Role of GPU in ALICE Online-offline reconstruction at LHC Run 3

M. Concas, D. Elia, F. Noferini, S. Piano, M. Puccio

Workshop sul Calcolo nell'I.N.F.N.
Paestum 23 - 27 maggio 2022

ALICE @ Run 3 and 4

$L_{Pb\text{-}Pb} = 7.0\ nb^{-1}$

Shutdown/Technical stop
Protons physics
Ions
Commissioning with beam
Hardware commissioning/magnet training

$L_{Pb\text{-}Pb} = 6.0\ nb^{-1}$

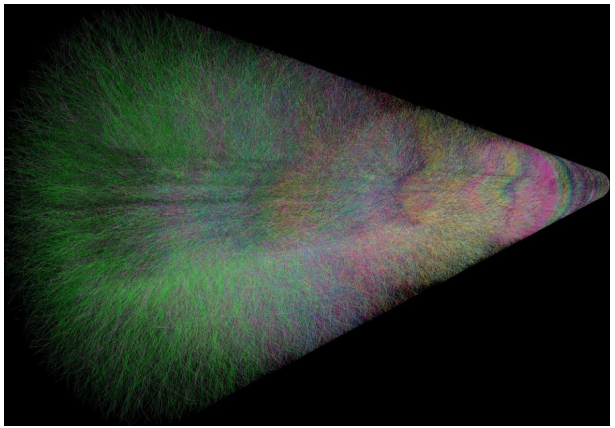

Continuous Readout

Pb-Pb@50 kHz
pp@1 MHz

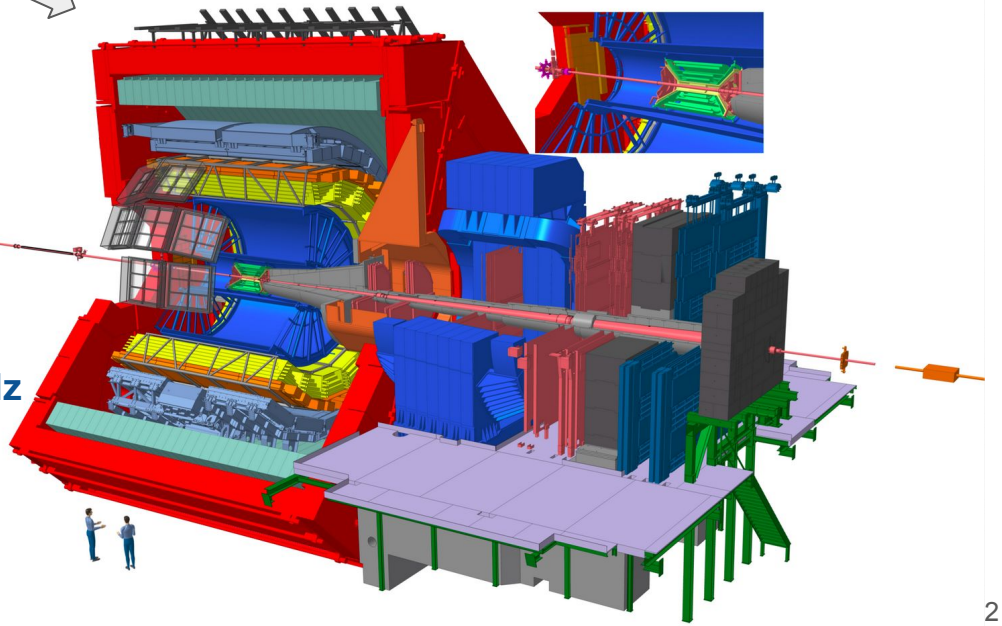

ALICE @ Run 3 and 4

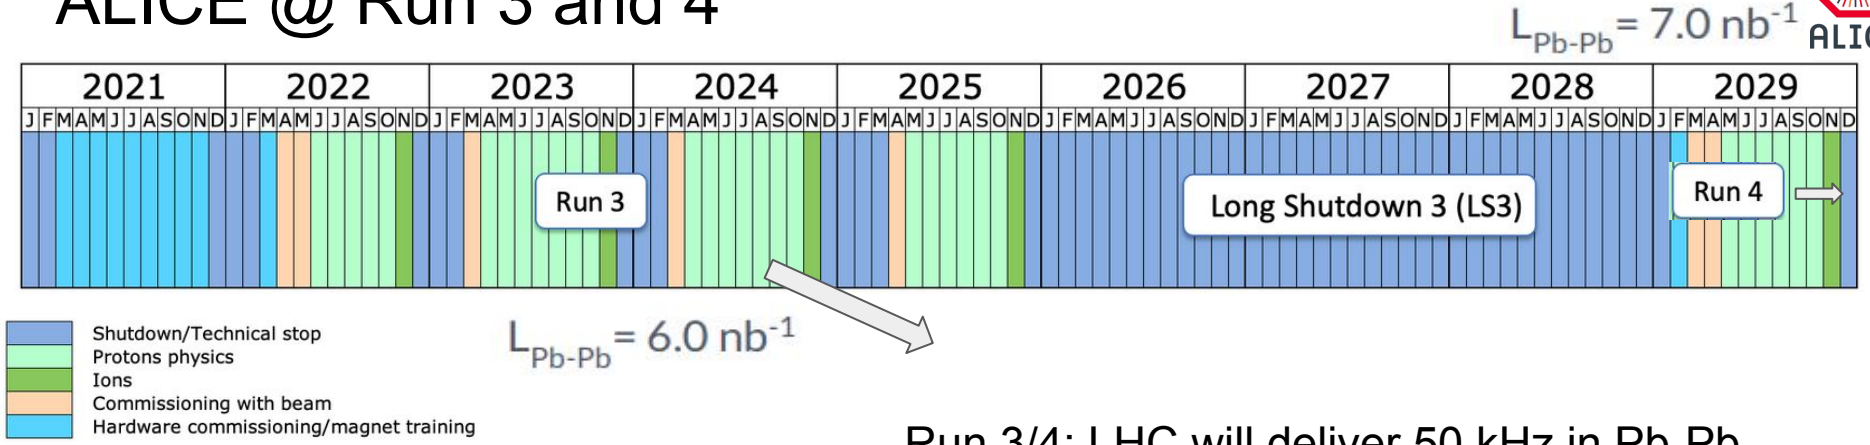

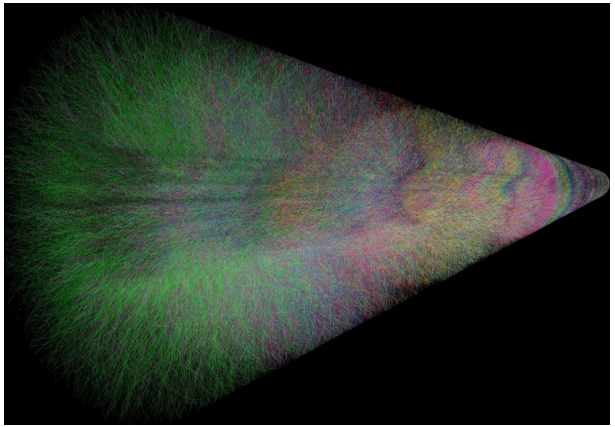

Continuous Readout

Pb-Pb@50 kHz
pp@1 MHz

Run 3/4: LHC will deliver 50 kHz in Pb-Pb collisions

- ALICE aims to record >10 nb^{-1} integrated luminosity, x50 times more minimum bias data wrt Run 2

