

Contribution ID: 566

Type: Parallel

The new ALICE asynchronous software trigger processing

During the LHC Run 3 data taking period, ALICE is reading out a factor of 600 more proton–proton collisions compared to Run 2, generating a data stream to the CERN T0 of over 30 GB/s —the highest among all LHC experiments. This dramatic increase was made possible through major upgrades to both the detector systems and the underlying data processing infrastructure.

The full data stream is continuously compressed and written to disk, then passed through the ALICE calibration and reconstruction pipeline. A flexible software trigger is applied afterward to identify and retain the relevant events for analysis.

Thanks to this efficient processing model, in 2024, ALICE collected and processed over 53 pb⁻¹ of data —more than ever before in the experiment's history. From an initial 200 PB of raw data, just 8 PB were retained for analysis. Remarkably, the full trigger chain —including final physics reconstruction —was completed within just eight weeks of data taking, with datasets already available for analysis. This enabled a broad range of results based on 2024 data to be presented at conferences that same year.

This presentation will provide an overview of ALICE's upgraded processing chain with a focus on the new asynchrounous software trigger strategy, and demonstrate how these innovations are maximizing the experiment's physics reach in Runs 3 and 4.

Secondary track

Author: COLLABORATION, ALICE

Session Classification: T12

Track Classification: T12 - Data Handling and Computing