

Tracking in high-multiplicity events

Maximiliano Puccio^{*†}

Universita e INFN Torino

E-mail: maximiliano.puccio@cern.ch

The ALICE experiment is preparing a major upgrade of its inner silicon tracker (the Inner Tracking System) and of its Online and Offline systems for the upcoming Run3 of the LHC starting in 2021.

During Run3, LHC will deliver Pb-Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV with a peak luminosity $\mathcal{L} = 6 \times 10^{27} \text{cm}^{-2}\text{s}^{-1}$ and an interaction rate of 50 kHz, to be compared to the 8 kHz design interaction rate currently delivered by the LHC. The aim of ALICE is to cope with such a high interaction rate improving at the same time the resolution and the efficiency of the silicon tracker. In this context, one of the requirements for a prompt calibration of external detectors and to speed up the offline data processing is to run online the reconstruction of tracks in the Upgraded Inner Tracking System.

A new algorithm based on Cellular Automata has been developed to tackle this issue. In this algorithm the tracking is split in multiple phases to profit from data locality. At first, hit points are organised in sectors of azimuthal angle and longitudinal coordinate; then the algorithm looks for track segments within these sectors of the detector, independently. Track segments with compatible track parameters are marked as neighbours. Neighbouring track segments are then merged at the final stage using a set of rules defined by the Cellular Automaton mechanism, somewhat similar to the set of rules used in the Conway's Game of Life.

The obtained computing and tracking performance are compliant with the requirements of ALICE, being able to reconstruct tracks of transverse momentum down to 100 MeV/c in events with high track density ($dN/d\eta$ up to 2000). The tracking and computing performance of this algorithm will be shown in the case of central Pb-Pb events at $\sqrt{s_{NN}} = 5.5$ TeV.

The 25th International workshop on vertex detectors

September 26-30, 2016

La Biodola, Isola d'Elba, ITALY

^{*}Speaker.

[†]On behalf of the ALICE collaboration

1. Introduction

Dealing with a high density track environment is one of the major challenges for the experiments running at the LHC. In a single Pb–Pb collision at the current LHC top energy ($\sqrt{s_{NN}} = 5.02$ TeV) the charged track multiplicity per pseudorapidity unity ($dN_{ch}/d\eta$) can be as high as 2000 in the 5% most central collisions between the lead nuclei [1].

In order to reconstruct the trajectories of all the particles generated in a heavy ion collision there are two main requirements:

- highly granular tracking detectors,
- smart algorithms able to handle high combinatorial complexity.

The focus of these proceedings is on the tracking algorithms used in the reconstruction of high multiplicity events. In particular we will describe the current strategies adopted by the ALICE collaboration, the plans for the data processing and reconstruction during the LHC Run3 and the tracking algorithm for the new silicon tracker of ALICE (the Inner Tracking System Upgrade) that is currently under development.

2. ALICE offline and online tracking during LHC Run2

The ALICE experiment is dedicated to the study of heavy ion collisions and it has been designed to cope with a high multiplicity environment. The granularity required to track particles in heavy ion collisions is provided by a silicon tracker – the Inner Tracking System (ITS) – and by a Time Projection Chamber (TPC) [2].

The ITS is composed by six layers of silicon detectors adopting three different technologies. The ITS layers are located at radii between 4 and 43 cm covering a pseudorapidity range of $|\eta| \leq 0.9$. The first two layers are made of Silicon Pixel Detectors (SPD), the two intermediate layers are composed of Silicon Drift Detectors (SDD) while the outermost layers are made of Silicon Strip Detectors (SSD). An overview of the ITS adopted technologies, operational experience and performance can be found in these proceedings (C. Deplano, "ALICE ITS Operational Experience").

The TPC is partitioned in 36 sectors, 18 azimuthal sections divided in 2 longitudinal halves by the central electrode. Each sector is segmented by pads organised in rows and the longitudinal coordinate is given by the drift time. The TPC covers a pseudorapidity range of $|\eta| \leq 0.9$ while the radial coverage is $85 < r < 247$ cm.

