Track Reconstruction in the ALICE TPC using GPUs for LHC Run 3

David Rohr for the ALICE Collaboration

drohr@cern.ch, CERN

Connecting the Dots 2018

21.3.2018

ALICE Upgrade for Run 3

- **ALICE uses mainly 3 detectors for tracking: ITS, TPC, TRD + (TOF)**

- **Luminosity will increase significantly in Run 3:**
  500 Hz – 1 kHz trigged → 50 kHz continuous min-bias Pb-Pb.
  - We want to adapt the HLT tracking for Run 3.
  - Ideally, create one tracker as fast as the HLT version with competitive resolution to the offline version.

- **50+ times as many events!**

50+ times as much data?

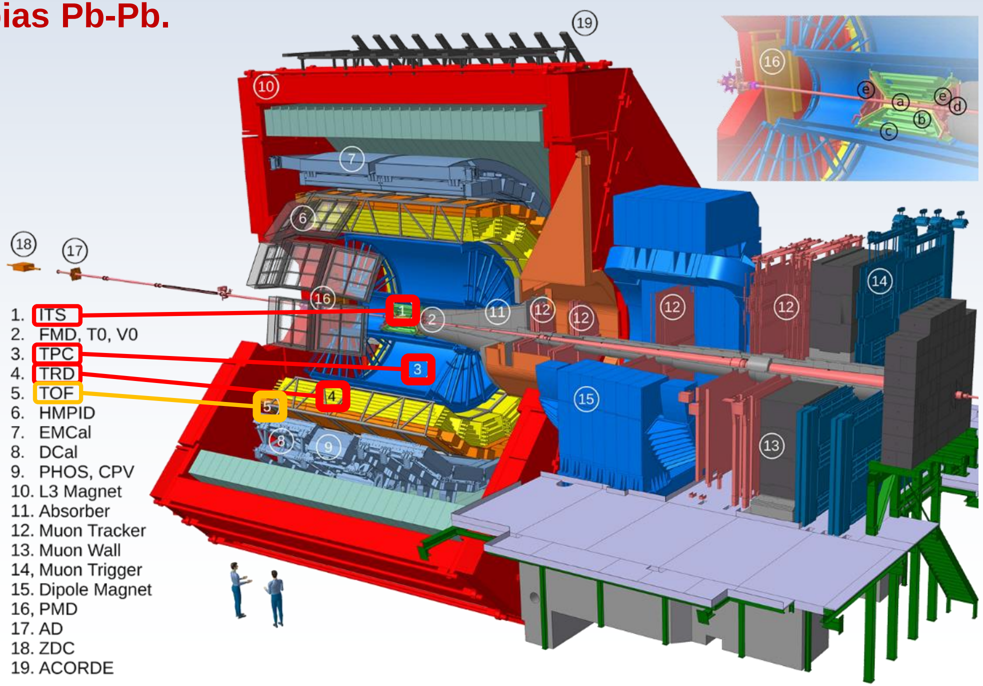

ALICE Upgrade for Run 3

- **ALICE uses mainly 3 detectors for tracking: ITS, TPC, TRD + (TOF)**

- **Luminosity will increase significantly in Run 3:**
  500 Hz – 1 kHz trigged → 50 kHz continuous min-bias Pb-Pb.
  - We want to adapt the HLT tracking for Run 3.
  - Ideally, create one tracker as fast as the HLT version with competitive resolution to the offline version.

- **50+ times as many events!**

50+ times as much data?

We need either a huge stack of disks…

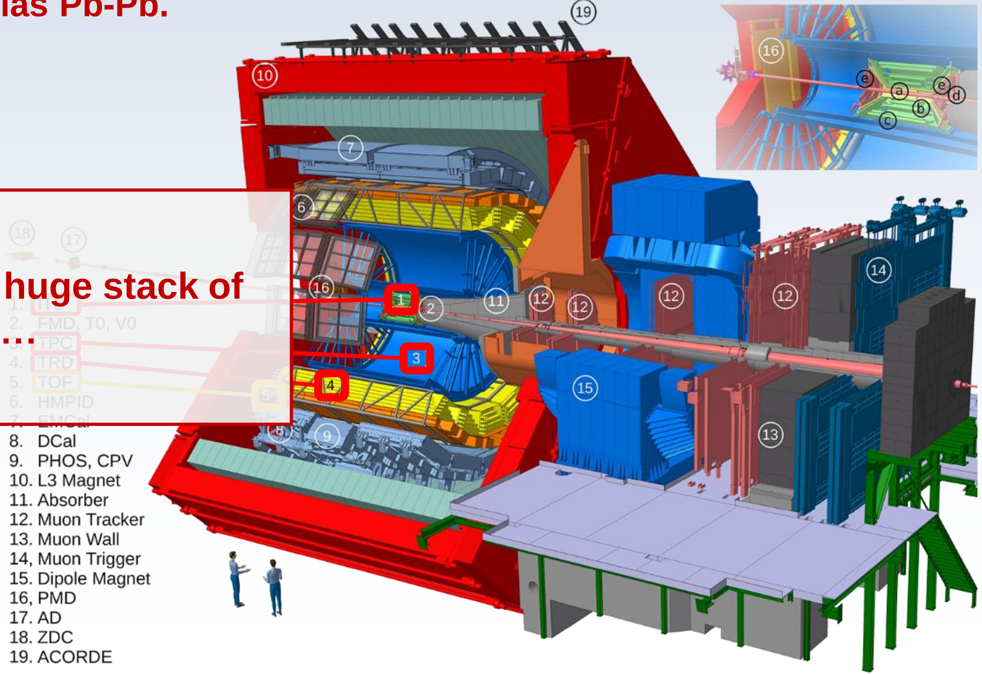

ALICE Upgrade for Run 3

- **ALICE uses mainly 3 detectors for tracking: ITS, TPC, TRD + (TOF)**

- **Luminosity will increase significantly in Run 3:**
  500 Hz – 1 kHz trigged → 50 kHz continuous min-bias Pb-Pb.
  - We want to adapt the HLT tracking for Run 3.
  - Ideally, create one tracker as fast as the HLT version with competitive resolution to the offline version.

- **50+ times as many events!**

50+ times as much data?

…or shrink the event.
- Same problem with computing

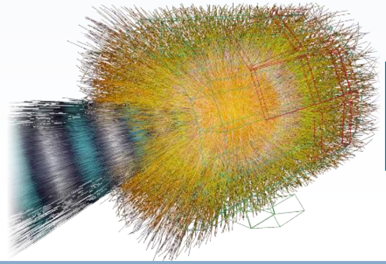

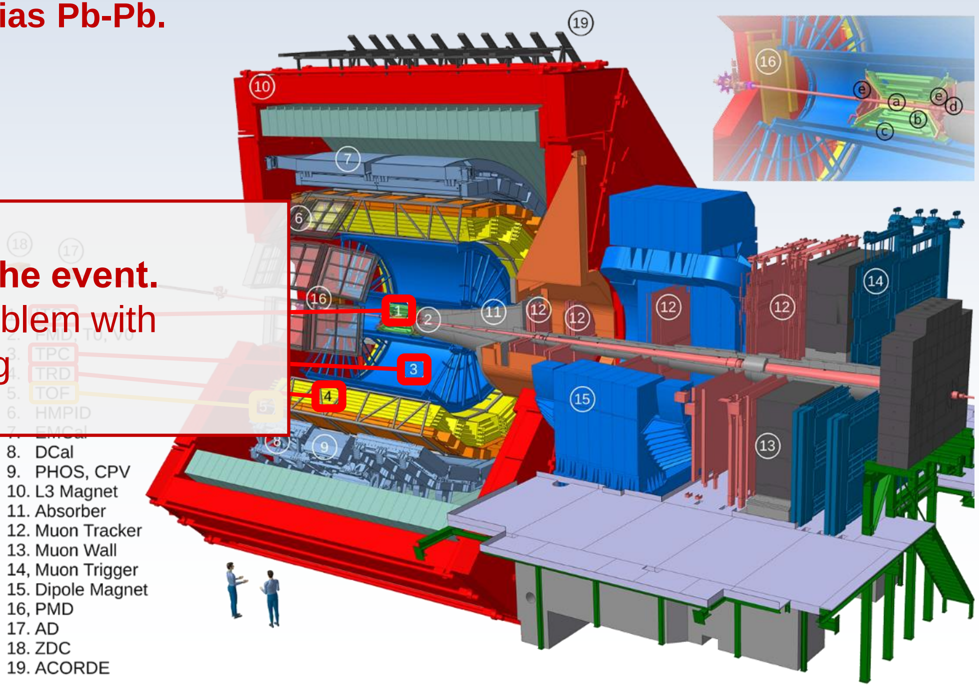

Implications of the ALICE Upgrade for Tracking

- **ALICE will perform a major upgrade for high luminosity data taking.**
  - This includes the entire computing infrastructure: ALICE Online Offline Computing Upgrade **O²**.

Change of paradigm:

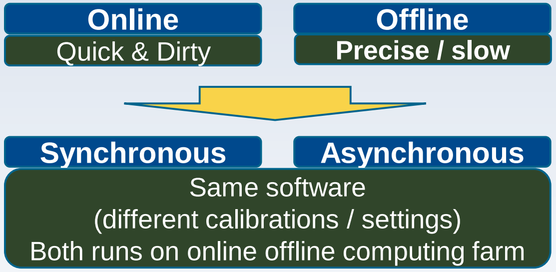

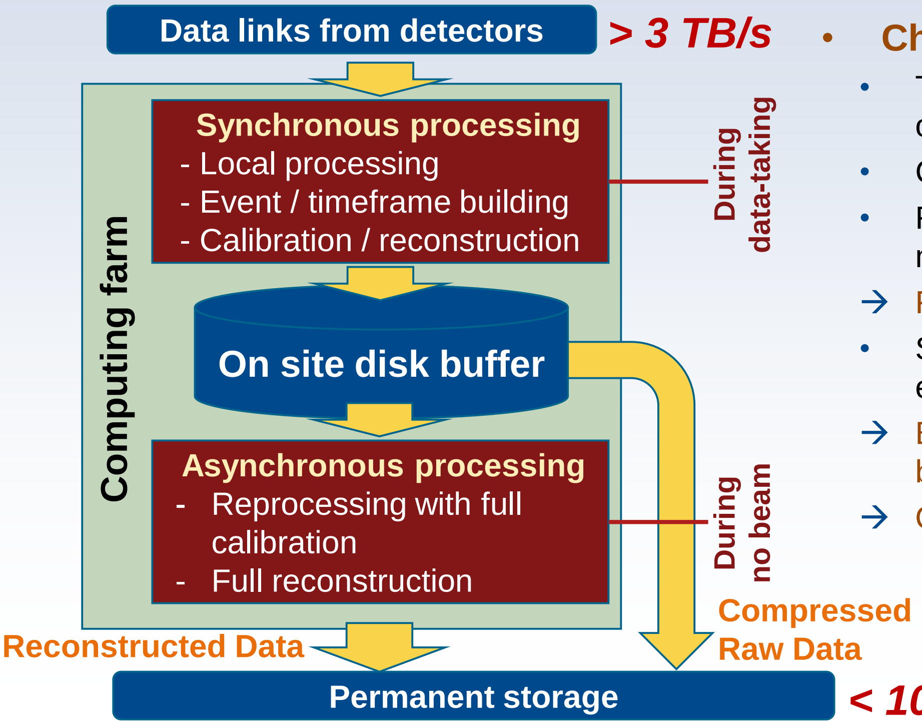

- **Challenges:**
  - Tracking time frames with overlapping events.
  - Continuous readout.
  - Reconstruct 50 times more events in same time.

  → Faster reconstruction
  - Store 50 times more events than in Run 2.

  → Better data compression based on tracks.

  → Online calibration.

From few 100 Hz single events to Pb-Pb time frames

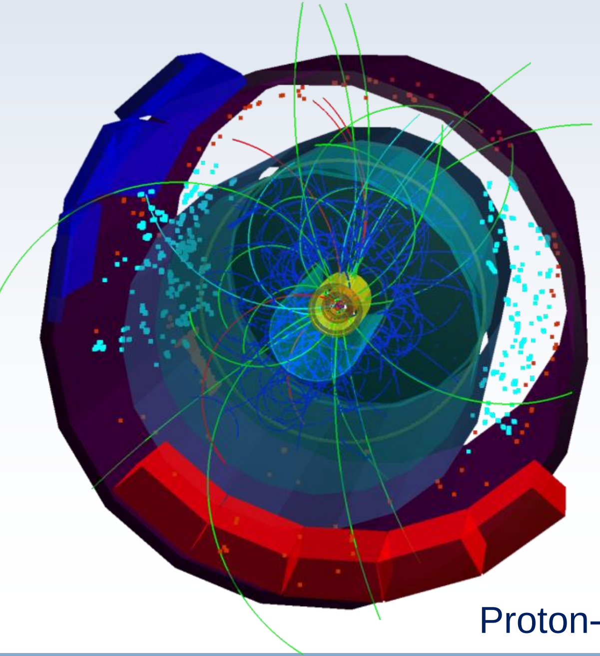

Proton-Proton

From few 100 Hz single events to Pb-Pb time frames

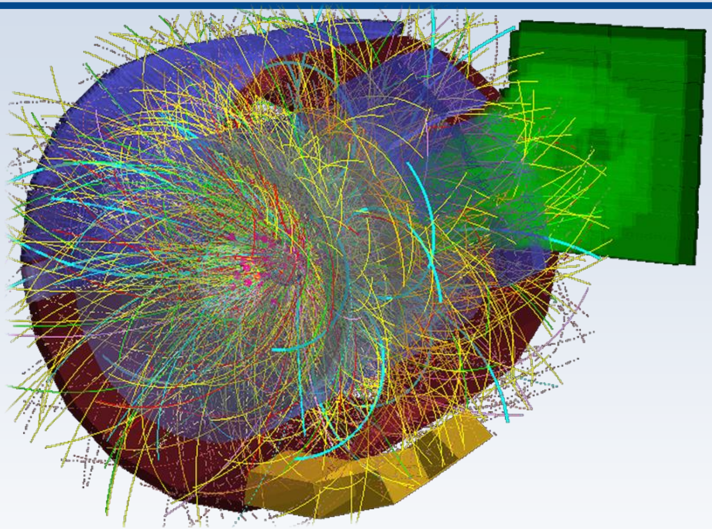

Proton-Lead

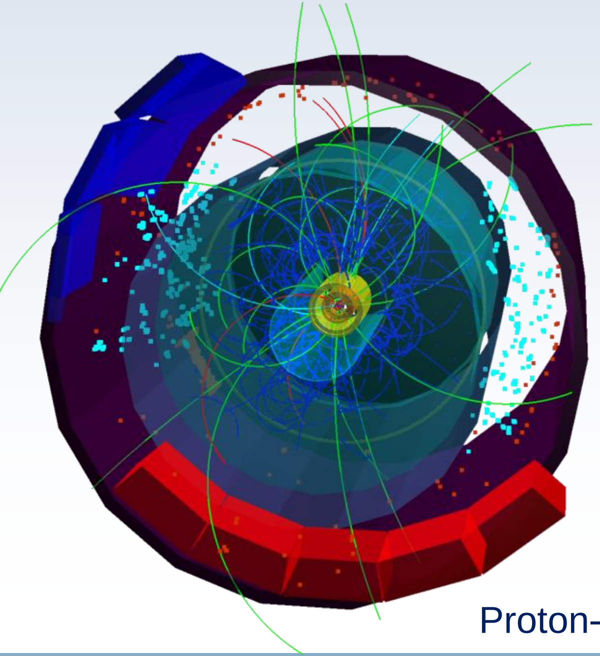

Proton-Proton

From few 100 Hz single events to Pb-Pb time frames

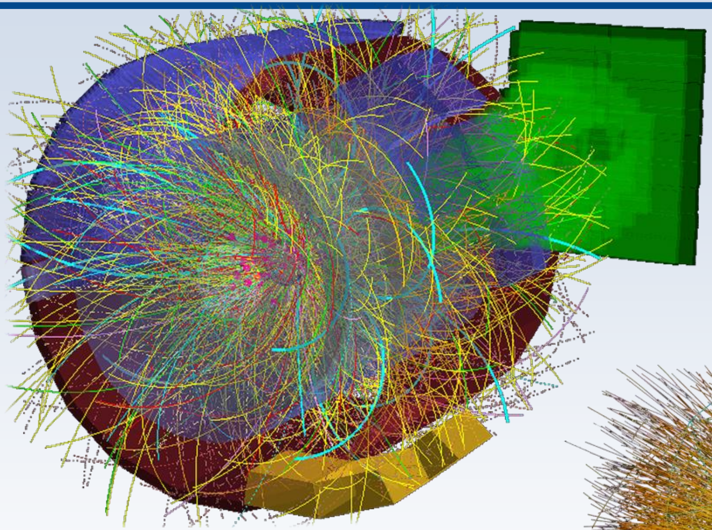

Proton-Lead

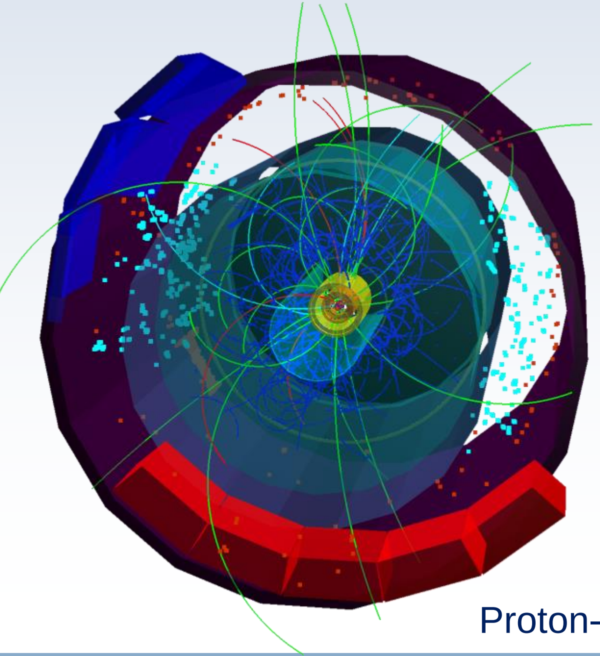

Proton-Proton

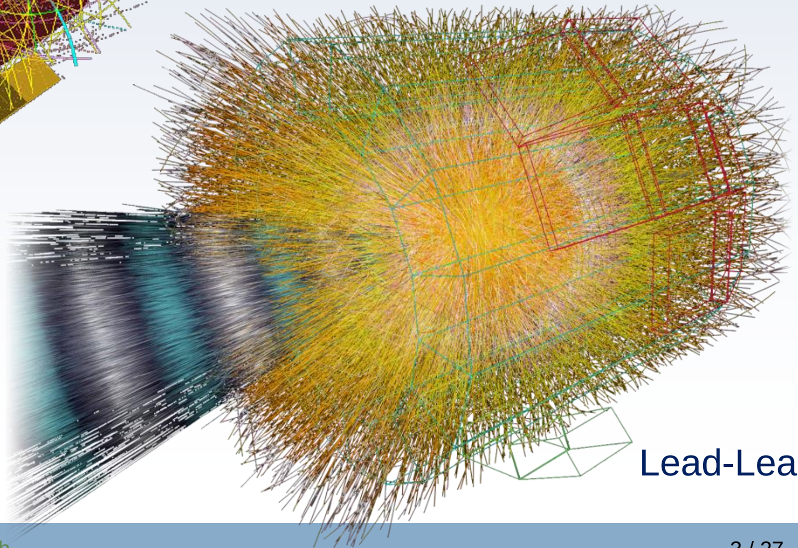

Lead-Lead

From few 100 Hz single events to Pb-Pb time frames

- Longitudinal crosssection of time frame in TPC, different collisions in different color.

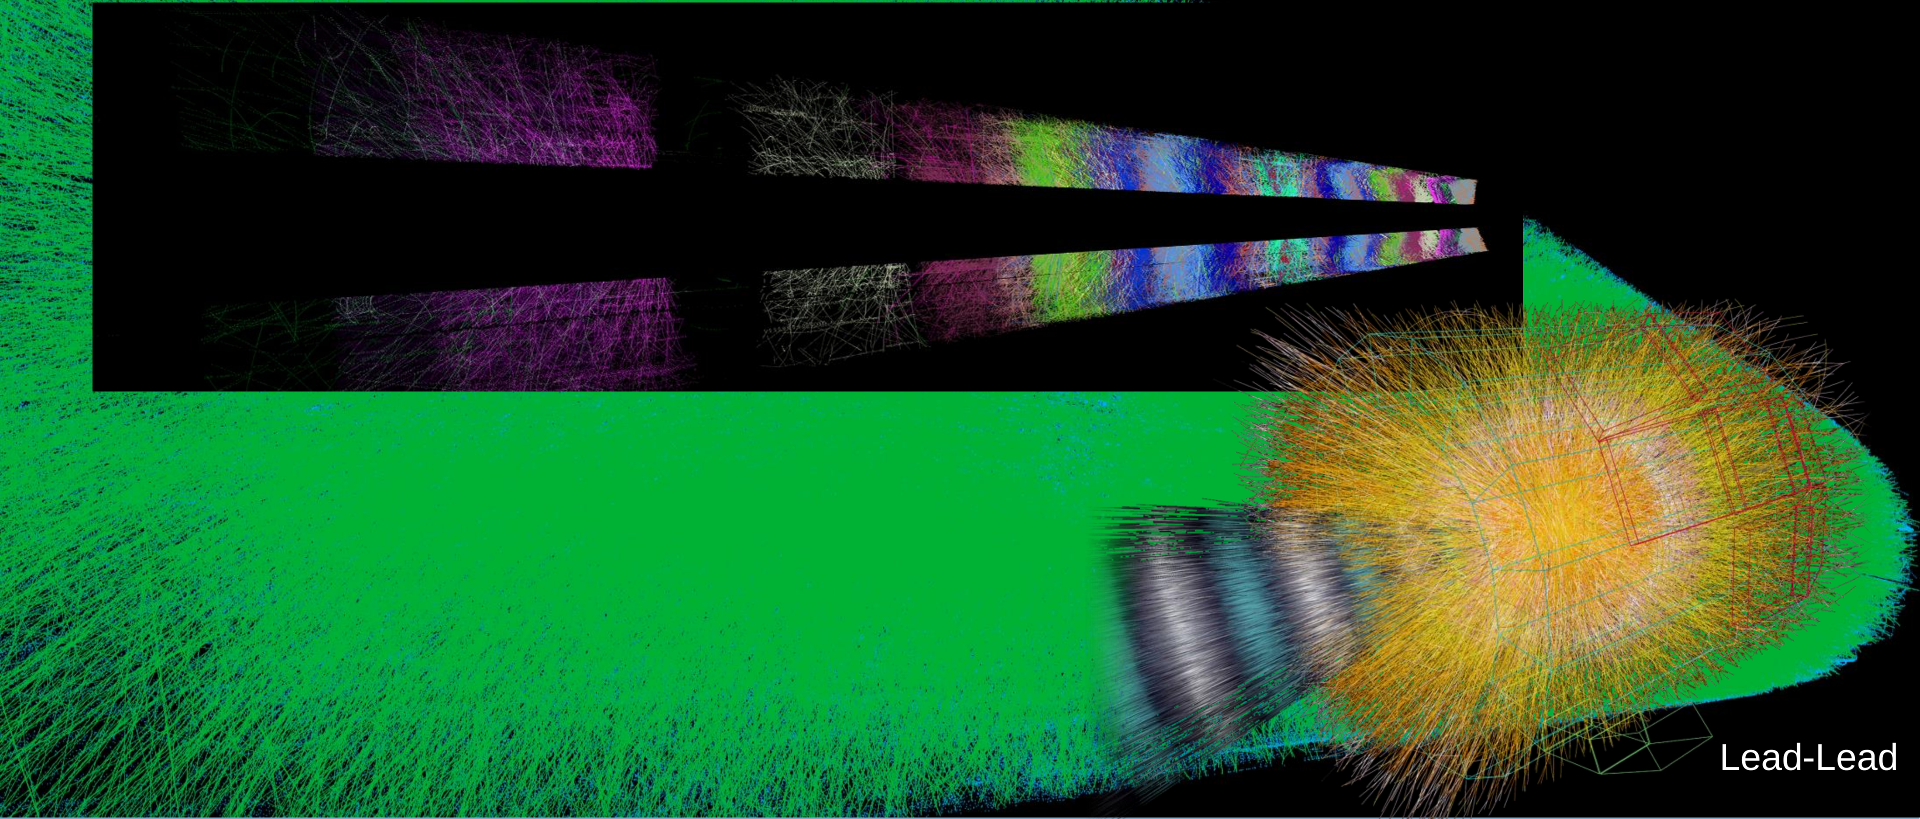

From few 100 Hz single events to Pb-Pb time frames

- Longitudinal crosssection of time frame in TPC, different collisions in different color.

