

ALICE data processing for Run 3 and Run 4 at the LHC

Chiara Zampolli for the ALICE Collaboration^{a,*}

^a*CERN,*

Espl. des Particules 1, 1211 Meyrin, Switzerland

E-mail: Chiara.Zampolli@cern.ch

During the upcoming RUN 3 and RUN 4 of the LHC, ALICE will take data at a peak Pb–Pb collision rate of 50 kHz. This will be made possible thanks to the upgrade of the main tracking detectors of the experiment, and with a new data processing strategy. In order to collect the statistics needed for the precise measurements that ALICE aims at, a continuous readout will be adopted. This brings about the challenge of handling unprecedented data rates. The 3.5 TB/s of raw data from the detectors will be reduced to about 600 GB/s on the First Level Processing (FLP) nodes, and sent to the Event Processing layer for further processing and reduction to ~100 GB/s of data to be stored permanently. This synchronous processing stage, which will include reconstruction, calibration and compression procedures, will be followed by an asynchronous one to account for final calibrations. Quality Control (QC) will be intensively used in all the processing stages. This contribution illustrates the processing flow for ALICE in RUN 3 and RUN 4, with emphasis on the components of the synchronous processing. The chosen software design will be described. An overview of the data analysis framework is included as well.

*40th International Conference on High Energy physics - ICHEP2020
July 28 - August 6, 2020
Prague, Czech Republic (virtual meeting)*

*Speaker

1. ALICE upgrade

A Large Ion Collider Experiment [1], [2] is the experiment at the CERN LHC dedicated to the study of heavy-ion collisions. In order to cope with the very large track density of such environment, which reaches ~ 2000 charged particles per unit rapidity in the most central collisions [3], ALICE is endowed with optimized and precise tracking detectors, covering the pseudorapidity region $|\eta| < 0.9$ and full azimuth, which allow for track reconstruction down to very low momenta, of the order of a few hundreds of MeV. The main ALICE tracking device in the central region is the TPC, a large Time Projection Chamber based on multi-proportional wire chambers, recording up to 159 reconstructed points for charged tracks. Thanks to the diverse particle identification capabilities of many of its detectors, ALICE identifies hadrons from the very low momenta up to a few GeV/c. During the very successful RUN 2 data taking, ALICE collected $\sim 1 \text{ nb}^{-1}$ Pb–Pb events, running at an Interaction Rate (IR) up to $\sim 10 \text{ kHz}$, limited by the Time Projection Chamber readout capability ($< 3.5 \text{ kHz}$) and the bandwidth of the optical links used by the online data acquisition of the experiment ($\sim 80 \text{ GB/s}$). The goal of RUN 3 and RUN 4 is to collect $\sim 10 \text{ nb}^{-1}$ of Pb–Pb collisions at the nominal magnetic field of 0.5 T, in addition to $\sim 3 \text{ nb}^{-1}$ at lower B, $B = 0.2 \text{ T}$. In order to be able to record this large volume of data, ALICE will run at a peak interaction rate of 50 kHz, and will undergo very important upgrades involving the data taking paradigm, the detectors, and the data processing [5], [4], [6]. The RUN 1 and RUN 2 trigger-based data taking will be replaced by a continuous readout mode both for pp and heavy-ion collisions, which imposes drastic changes in some of the detectors. The upgrade of the ALICE Inner Tracking System (ITS) [7], the TPC [8], [9] and the new Muon Forward Tracker (MFT) [10] are discussed in the contribution [11]. Never the less, here we are going to briefly summarize the main changes that the ALICE TPC will go through, since these are the key elements to understand the data processing chain of the experiment in RUN 3 and RUN 4.

In order to be operated in continuous data taking mode, the ALICE TPC will replace the multi-proportional wire chambers (having the strong limitation of applying a gating grid on a trigger signal to prevent ions produced during the amplification process to enter the gas volume) with Gas Electron Multiplier (GEM) readout chambers. The final configuration of the TPC was defined after intense R&D studies, and led to GEM stacks made of 4 layers working under highly optimized high voltage settings [12], which ensure an ion back flow, a tracking performance and energy loss determination as defined in the ALICE upgrade requirements [8].