As extensively explained in [3], ALICE developed different approaches for the Offline and the Online track reconstruction:

- Offline reconstruction: a classical Kalman Filter track following algorithm is used to find the tracks in the TPC. Found tracks are then propagated to the ITS updating the track parameters with the information of the ITS reconstructed hits (*recpoints*) being attached to the track. All the ITS recpoints attached to the tracks are marked as used. A ITS standalone algorithm – also based on Kalman Filter – further reconstructs tracks with those clusters not used in the previous step. Thanks to the ITS standalone algorithm, ALICE is able to reconstruct

tracks down to $p_T \sim 80 \text{ MeV}/c$. In the processing of low multiplicity events – such as pp or p–Pb collisions – the ITS tracking is used also to reconstruct tracks independently from the ITS+TPC algorithm using all the ITS recpoints. In high multiplicity events, the ITS standalone algorithm requires tens of seconds of CPU time and therefore it is not run using all the ITS recpoints.

- Online reconstruction: a faster version of the TPC standalone tracking is used to reconstruct tracks in the High Level Trigger (HLT) reconstruction chain [4]. A simplified ITS tracking is also used to reconstruct high p_T tracks that are used to determine the luminous region in ALICE and that can be used to calibrate the drift velocity in the TPC volume.

3. ALICE Upgrade of the Offline and Online reconstruction

A new approach will be used by ALICE during the Run 3 of the LHC, moving Online most of the present Offline reconstruction steps.

The foreseen data throughput is up to 1.1 TB/s of data (Table 1), most of which coming from the upgraded Time Projection Chamber [5]. Such a high data volume can not be written on tape, therefore a strong data volume reduction will be necessary. The estimations of the data volume coming from the ITS upgrade and the Muon Forward Tracker (MFT) are still uncertain since the development of the readout chip still ongoing. In Table 1 the values for ITS upgrade and MFT refers to the baseline of the ITS Technical Design Report (TDR) [6] assuming a random Gaussian noise of the order of 10^{-5} hits per pixel per readout cycle. The TDR estimate revealed to be very pessimistic as the most recent tests show a rate of random hits well below 10^{-6} over a large operational range (see in these proceedings: S. Beolè, "The upgrade of the ALICE ITS").

Detector	Event Size	Pb–Pb at 50 kHz
TPC	20.7 MB	1012 GB/s
ITS	0.8 MB	40 GB/s
TRD	0.5 MB	20 GB/s
MFT	0.2 MB	10 GB/s
Others	0.3 MB	12.2 GB/s

Table 1: Data size per Pb–Pb event and data throughput foreseen for ALICE detectors during LHC Run 3.

Online processing of the data is mandatory to achieve the level of compression necessary and it is a strong requirement of the Online and Offline upgrade program (O²) [7]. The O² project will fulfil this requirement by integrating Data Acquisition (DAQ), HLT and Offline in a single infrastructure that will process online the collected data. The O² infrastructure will be a High Performance Computing (HPC) facility hosting heterogeneous hardware, including:

- 250 First Level Processor worker nodes equipped with FPGAs,
- 1500 Event Processing Nodes equipped with GPUs.

Step	Method	Time fraction	Device
Seeding	Cellular Automaton	$\sim 30\%$	CPU or GPU
Track Following	Kalman Filter	$\sim 60\%$	CPU or GPU
Track Merging	Combinatorics	$\sim 2\%$	CPU
Track fit	Kalman Filter	$\sim 8\%$	CPU

Table 2: Computing time fractions spent in the various step of the TPC online tracking.

As illustrated in [7], the compression factor needed is about 20 and it can be achieved through zero suppression of TPC data and online reconstruction of the TPC hits (compression factor 2.5), removing all the TPC hits not associated to particle tracks (compression factor 6) and using an optimised data format (compression factor 1.35). Therefore ITS+TPC tracking will be required in order to reduce the data size; moreover the ITS tracking and matching to external detector as the Transition Radiation Detector (TRD) and the Time Of Flight (TOF) are necessary to calibrate online the TPC.

3.1 TPC online tracking

The online TPC tracking has been already exercised in the ALICE HLT reconstruction chain during Run 1 and Run 2 of the LHC. This algorithm represents the most relevant example of high performance computing in ALICE and it is based on Cellular Automaton [8].