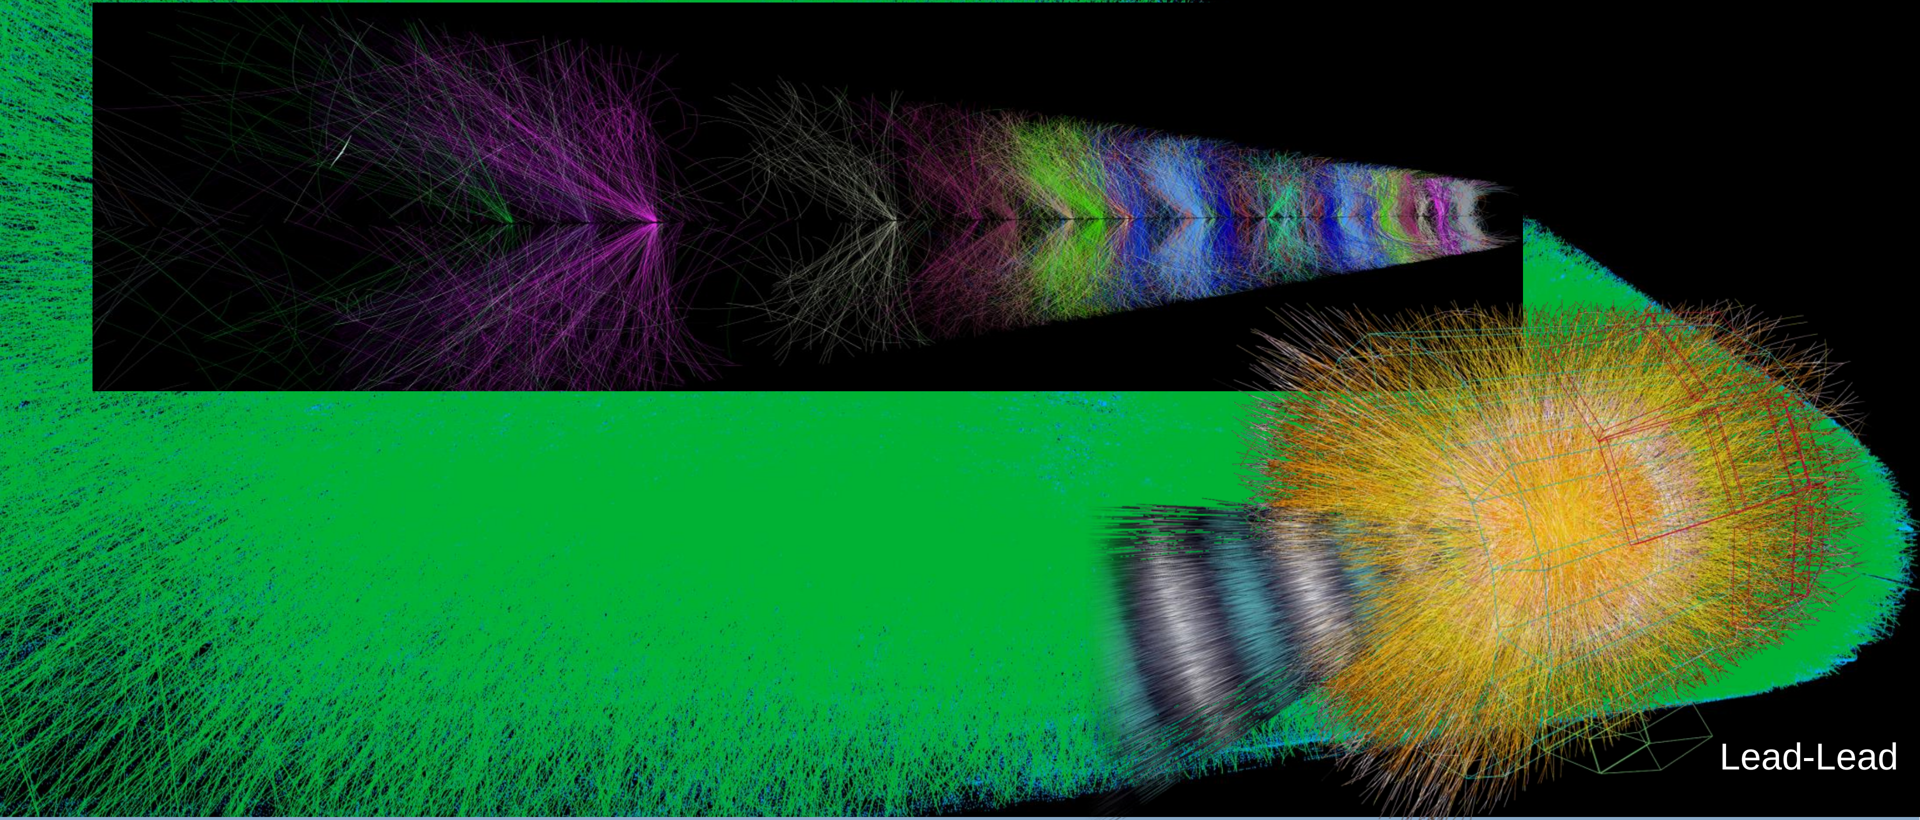

From few 100 Hz single events to Pb-Pb time frames

- Longitudinal crosssection of time frame in TPC, different collisions in different color.

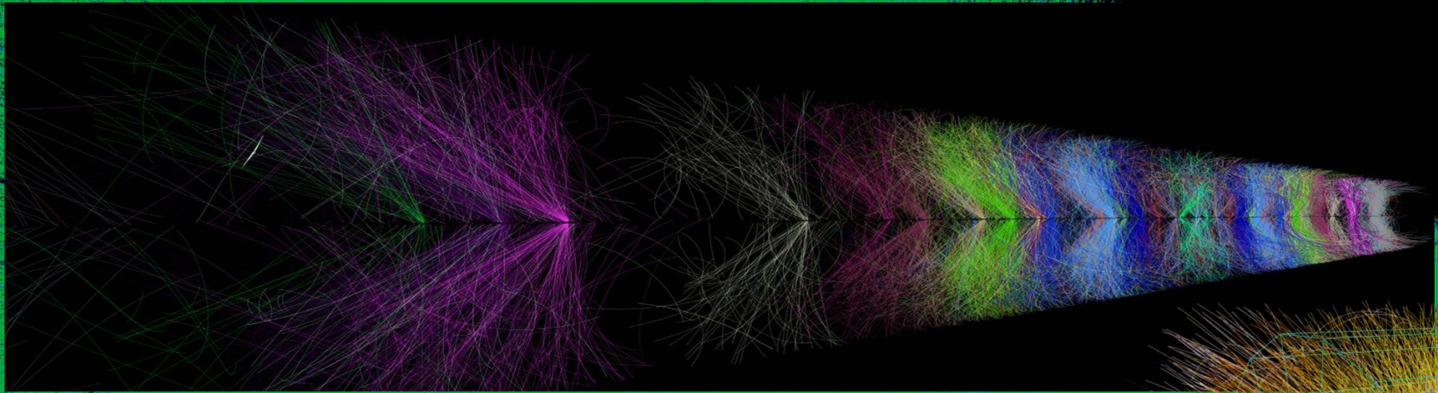

  - Time frame with pp collisions.

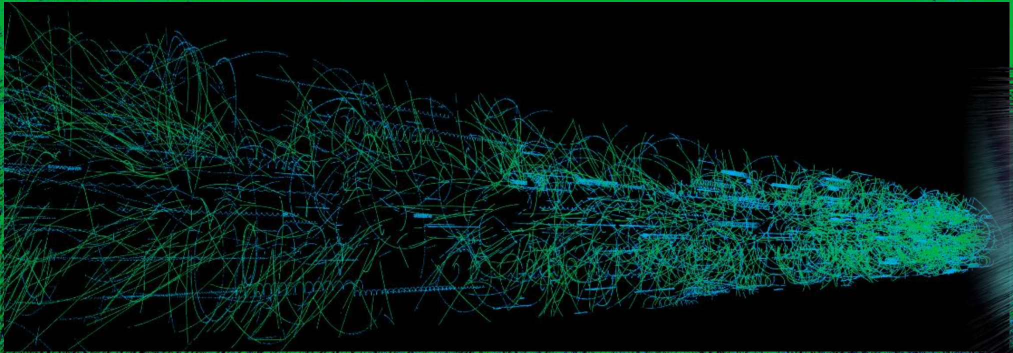

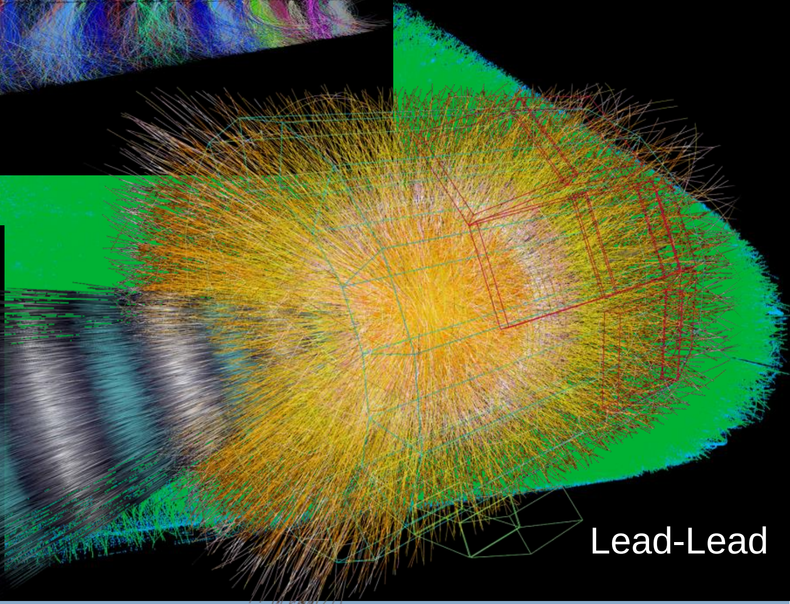

From few 100 Hz single events to Pb-Pb time frames

- Lead-Lead time frame at 50 kHz collision rate

- Outline of this talk (Tracking plans of ALICE):
  - Use fast online algorithms for O².
  - Focus on TPC Tracking (short summary for other detectors).
  - Improve HLT tracking to reach Offline Quality.
  - Adapt to the settings of O2.
    - Calibration
    - Tracking of time frames
    - Data compression

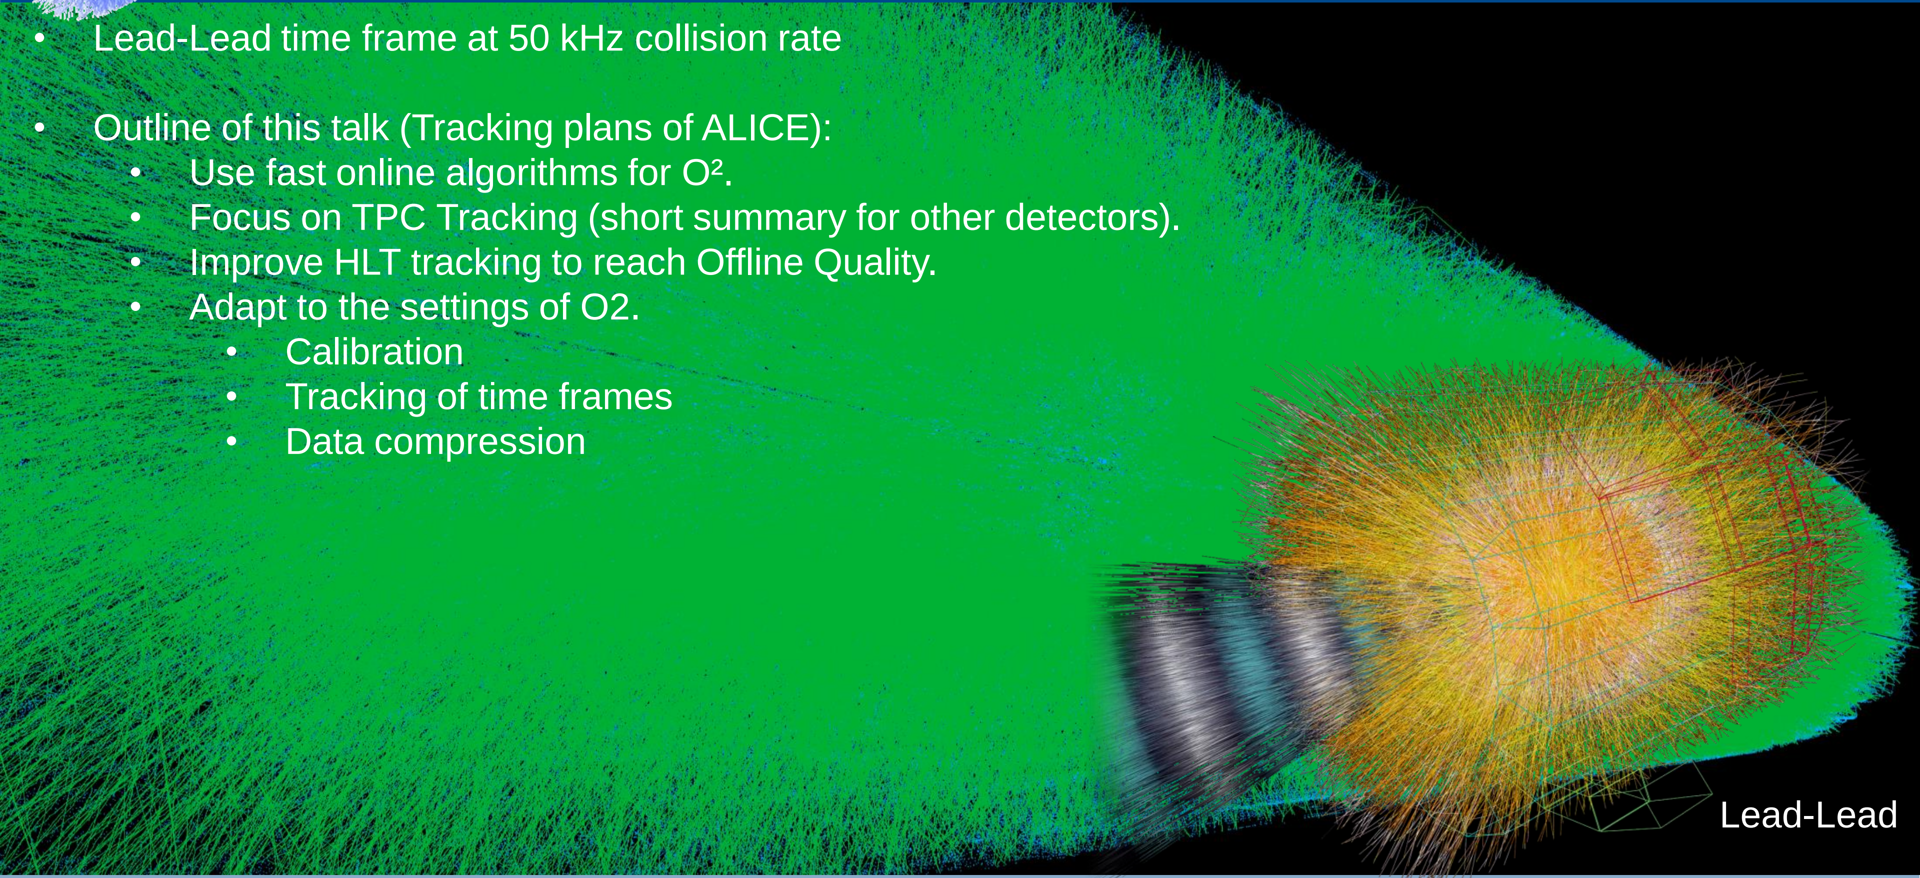

Tracking in ALICE

- **ALICE uses mainly 3 detectors for tracking: ITS, TPC, TRD + (TOF)**
  - **7 layers ITS** (Inner Tracking System – silicon tracker)
  - **152 pad rows** TPC (Time Projection Chamber)
  - **6 layers TRD** (Transition Radiation Detector)
  - **1 layer TOF** (Time Of Flight Detector)

- **Today, ALICE has two tracking implementations for the TPC:**
  - A fast online tracker in the HLT.
  - The offline TPC tracker as reference.
  - Both use the Kalman filter for track following and track fitting.
  - The HLT uses a Cellular Automaton for seeding.

- **Other detectors:**
  - HLT and offlline tracker for ITS (TPC prolongation).
  - No HLT TRD tracking so far.
  - TOF can be used as single last tracking layer.

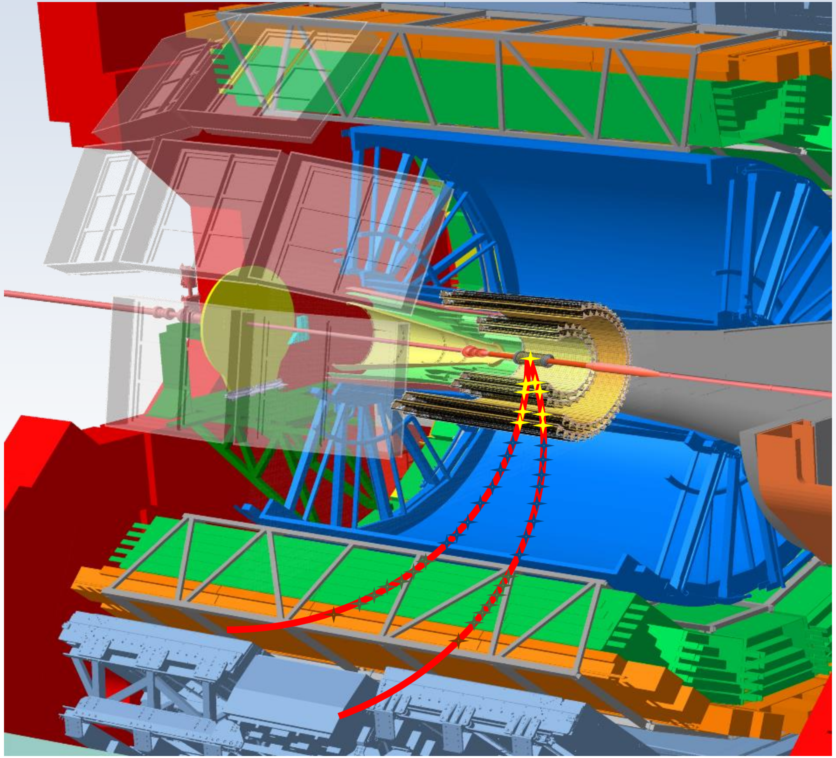

ONLINE TRACKING WITH OFFLINE QUALITY

Run 3 Tracking derived from current tracking in ALICE HLT

- **TPC Volume is split into 36 sectors.**
  - The tracker processes each sector individually.
  - Increases data locality, reduce network bandwidth, but reduces parallelism.
  - Each sector has 159 read out rows in radial direction.
  - Tracking runs in 2 phases:

- **1. Phase: Sector-Tracking (within a sector)**
  - Heuristic, combinatorial search for track seeds using a **Cellular Automaton**.
    - A) Looks for three hits composing a straight line (**link**).
    - B) Concatenates links.
  - Fit of track parameters, extrapolation of track, and search for additional clusters using the **Kalman Filter**.

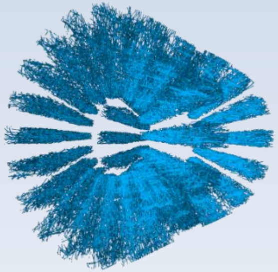

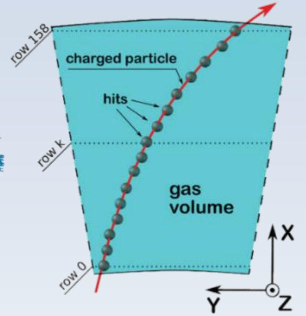

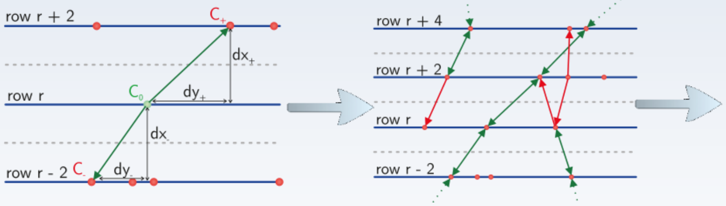

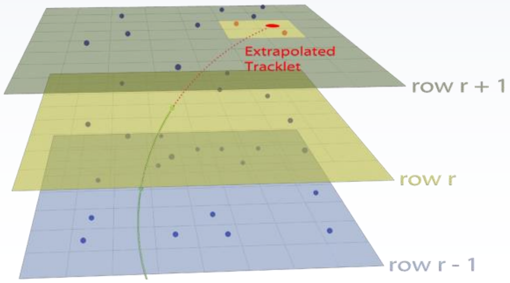

Run 3 Tracking derived from current tracking in ALICE HLT

- **TPC Volume is split into 36 sectors.**
  - The tracker processes each sector individually.
  - Increases data locality, reduce network bandwidth, but reduces parallelism.
  - Each sector has 159 read out rows in radial direction.
  - Tracking runs in 2 phases:

- **1. Phase: Sector-Tracking (within a sector)**
  - Heuristic, combinatorial search for track seeds using a **Cellular Automaton**.
    - A) Looks for three hits composing a straight line (**link**).
    - B) Concatenates links.
  - Fit of track parameters, extrapolation of track, and search for additional clusters using the **Kalman Filter**.

- **2. Phase: Track-Merger**
  - Combines the track segments found in the individual sectors.

- **Phase 1 track finding implemented in a common generic source code, that runs on CPU and GPU, supporting CUDA, OpenCL, and OpenMP.**

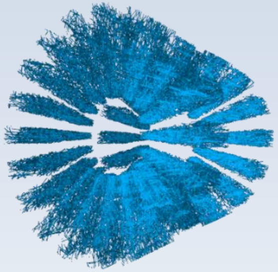

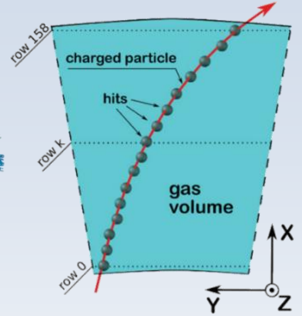

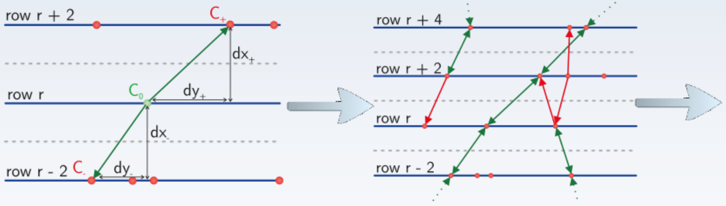

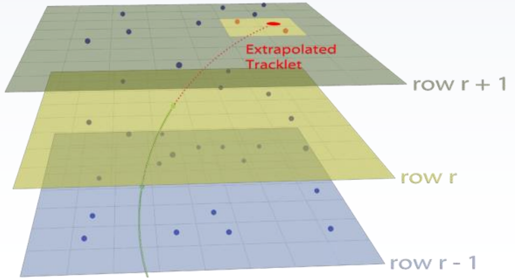

Tracking efficiency (Run 2, HLT v.s. Offline – Pb-Pb)

- Offline and HLT feature good efficiency.
  - In order to compare the tracking algorithm, the offline calibration and error parameterization is applied in both cases.
  - Absolutely same result in pp.

- All plots are Monte Carlo.
- All plots are TPC only.
- Resolutions at inner end of TPC.
- Findable tracks: min 70 TPC hits.
- Others: min 1 TPC hit.
- Other offline features (dE/dx, …) disabled.

