

Abstract

In 2021 the Large Hadron Collider at CERN will start its third data taking period, so-called Run 3, which will last until the 2023. During the Long Shutdown 2, the pause between Run 2 and Run 3, the four major experiments at the accelerator (ALICE, ATLAS, CMS, LHCb), are undergoing through major upgrades and maintaining of their layouts. Notably among them, the A Large Hadron Collider Experiment (ALICE) experiment, to improve its measurement capabilities, is not only performing critical improvements on its detector hardware, but also on the software machinery used for the data acquisition and processing. The ALICE experiment aims at studying the Quark-Gluon Plasma, a state of matter where quarks and gluons are not confined into hadrons and that can be inspected and characterised using high-energy ion collisions. Plans for ALICE in Run 3 include the collection of 10 nb⁻¹ of heavy-ion collisions, with an instantaneous luminosity up to $6 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ and a collision rate of 50 kHz at 5.5 TeV, corresponding to a total of 10^{11} interactions recorded. This is the minimum required rate to address the proposed physics programme that focuses on rare probes both at low and high momentum. For the proton collisions programme the experimental apparatus will acquire data with a rate up to 400 kHz of interactions at 13 TeV, to obtain a solid data reference instrumental to the study of Ion collisions. ALICE will also move to a brand new paradigm for what concern the data acquisition for many of its sub-detectors, the so-called continuous readout. It is a trigger-less mode that consists on continuously acquiring readout signals and to process and reconstruct them as a stream of data in an on-line fashion, meaning during the data taking times. Such conditions impose stringent constraints to the detector performances in terms of acquisition rate and challenges for what concerns the output data bandwidth, forcing a general upgrade of many of the aspects of the experiment.

To comply with this ambitious scenario, many interventions are being operated on sub-detectors of the ALICE apparatus, mostly dedicated to the readout electronics and to sensor readout capabilities, to provide a faster acquisition rate compared to working conditions attained during Run 2. Also, two brand new detectors will be deployed in the setup aiming at improving the precision of the measurements: a Muon Forward Tracker (MFT) and a completely renovated Inner Tracking System (ITS) will be placed inside the inner barrel, the innermost part of the detector, in the central barrel volume of the ALICE apparatus. The latter is cylindrical detector built up with seven concentric layers of silicon pixels chips adopting a Monolithic Active Pixels Sensors (MAPS) layout, a new technology for faster and more precise measurements. The upgraded layout is characterised by finer pixel granularity, inner barrel closer to the beam with respect to its predecessor and counts one more layer. These features are translated in a better resolution of the interaction point (tridimensional space point where beams collide) and better resolution for tasks such the reconstruction of the trajectories of charged particles generated in the collisions.

The continuous readout set also critical challenges in terms of online data reduction, compression and reconstruction. It is estimated that the whole online reconstruction workflow will be able to cope with a 3.5 TB/s bandwidth of data in input and will write ~ 100 GB/s of data on persistent storage. To enable the processing of continuous readout data stream it has been developed a completely renewed software stack that includes parts to

be run both online, that is during the data taking, and offline on reconstructed data. It is named after Online-Offline (O^2) framework. Such a new platform is thought to be deployed both on large computing clusters like the Event Processing Node farm (EPN), the main computing centre used for online reconstruction, located on-site close to the ALICE experimental apparatus, but also on smaller computing resources, down to the single workstation of a scientist. It presents a new operational design, compared with former software used in ALICE. It based on workflows operated by generic logical "devices": computing processes that aims at continuously execute some routine depending on designed roles. Devices are then *plugged* each others to create actual pipelines or topologies, where they can communicate and cooperate to execute more complex tasks. The O^2 embodies a multi-process paradigm where each device is represented by a process on a computing resource and the workflows are described by topologies connecting devices following a data-flow models. More in details, each device is responsible to perform a piece of a large routine and the output data are exchanged across devices by mean of messages that can be implemented using different technologies, depending on the underlying computing infrastructure. Every type of workflow in the O^2 is described and implemented using this design, from data simulation via Monte Carlo to the online data reconstruction and even the analysis framework is being developed using the same core approach. The schema os flexible and scalable, each entity can be, if required, replicated up to cope with specific demands and the deployment of a workflow can be done a laptop as well as on a High-Performance Computing cluster, in some cases. These features are completely transparent to the final user, which will be able to agilely work with both of the resources with minimal techincal effort. Ultimately, the framework does support heterogeneous architectures, allowing devices to offload payloads on computing accelerators like Graphic Processing Units (GPUs) and Fully Programmable Field Arrays (FPGAs), to obtain high-throughput computations with a fully integrated data model to support them.

The work in this thesis will present the design, development and implementation of a primary vertex reconstruction algorithm with ITS-only data that will be used by ALICE in the online reconstruction phase during the Run 3. The estimation of the primary vertex position is instrumental for calibrating detectors because it provides useful information on the beam position. More importantly, it is a critical information for some online reconstruction processes such as the ITS tracking, the process that reconstruct the trajectories of charged particles, as a critical starting point for the process. The work presented is integrated in the O^2 framework, and provides both a CPU and a GPU-accelerated parallel version of the algorithm. It is able to identify also the so-called pile up of events, when more than one collision data are present in the same input dataset, and provide their position with a resolution compliant with upgrade requests. Concerning the GPU version, two implementations are presented: on using Nvidia GPUs and one using Advanced Micro Devices (AMD) GPUs. Two version are coded using the two corresponding development frameworks, the Computing Unified Device Architecture (CUDA) for Nvidia and Heterogeneous-Computing Interface for Portability (HIP) from AMD. Consistency checks among three implementations on obtained results are performed, together with performance benchmarks. Time measurements are reported to compare three implementations. The primary vertex reconstruction is proven to be compliant with the O^2 requirements in terms of resolution and time performances.

Knowledge acquired in working on this will be used in future also to further extend the

dominion of the GPU-accelerated workflows in the O^2 context.