



Contribution ID: 41

Type: Oral

Overview and 2022 data taking the experience of ALICE online and offline processing in Run 3

ALICE has upgraded many of its detectors for LHC Run 3 to operate in continuous readout mode recording Pb-Pb collisions at 50 kHz interaction rate without trigger.

The computing infrastructure was upgraded accordingly, and new reconstruction software has been developed to process the data in real time at rates 50 times higher than during Run 2.

ALICE has put a new computing scheme in place, using the same software, and the same computing farm for synchronous online reconstruction during data taking and for asynchronous reconstruction when there is no beam in the LHC.

The synchronous reconstruction performs mostly detector calibration, QA, and raw data compression, while the asynchronous reconstruction yields the AODs with the final reconstruction output.

To handle the large data rates of Run 3, ALICE leverages heavily the compute power of GPUs.

This synchronous processing is dominated by the TPC, which produces by far the largest data volume, and TPC reconstruction is fully running on GPUs.

When the online computing farm is not used for synchronous processing while there is no stable beam, it runs asynchronous processing jobs.

Since the majority of the compute performance of the online farm is in the GPUs, and since the asynchronous processing is not dominated by the TPC in the way the synchronous processing is, there is an ongoing effort to offload a significant amount of compute load from other detectors to the GPU as well.

The talk will give an overview of the ALICE computing scheme and the experience from 2022 data taking.

We give an overview of the current state and the plans for the asynchronous reconstruction, and the current performance of synchronous and asynchronous reconstruction with GPUs for pp and Pb-Pb data.

Primary author: Dr ROHR, David (CERN)

Presenter: Dr ROHR, David (CERN)