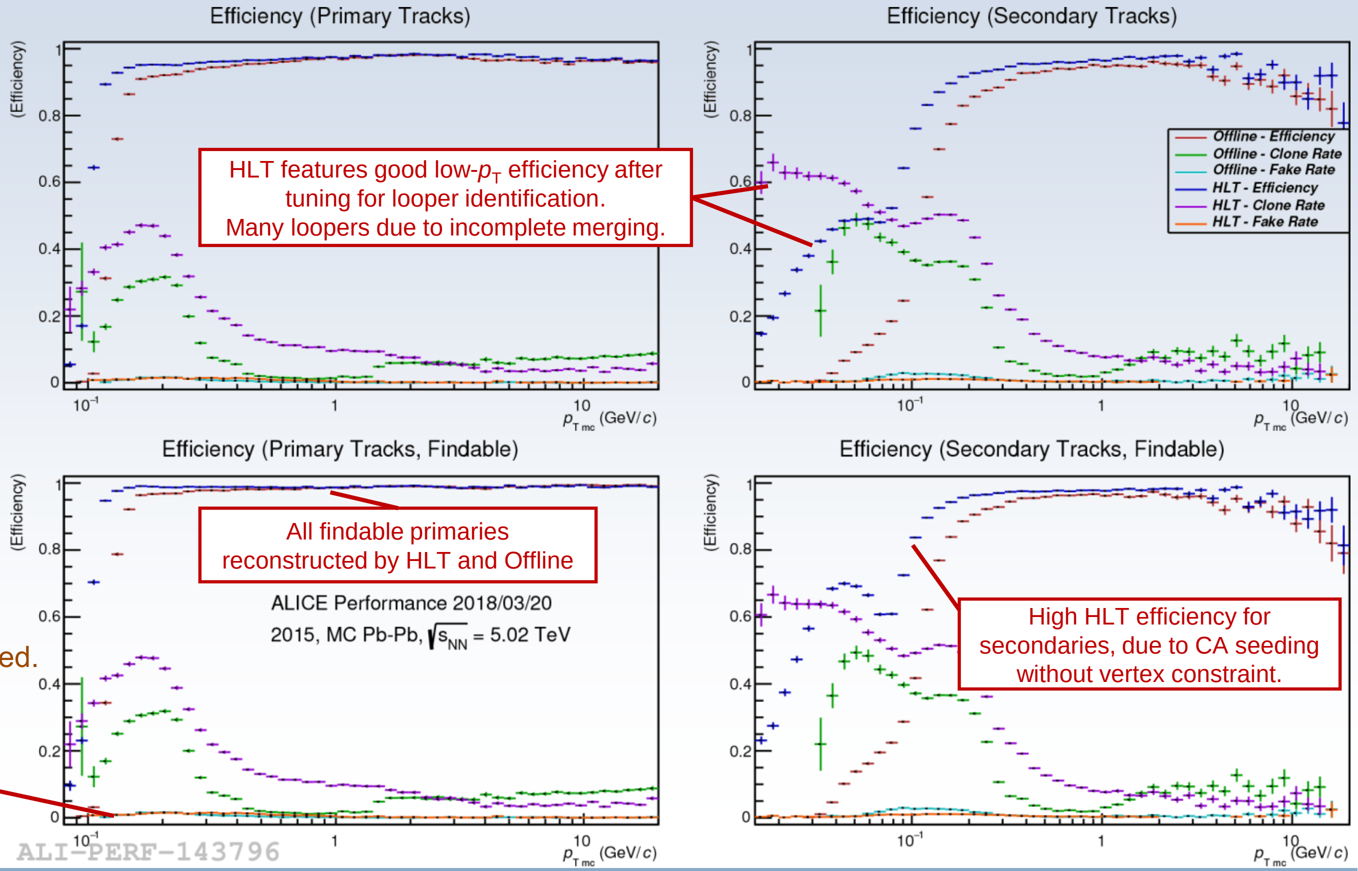

Practically zero fake rate for both trackers

Track resolution (Run 2, Pb-Pb, no space-charge distortions)

- HLT / Offline resolution practically identical (no space-charge distortions).

- Improvements in HLT tracking:
  - Propagation using polynomial approximation of 3D B-field.
  - Outlier cluster rejection during refit.
  - Improved cluster error parameterization, depending on flags set by clusterizer. (edge, deconvoluted, …)
  - 3-way fit. (inward, outward, inward)

- Absolutely same result for pp. (see backup)

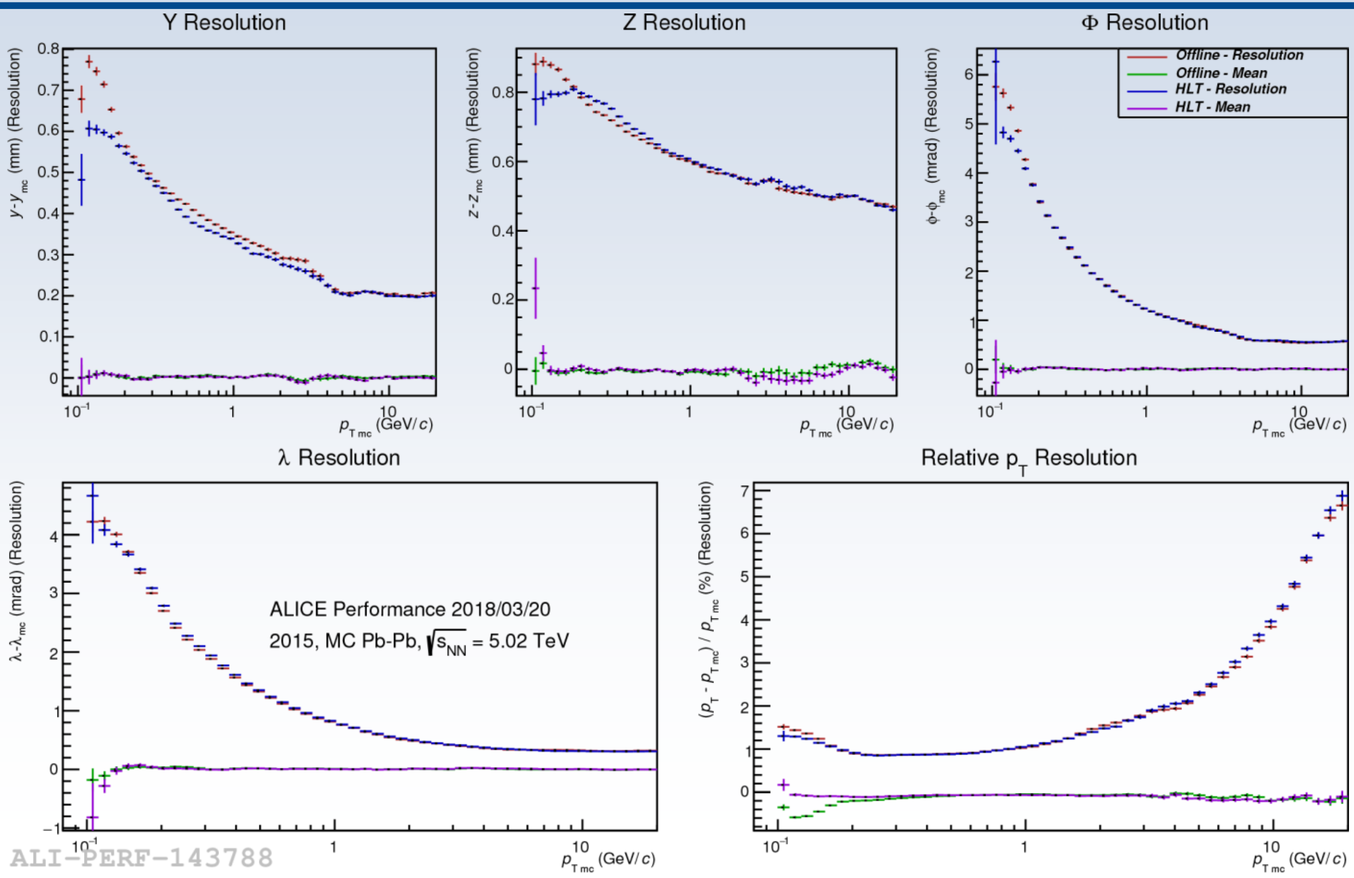

Track resolution (Run 2, Pb-Pb, with space-charge distortions)

- Small differences with space-charge distortions.
- Similar structure in y-resolution.

- HLT/O² has not been tuned for distortions so far.
- Only using systematic cluster error parameterization obtained from offline distortion map residuals.

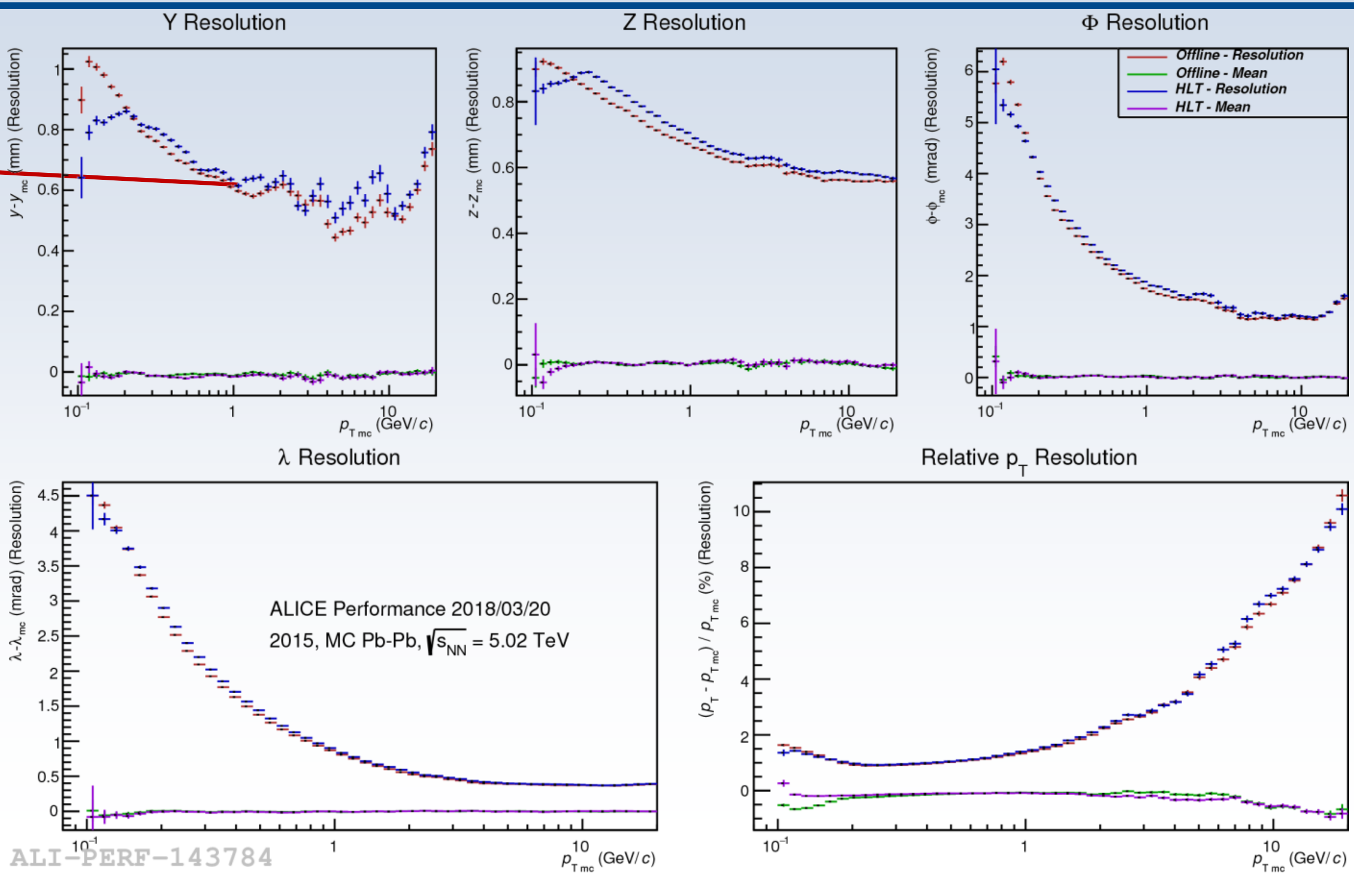

RECONSTRUCTION AND CALIBRATION IN RUN 3

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

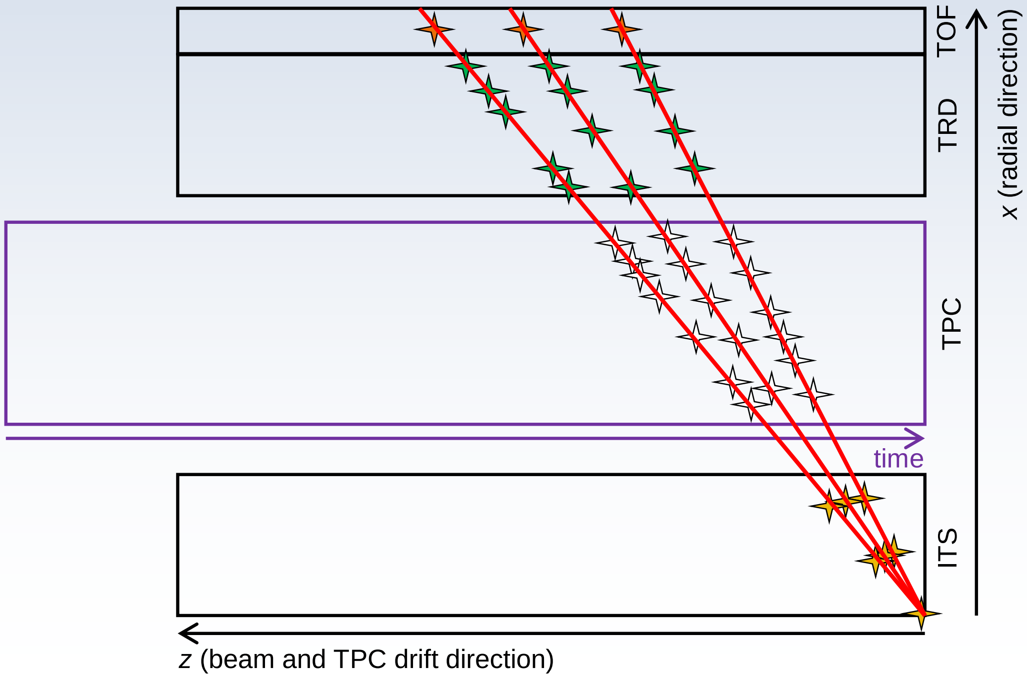

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**

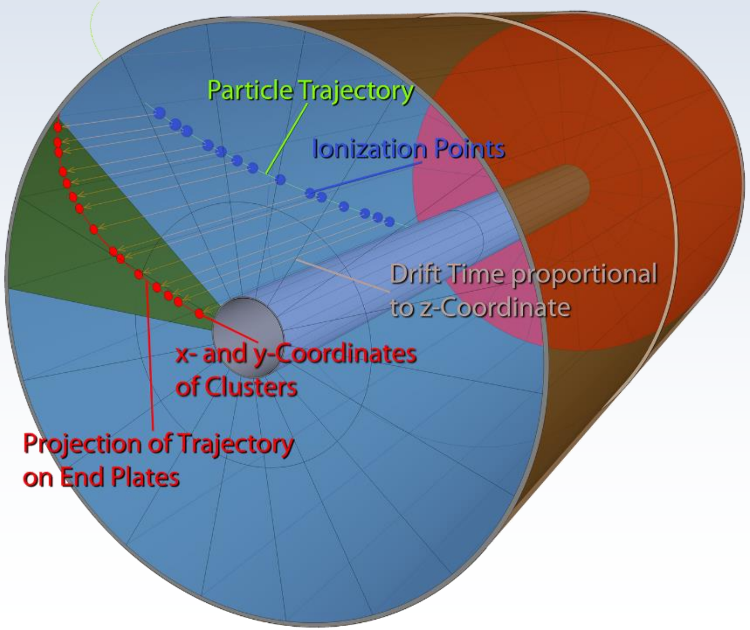

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

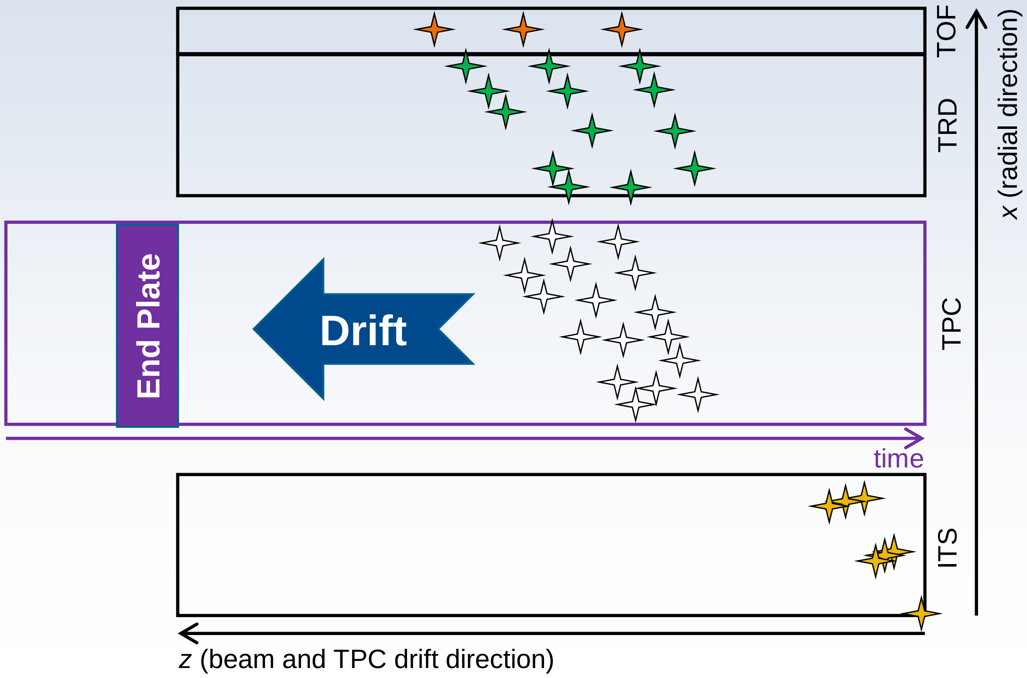

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**

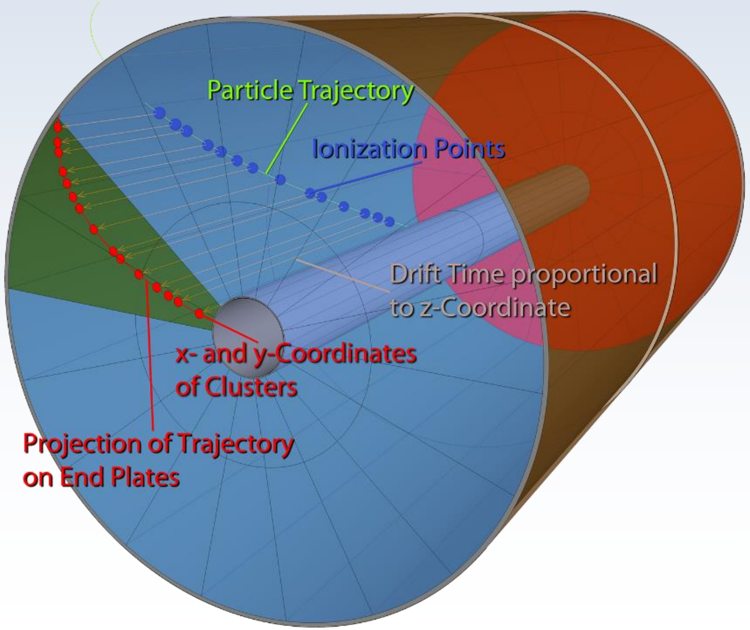

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

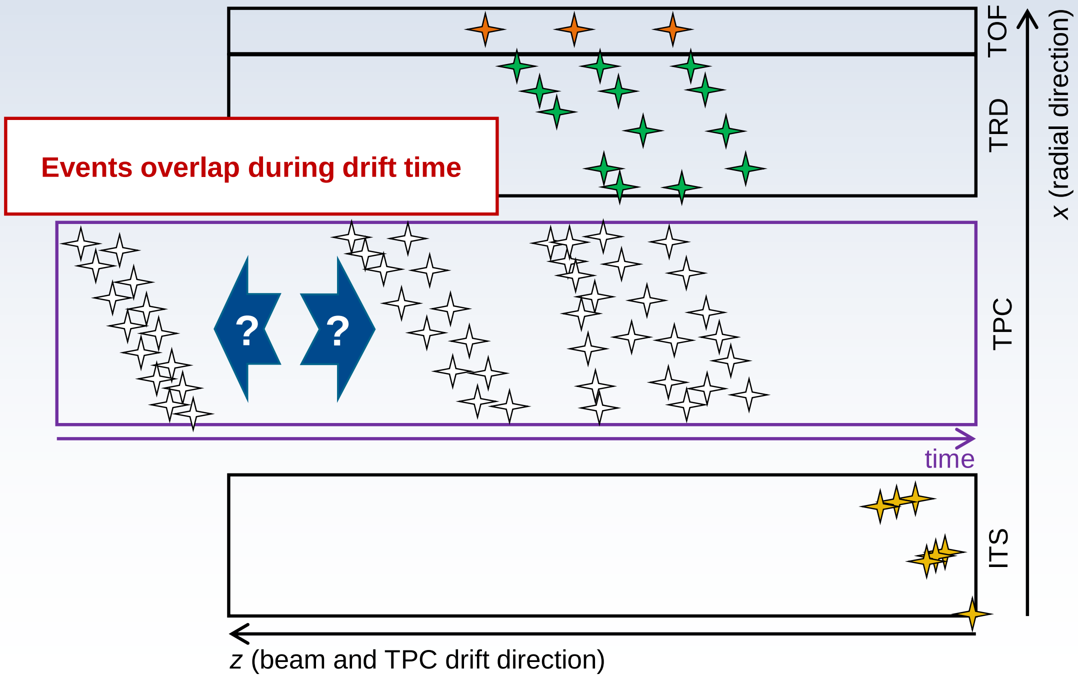

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**

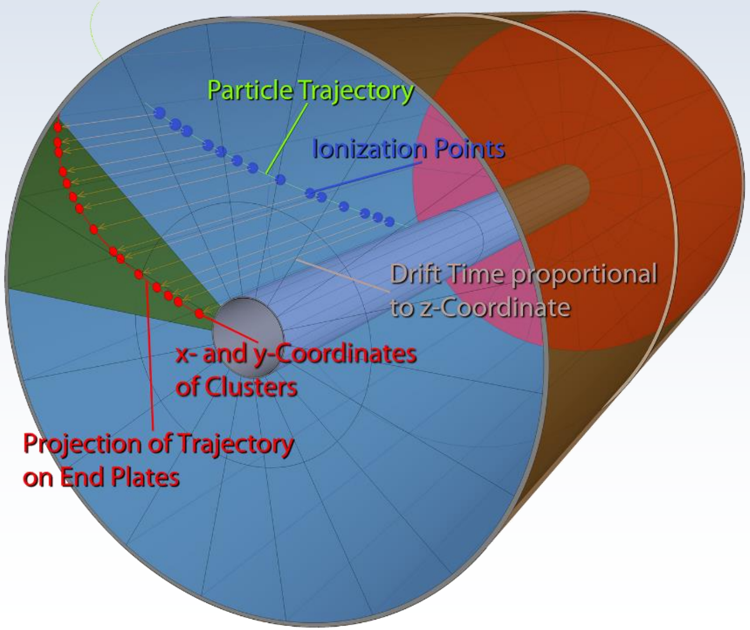

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

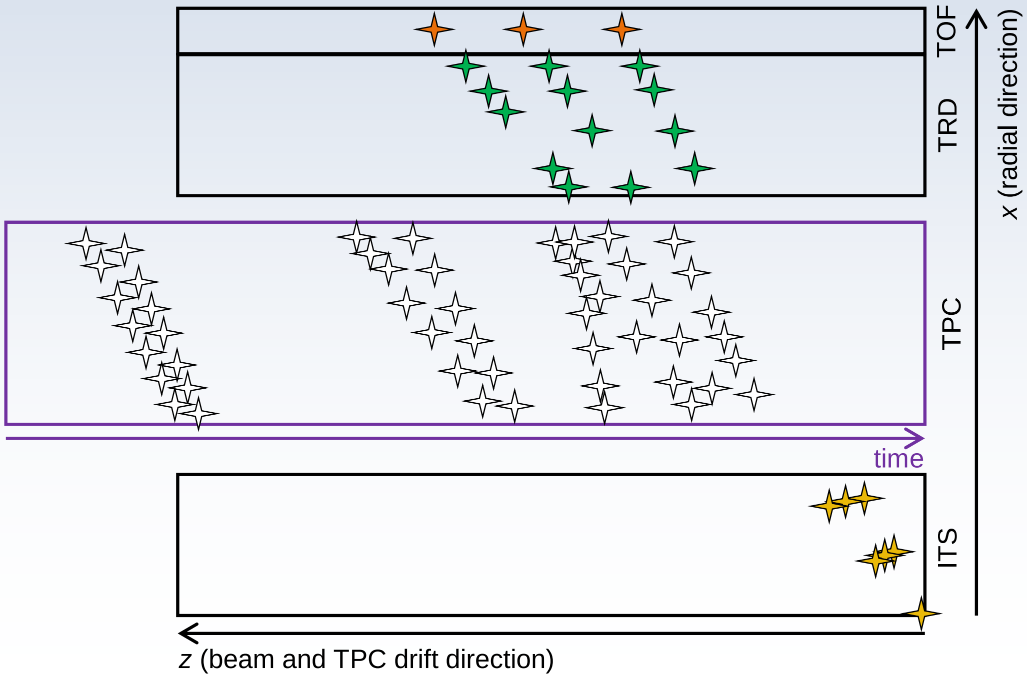

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**

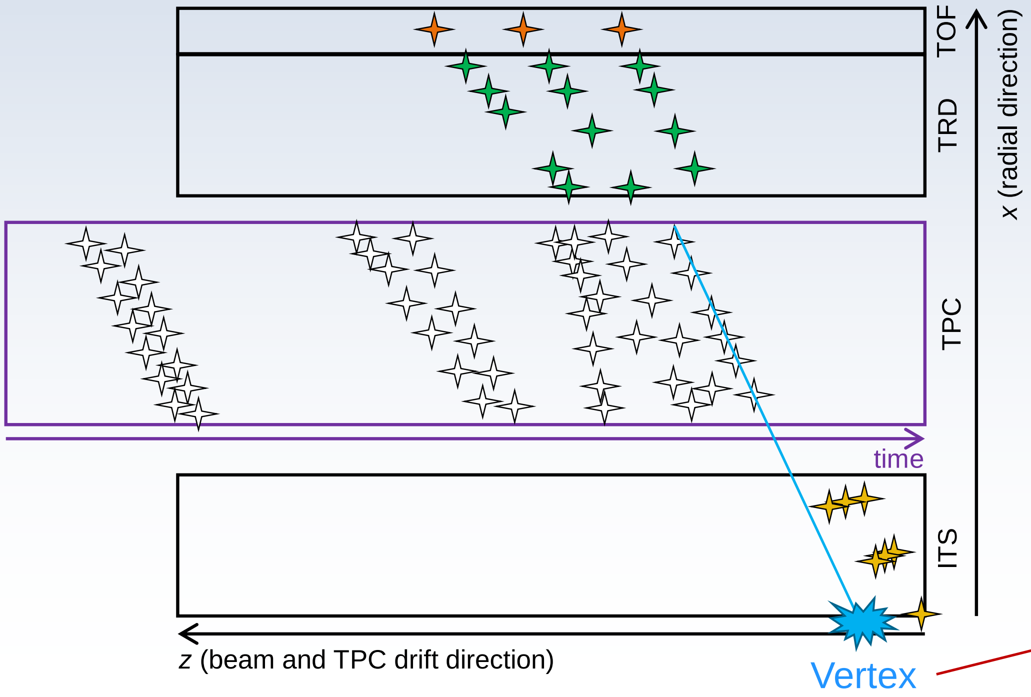

$z \sim t - t_{\text{Vertex}}$

→ Need to identify the primary vertex, before assigning final *z* to cluster.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