The Cellular Automaton tracking exploits the native partition in sectors of the TPC. Track segments are reconstructed in each sector – taking advantage from good data locality – and then the track segments are merged and refitted using the Kalman Filter.

The tracking is performed in few steps:

- neighbours finding (figure 1a): for each row of pads k pairs of hits from row $k+2$ and $k-2$. The three hits are required to lay on a straight line. The best pair for each cluster is found and the hit-to-hit links are stored;
- evolution step: the concatenation of subsequent links defines a tracklet (figure 1b) and all the orphan links are removed;
- track segments are created combining tracklets according to compatibility criteria and a simplified Kalman filter is used to fit geometrical trajectories and in case of overlapping track segments the longest is kept. Using the parameters estimated with the Kalman Filter the trajectory is extrapolated to adjacent TPC rows in order to catch the leftover clusters using a classical track following algorithm;
- track merging: tracks in different sectors are merged if their parameters are in agreement;
- track fitting: tracks are refitted and parameters are recomputed using Kalman Filter using the information of all the reconstructed hits attached to the tracks.

In Table 2 the time fractions spent for the various steps are shown. For the computing hot spots – the seeding and the track following – the algorithms have been implemented both for CPU

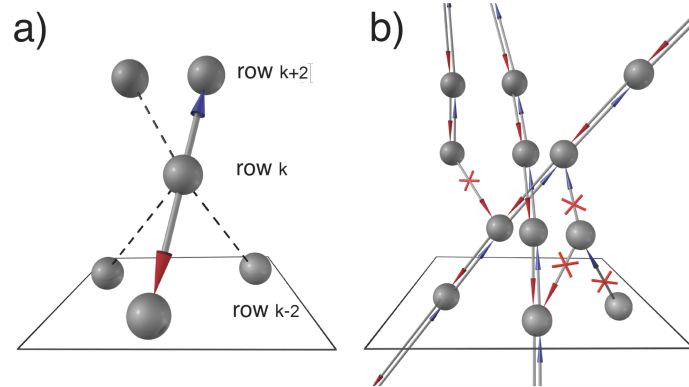


Figure 1: Neighbours finding (a) and evolution (b) steps of the TPC Cellular Automaton tracking.

and for GPU. The CPU version uses OpenMP [9] to enable multithreading while the GPU version has been implemented both in OpenCL [10] and in CUDA [11]. Most of the source code is shared between the different implementations and specific wrappers are used depending on the device in use.

In absolute terms, the GPU implementation of the tracking takes 145 ms to reconstruct a central Pb–Pb events on a AMD S9000 graphics card and using 3 CPU cores to pre-process the data for each sector before the tracking.

3.2 ITS online tracking

A similar algorithm to that used in the online TPC tracking is under development for the tracking of the new ITS. ALICE collaboration will install a complete new silicon tracker during the Long Shutdown 2 of the LHC. The new tracker features 7 silicon pixel layers covering 3 units of pseudorapidity for events with a primary collision vertex located within 10 cm on the longitudinal axis with respect to the nominal collision spot. With respect to the current ITS there is an extension of the reconstruction efficiency at low p_T and a gain of a factor 5 in terms of pointing resolution thanks to the reduced pixel size (from $50 \times 425 \mu\text{m}^2$ to $O(30 \times 30 \mu\text{m}^2)$, the reduced material budget per layer (e.g. from $1.1\% X_0$ to $0.3\% X_0$ for the innermost layers) and an innermost layer closer to the interaction point (from $r = 39$ mm to $r = 23$ mm).

The online tracking of the ITS Upgrade will be required to perform an online calibration of the space charge distortion of the TPC: looking to track point residual between ITS tracks and TPC clusters is possible to reconstruct a distortion map that allows the cluster displacement to be corrected. This method has been already developed and shown to be working offline during LHC Run2, and there are developments ongoing to bring this calibration online.

