



Brief recap from ALICE on R&D activities

Stefano Piano, Andreas Morsch and Peter Hristov



Update from October 2020

see:

<https://indico.cern.ch/event/919840/>

Visualization of 2 ms of 50 kHz Pb-Pb data as expected in the ALICE TPC in LHC Run 3
Basic processing unit: **Time Frame** (~10 ms of data ~500 collisions @ 50 kHz Pb-Pb)

Usage of GPUs in ALICE

- **Synchronous Reconstruction:**

- Fully dominated by the TPC (99% CPU time): already fully runs on the GPU (GPU-bound)
- Since July global commissioning with the synchronous workflow on EPN

- **Asynchronous Reconstruction:**

- When the EPN farm is not (fully) used for synch. processing, it will be used for asynchronous processing of the raw data stored on the disk buffer
- This ensures constant 100% duty cycle of the EPN farm
- More than 80% is already on the GPU (baseline scenario) but still CPU-bound
- Keep working to offload more tasks to GPU in asynchronous reconstruction ongoing:
 - First benchmark on EPN of asynchronous workflow with TPC tracking on GPU (2x)
 - It will reach higher GPU speedup once optimistic scenario is fully implemented

- **Simulation:**

- Plan to develop digitization on GPU (at least for TPC)

HPC and cloud resources

- **Thanks to the new O2 simulation and asynchronous reconstruction code (Run 3) possible to fully exploit the multi process features:**
 - Event generator and Transport via VMC Interface:
 - Sub-event parallel simulation across multiple transport workers
 - Digitization:
 - Implemented in the ALICE Run 3 framework (DPL)
 - Parallel across and within detector components
 - New graph-pipeline for multicore scheduling to achieve good multi-core efficiency
- **Significant progress has been made to incorporate HPC and cloud resources in the standard ALICE Grid workflows:**
 - Multicore queues, Intel based HPC and cloud resource delivery at CERN
- **Integration of non-x86 platform (IBM+Nvidia, ARM):**
 - Needed to exploit a wider range of computing resources

Analysis

- Novel O^2 analysis framework available
- New organized distributed analysis “**hyperloop**”
- Performed benchmark of real analysis in distributed environment
 - Strong optimization of event throughput:
 - Factor 10 wrt Run 2 analysis framework
 - No optimization for optimal core usage done yet
 - But multi core shows speed up
- Analysis facilities (GSI and Wigner)
 - Collect ~10% of AODs on a few dedicated sites (AFs) that are capable of quickly processing locally large data volumes
 - Goal: fast turnaround for cut tuning and task validation
- Interested in high performance interactive machines for analysis (ML/DL)