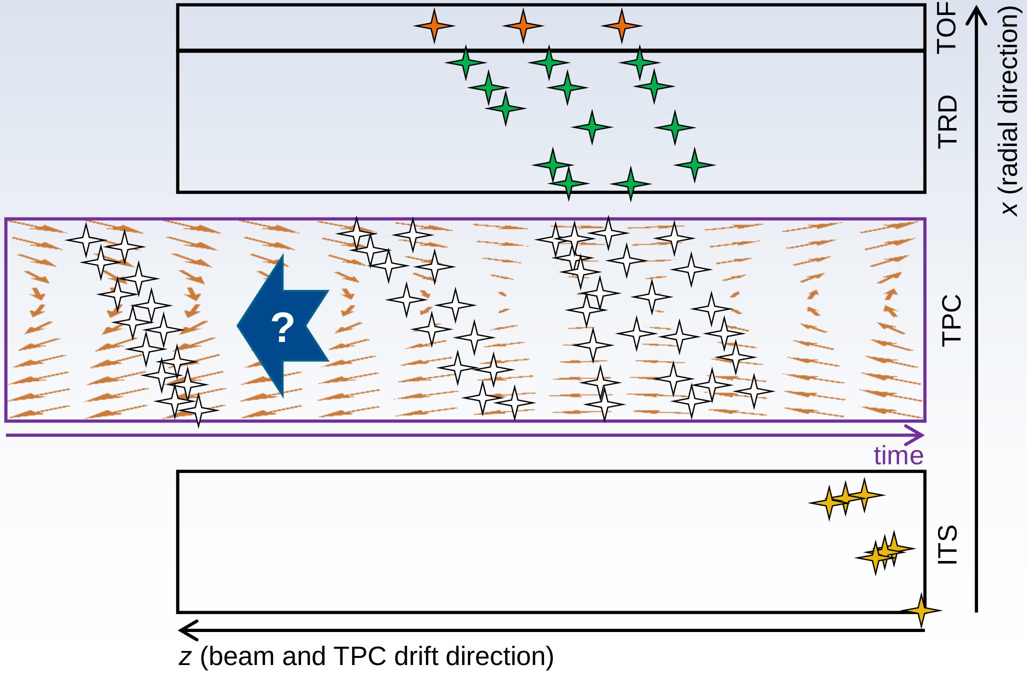

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

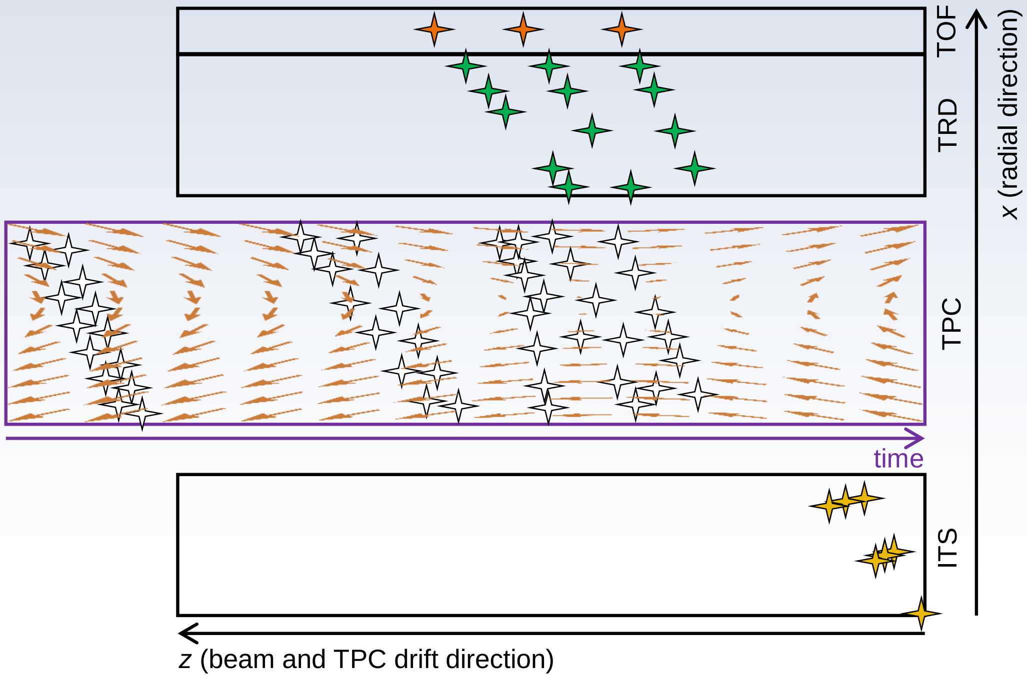

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

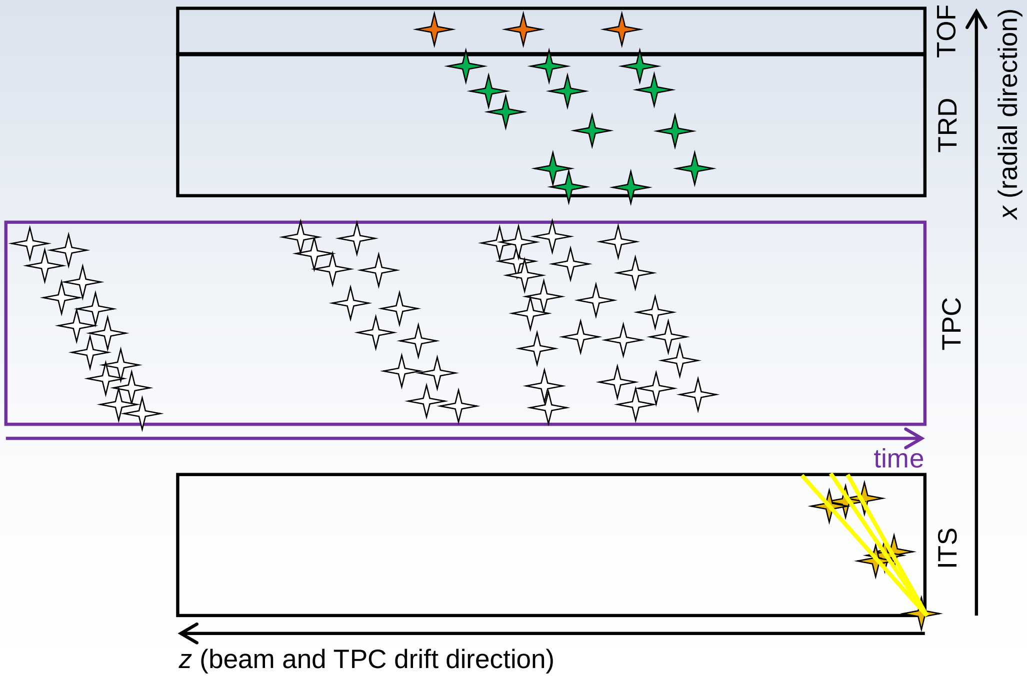

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

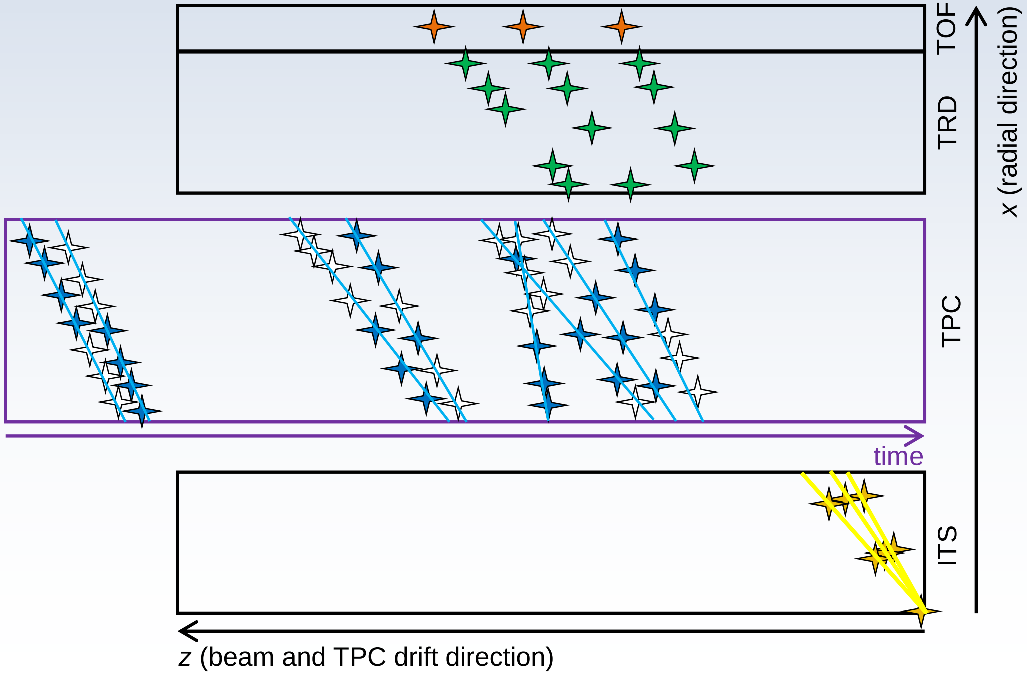

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

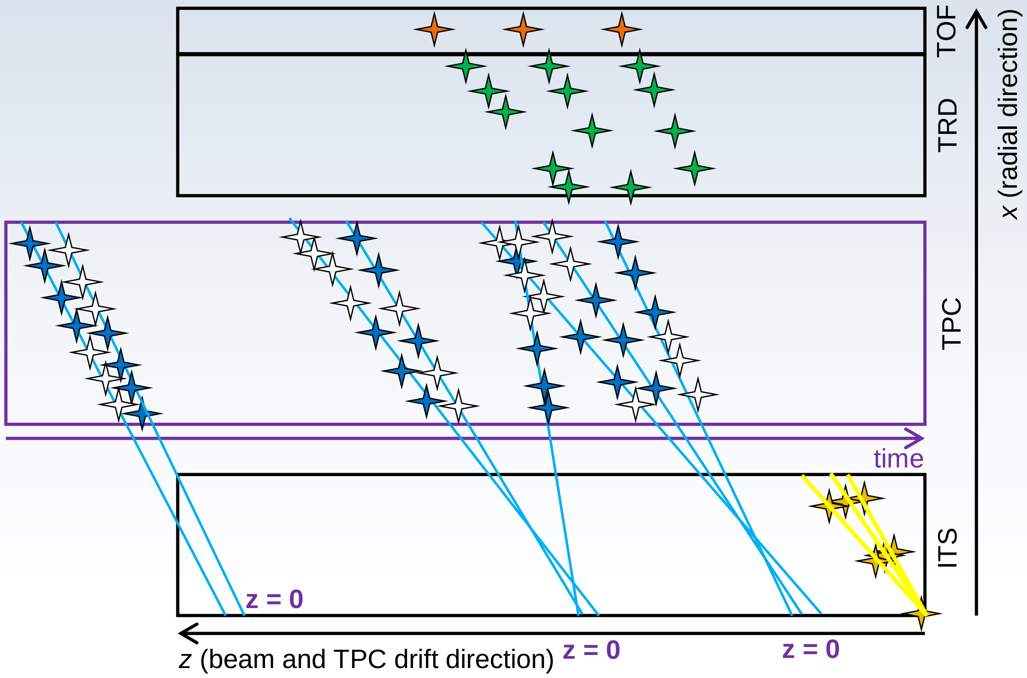

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.

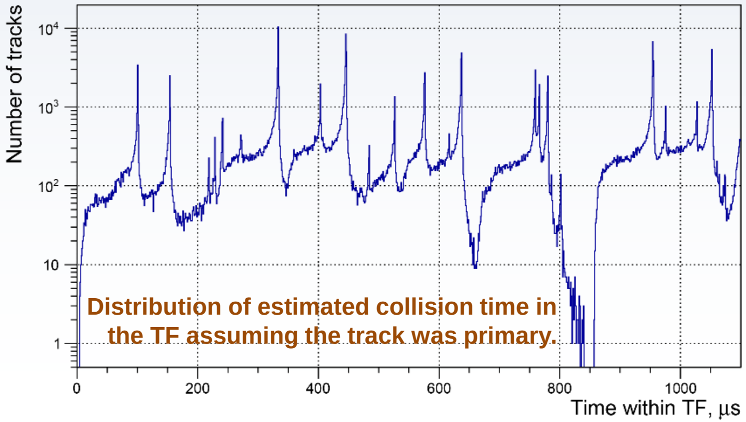

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

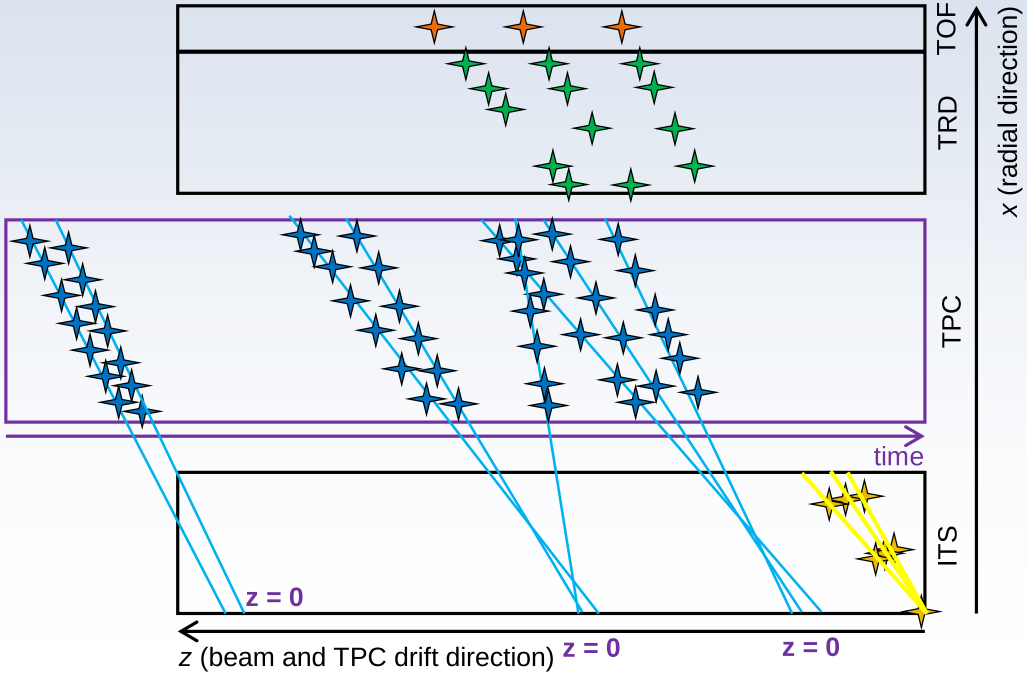

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

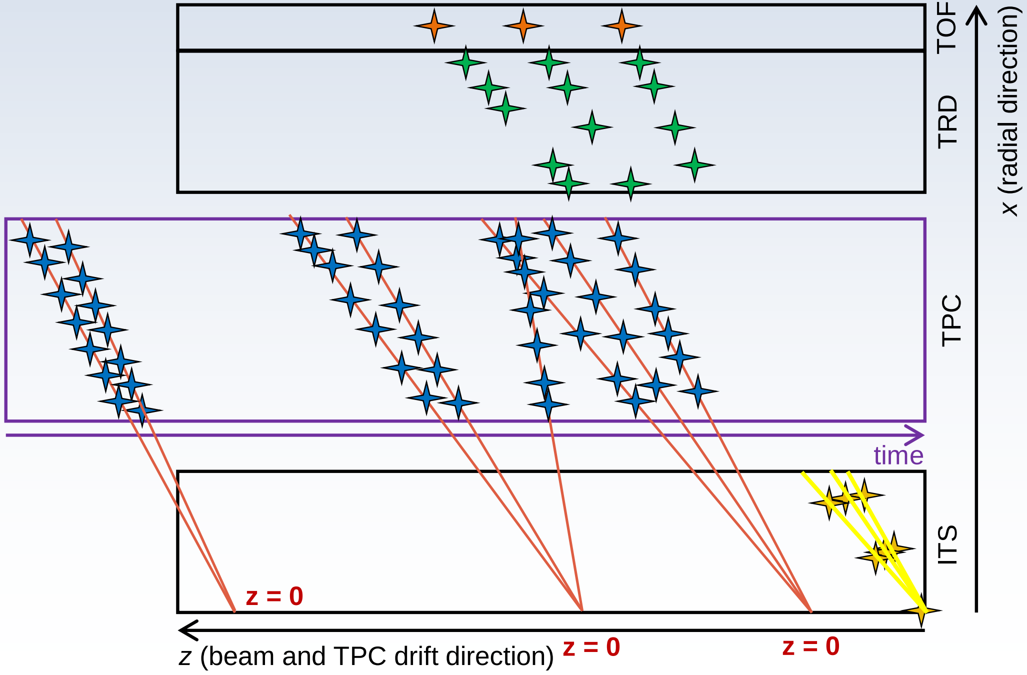

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

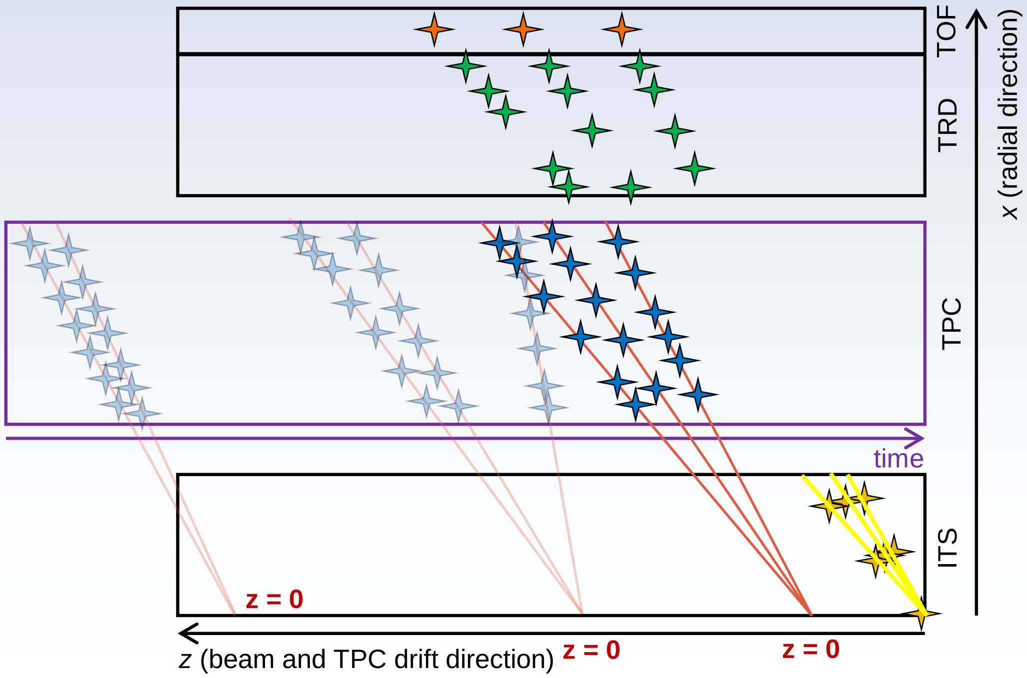

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision
- For the tracks seen in one ITS read out frame, select all TPC events with a matching time (from z = 0 estimate).

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

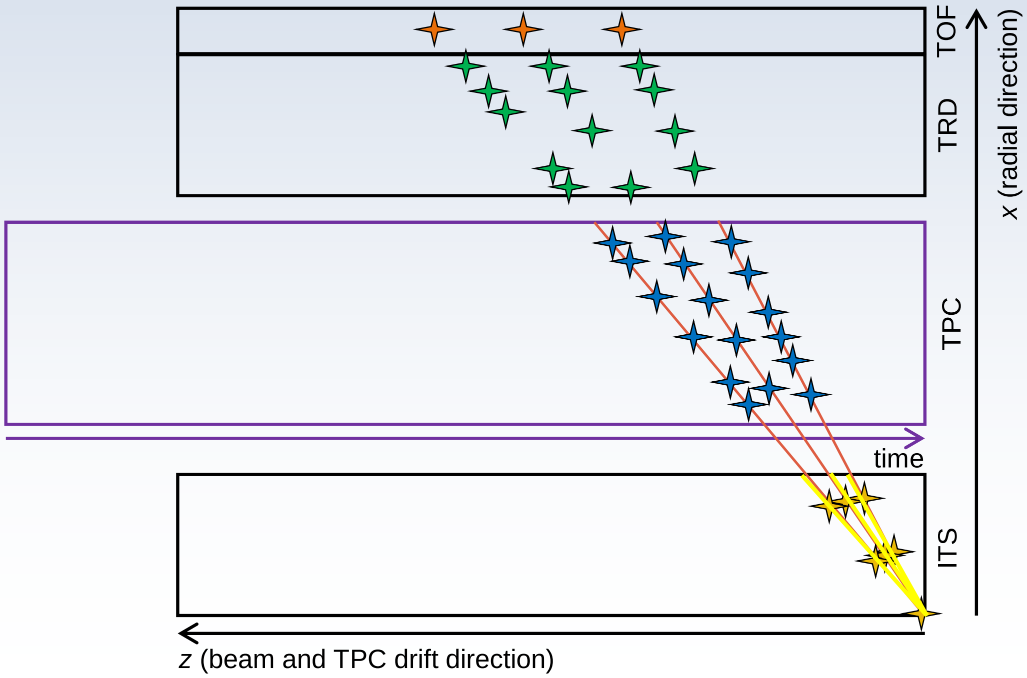

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision
- For the tracks seen in one ITS read out frame, select all TPC events with a matching time (from z = 0 estimate).
- Match TPC track to ITS track, fixing the time and thus the z position of the TPC track.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

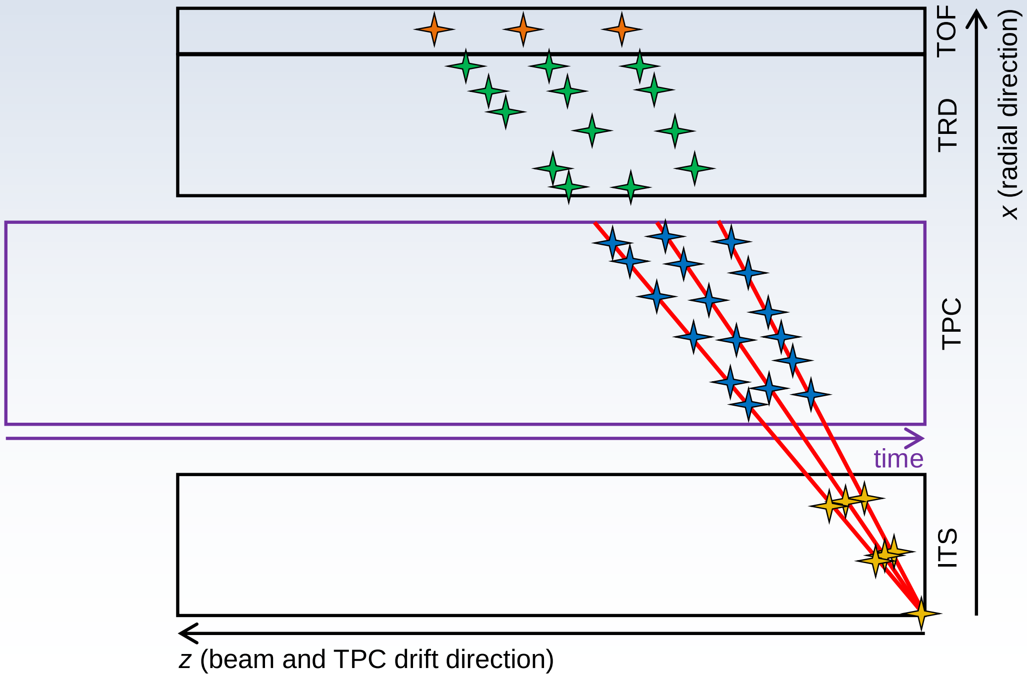

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision
- For the tracks seen in one ITS read out frame, select all TPC events with a matching time (from z = 0 estimate).
- Match TPC track to ITS track, fixing the time and thus the z position of the TPC track.
- Refit ITS + TPC track outwards.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

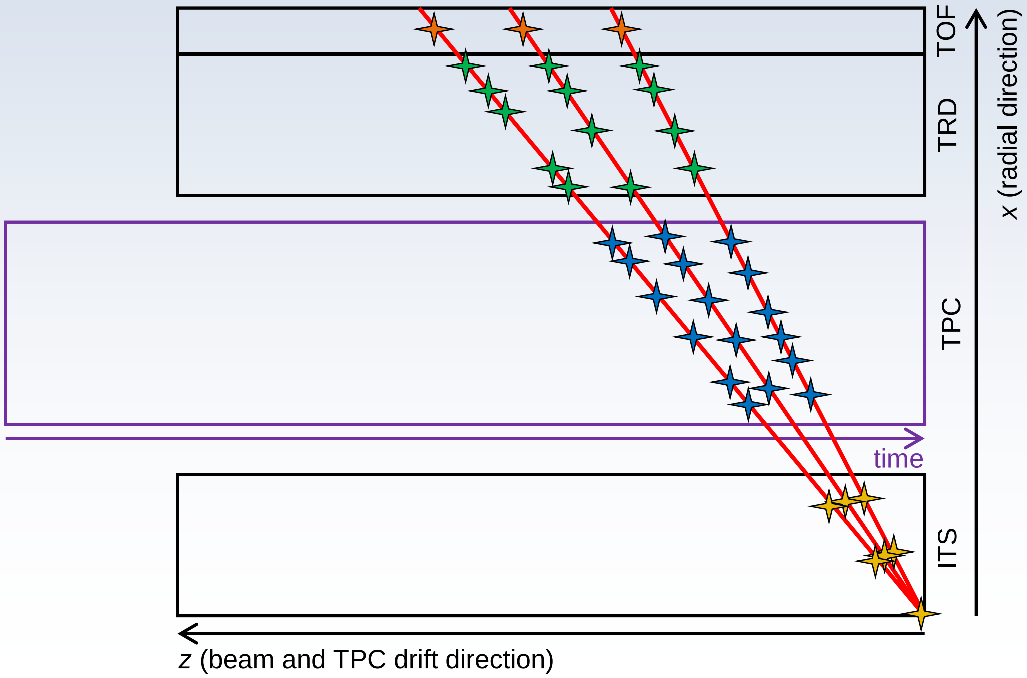

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision
- For the tracks seen in one ITS read out frame, select all TPC events with a matching time (from z = 0 estimate).
- Match TPC track to ITS track, fixing the time and thus the z position of the TPC track.
- Refit ITS + TPC track outwards.
- Prolong into TRD / TOF.