The baseline for the ITS standalone tracker is an algorithm based on the Cellular Automaton, and it is been already coded for CPU. The new tracking in the ITS upgrade is performed in 4 steps:

- first the position of the primary vertex is found by a fast algorithm,
- a pattern recognition method is used to find track candidates,

- then track candidates are fitted in three passes using the Kalman Filter: from the outermost layer inward, then outward and then refitted inward,
- in case of two or more candidates sharing reconstructed hits only the one with the best χ^2 is kept.

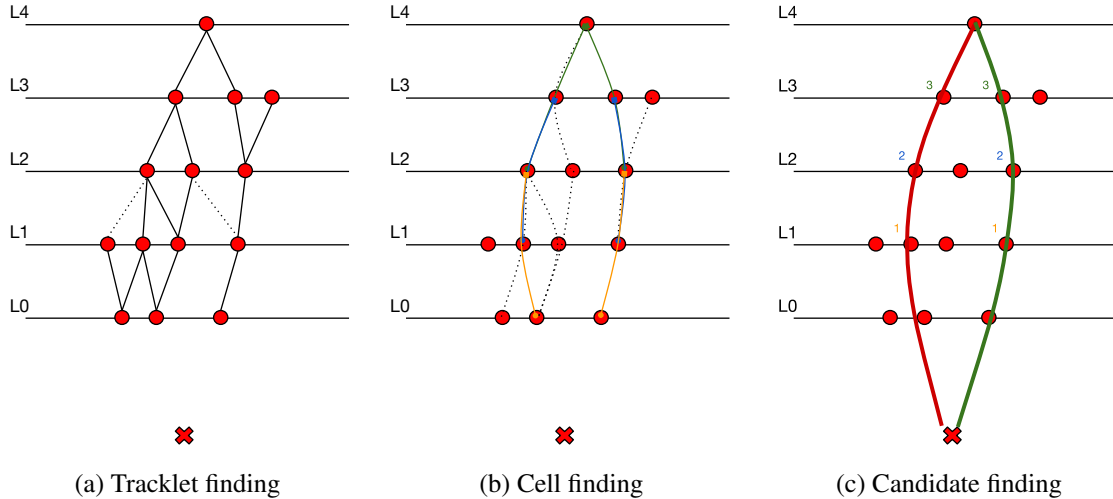


Figure 2: Reconstruction steps of the Cellular Automaton pattern recognition used for the ITS Upgrade tracking. The red cross represent the reconstructed position of the primary collision while the red dots are the reconstructed hits on the ITS Upgrade layers.

This procedure uses standard ALICE methods and algorithms to determine the primary vertex and to fit the track candidates using the Kalman Filter. In the following the details about the pattern recognition are explained. The pattern recognition used is based on the Cellular Automaton algorithm. The tracking algorithm has been implemented on CPU and it is able to reconstruct a central Pb–Pb event simulated with the full geometry of the ITS Upgrade in approximately 0.7 seconds.

The first step before the pattern recognition is the pre-processing of the input data. The reconstructed ITS hits are contained in structure of arrays sorted according the azimuthal coordinate of the hit and then on the z coordinates. Thanks to an index table during the pattern recognition it is possible to access quickly the reconstructed hits in a region of interest of the detector. This approach increases the data locality in the memory and it enables the possibility of parallelisation.

Figure 2 shows the main steps in the tracking algorithm:

- For each cluster on each layer a two dimensional window (in azimuth and z) is opened, then the clusters are joined with those on the next layer within the window forming the so-called *doublets*. Here the z position of the primary vertex is necessary to use the cut on the longitudinal coordinate rejecting part of the combinatorial background;
- doublets spanning on consecutive layers with one cluster in common and with compatible directions are considered to be subsequent and are combined into *cells* (3 points seeds) whose

track parameters are computed. Two cells with compatible track parameters are considered to be connected;

- a continuous sequence of compatible cells represents a track candidate. If two or more candidates share an hit, then the longest is kept, while if they have the same length, then the discrimination is done at the level of track fitting.

Figure 3 shows the tracking efficiency of the ITS Upgrade tracking for particles with one hit on each layers. While the efficiency saturates to almost 100% for p_T higher than 0.6 GeV/c it drops at lower transverse momentum values down to 40% efficiency for 0.1 GeV/c. One third of the reconstructed tracks with momentum 0.1 GeV/c have at least one wrongly associated cluster and it is considered to be fake. The rate of fake tracks drops rapidly to become negligible for p_T higher than 1 GeV/c.