ALICE Raw Data Flow in Run 3

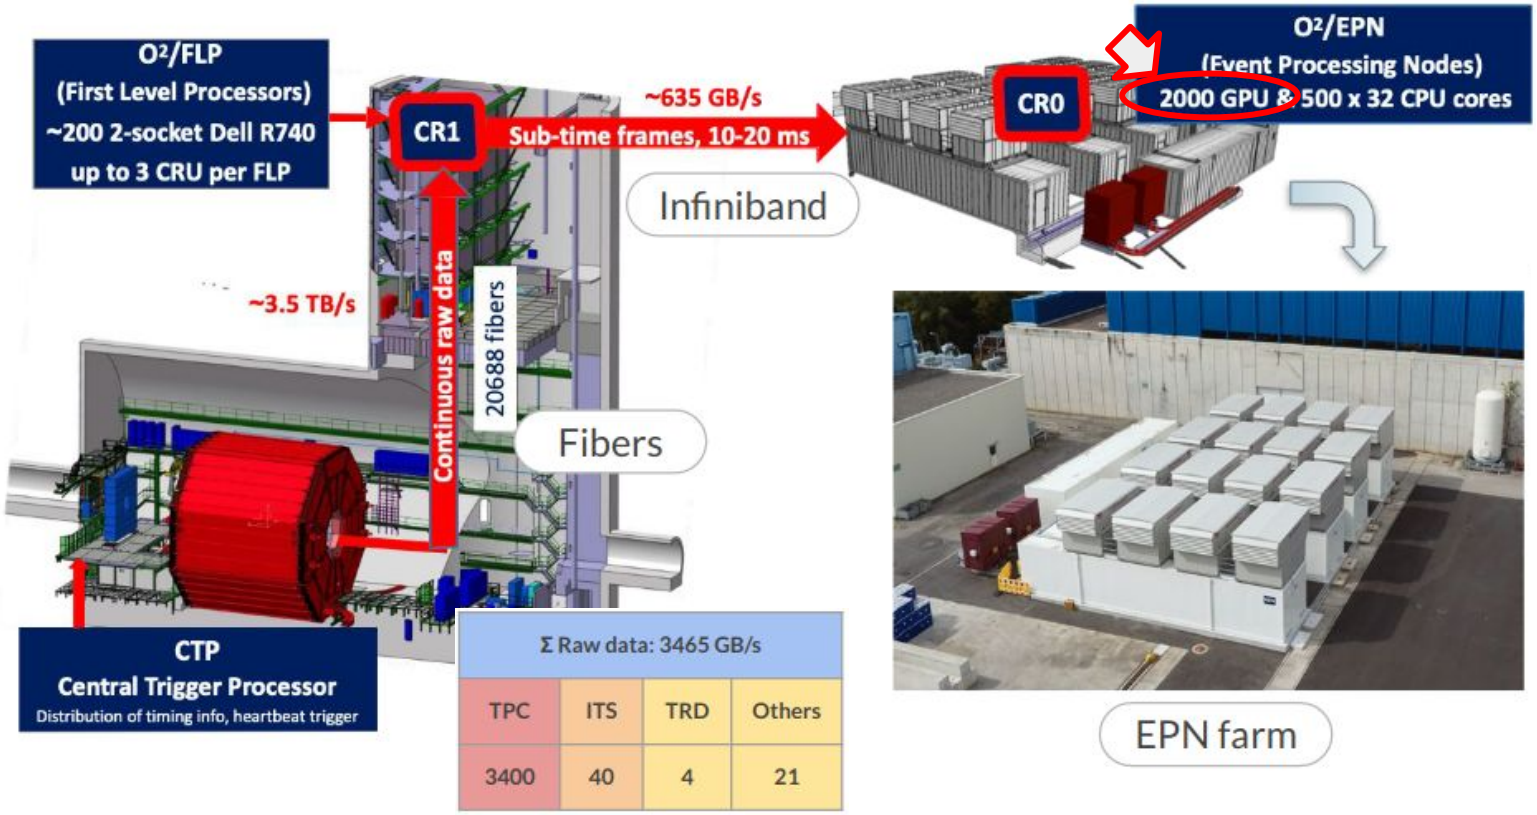

Σ Raw data: 3465 GB/s			
TPC	ITS	TRD	Others
3400	40	4	21

ALICE O^2 computing model

Reconstruction
Calibration
Archiving
Analysis

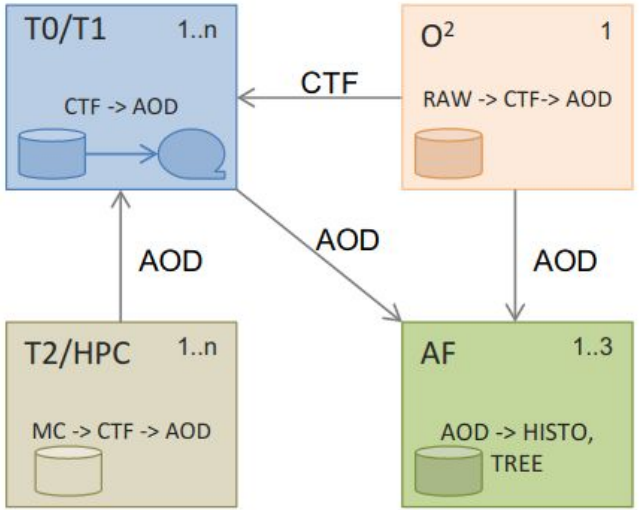

Compression
Reconstruction
Calibration

Simulation
Analysis

Analysis

CPU share	O^2	T0	T1s	T2s	AF
Synch.	100%	0%	0%	0%	0%
Asynch.	33%	33%	33%	0%	0%
MC	0%	0%	0%	100%	0%
Analysis	0%	30%	20%	0%	50%

- Subject to fine tuning
- MC can be run as a backfill

2/3s of CTFs processed by O^2 + T0 and archived at T0;