The tracking challenge

- **Tracking continuous data…**
  - The TPC sees **multiple overlapped collisions** (shifted in time).
  - Other detectors know the (rough) time of the collision.

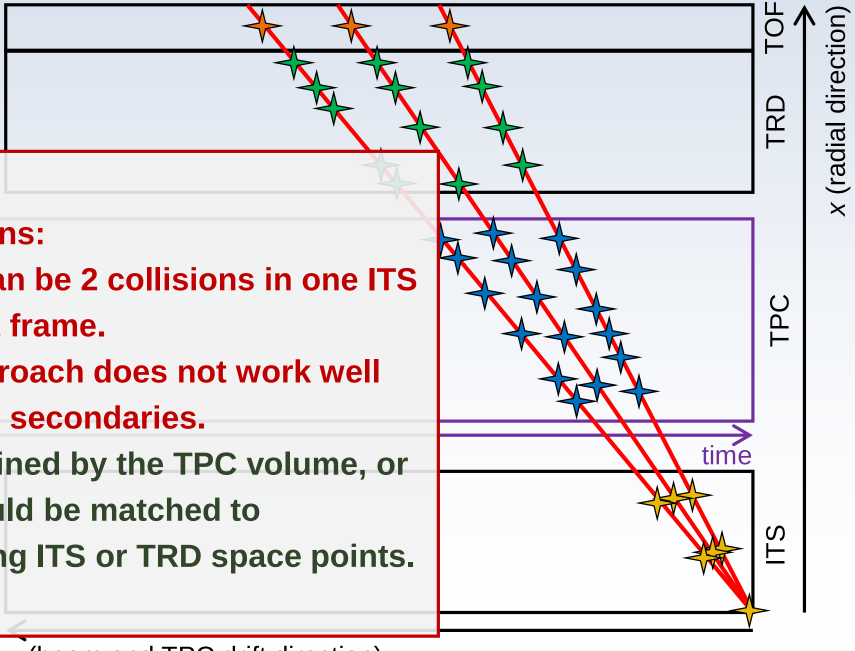

Complications:

- **There can be 2 collisions in one ITS read out frame.**
- **The approach does not work well for deep secondaries.**
- **Constrained by the TPC volume, or they could be matched to remaining ITS or TRD space points.**

- **Problem: TPC clusters have no defined z-position but only a time. They can be shifted in z arbitrarily.**
- **GEM amplifications produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute z position.**

- Standalone ITS tracking.
- Standalone TPC tracking, scaling t linearly to an arbitrary z.
- Extrapolate to x = 0, define z = 0 as if the track was primary.
- Track following to find missing clusters. For cluster error parameterization, distortions, and B-field, shift the track such that z = 0 at x = 0.
- Refine z = 0 estimate, refit track with best precision
- For the tracks seen in one ITS read out frame, select all TPC events with a matching time (from z = 0 estimate).
- Match TPC track to ITS track, fixing the time and thus the z position of the TPC track.
- Refit ITS + TPC track outwards.
- Prolong into TRD / TOF.

Tracking with GEM TPC, plan for Run3

- **Run 2 tracking needs on thy fly calibration…**
  - …and cluster conversion from row, pad, time to x, y, z.
  - Online calibration exercised in the HLT for TPC drift velocity.
    - No difference to Offline drift velocity calibration.
  - Online calibration with feedback loop: Calibration produced during n seconds / minutes used for the following n seconds / minutes.

  - Two space-charge distortion calibration algorithms foreseen.
    1. ITS – TRD interpolation (see next slide).
    2. Integrated digital currents.
  - The first is already running for Run 2 offline, the second requires continuous read out.

- **Work ongoing to perform most of the distortion calibration work online.**

- **A related problem imposed by the distortions:**
  - The tracking must perform the Kalman update at the correct radius even when clusters are seen at the wrong pad row.

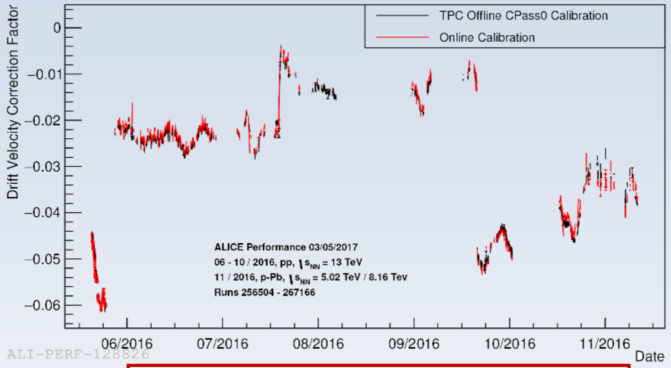

Residuals of cluster *z*-position in HLT:
Before online calibration: **30 mm**
With online calibration: **0.5 mm**

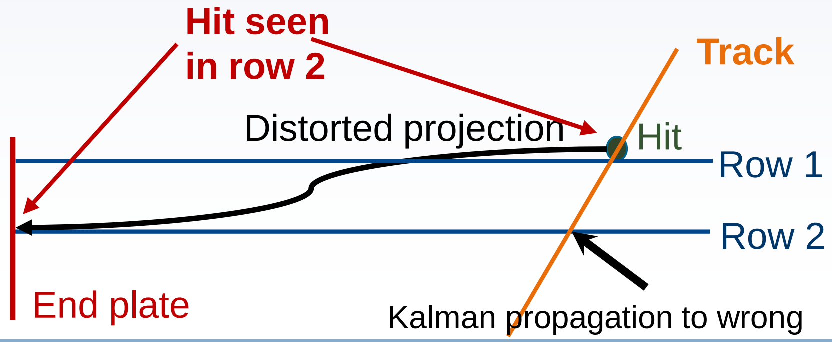

Space-charge Distortions corrections (in Pb-Pb Run 2)

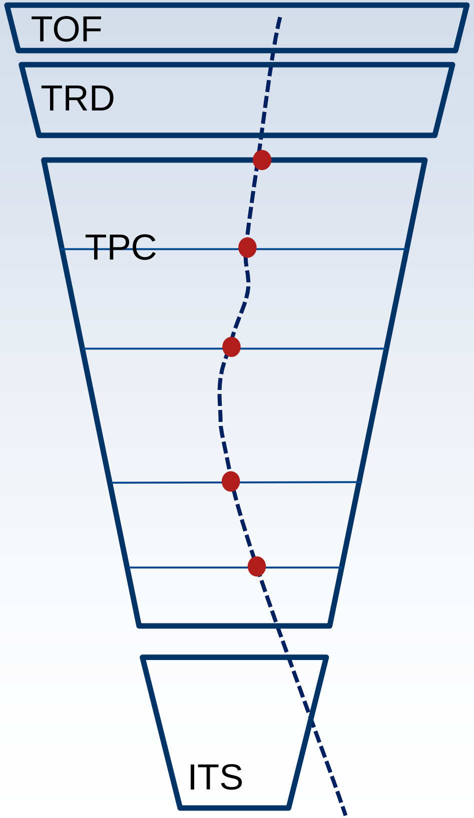

- Reconstruct TPC tracks with relaxed tolerances (applying “default distortion maps” if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

Space-charge Distortions corrections (in Pb-Pb Run 2)

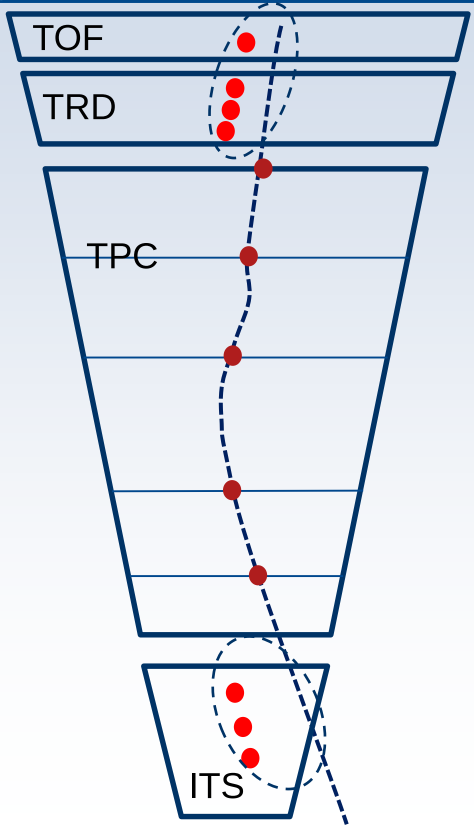

- Reconstruct TPC tracks with relaxed tolerances (applying “default distortion maps” if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

Space-charge Distortions corrections (in Pb-Pb Run 2)

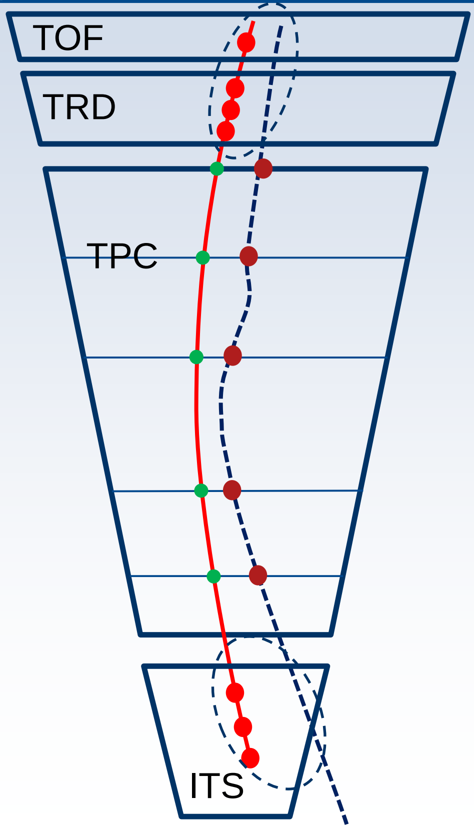

- Reconstruct TPC tracks with relaxed tolerances (applying “default distortion maps” if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

Space-charge Distortions corrections (in Pb-Pb Run 2)

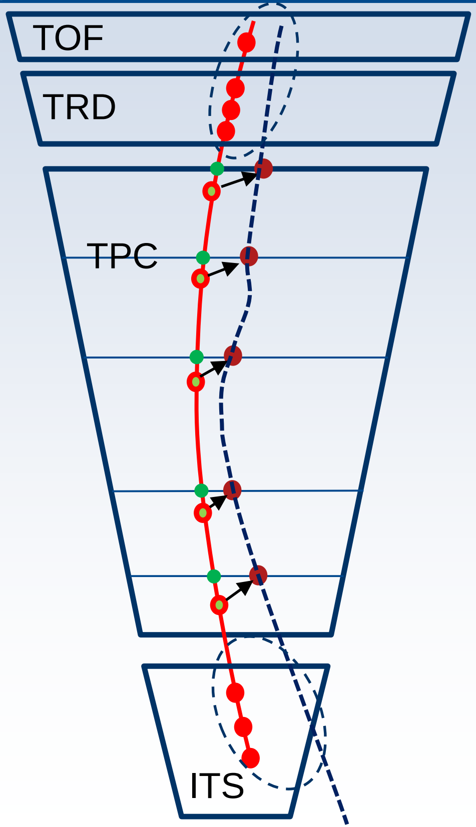

- Reconstruct TPC tracks with relaxed tolerances (applying “default distortion maps” if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

Space-charge Distortions corrections (in Pb-Pb Run 2)

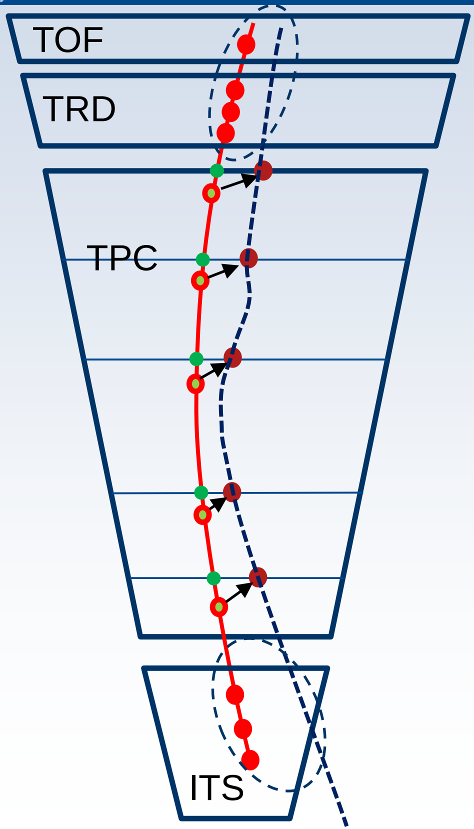

- Reconstruct TPC tracks with relaxed tolerances (applying “default distortion maps” if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

Space-charge Distortions corrections (in Pb-Pb Run 2)

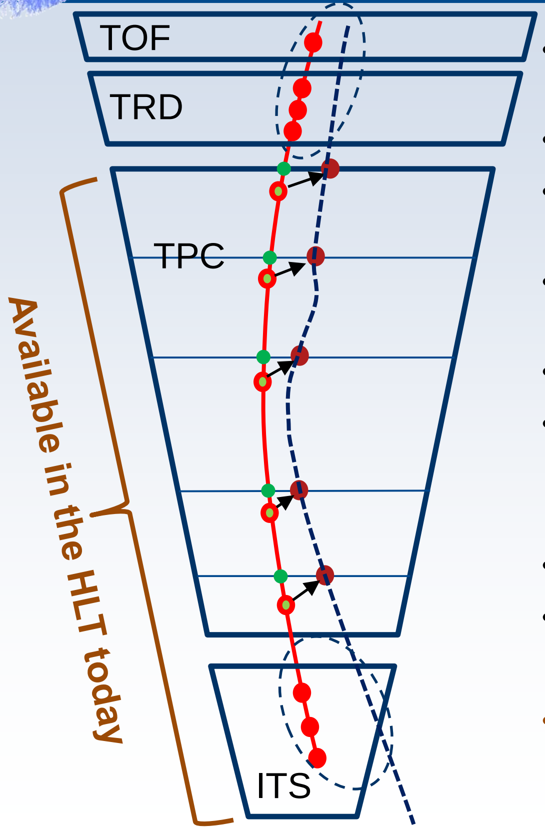

- Reconstruct TPC tracks with relaxed tolerances (applying "default distortion maps" if available)
- Match to ITS and TRD/TOF with relaxed tolerances
- Refit ITS-TRD-TOF part and interpolate to TPC as a **reference** of the **true track position** at every pad-row
- Collect Y, Z differences between **distorted clusters** and **reference** points in sub-volumes (voxels) of TPC
- Extract 3D vector of distortion in every voxel
- Create smooth parameterization (fast interpolation by Chebyshev polynomials) to use for correction during following reconstruction

- Distortions change with time: 40 min intervals. (min 15-20 min for statistics)
- 15 (in Y/X) x 5 (in Z/X) voxels per padrow ⇒ ~430K in total

- **For Run 3: In addition, integrate the digital currents at the TPC pads to account for fluctuations in the distortions.**

- **Running offline**
- **To be implemented for HLT during Run 2 as prototype for Run 3**
- **Needs TRD and TOF reconstruction in the HLT (not available yet)**

TPC TIME FRAME TRACKING IN O^2

David Rohr, drohr@cern.ch

Time frame tracking / robustness

- **The following plots…**
  - Compare the HLT / O² tracking for different scenarios (pile-up, time frame length, interaction rate, …)
  - Exact same data as before (same events, but reshuffled and arranged in time frames).
  - Raw QA of all tracks, no cuts except for $|\eta| < 0.9$.
  - Full time frame border simulated. Collisions with incomplete drift time ignored on QA (Quality Assurance) level.
  - In order to compare only the tracking algorithm, events are overlapped on the level of clusters → no clusterization effects.

Multiplicity / event pile-up (pp)

- Overlaying up to $\mu = 100$ pp TPC events (in-bunch pile-up) has absolutely no impact on efficiency, minimal impact on fake rate.
- At 300 overlaid pp events, one starts to see a small deterioration in the efficiency below 120 MeV/c.
- Above (at μ = 1000), there is a significant effect, but the tracking still works.
- Pile-up has does not affect resolution at all.

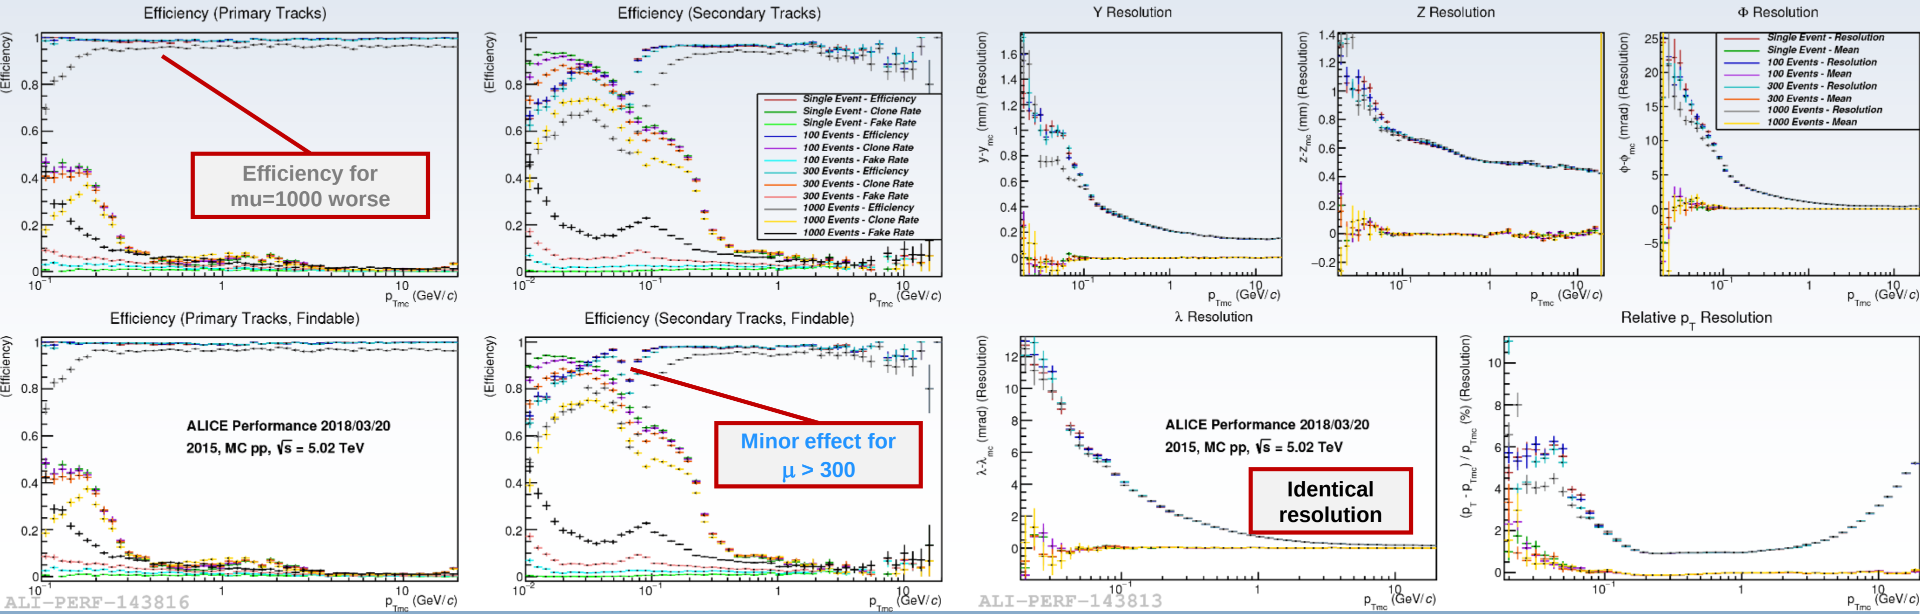

Normal tracking / z-independent tracking

- In continuous tracking, the absolute z-position of the track is not known, but estimated from the assumption that the track is primarily pointing towards the origin (B-field and cluster errors are computed under this assumption).
- Naturally, secondary tracks suffer a bit, while primaries are mostly unaffected.
- **No significant difference between Run 2 tracking and z-independent Run 3 tracking.**

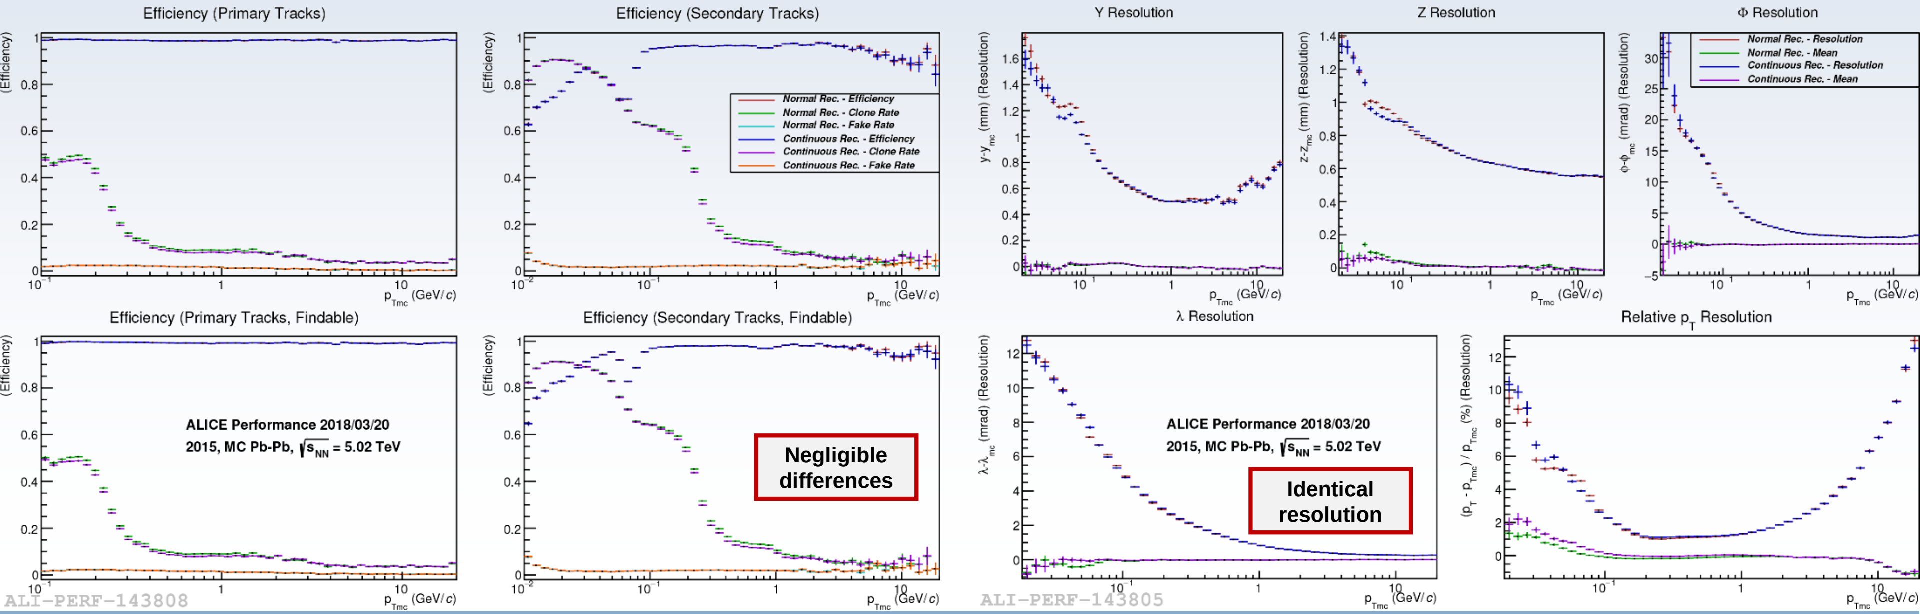

Length of time frame

- **Identical result independent of length of time frame.**
  - No problem with floating point precision / representation (thanks to shift in z).

