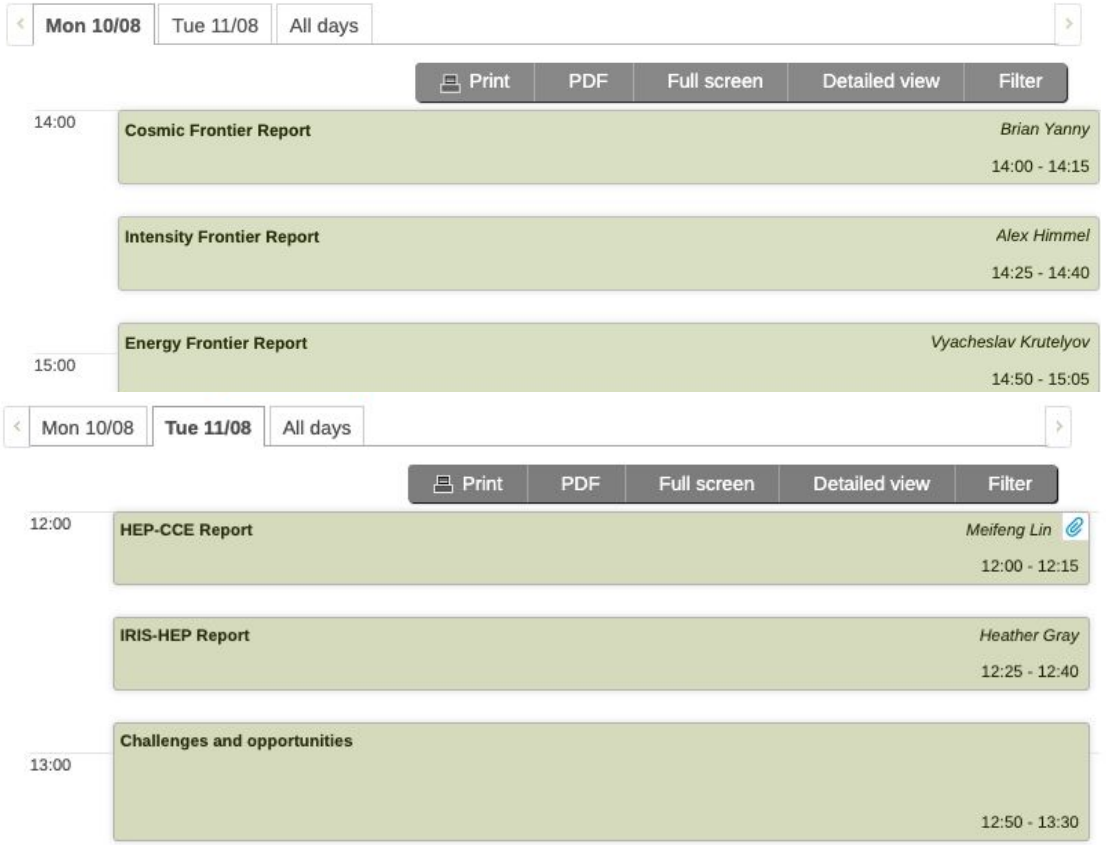

Cross-frontier themes

- Many different needs, from different experiments or different algorithms
  - difficult to have a one-fits-all solution, even within a frontier
  - possible exceptions: accelerated FFTs in Cosmic and Neutrino frontier, real time processing (trigger/broker applications)
- Transitioning from HTC to HPC (or using both)
  - evolution of the programming model
- (Optimal) use of heterogeneous resources
  - how to keep the GPUs busy?
- Interplay between ML and traditional reco algorithms
  - switch to ML approaches vs rewriting algorithms
  - avoid separate workflows, ensure feedback between the two

Cosmic frontier report takeaways / Brian Yanny

- Wide range of experiments (CMB surveys, optical surveys, dark matter detection experiments) that all have different needs with regard to experimental algorithm development
- Future surveys will stress computing resources at new level and experimental algorithms will become bottleneck if not addressed appropriately
- Challenges with regard to use of HPC vs HTC resources
- Data from optical surveys will be reprocessed several times and experimental algorithms might need to be improved to enable handling the high quality data during the lifetime of the survey
- ML approaches will play important role for identifying artifacts as well as new transients, variables, and moving objects of interest
- Cross-frontier challenges: Fast FFTs, search for rare events, ML approaches, challenges with regard to using heterogeneous architectures efficiently

Intensity frontier takeaways / Alex Himmel

- Covered DUNE, not necessarily applicable to other neutrino experiments. Definitely not applicable to RPF
  - Variety of detector technologies, even within the same experiments. How to have solutions common?
- Differently from LHC, the main problem of DUNE is not event rate but volume rates at raw data level, so solutions may not be the same. Need to process in parallel separate detector volumes
- Different competing approaches, often using separate processing resources/workflows are difficult to combine and get the most out of them
- Usage of simulation tools within reconstruction?
- How to make effective usage of heterogeneous resources?

Energy frontier takeaways / Slava Krutelyov

- Efforts to parallelize algorithms start within experiments and different R&D venues bring projects together.
- How much time is needed to adapt algorithms for GPUs (and other resources) vs training an ML algorithm?
- Tracking is the largest contributor to the processing time, and where most of R&D activities focus on. After tracking, no one algorithm dominates resource usage. How to prepare many algorithms for heterogeneous resources with limited expertise and person power?
- Teams started porting to new architectures (e.g. GPUs), planning to focus later on portability
    - Portability solutions are not that mature and you need to start somewhere

HEP-CCE Report / Meifeng Lin

DOE initiative that aims to promote excellence in using HPC for data-intensive applications. Four thrusts: Portable Parallelization Strategies (PPS), I/O and Storage (IOS), Event Generators (EG), Complex Workflows (CW)

Started with three self-contained use cases from ATLAS, CMS, and DUNE, all with an initial CUDA implementation. Investigating different portable programming models for each use case, ranging from libraries with different backends to compiler directives.

Initial goal prioritizes portability over performance, although initial results show Kokkos with CUDA backend is close to CUDA standalone.

Deliverable of 3-year project is a set of recommendations for portability

IRIS-HEP report / Heather Gray

NSF-funded software institute, activities within IRIS-HEP that are relevant for CompF01 are in innovative algorithms group. Focusing is on LHC challenges

Main activities include development of tracking algorithms, preparation of algorithms for GPUs, exploiting ML solutions. Each project spans multiple if not all of these activities. Projects are fairly mature, with papers published and experiments adopting the proposed solutions.

Institute focuses both on specific and common solutions. Many algorithms developed by different groups within one experiment, so the main goal is to find specific solutions but it is also useful to find commonalities. Other developments start as common solution that can be leveraged by different experiments.

Challenges discussion / All

- Definition of metrics for experimental algorithms and how to weigh them
  - Ease of implementation and portability
  - Performance
  - Longevity (design/optimize for machine available now or try to develop long lasting implementation)
- Traditional grid resources vs HPC centers
  - Main limitation of the grid: cannot point code to machine with a particular set of features
  - Using a variety of centers makes verification challenging given the diverse set of resources
- Training
  - Have to improve teaching of software development to be able to face future challenges
  - Most students focus on high-level programming, while efficient algorithm development often needs deeper knowledge of programming paradigms

Thank you!

Stay connected:

- Mailing list: snowmass-compf01-expalgos[at]fnal.gov
- Slack Channels: #compf01-expalgos, #comp_frontier_topics
- Wiki page: https://snowmass21.org/computational/algorithms