1/3 of CTFs exported, archived and processed on T1s;

One calibration (sync.) and two reconstruction passes (async.) over raw data each year, with processing on EPN farm (> 2/3)

The goal is to minimize data transfer and optimize processing efficiency
10% of AODs sampled and sent to the Analysis Facility for quick analysis and cut tuning;

Analysis of full data sample across T0/T1s only performed upon Physics Board approval.

Processing plan for Pb-Pb

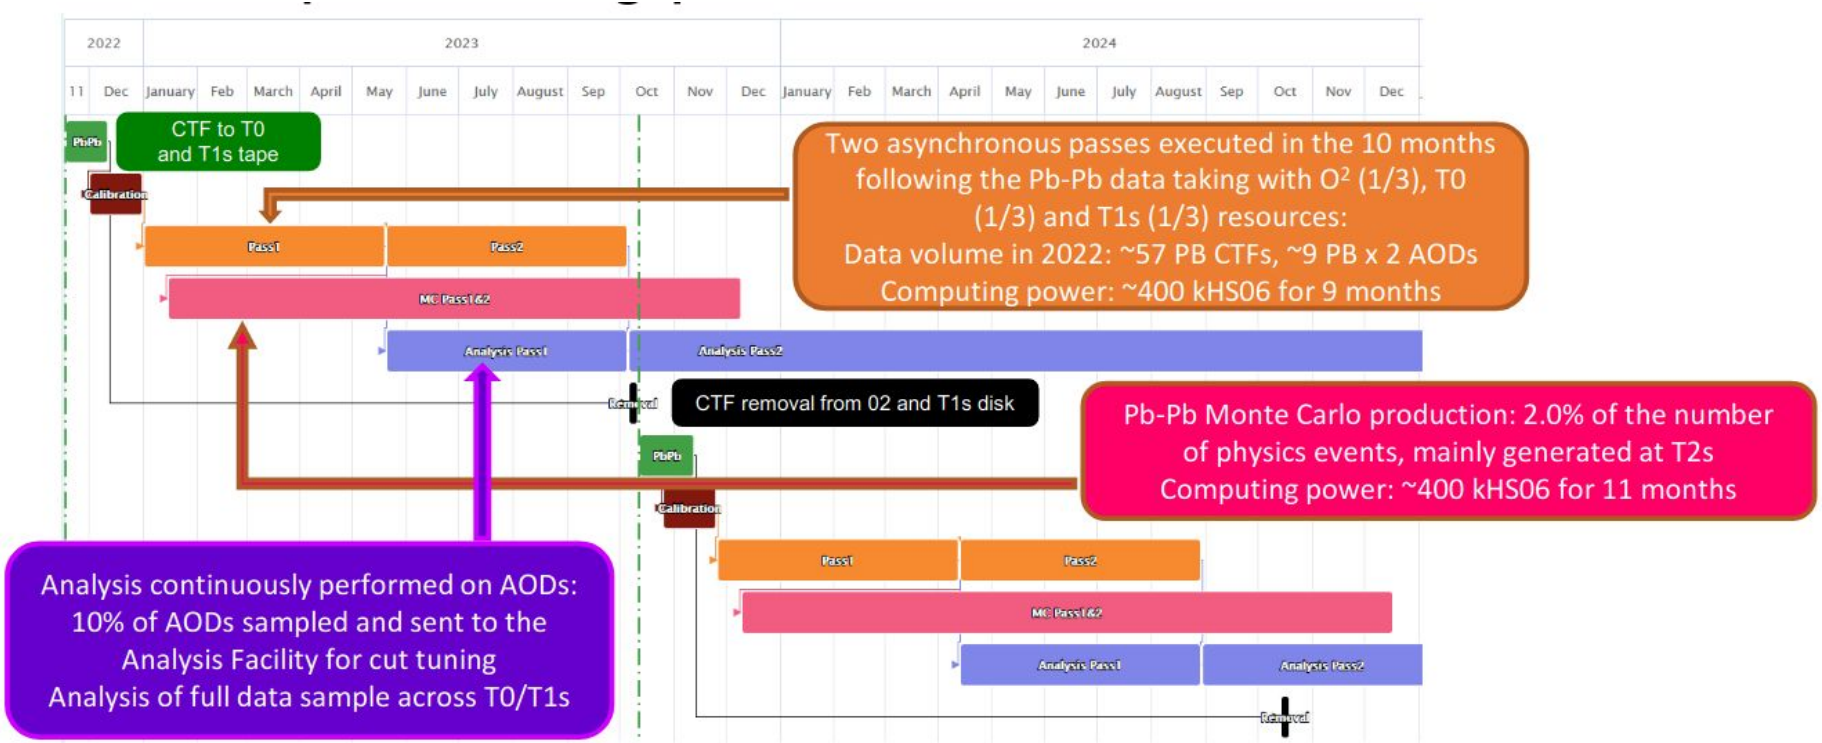

ALICE O^2 activities

ALICE upgrade to continuous readout required a new Online-Offline (O^2) framework