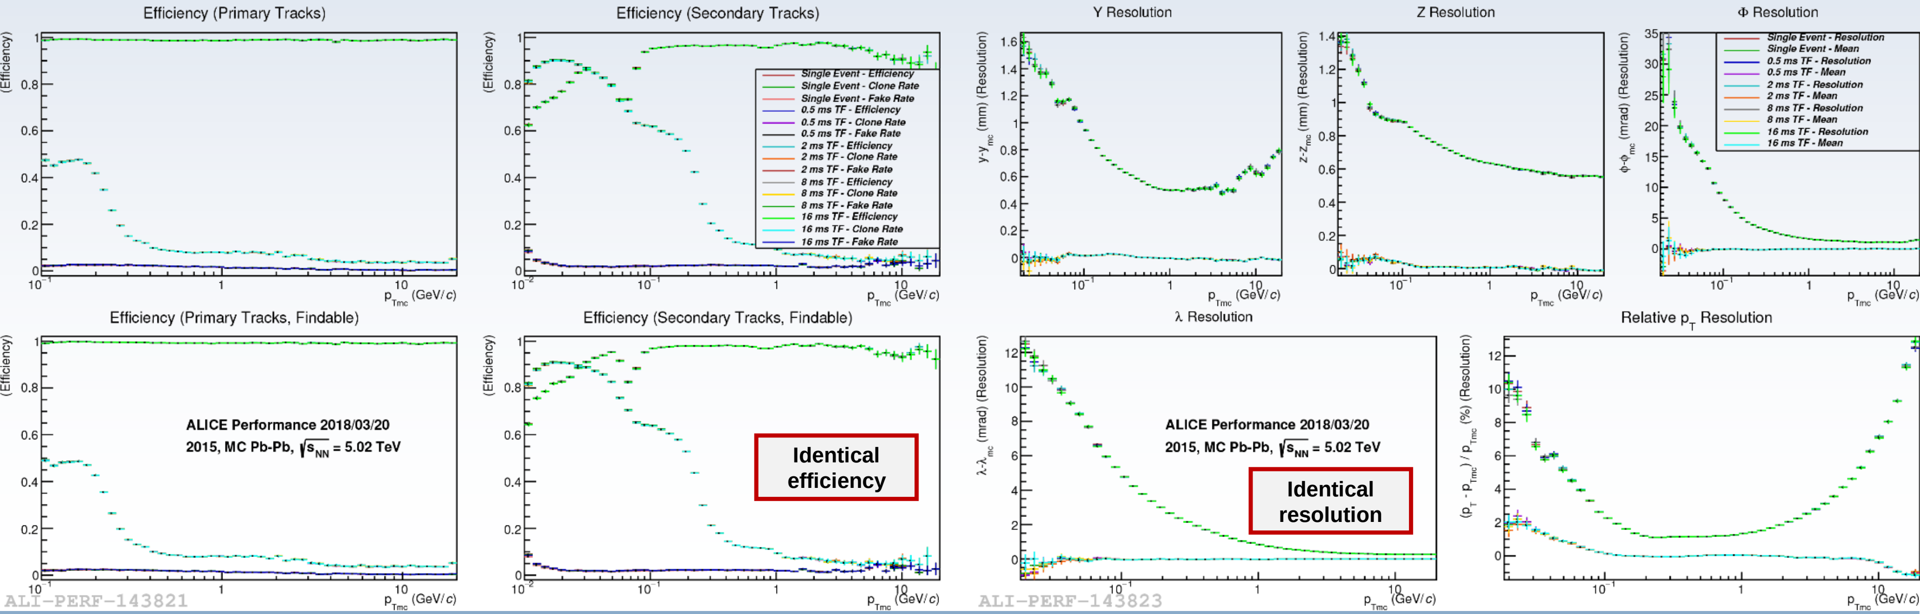

Tracking time frames at different interaction rates

- Simulation uses correct bunch structure as expected for Run 3 Pb-Pb(from ALICE TPC upgrade TDR).
- Practically no deterioration of resolution, even at 50 kHz.
- Minor efficiency decrease below 150 MeV/c.
- Still, fake rate increases with interaction rate (in particular for low p_T) – Should improve with better merging.

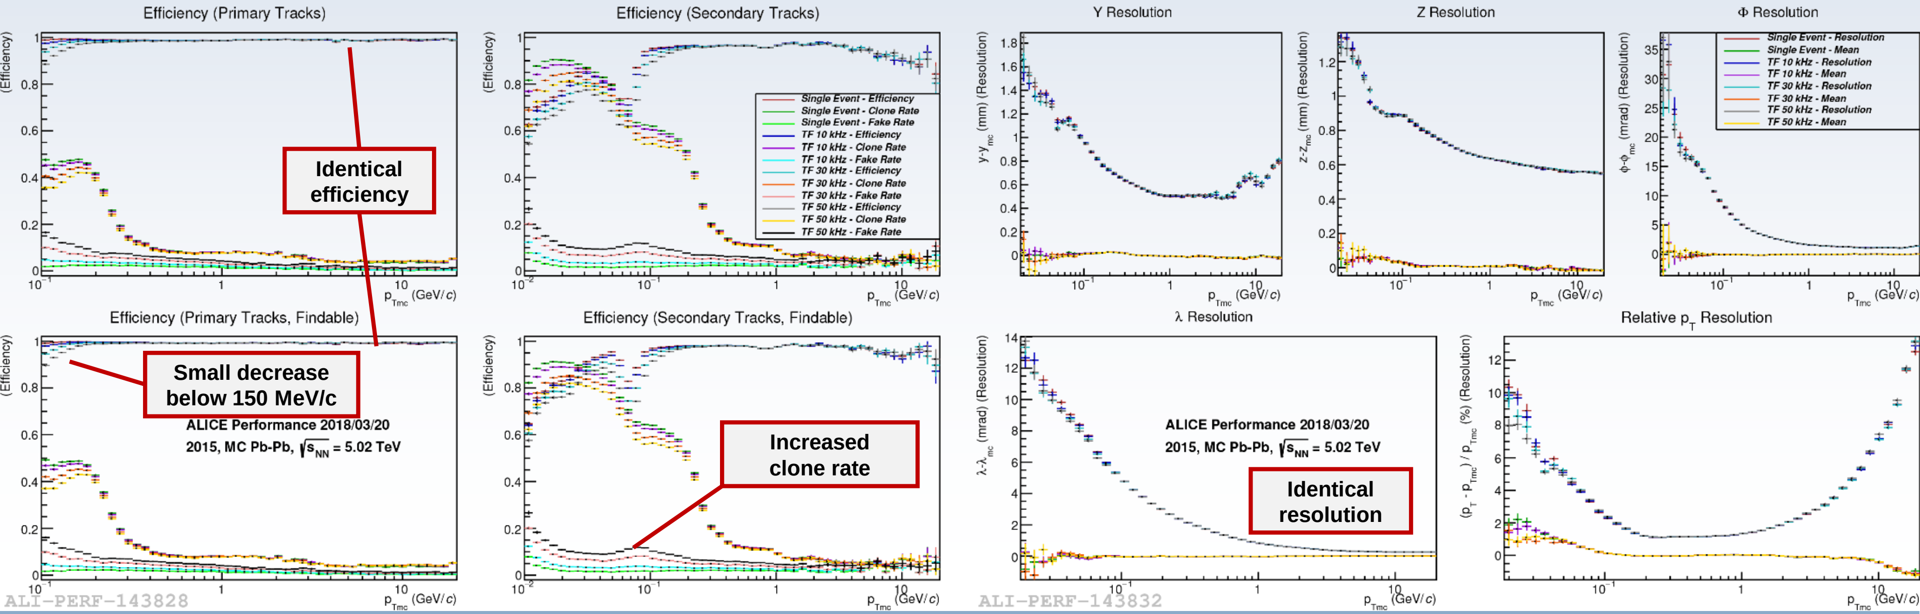

Tracking time (Run 2 / Pb-Pb)

- **Tracking time is linear with event size in all cases.**
- About 50 ms for largest events in Run 2. (Central Pb-Pb with pile-up.

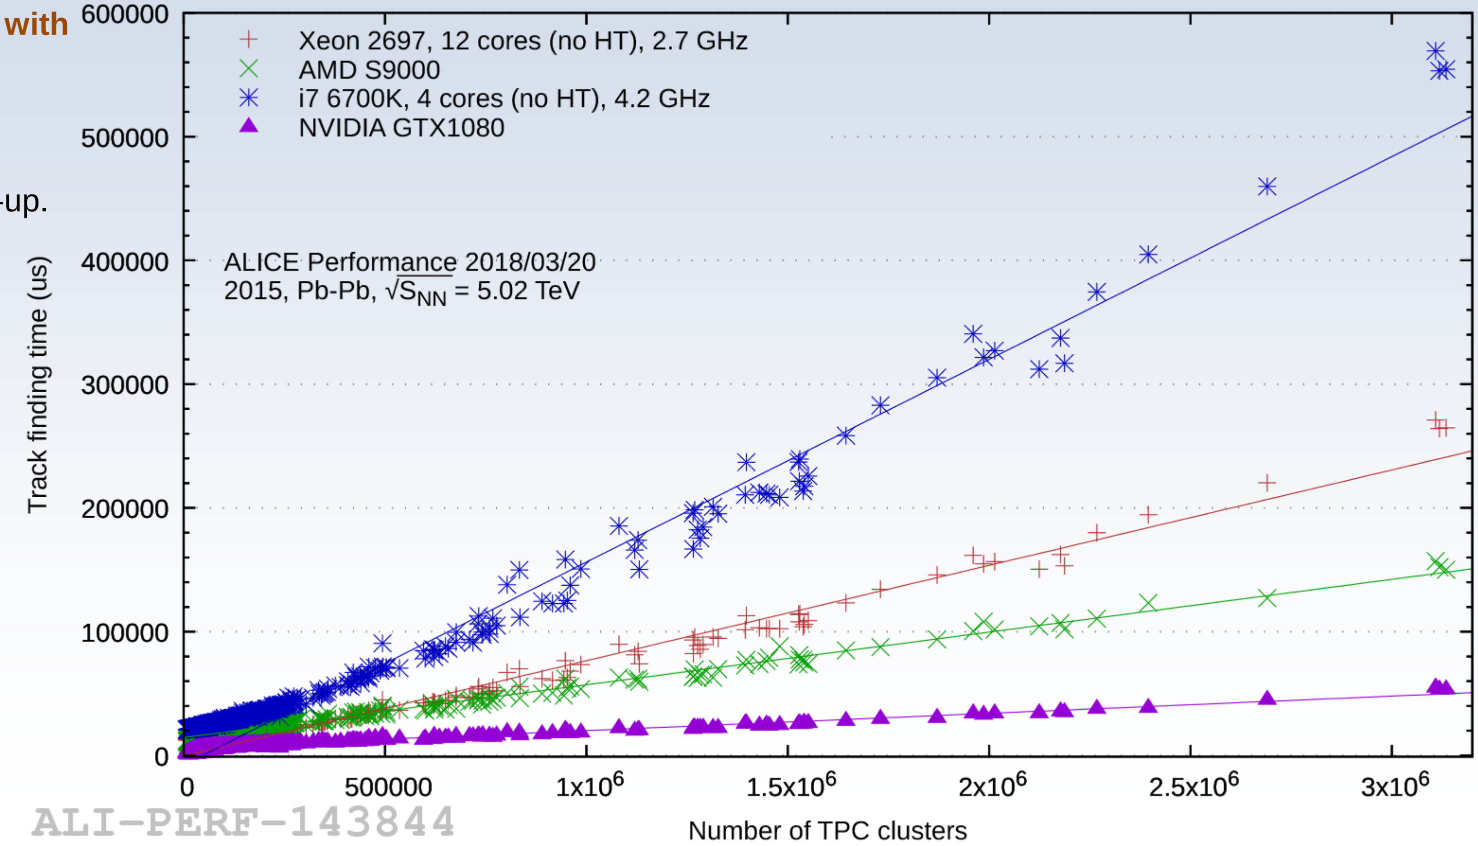

Tracking time

- Speed-up normalized to single core.
  - For the blue curve, this is exactly the speed-up.
  - For the GPU curves, this is corrected by the CPU resources required for GPU pre- and postprocessing.
    - How many cores does the GPU replace.
- Significant (>20x) speed-up compared to offline.
- A modern GPU replaces about 40 cores @4.2 GHz.
- Significant gain with new GPU models.
- One GPU replaces >800 CPU cores running Offline tracking. (at same efficiency / resolution)

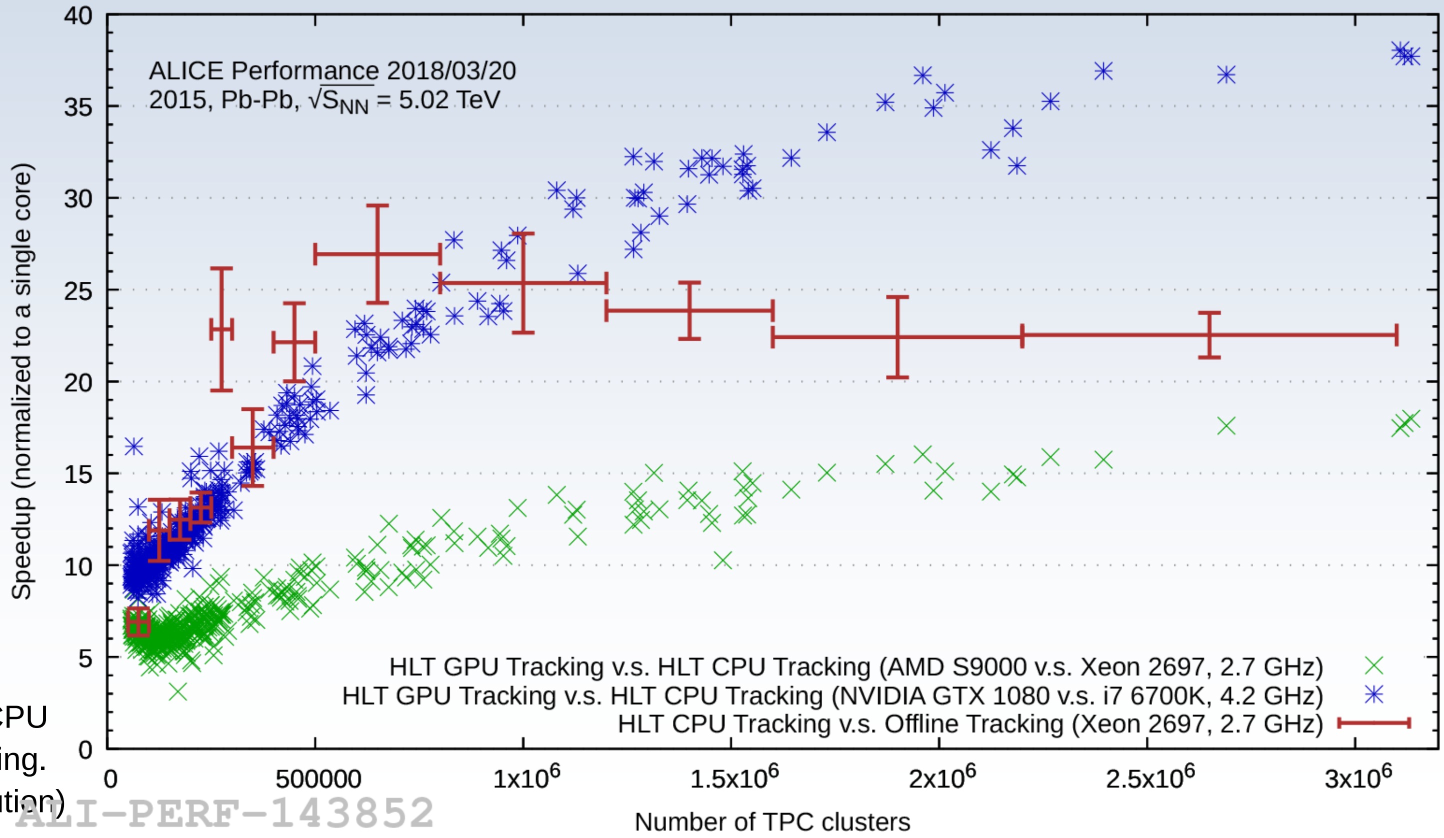

DATA COMPRESSION & LOOPER REJECTION

David Rohr, drohr@cern.ch

Data Compression

- **Data compression steps:**
  - Online cluster finding (in hardware on FPGA).
    - Enables better entropy compression.
    - Can perform noise suppression.
  - Entropy reduction steps.
    - Fixed point integer format.
    - Store differences instead of absolute positions.
    - Use track model to reduce position / charge entropy.
  - Entropy encoding.
    - Huffman encoding (Run 1 – 2) / Arithmetic encoding / ANS / etc.

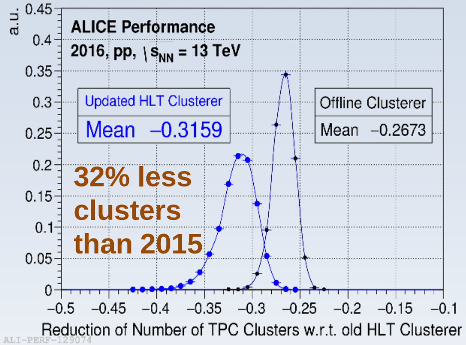

Data Compression

- **Data compression steps:**

Factor 7.2

- Online cluster finding (in hardware on FPGA).
  - Enables better entropy compression.
  - Can perform noise suppression.
- Entropy reduction steps.
  - Fixed point integer format.
  - Store differences instead of absolute positions.
  - Use track model to reduce position / charge entropy.
- Entropy encoding.
  - Huffman encoding (Run 1 – 2) / Arithmetic encoding / ANS / etc.

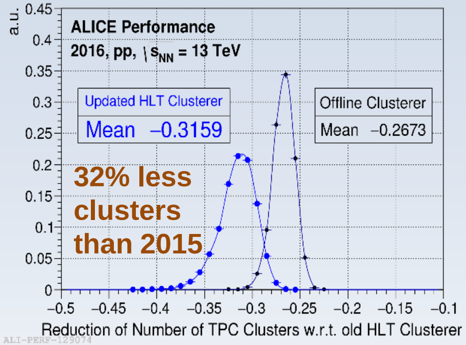

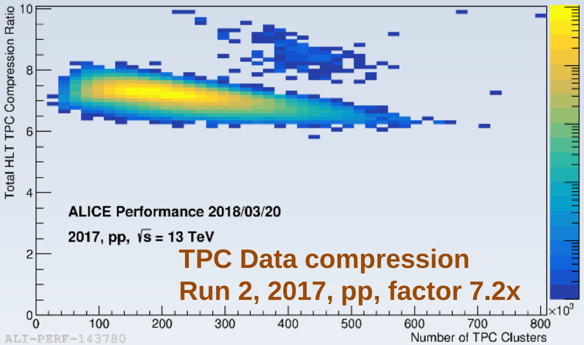

- Remove clusters of tracks not used for physics:
  (looping tracks below 50 MeV/c, additional legs of tracks below 200 MeV/c, track segments with high inclination angle)
  - Normal TPC tracking finds low-p_T tracks down to 10 MeV/c.
  - “Afterburner” (Hough transform, machine learning, etc.) finds what is left.

Work in Progress

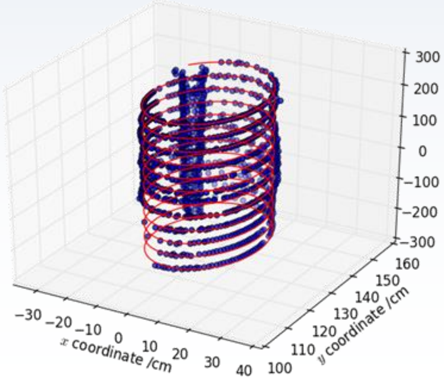

TPC Data Compression

- **Overview of compression achieved per TPC compression step**

Configuration							
Data	2013	2016 pp	2015 Pb–Pb	2017 pp	2017 pp	*2016 pp*	*2015 Pb–Pb*
TPC gas	neon	argon	argon	neon	neon	*argon*	*argon*
RCU version	1	2	1	2	2	*2*	*2*
Cluster finder version	run 1	old	old	old	improved	*improved*	*improved*
Compression version	run 1 / 2	run 1 / 2	run 1 / 2	run 1 / 2	run 1 / 2	*run 3 prototype*	*run 3 prototype*
Compression step							
Cluster finder	1.20x	1.50x	1.28x	1.42x	1.81x	*1.72x*	*1.70x*
Branch merging	1.05x	-	1.05x	-	-	-	-
Integer format	2.50x	2.50x	2.50x	2.50x	2.40x	*2.40x*	*2.40x*
(bits per cluster)	77 bits	77 bits	77 bits	77 bits	80 bits	*80 bits*	*80 bits*
Entropy reduction							
Position differences	-	2% / -1.2 bits	16% / -7.2 bits	2% / -1.0 bits	2% / -1.0 bits	*-1.0 bits*	*-4.5 bits*
Track model	-	-	-	-	-	*-14.5 bits*	*-14.3 bits*
Track model + differences	-	-	-	-	-	*-8.0 bits*	*-8.41 bits*
Logarithmic precision	-	-	-	-	15% / -6.6 bits	*-7.3 bits*	*-7.3 bits*
Entropy encoding							
Huffman coding	1.36x	1.49x	1.75x	1.46x	1.68x	*2.08x*	*2.12x*
Arithmetic coding	-	-	-	-	-	*2.18x*	*2.22x*
Total compression	4.26x	5.58x	5.89x	5.18x	7.28x	*9.00x*	*9.10x*
(bits per cluster)	56.6 bits	51.7 bits	44.0 bits	52.8 bits	47,7 bits	*36,7 bits*	*36,0 bits*

TPC Data Compression

- **Overview of compression achieved per TPC compression step**

Configuration							
Data	2013	2016 pp	2015 Pb–Pb	2017 pp	2017 pp	*2016 pp*	*2015 Pb–Pb*
TPC gas	neon	argon	argon	neon	neon	*argon*	*argon*
RCU version	1	2	1	2	2	*2*	*2*
Cluster finder version	run 1	old	old	old	improved	*improved*	*improved*
Compression version	run 1 / 2	run 1 / 2	run 1 / 2	run 1 / 2	run 1 / 2	*run 3 prototype*	*run 3 prototype*
Compression step							
Cluster finder	1.20x					*1.72x*	*1.70x*
Branch merging	1.05x					-	-
Integer format	2.50x					*2.40x*	*2.40x*
(bits per cluster)	77 bits					*80 bits*	*80 bits*
Entropy reduction							
Position differences						*-1.0 bits*	*-4.5 bits*
Track model						*-14.5 bits*	*-14.3 bits*
Track model + differences						*-8.0 bits*	*-8.41 bits*
Logarithmic precision						*-7.3 bits*	*-7.3 bits*
Entropy encoding							
Huffman coding	1.36x					*2.08x*	*2.12x*
Arithmetic coding						*2.18x*	*2.22x*
Total compression	4.26x	5.58x	5.89x	5.18x	7.28x	*9.00x*	*9.10x*
(bits per cluster)	56.6 bits	51.7 bits	44.0 bits	52.8 bits	47,7 bits	*36,7 bits*	*36,0 bits*

Run 3 compression prototype: factor 9.1.

