

Fast Entropy Coding for ALICE Run 3

Thursday, 30 July 2020 09:20 (20 minutes)

In LHC Run 3, the upgraded ALICE detector will record 50kHz Pb-Pb collisions using continuous readout. The resulting stream of raw data at ~3.5TB/s - a fiftyfold increase over Run 2 - must be processed with a set of lossy and lossless compression and data reduction techniques to decrease the data rate to storage to ~100GB/s without affecting the physics. This contribution focuses on lossless entropy coding for ALICE Run 3 data which is the final component in the compression stage. It describes the challenges imposed by the source data with a non-standard, patchy distribution of up to 25Bit per symbol. It then explains how we use rANS, a variant of Asymmetric Numeral System coders for compression with special emphasis on our adaptations and additions to the algorithm. Finally, we will show the first production ready rANS implementation for both CPU and GPU as part of the ALICE O² framework and demonstrate its capability of compressing data close to entropy at rates of up to 750 MB/s.

I read the instructions

Secondary track (number)

Primary author: LETTRICH, Michael (Technische Universitat Munchen (DE))

Presenter: LETTRICH, Michael (Technische Universitat Munchen (DE))

Session Classification: Computing and Data Handling

Track Classification: 14. Computing and Data Handling