The data processing strategy adopted by ALICE in RUN 1 and RUN 2 won't fit anymore the data taking conditions of the experiment. The extremely large data size that will be read out by the detectors ($\sim 3.5 \text{ TB/s}$) dictates the need to adopt data reduction techniques during the data taking itself, synchronously. A total compression factor of 35 will be needed in order to be able to store on a 60 PB buffer disk the data collected during a Pb–Pb data taking period (few weeks per year), achievable only if, in addition to data format optimization techniques strongly improved and adapted with respect to RUN 1 and RUN 2 (e. g. zero suppression, entropy reduction . . . [13]), the output of online reconstruction and calibration will also be used. The ALICE Online-Offline (O^2) processing is the synergy of the two online and offline worlds that existed as separate entities in RUN 1 and RUN 2, and will work as a whole during RUN 3 and RUN 4.

2. Synchronous versus asynchronous data processing

The ALICE RUN 3 and RUN 4 data processing flow can be divided in two phases: a synchronous processing, occurring during the data taking, and an asynchronous one, happening with a few weeks delay with respect to the data taking. The ultimate goal of the synchronous processing will be to reduce the size of the incoming data from the detectors' readout links such to be able to store them on disk.

Data will come at a speed of 3.5 TB/s to the First Layer Processing (FLP) nodes, which constitute the first processing level of the O² farm. Here, zero suppression will happen inside FPGA-based readout cards (Common Readout Unit, CRU, [5], [14]). The data coming continuously from the detectors will be chopped on each FLP into 128 – 256 orbits (corresponding to ~10–20 ms) data packets called sub-Time Frames (TF). The name refers to the fact that they will contain the data coming from only a subset of the readout links of the experiment. At this stage, the data rate will be reduced by a factor of < 6, with ~ 0.6 TB/s reaching the second processing level, happening on the Event Processing Nodes (EPN). On the EPNs, the full TF will be built, merging the data from the different sub-time frames. Synchronous reconstruction, calibration, and data compression will be performed, resulting in the Compressed Time Frames (CTFs) stored on a 60 PB disk buffer at a data rate of < 0.1 TB/s. The total reduction during synchronous processing results in a factor of 35.

The O² farm will consist of 200 FLPs, and 250 EPN servers hosting each 8 GPUs and 64 CPU cores. Such computing power was estimated to be sufficient to cope with the peak data taking rate of 50 kHz expected during the Pb–Pb data taking.

After a time interval of between 2 and 4 weeks, during which refined calibrations not crucial for data reduction will happen, a second, and eventually third reconstruction step will take place, in the so called asynchronous processing phase. The reconstruction load will be shared equally between the O² farm and the Tier 0 (T0) and Tier 1 (T1) grid computing nodes. The output of the asynchronous reconstruction will be stored in Analysis Object Data (AODs) that will be saved to permanent storage and will be the input for physics analysis, while CTFs, previously transferred to the T0 and T1's, will be deleted from the disk buffer to make space for new incoming data. Figure 1 shows schematically the ALICE processing stages for RUN 3 and RUN 4.

It is worth to note that the same type of processing will be adopted for both pp and heavy-ion collisions, with the additional implementation of software, physics-oriented triggers during the asynchronous reconstruction in the case of pp, to be able to cope with the expected data sizes and rates. This means that, for pp data, CTFs will be rewritten to contain only the interesting events at the end of the asynchronous stage, when also AODs will be produced.

2.1 TPC calibration and reconstruction

As mentioned before, the goal of the synchronous reconstruction will be to reach a factor of 35 in data reduction. The main concerned detector will be the TPC, which will have to go down from ~3.4 TB/s to ~70 GB/s. The TPC data compression will consist of different ingredients, among which zero suppression, clusterization and optimized cluster data format, entropy reduction and encoding [13], and tracking (aimed at removing clusters not associated to signal tracks and further improve the data format used for clusters). This challenging goal will be attained through