- **Track model compression stores cluster to track residuals instead of absolute positions.**
- **Arithmetic encoding (small improvement of 4-5%.**
- **General improvements.**

Next step: reject clusters not used for physics.
Goal: 20x total compression factor.

Track model compression

Work in Progress

- **Idea: store only residuals of clusters to extrapolated track.**
  - **Smaller entropy** than absolute (differential) coordinates → **Better Huffman compression**.
  - **Constraint**: Clusters shall be stored in native TPC coordinates (Row, Pad, Time), independent from calibration.

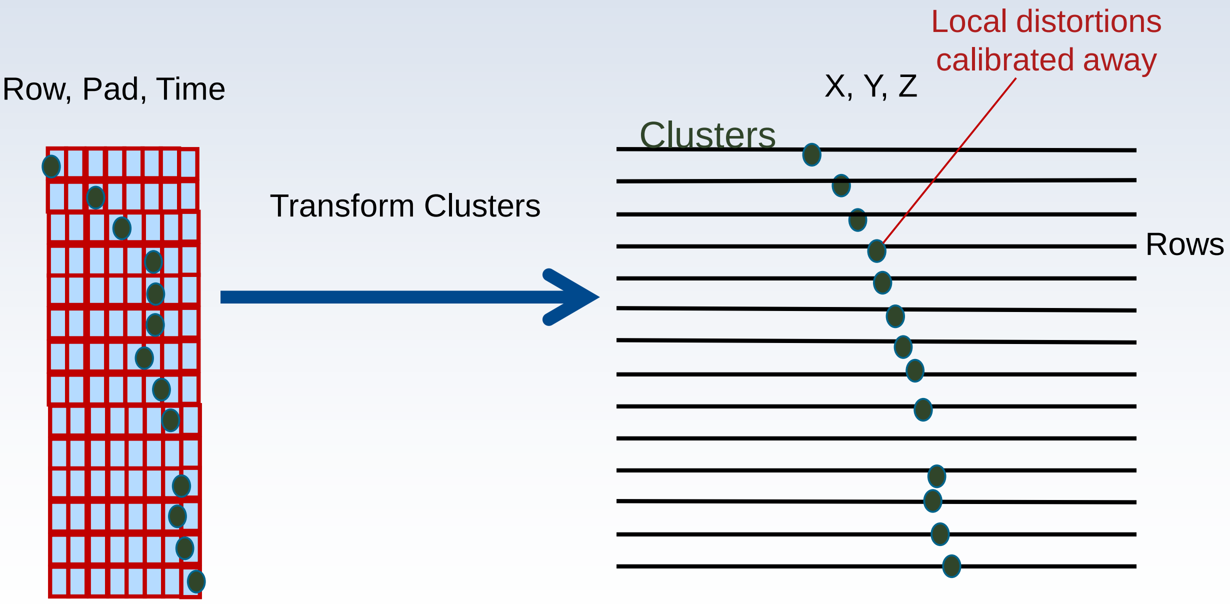

Track model compression

Work in Progress

- **Idea: store only residuals of clusters to extrapolated track.**
  - **Smaller entropy** than absolute (differential) coordinates → **Better Huffman compression**.
  - **Constraint**: Clusters shall be stored in native TPC coordinates (Row, Pad, Time), independent from calibration.

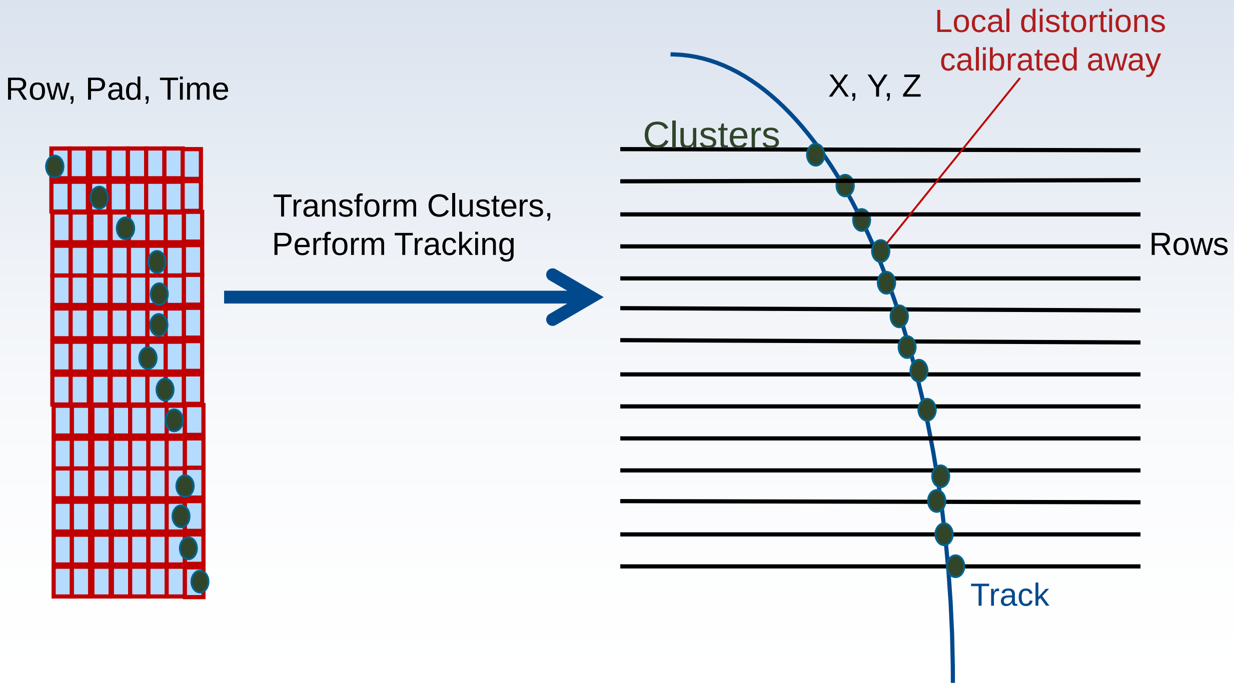

Track model compression

Work in Progress

- **Idea: store only residuals of clusters to extrapolated track.**
  - **Smaller entropy** than absolute (differential) coordinates → **Better Huffman compression**.
  - **Constraint**: Clusters shall be stored in native TPC coordinates (Row, Pad, Time), independent from calibration.

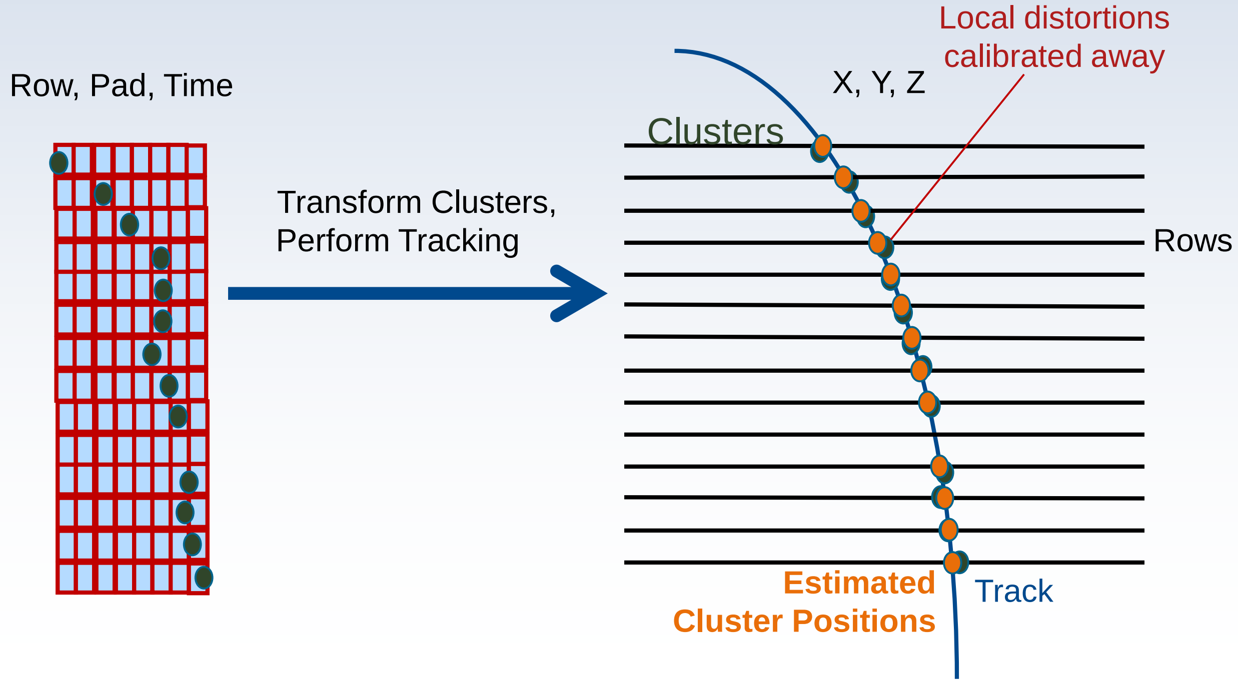

Track model compression

Work in Progress

- **Idea: store only residuals of clusters to extrapolated track.**
  - **Smaller entropy** than absolute (differential) coordinates → **Better Huffman compression**.
  - **Constraint**: Clusters shall be stored in native TPC coordinates (Row, Pad, Time), independent from calibration.

- **Problems**:
  - Helix prolongation yields **large residuals** → inefficient compression.
  - Linear back-transformation **cannot revert transformation** based on full calibration.
  - Helix **cannot accommodate space charge distortions** in the TPC.

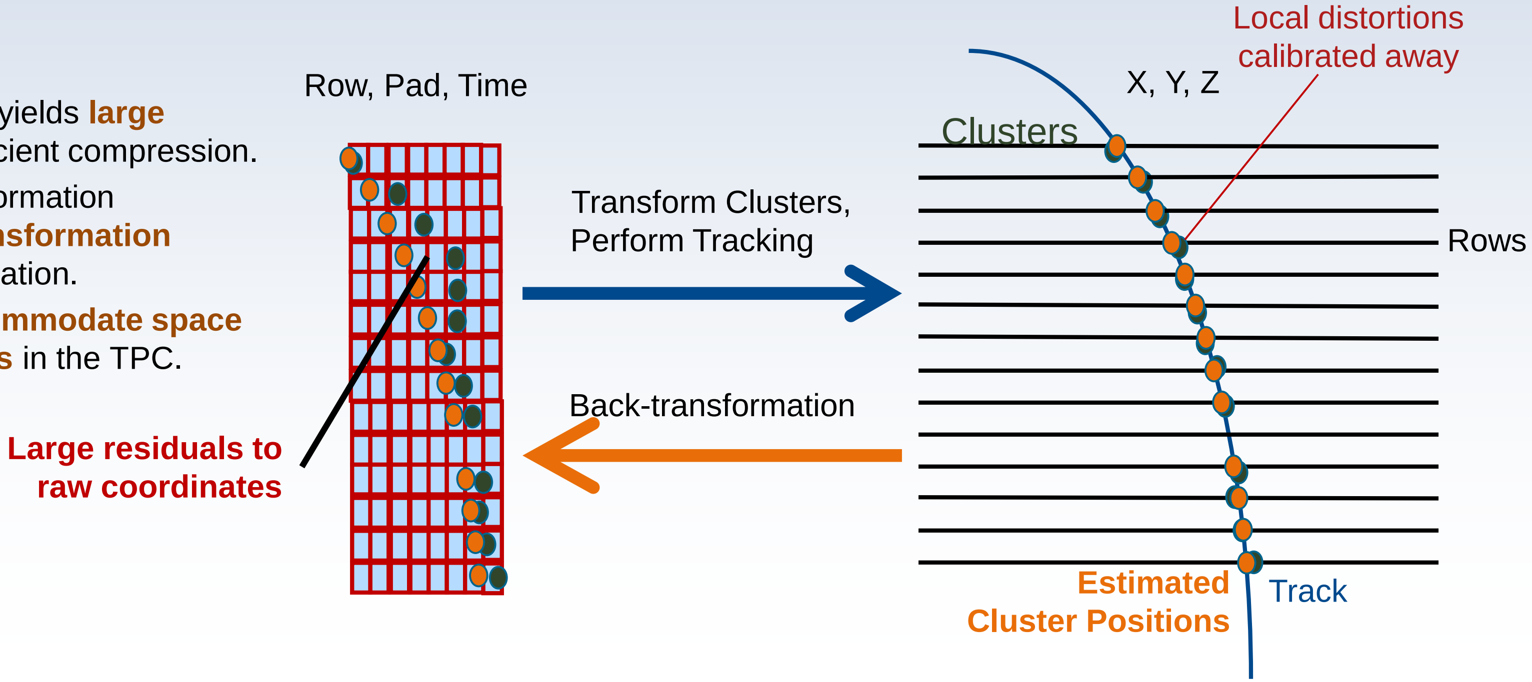

Track model compression

Work in Progress

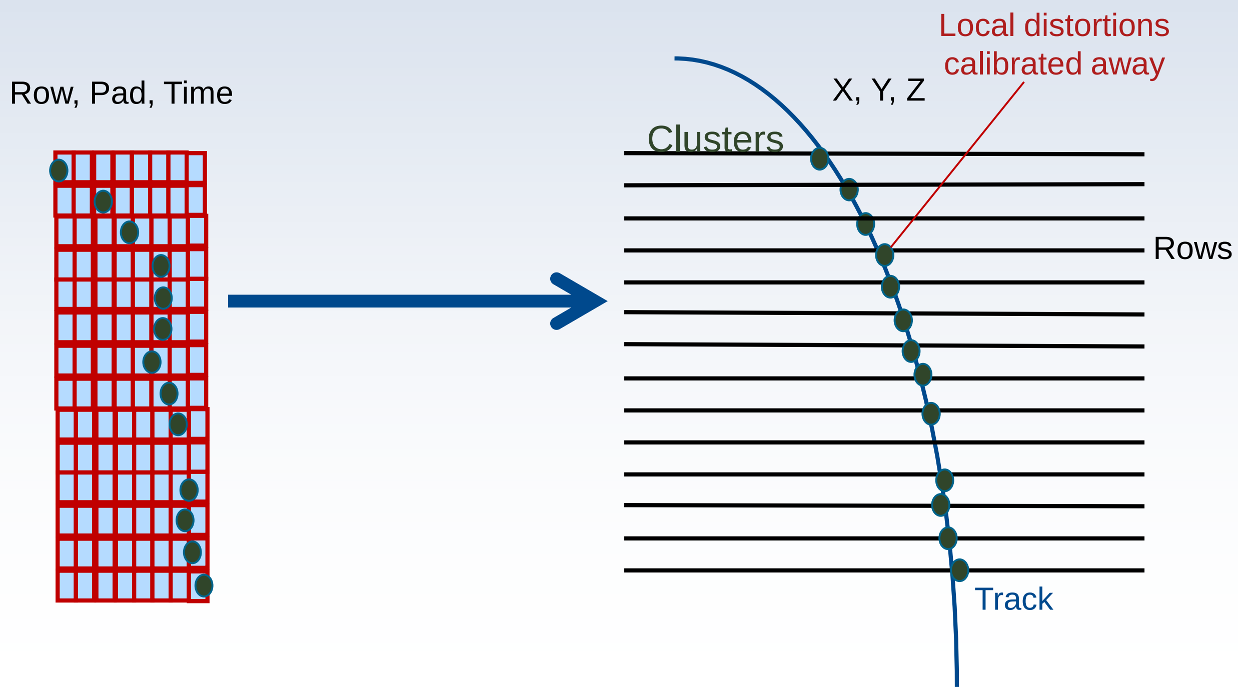

Track model compression

Work in Progress

- **Solution:**
  - Employ **fast, reversible polynomial approximation**. *(In principle, every transformation works, but the closer the better!)*

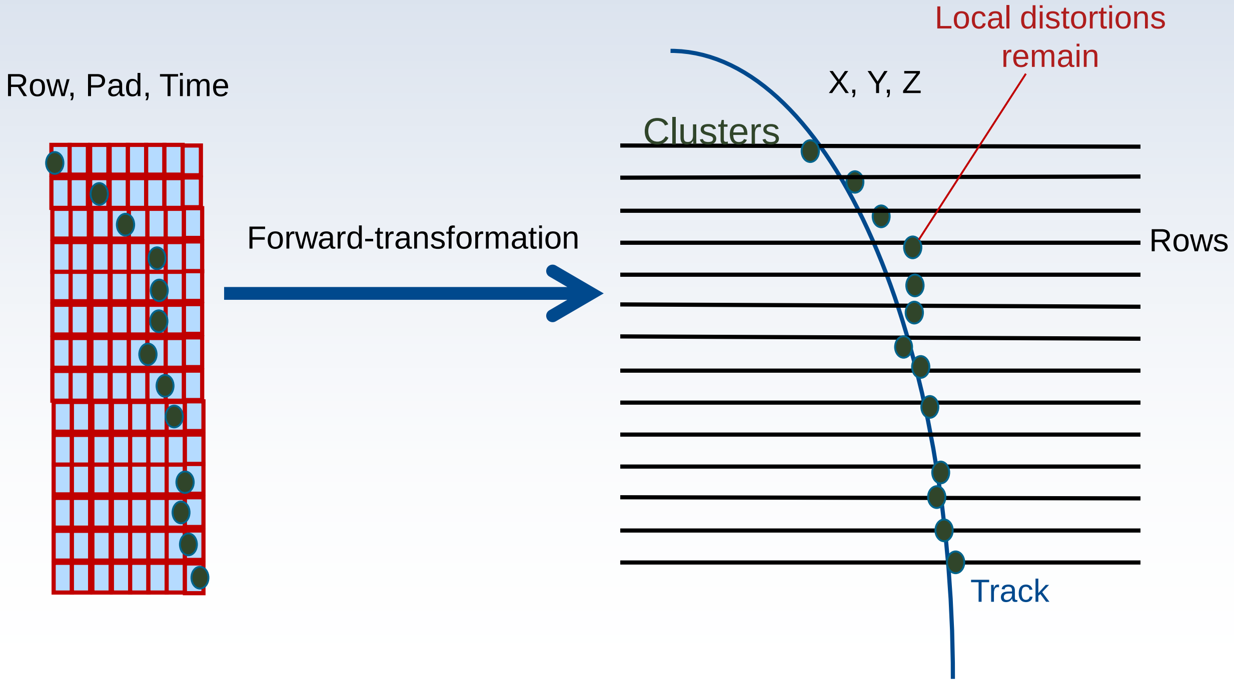

Track model compression

Work in Progress

- **Solution:**
  - Employ **fast, reversible polynomial approximation**. *(In principle, every transformation works, but the closer the better!)*
  - **Refit track in distorted coordinate system**.

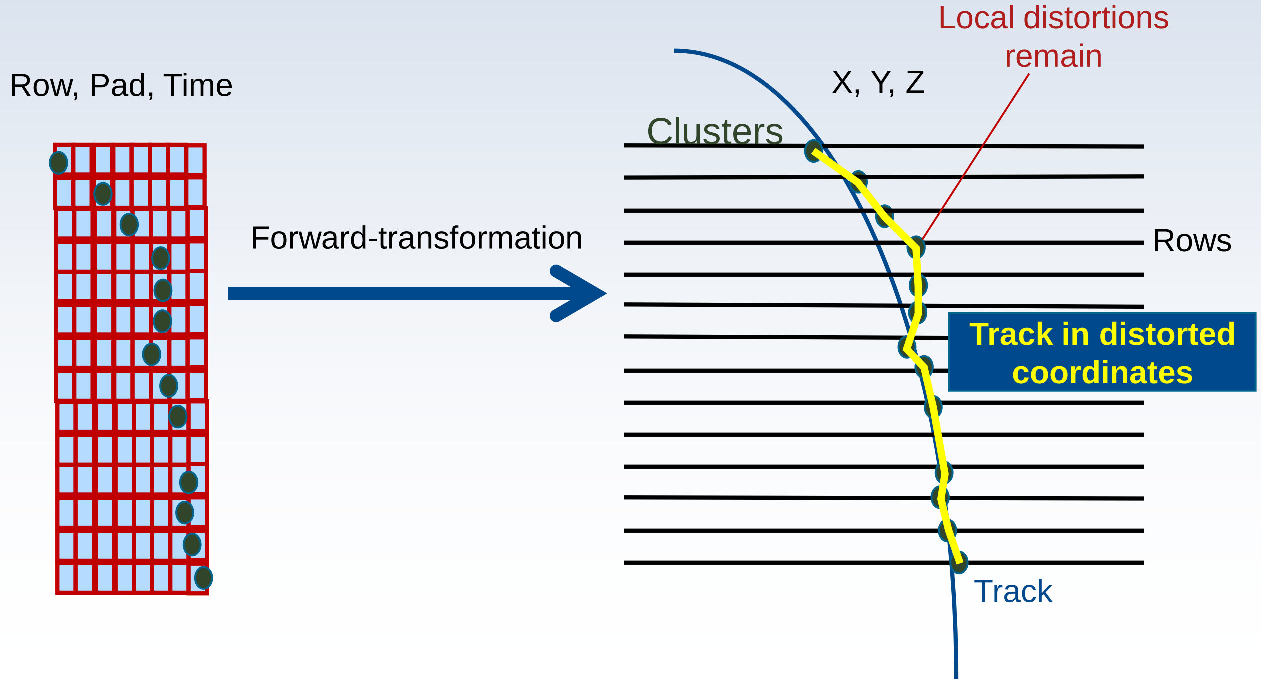

Track model compression

Work in Progress

- **Solution:**
  - Employ **fast, reversible polynomial approximation**. *(In principle, every transformation works, but the closer the better!)*
  - **Refit track in distorted coordinate system**.

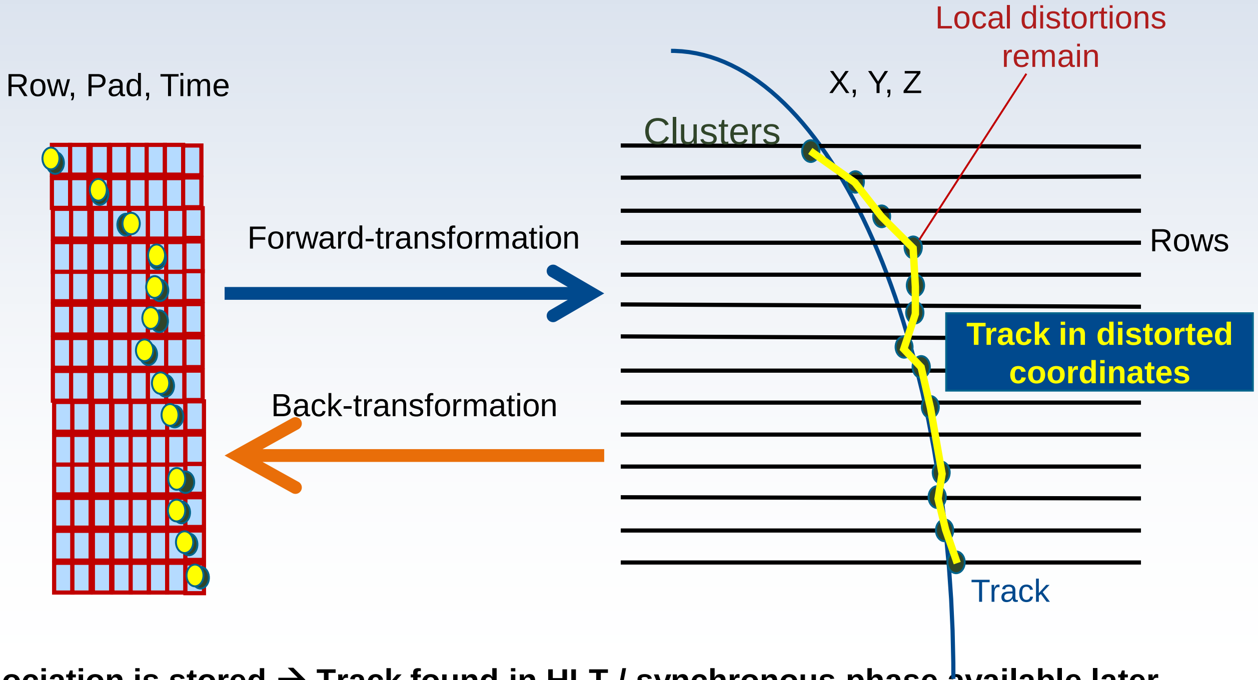

- **Non-associated clusters still compressed with differences scheme.**
- **Additional benefit: Cluster to track association is stored → Track found in HLT / synchronous phase available later.**

Tracking of low-p_T tracks

- **Track-merging and fit improved for low-p_T:**
  - Most legs reconstructed correctly (green)
  - Refit fails rarely (only 1 white leg left).
  - Most legs merged on at least one side.
    - Cannot merge on both sides right so far.
  - Some seeds left, which do not make it to tracks.
    - Track fit fails for seeds below 10 MeV/c.
    - Some track-fit failures remain to be understood.