- message-queues based processing by separate device(processes)
- Data Processing Layer (DPL)
- Workflows are built for group of Devices by automatic matching of their Inputs and Outputs

G. Eulisse, R. Shaoyan (CHEP 2019)

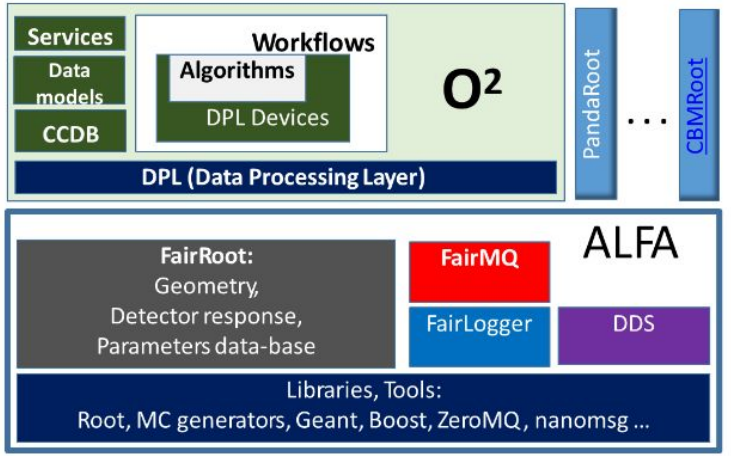

Simulation -> High parallelization (GRID: 8-core queue since last year)

Reco -> sync + async reco with CPU+**GPU** -> EPN farm (GRID with CPU)

Analysis -> O^2+ Apache Arrow Table

M. Concas (CCR 2021)

Using GPU in Run 3 (reconstruction) with ALICE O^2

Processing on dedicated farm at experimental site

- 250x Event Processing Nodes (EPNs) with 2x32 core CPUs
- 8x AMD Graphic Processing Units (GPUs)
- ~1600 GPUs required to process 50 kHz Pb-Pb collisions

→**GPU usage is mandatory for sync reconstruction and calibration**

➢ All GPU software written in a generic way

➢ Same software runs on GPUs of different vendors and on the CPU

Generic software and GPU Benchmarks

Vendor/architecture-independent software:

- All algorithms are written in generic C++, and can be dispatched to HIP, CUDA, OpenCL on GPUs or OpenMP on CPUs using small wrappers → good code maintainability
- GPU libraries linked dynamically on demand → can distribute same binary software to CPU and GPU nodes

Benchmarking of the synchronous software completed in August 2020:

- GPU performance @ 50kHz Pb-Pb
  - ~1600 AMD MI50 and ~1100 NVIDIA Quadro RTX 6000
  - Compatible with our previous estimates <2000 GPU including 20% margin
- GPU Memory optimization
  - 128 orbit TF (~ 11 ms) needs 24 GB
- EPN Full System Tests performed with 70 orbit TF
  - Validated processing rate of 1/230 of assumed rate at 50 kHz Pb-Pb (nominal 1/250)
  - Max. server memory consumption 280 GB and CPU load 44 cores (+20% in the final setup)

ALICE reconstruction

SYNC

- Rough corrections/calibrations for all detectors
- Full reconstruction of TPC (data reduction on GPU + space distortion corrections)
- TPC – ITS tracks matching (for a small subsample)
- Tracks propagation to outer detectors (TRD, TOF)
- Global track fits
- Primary and secondary vertices
- PID hypothesis
- **CTF and calibrations as output**

When the EPN farm is not (fully) used for synch. processing, it will be used for asynch. processing of the raw data stored on the disk buffer
EPN will perform ~1/3 of the Pb-Pb asynchronous processing

ASYNC

- Full correction of TPC distortions (nominal resolution), full calibration for all detectors
- TPC – ITS tracks matching
- Tracks propagation to outer detectors (TRD, TOF)
- Global track fits
- Primary and secondary vertices
- PID hypothesis
- **Calibration/QC and AOD as output**

➢ Different relative importance of GPU / CPU algorithms compared to synchronous processing

➢ TPC part faster than in synchronous processing (less hits, no clustering, no compression

Reconstruction steps for GPU-offload

Sync and Async reconstruction
mainly focused on central barrel → TPC reconstruction (based scenario)
→optimizing resource usage is mandatory also in async!