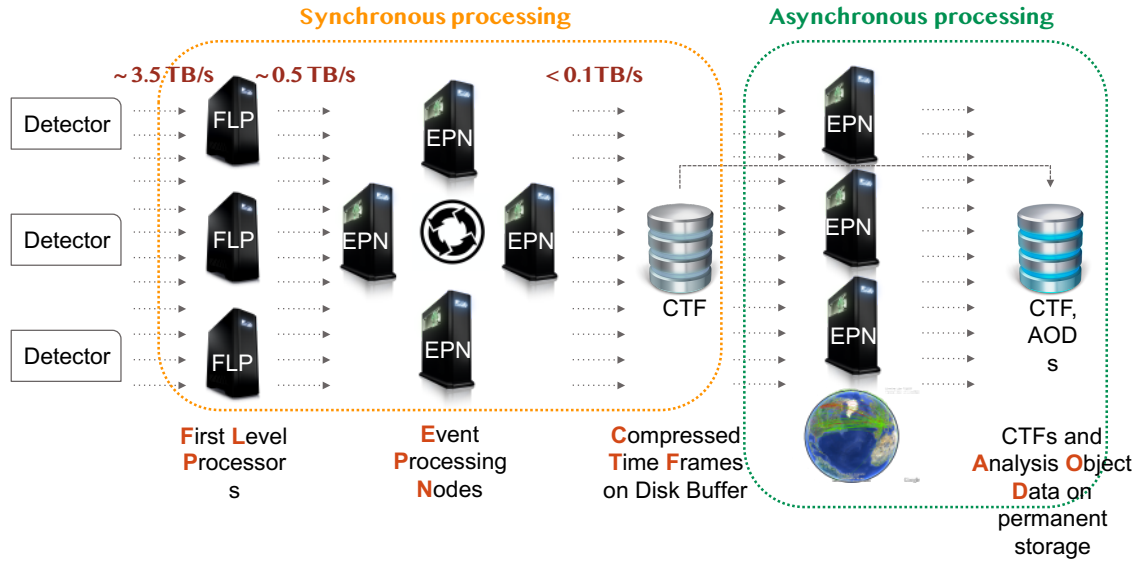


Figure 1: Synchronous and asynchronous processing stages for the ALICE data taking in RUN 3 and RUN 4. Data rates refer to the peak interaction rate foreseen in Pb–Pb collisions, 50 kHz.

an intense use of GPUs, which ensure the necessary computing power at limited costs to process the data synchronously [15], and which will also be exploited during the asynchronous stage.

Despite of the fact that the design and configuration of the TPC GEMs will reduce the ion back flow to the minimum, positive charge will never the less accumulate and move in the TPC drift volume, modifying the electric field, and resulting in space point distortions affecting the TPC clusters. This effect was seen already in RUN 1 and RUN 2 [8], even though with different magnitudes and patterns in the r and z directions and $dr\phi^1$ from those expected in RUN 3 and RUN 4. The main ingredient to the calibration strategy will be the same as in RUN 2 [16], i.e. the ITS-TPC-TRD (Transition Radiation Detector)-TOF (Time Of Flight) track interpolation. Such processing will happen during the synchronous phase on a small fraction ($< 1\%$) of the tracks as a preparatory step for the calibration to be applied during the final reconstruction, and it will rely on the tracking in the central barrel of ALICE, which will be one of the main challenges in RUN 3 and RUN 4 for the experiment: at the IR of 50 kHz five collisions will overlap in the TPC drift time ($\sim 100 \mu\text{s}$), and clusters will be produced in continuous readout mode with an absolute time without a well defined associated z coordinate, and will suffer from the aforementioned position-dependent distortions. The description of the track reconstruction strategy for RUN 3 and RUN 4 is detailed in [17]. While the track-based interpolation between detectors, requiring to collect and process statistics beyond the synchronous timescale, will serve as a correction for the average effects of the space point distortions at the asynchronous stage, during the synchronous phase, a reference space point distortion map (from previous reference data) will be scaled by the actual luminosity to provide the needed corrections. Additionally, fluctuations originating from variations in the number of effective instantaneous event and track multiplicity, will not be caught by neither procedure, but will need to be corrected for. This will be critical during the asynchronous processing to guarantee

¹The r direction is the radial one, while z is the longitudinal direction along the beam.

the optimal performance of the TPC, with an intrinsic resolution of $\sim 200 \mu\text{m}$. To account for fluctuations, the history of the measured charge at the readout plane (digital currents) will be used with different granularity for the synchronous and asynchronous reconstruction stages, as it can be translated into number and position of ions, and, in turn, space charge and distortions.

3. O² processing model and analysis framework

No matter what will be the processing step, the same software framework will be used during RUN 3 and RUN 4, developed in collaboration with the FAIR Software Group at GSI, and inspired by the ALICE RUN 1 and RUN 2 HLT [18]. The software framework is based on three layers. The first one is the Transport Layer, which defines that the data will be transported as of FairMQ messages through the different devices building the ALICE software pattern (topology). This guarantees a low level generalization of the software architecture which will be flexible and suitable for different network configurations, with the advantage of a shared memory to be used if more devices sit on the same node. The second layer is the Data model, and it defines the detailed description of the messages to be passed, which is computer language agnostic, and extendible. It allows for multiple data formats and serialization methods, like custom data structure that may be GPU oriented (useful e.g. for reconstruction), ROOT objects (for Quality Control, QC), and Apache Arrow (for analysis). Finally, the Data Processing Layer (DPL) serves as a sort of translator of the user's computational problem into a low level topology of devices exchanging messages.