- Green: Final tracks
- Blue: Unused clusters
- Purple: Segments found in first CA seeding phase but track prolongation did not find good track.
- White: Track prolongation found track, but the track is rejected later.

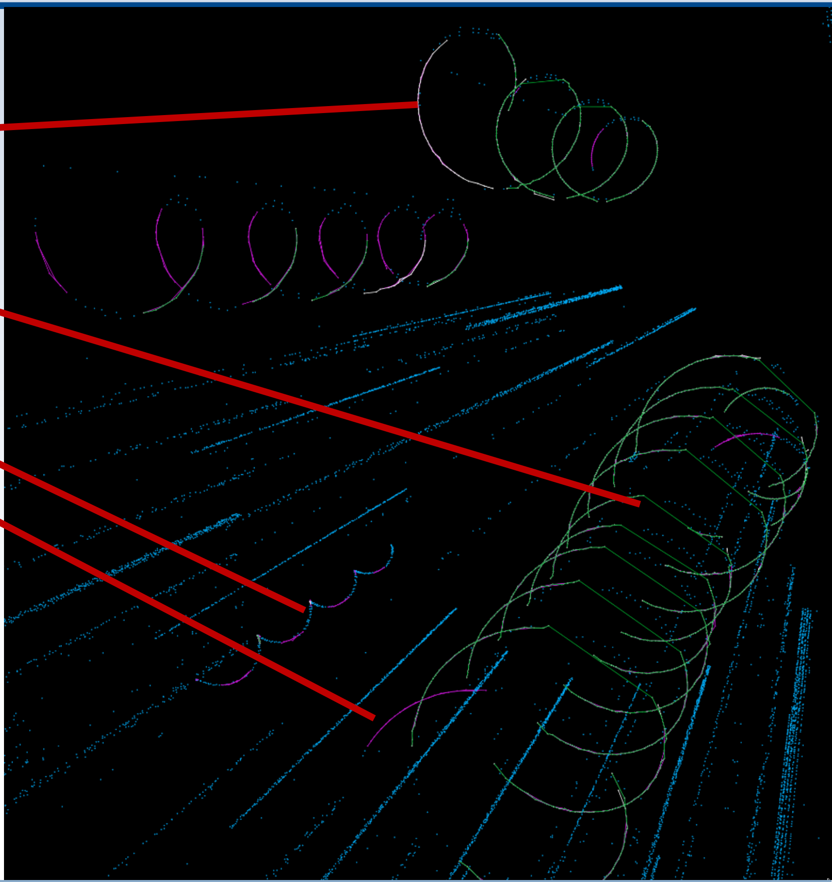

Low-p_T looper finding

- Very good efficiency down to 10 MeV/c.
- Practically zero fake-cluster attachment for single Pb-Pb event processing.
- Small but non-zero fake attachment at 50 kHz.
- Small difference in fraction of attached hits.
- With additional feature to attach all adjacent hits, attachment rate should be close to the blue / orange curves.

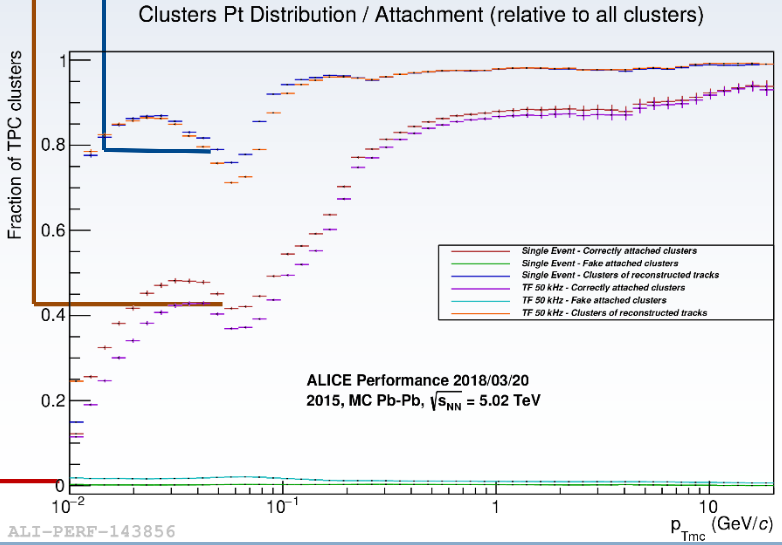

Efficiency, Pb-Pb, 50kHz

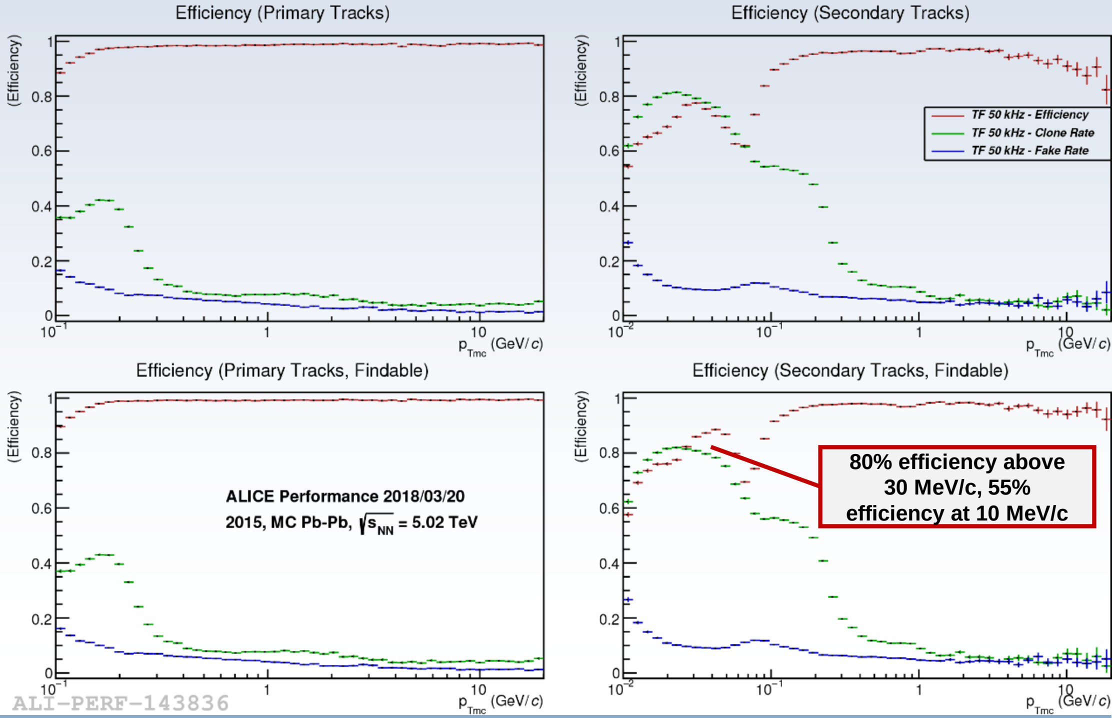

Cluster removal for O^2

- Strategy: Track as low in p_T as possible with relaxed cuts (number of clusters, etc.), merge legs to get sufficient track quality.
  - Attach all adjacent clusters using inter- / extrapolation.
- Remove all clusters assigned to
  - Tracks below 50 MeV/c.
  - Additional legs of tracks below 200 MeV/c.
  - High inclination-angle track segments.
- Use other method to remove what is left after tracking.
- Cluster statistics:
  - Purple: all clusters
  - Red: clusters attached to the correct track.
  - Green: clusters attached to wrong track.
  - Blue: All clusters (if attached or not) of a reconstructed track.
- *Shared clusters and multiple-attached clusters are shown multiple times weighted correctly, so that the integral yields the total number of clusters.*

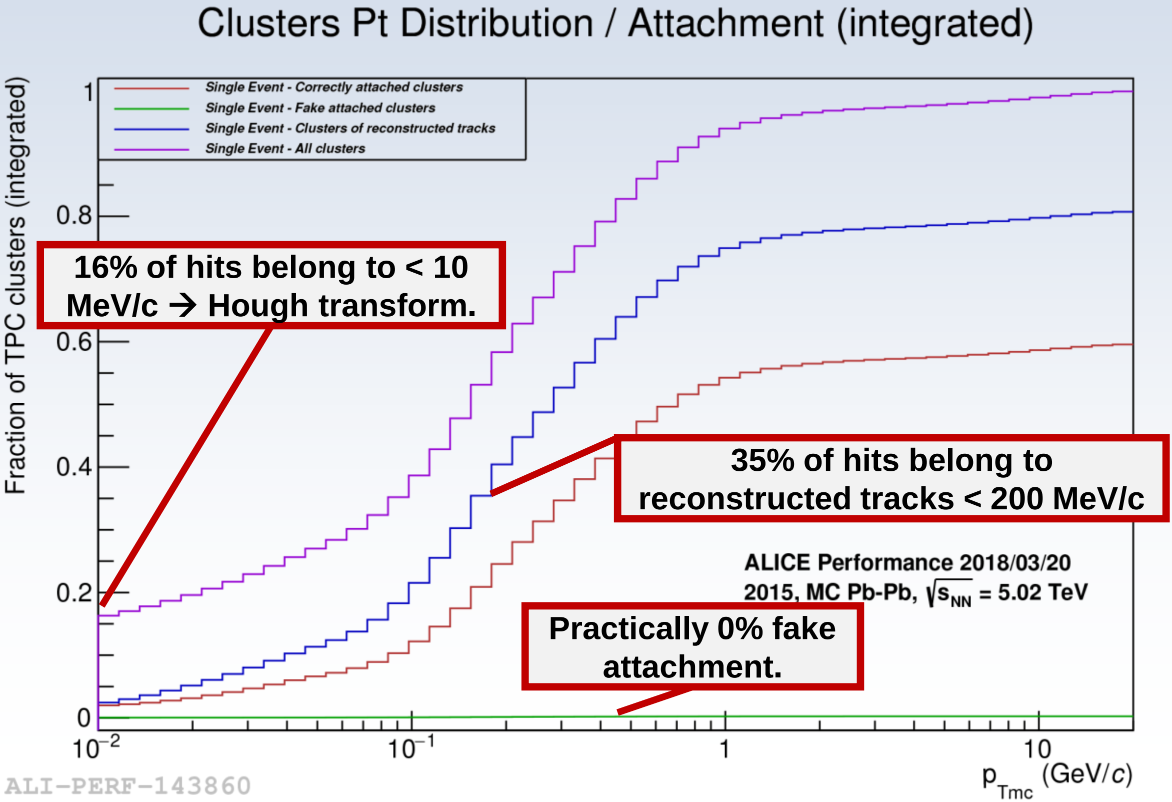

OTHER DETECTORS & THE GLOBAL PICTURE

Online TRD tracking in the HLT and for Run 3

Work in Progress

- **Space-Charge Distortion Calibration requires online TRD tracking.**
  - Offline tracking uses **offline tracklets** created from hits around TPC tracks.
  - **Not available in Run 3**, use **online tracklets** created by TRD readout.
- **Same matching efficiency for online and offline tracklets > 1.5 GeV/*c*.**
  - Efficiency drop at low p_T is caused by the absence of online tracklets (currently optimized for electron trigger at p_T > 3 GeV/*c*)
- **Algorithm: Extrapolate TPC track through TRD layers, find closest tracklet.**
- **Next steps:**
  - Decision-tree based algorithm foreseen if needed for high occupancy.
  - Extending matching to p_T ~ 0.6 GeV/*c* is important both for:
    - Disentangling between radial and rϕ distortions.
    - Bridging TPC tracks to TOF.
  - Add TOF detector as well.

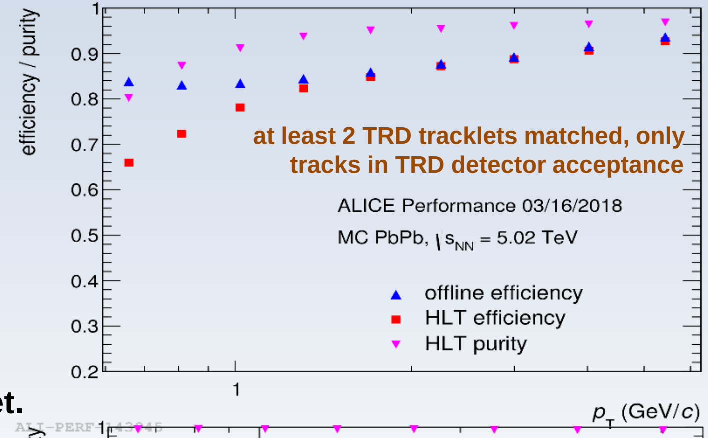

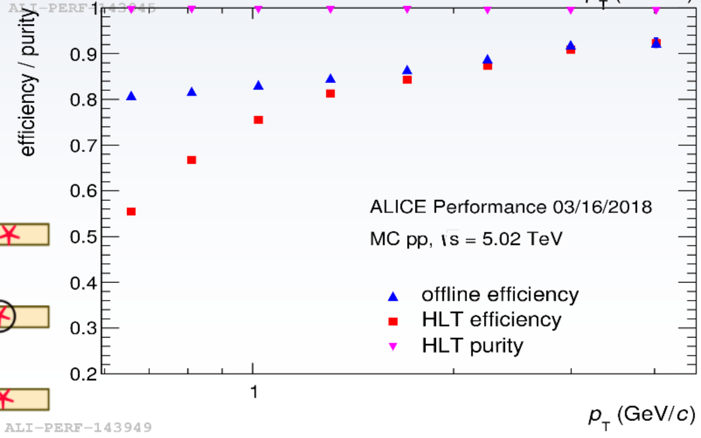

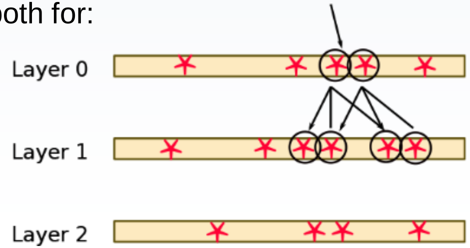

TPC-ITS matching performance

Work in Progress

- **TPC-ITS matching for O² implemented and principally working.**
  - Currently, significant amount of fake matches.
  - Unmerged loopers cause many fake matches.
  - Still working with loose time bracketing to avoid problems with inaccurate TPC vertex time estimate.

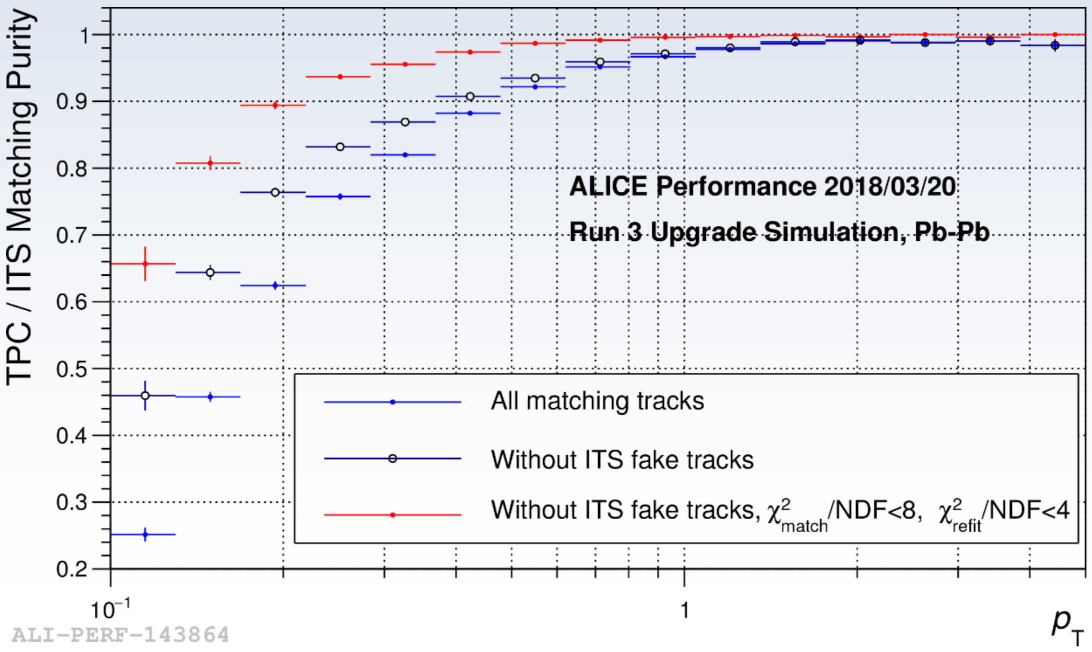

Summary

- **ALICE Run 3 upgrade is a major challenge for online computing / tracking / data compression.**
  - Improved HLT tracking used for O².
  - Identical resolution, equal or better efficiency compared to Run 2 offline tracking.
- **TPC tracking very robust.**
  - Length of time frame plays no role.
  - No effect until μ = 100, negligible effect at μ = 300, operational at reduced efficiency at μ = 1000.
  - No significant problem with interaction rates up to 50 kHz (only fake rate increases, should be mitigated by better merging).
  - Same resolution and efficiency for z-independent time frame tracking.
- **GPU replaces about 40 CPU cores @ 4.2 GHz (GTX 1080 v.s. Core i7 6700K).**
  - Tracking time is completely **linear** with number of TPC clusters.
  - > 20x faster than Run 2 offline tracker (on the CPU).
- **Good low-p_T tracking efficiency for secondaries / loopers down to 10 MeV/c (55% at 10 MeV/c, > 80% above 30 MeV/c).**
  - Will allow rejection of significant fraction of clusters attached to tracks not used for physics.
  - Remaining clusters compressed with factor **9.1** using track model and entropy encoding, aiming for 20x total compression.
- **TPC and ITS tracking / merging available in O² software.**
  - TRD tracking and TPC data compression being commissioned.
  - Online calibration scheme tested partially in Run 2 HLT.

APPENDIX

Tracking efficiency / resolution (Run 2, HLT v.s. Offline – pp)

- **For reference (same situation for pp).**
  - Identical resolution.
  - Same efficiency for primaries.
  - Better efficiency for secondaries / low p_T.

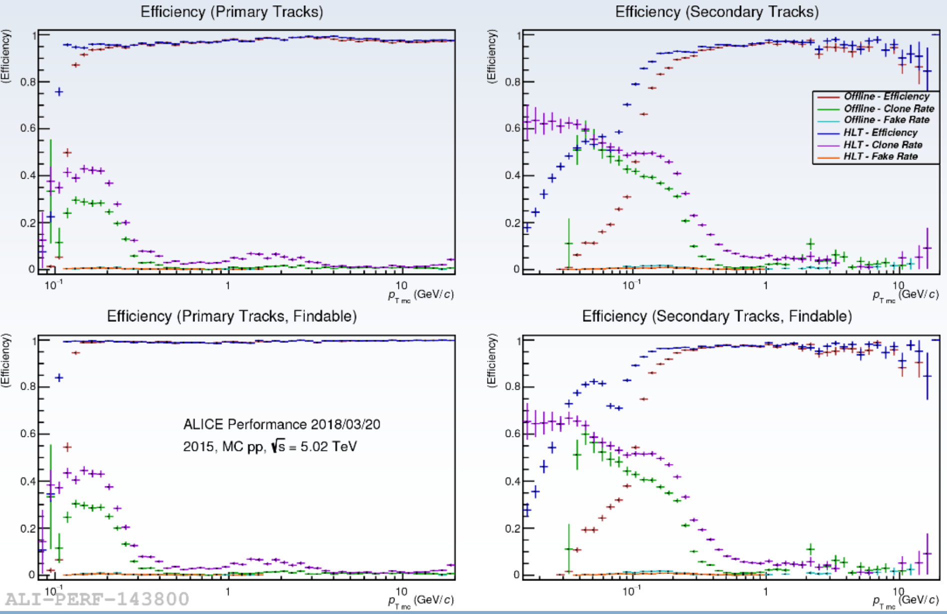

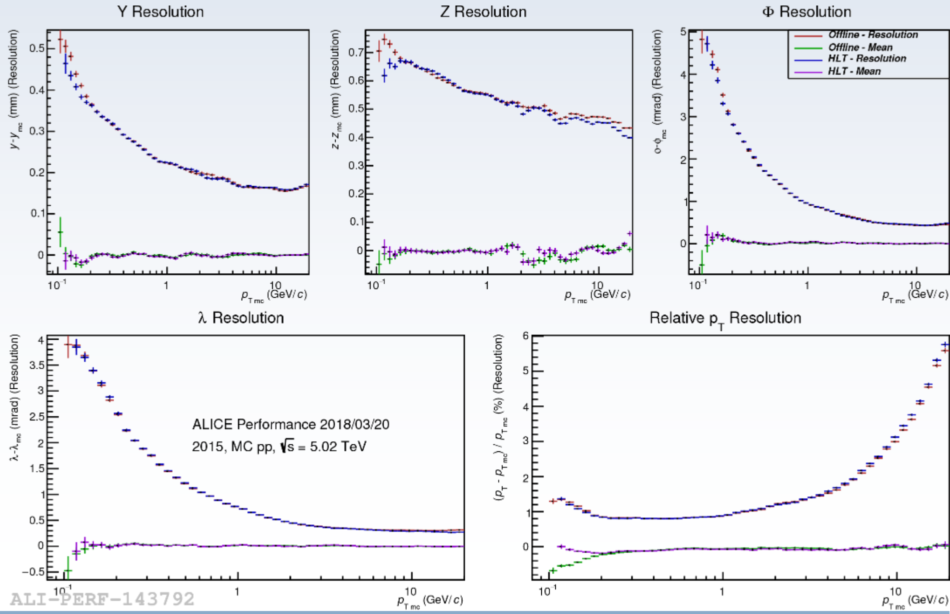

TPC Track Reconstruction Status

- **List of tracking improvements**
  - Some improvements to tracking efficiency and resolution, in particular low-p_T
  - Finally implemented fully z-independent tracking for O² (can shift clusters in Z, assuming it is a primary pointing to the vertex at Z = 0).
  - Finally performed tuning of cluster rejection.
  - Using offline error parameterization.
  - Some features needed for O² TPC-ITS matching.
  - Finalized work on 3-way track fit with full 3d magnetic field parameterization.
  - Full O2 / HLT tracking available in official ALICE repositories and default build.
  - Outer TPC parameters available in HLT farm for TRD tracking.
- **Still missing:**
  - Full low-p_T merging still not implemented (clone rate higher than before due to better efficiency).
  - Prolongation and merging across central electrode.
  - Attaching all adjacent clusters to tracks – for looper rejection and better track model compression.
  - Test HLT tracking with run 3 distortions.
  - Update of compression prototype / retest compression in run 3 scenario.
- **Related:**
  - Work on calibration / transformation framework (O² needs tracking in pad, row, time → transformation during the tracking on GPU).
  - TPC Online Calibration could now provide histograms for TPC online QA.

TPC Track Reconstruction Status

- **In summary: what has changed:**
  - With all the recent tuning, in particular cluster rejection, and usage of offline cluster error parameterization, HLT / O² TPC tracking reaches equal resolution as offline track reconstruction (in overall plot, special cases like tracks going through the dead zone might still be different).
  - Track fit time increases significantly (3-way fit, etc.), but still faster than track finding – and room for tuning.
  - Full TPC tracking available in O² in official repository, used for TPC / ITS matching.
  - All major apparent issues solved.

  - I'd consider this a good baseline now (having the same resolution as offline in Run 2).
    - From now, we go on and optimize for Run 3 – we only need to maintain the resolution and we'll be as good as today.

- **The following plots…**
  - …basically repeat all of the comparisons shown during last offline weeks / tracking presentations / conferences, using the most recent tracking version.
  - Many artifacts seen before have vanished.
  - Findable TPC tracks for efficiency are required to have at least 70 TPC hits (other tracks need to have at least 1 TPC hit).
  - Cuts are $|\eta| < 0.9$, $p_T > 200$ MeV/c.
  - The comparison plot to offline applies in addition to the standard cuts from AliPerformanceRes/AliPerformanceEff.
  - All plots use same calibration / error parameterization from Offline also for HLT.
  - Comparing pure tracking algorithm.
  - dE/dx and related features (x-talk / ion tail) disabled in offline for direct comparison (affects only dE/dx and timing, but not resolution.)