Synchronous reconstruction

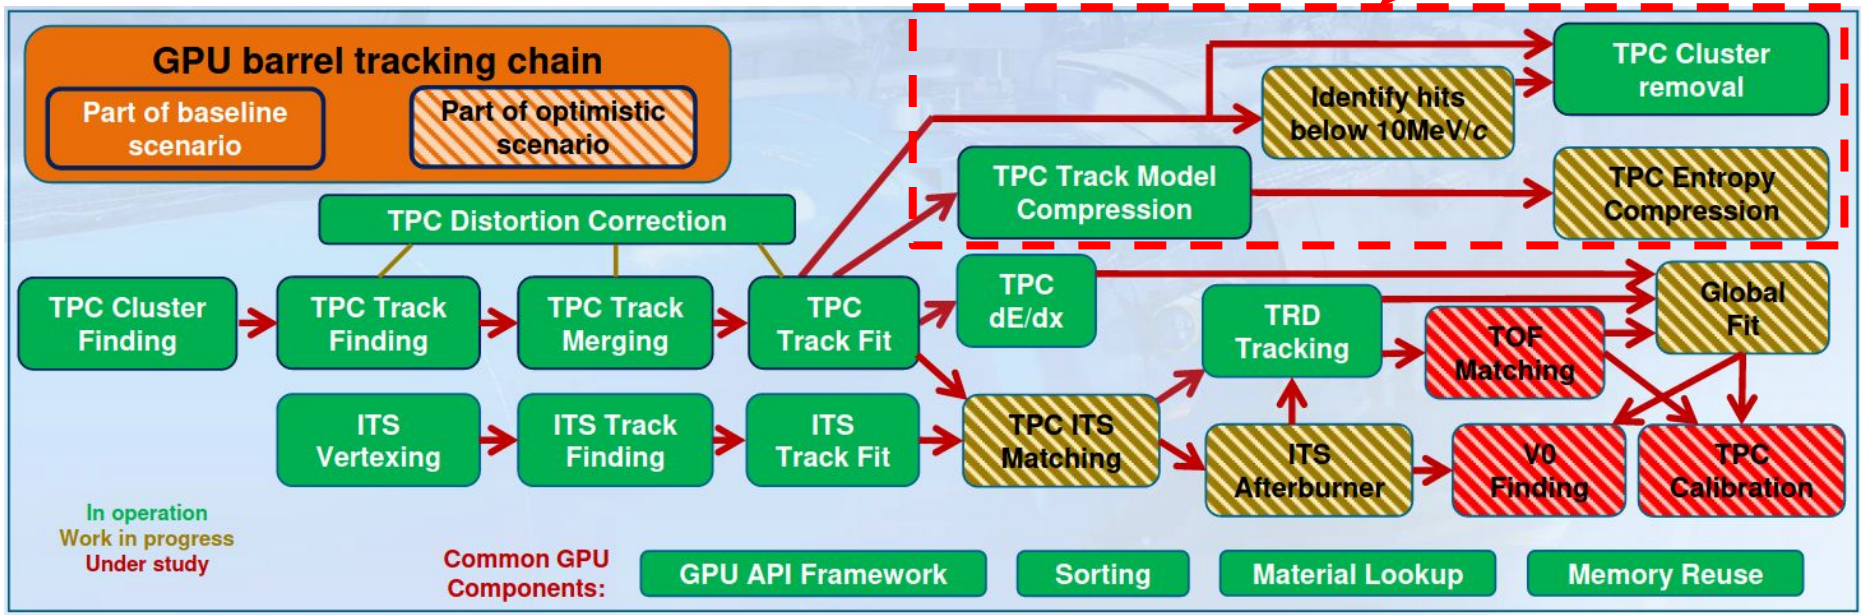

D. Rhor (CHEP 2021)

Reconstruction time covered by GPUs

room of improvements in async reconstruction

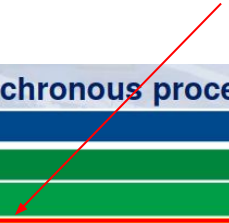

Synchronous processing

Processing step	% of time
TPC Processing	99.37 %
EMCAL Processing	0.20 %
ITS Processing	0.10 %
TPC Entropy Coder	0.10 %
ITS-TPC Matching	0.09 %
MFT Processing	0.02 %
TOF Processing	0.01 %
TOF Global Matching	0.01 %
PHOS / CPV Entropy Coder	0.01 %
ITS Entropy Coder	0.01 %
FIT Entropy Coder	0.01 %
TOF Entropy Coder	0.01 %
MFT Entropy Coder	0.01 %
TPC Calibration residual extraction	0.01 %
TOF Processing	0.01 %

Asynchronous processing

Processing step	% of time
TPC Processing	72.01 %
TRD Tracking	12.69 %
TOF-TPC Matching	9.94 %
MFT Tracking	1.69 %
ITS Tracking	0.78 %
TPC Entropy Decoder	0.73 %
Secondary Vertexing	0.69 %
ITS-TPC Matching	0.56 %
Primary Vertexing	0.14 %
TOF Global Matching	0.11 %
PHOS / CPV Entropy Decoder	0.10 %
FIT Entropy Decoder	0.10 %
ITS Entropy Decoder	0.06 %
TOF Entropy Decoder	0.05 %
MFT Entropy Decoder	0.05 %

Running on GPU in baseline scenario

Running on GPU in optimistic scenario

Preliminary numbers: some algorithms not yet complete or not optimized!

D. Rhor (CHEP 2021)

Reconstruction time of O2 GPU Synchronous Reconstruction on GPU

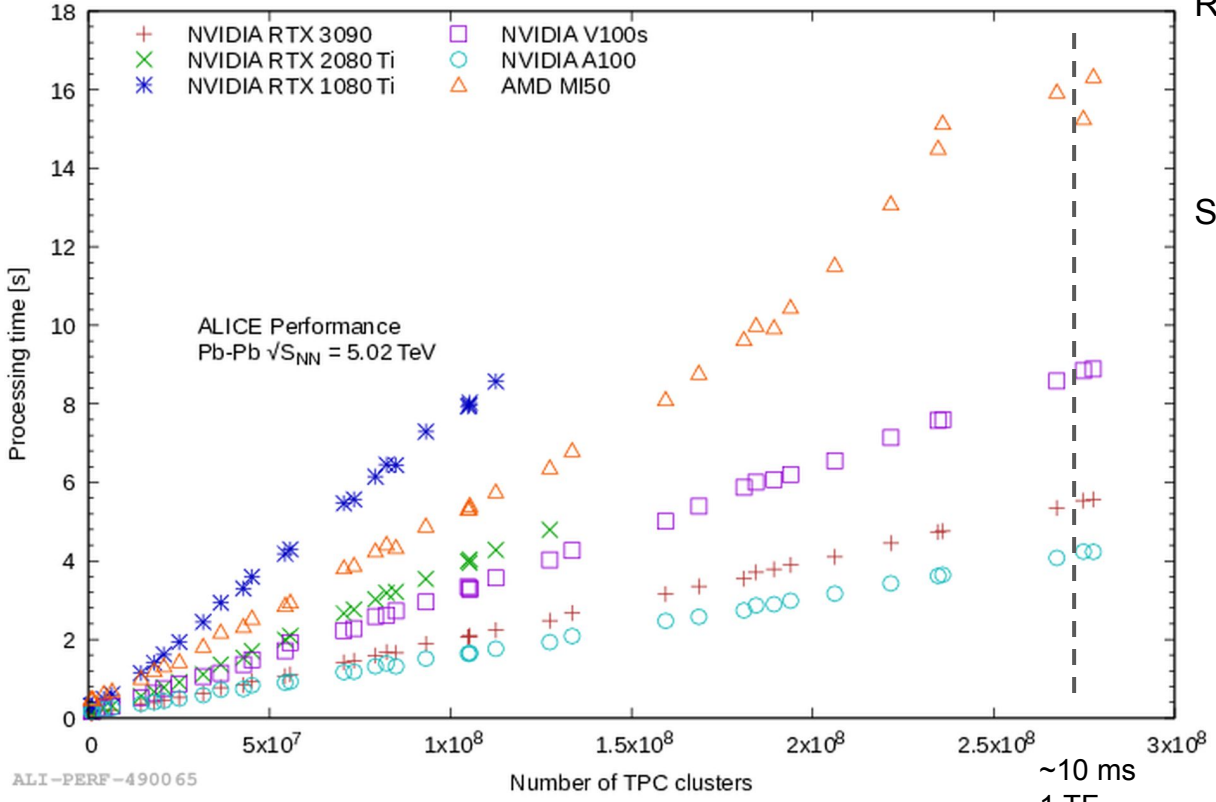

~10 ms
1 TF

Runs only EPN servers:

- 1x Supermicro server, 512GB of memory
- 2x 32cores Rome CPUs
- 8x AMD MI50 GPUs

Standalone tests for GPU performance:

- **Time linearly scales with the size of processed data**

Speedup of O2 GPU synchronous reconstruction versus CPU

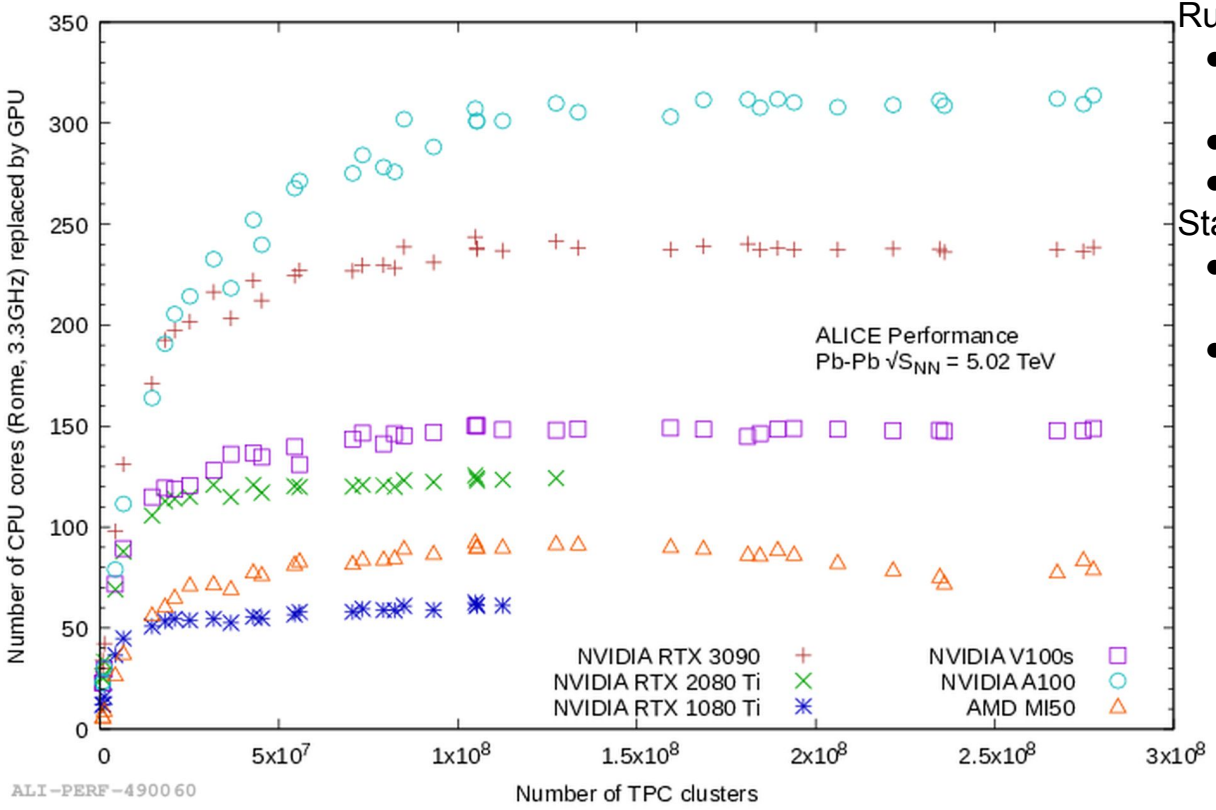

Runs only EPN servers:

- 1x Supermicro server, 512GB of memory
- 2x 32cores Rome CPUs
- 8x AMD MI50 GPUs

Standalone tests for GPU performance:

- Time linearly scales with the size of processed data
- **1x AMD M150 GPU replaces up to 80 Rome cores**