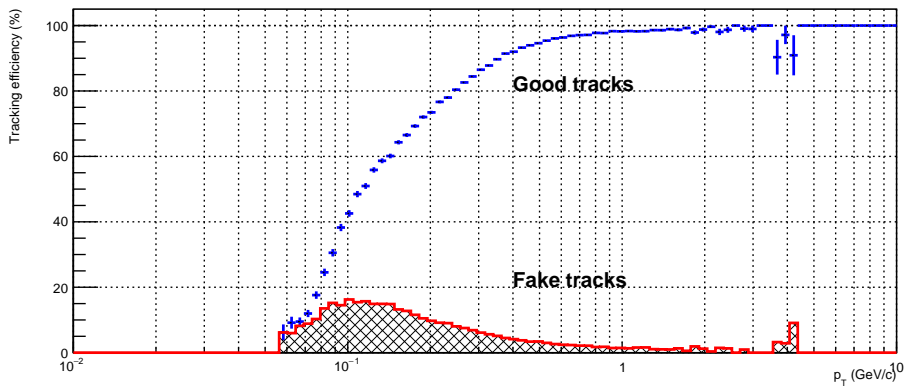


Figure 3: Efficiency and fake track fraction for the ITS upgrade tracking.

4. Conclusions

ALICE experiment already employs high performance computing strategies in its High Level Trigger to reconstruct online events with multiplicity up to 2000 charged particles per pseudorapidity unity. Currently the tracking is performed online only in the Time Projection Chamber and in the Inner Tracking System, but only for high p_T tracks in the latter.

For LHC Run3, the online reconstruction and the tracking will be the key elements to get a sustainable data rate to tape and to have a quick turnaround of the physics analyses. The development of high performance algorithms for the reconstruction of tracks in the ITS Upgrade already started and it is now moving towards the usage of heterogeneous resources such as GPUs.

References

- [1] J. Adam *et al.* [ALICE Collaboration], Phys. Rev. Lett. **116** (2016) no.22, 222302
doi:10.1103/PhysRevLett.116.222302 [arXiv:1512.06104 [nucl-ex]].
- [2] K. Aamodt *et al.* [ALICE Collaboration], JINST **3** (2008) S08002.
doi:10.1088/1748-0221/3/08/S08002
- [3] B. B. Abelev *et al.* [ALICE Collaboration], Int. J. Mod. Phys. A **29** (2014) 1430044
doi:10.1142/S0217751X14300440 [arXiv:1402.4476 [nucl-ex]].
- [4] M. Krzewicki, D. Rohr, S. Gorbunov, T. Breitner, J. Lehrbach, V. Lindenstruth and D. Berzano, J. Phys. Conf. Ser. **664** (2015) no.8, 082023. doi:10.1088/1742-6596/664/8/082023
- [5] B. Abelev *et al.* [ALICE Collaboration], CERN-LHCC-2013-020
- [6] B. Abelev *et al.* [ALICE Collaboration], J. Phys. G **41** (2014) 087002.
doi:10.1088/0954-3899/41/8/087002
- [7] Buncic, P and Krzewicki, M and Vande Vyvre, P, CERN-LHCC-2015-006. ALICE-TDR-019.
- [8] D. Rohr *et al.* [ALICE HLT Collaboration], J. Phys. Conf. Ser. **396** (2012) 012044.
doi:10.1088/1742-6596/396/1/012044
- [9] OpenMP Architecture Review Board, OpenMP Application Program, Interface Version 3.0,
<http://www.openmp.org/mp-documents/spec30.pdf>
- [10] Khronos OpenCL Working Group, The OpenCL Specification, Version 1.2,
<https://www.khronos.org/registry/cl/specs/opencl-1.2.pdf>
- [11] John Nickolls, Ian Buck, Michael Garland, Kevin Skadron, Scalable Parallel Programming with CUDA, ACM Queue, vol. 6 no. 2