As all the other processing elements (reconstruction, calibration, QC, simulation), the ALICE analysis framework will be based on the O² processing model. The AOD data will be organized in flat tables (using the Apache Arrow layout behind a classic C++ API) to minimize I/O costs, and allow for exploitation of vectorization and parallelization. The analysis core will be expressed in the form of a task, translated by DPL into FairMQ topologies and devices, where filters, selections and data manipulation will be possible, and table will be concatenated depending on the user needs. Finally, the analysis output will be based on ROOT.

Due to the very high number of data that will be collected (100 more collisions to be analyzed with respect to RUN 1 and RUN 2), new analysis strategies with respect to the present ones will be adopted in RUN 3 and RUN 4. Physics-oriented skimming campaigns will allow to selectively store collisions and tracks and reduce the size of interesting data. Besides, an Analysis Facility of 20000 cores will be used for fast analysis validation before the full submission on the entire data sample on the Grid. An organized approach like in RUN 1 and RUN 2 will be taken, with analyses organized in "trains" (grouping together analyses running on the same data samples) within the ALICE physics working groups.

4. Conclusions

ALICE is getting ready to a major upgrade for the upcoming RUN 3 and RUN 4, touching not only the data taking paradigm and detectors, but also the data processing at all stages, from reconstruction to analysis. Data will be first processed synchronously, to achieve an impressive size reduction factor of 35, allowing to store the data on a disk buffer for the physics-oriented

reconstruction stage. Online reconstruction and calibration will play a key role, and the TPC will be the crucial detector.

The O² processing framework based on three layers (data transport, data model, and data processing) will be used in all processing components, from reconstruction, to calibration, to analysis. The analysis framework for RUN 3 and RUN 4 is being built around the computing and storage needs imposed by the very large data samples that will be collected, ready to open the door to a new, beautiful physics.

References

- [1] K. Aamodt *et al.* [ALICE], JINST **3** (2008), S08002 doi:10.1088/1748-0221/3/08/S08002.
- [2] B. B. Abelev *et al.* [ALICE], Int. J. Mod. Phys. A **29** (2014), 1430044 doi:10.1142/S0217751X14300440.
- [3] J. Adam *et al.* [ALICE], Phys. Rev. Lett. **116** (2016) no.22, 222302 doi:10.1103/PhysRevLett.116.222302.
- [4] B. Abelev *et al.* [ALICE], J. Phys. G **41** (2014), 087001 doi:10.1088/0954-3899/41/8/087001.
- [5] P. Antonioli *et al.* [ALICE], CERN-LHCC-2013-019.
- [6] P. Buncic, M. Krzewicki and P. Vande Vyvre, CERN-LHCC-2015-006.
- [7] B. Abelev *et al.* [ALICE], J. Phys. G **41** (2014), 087002 doi:10.1088/0954-3899/41/8/087002.
- [8] C. Lippmann [ALICE], CERN-LHCC-2013-020.
- [9] CERN-LHCC-2015-002.
- [10] ALICE Collaboration, CERN-LHCC-2015-001.
- [11] S. Panebianco for the ALICE Collaboration, these proceedings.
- [12] R. Münzer [ALICE], Nucl. Instrum. Meth. A **958** (2020), 162058 doi:10.1016/j.nima.2019.04.012.
- [13] M. Lettrich for the ALICE Collaboration, these proceedings.
- [14] J. Mitra, S. A. Khan, S. Mukherjee and R. Paul, JINST **11** (2016) no.03, C03021 doi:10.1088/1748-0221/11/03/C03021.
- [15] M. Concas for the ALICE Collaboration, these proceedings.
- [16] M. O. Schmidt [ALICE], EPJ Web Conf. **245** (2020), 01003 doi:10.1051/epjconf/202024501003.
- [17] D. Rohr *et al.* [ALICE], [arXiv:1811.11481 [physics.ins-det]].
- [18] D. Rohr *et al.* [ALICE], J. Phys. Conf. Ser. **898** (2017) no.3, 032031 doi:10.1088/1742-6596/898/3/032031.