Number of GPUs required to run O2 GPU synchronous reconstruction at 50 kHz Pb-Pb

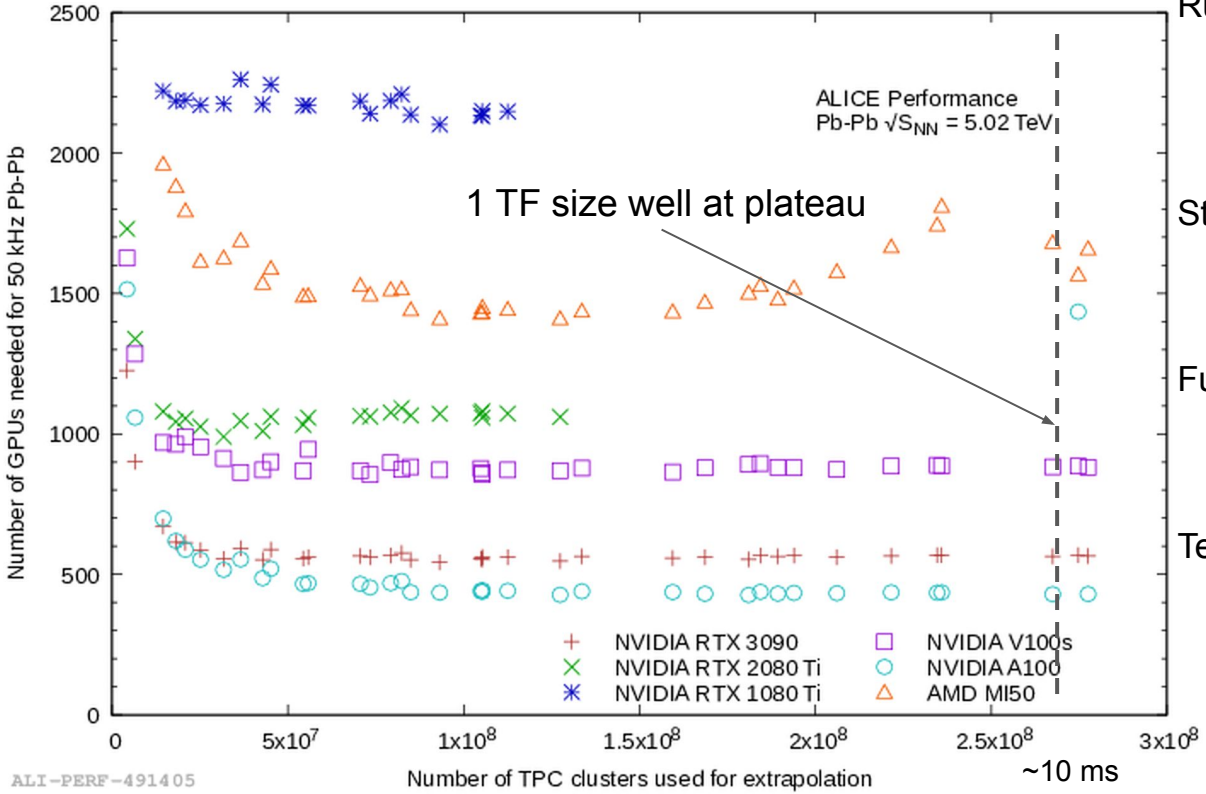

ALI-PERF-491405

~10 ms
1 TF

Runs only EPN servers:

- 1x Supermicro server, 512GB of memory
- 2x 32cores Rome CPUs
- 8x AMD MI50 GPUs

Standalone tests for GPU performance:

- 1x GPU replaces up to 80 Rome cores
- Time linearly scales with the size of processed data

Full system test is started for synchronous

- Replay of fraction of the rate expected in PbPb
- Server sustains the rate

Tested the expected rate for 230 servers

- **capable to process Pb-Pb data at 50 kHz with 20% margin (assuming 2000 GPUs)!**

HPC and cloud resources (simulation)

- Thanks to the new O^2 simulation and reconstruction code (Run 3) possible to fully exploit the multi process features
- Significant progress has been made to incorporate HPC and cloud resources in the standard ALICE Grid workflows
  - Multicore queues at CERN used to test and benchmark the O^2 MC code
  - Intel based HPCs (Marconi @ CINECA, Cori and Lawrencium @ LBNL) were used for the O^2 MC challenge
  - Cloud resources delivery at CERN - direct integration of Azure cloud as a Grid node
- Next steps:
  - porting the O^2 code to Power 9 and ARM platform

O^2 simulation

We definitely entered in the Run 3 phase: O^2 simulations replaced Run 2 ones in the GRID.

We recently simulate 1 billion pp@13.6 TeV events. Even if we are still collecting metrics the improvement with respect to Run 2 data was confirmed to be better than a factor 3 → impact on resources per simulated event is lower than Run 2 case both for Wall Time and disk usage.

Additional optimizations in place to speed up simulations: e.g. embedding in digitization → S. Wenzel (CHEP 2018)

Usage of GPUs not feasible → Geant4 transportation code is still the dominant component

Summary

ALICE will record 50 kHz Pb-Pb minimum bias data in Run 3 without trigger.

- Continuous TPC readout, time frames of ~10 (or ~20) ms instead of events.

Full online data processing on GPUs.

- Computing farm consists of 250 servers, with 8 AMD MI50 GPUs, 2 32-core Rome CPUs, and 512 GB RAM each.
- Currently 230 servers are sufficient for processing 50 kHz Pb-Pb (peak load).
- MI50 GPU replaces ~80 CPU cores (sync reco) or ~55 CPU cores (async reco)

All GPU software written in generic way, can run on different GPUs and on the CPU.

Processing farm used for synchronous (online) and asynchronous processing (periods without beam).

- Full baseline scenario with synchronous GPU processing ready.
- Planning to use GPUs as much as possible also in asynchronous processing.
    - In the optimistic scenario, we will be able to offload ~95% of the workload to the GPU.

Usage of GPUs in GRID nodes not yet explored (reco is performed at Tier-1s), is it something foreseen in INFN plans?

backup

Data buffer for Online-Offline (O^2) facility

- 100 PB raw capacity, RS 10+2 erasure coded (level of security to be defined, 84 PB usable space)
- Based on cheap JBODs, SATA drives, EOS managed

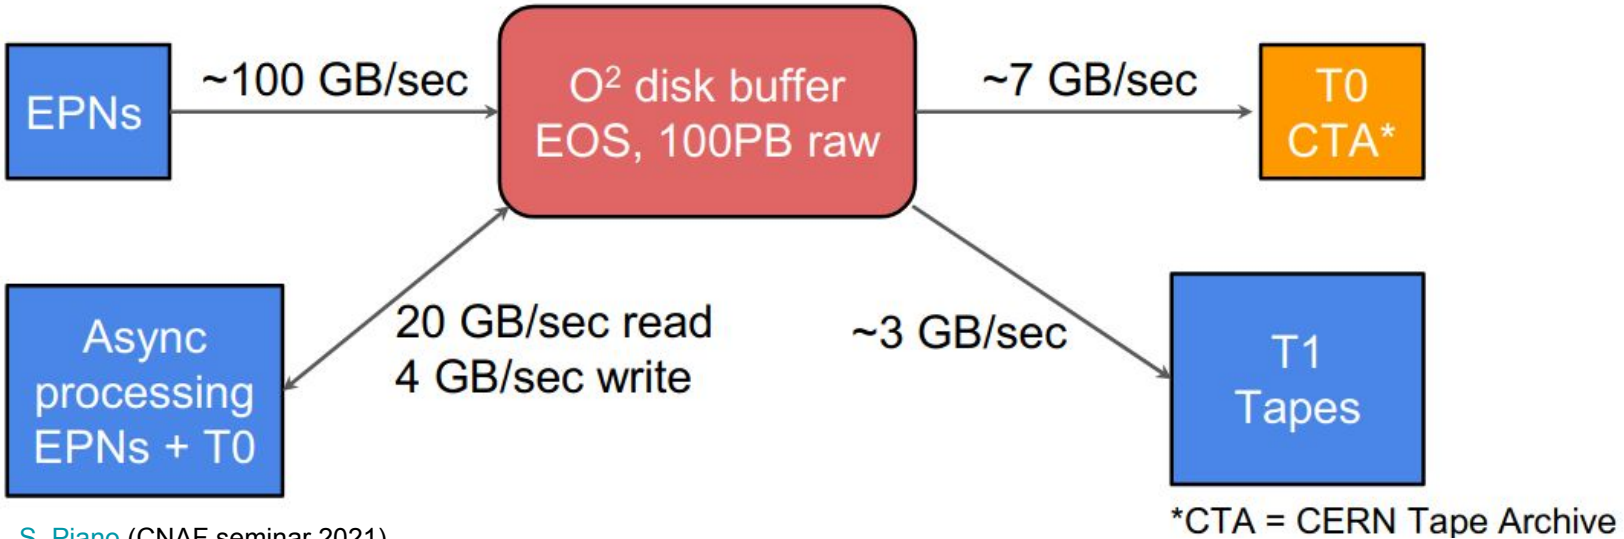

S. Piano (CNAF seminar 2021)

Generic software and GPU Benchmarks

Specifications of EPN defined (EPN PRR on Aug. 21st 2020)

Server choice SuperMicro AS-4124GS-TNR with

- **Eight GPUs AMD MI50 – 32 GB** / Two CPUs 32 core AMD Rome 7452
- 512 GB 3200 Hz Main Memory, 200 GB SSD plus 4 TB SSD
- 100 Gb/s InfiniBand HCA

(EPN delivery on Dec. 18th 2020)

250 servers totalling 2000 GPUs and 16000 physical cores

Benchmarking of the synchronous software completed in August 2020:

- GPU performance @ 50kHz Pb-Pb
  - ~1600 AMD MI50 and ~1100 NVIDIA Quadro RTX 6000
  - Compatible with our previous estimates <2000 GPU including 20% margin
- GPU Memory optimization
  - 128 orbit TF (~ 11 ms) needs 24 GB
- EPN Full System Tests performed with 70 orbit TF
  - Validated processing rate of 1/230 of assumed rate at 50 kHz Pb-Pb (nominal 1/250)
  - Max. server memory consumption 280 GB and CPU load 44 cores (+20% in the final setup)

Reconstruction time of O2 GPU Synchronous Reconstruction on GPU and CPU

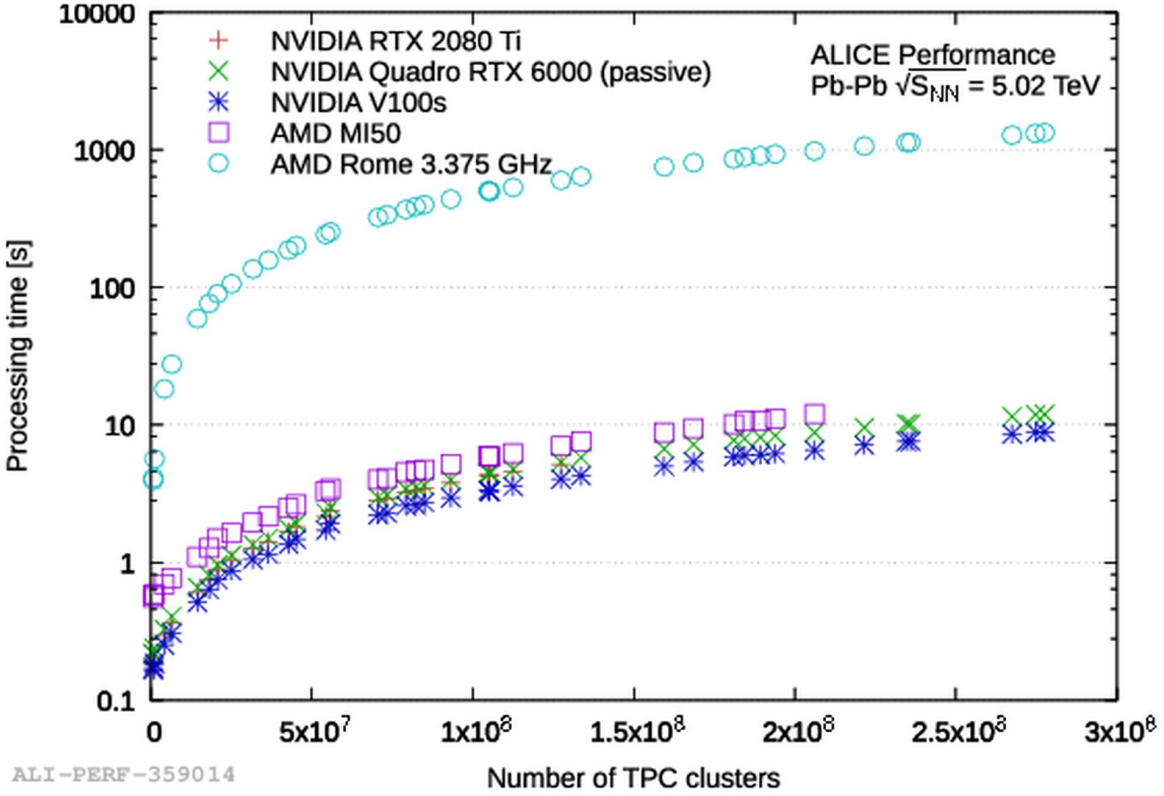

ALI-PERF-359014

Runs only EPN servers:

- 1x Supermicro server, 512GB of memory
- 2x 32cores Rome CPUs
- 8x AMD MI50 GPUs

Standalone tests for GPU performance:

- Time linearly scales with the size of processed data
- **GPUs are much faster than CPUs**