GPU-based Online Track Reconstruction in LHC Run 3 for the ALICE TPC with Continuous Read Out

David Rohr for the ALICE Collaboration

drohr@cern.ch, CERN

CHEP 2018

10.7.2018

Run 2: O(1) kHz single events…

- **What data will we see with the ALICE Upgrade…?**

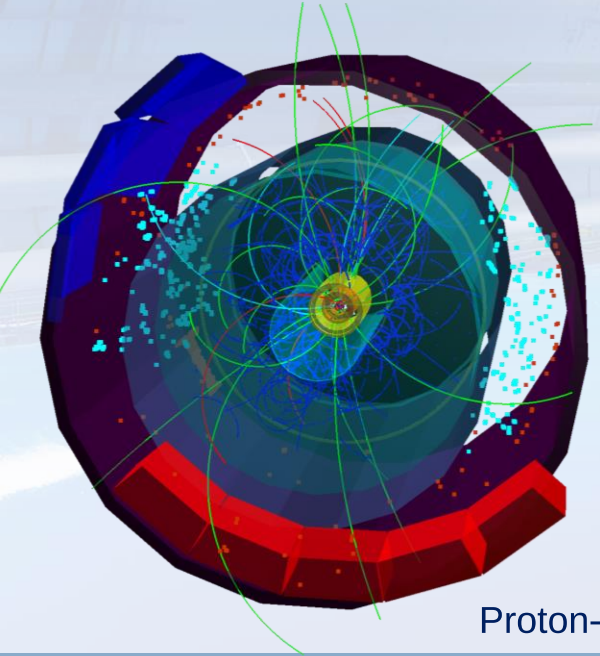

Proton-Proton

Run 2: O(1) kHz single events…

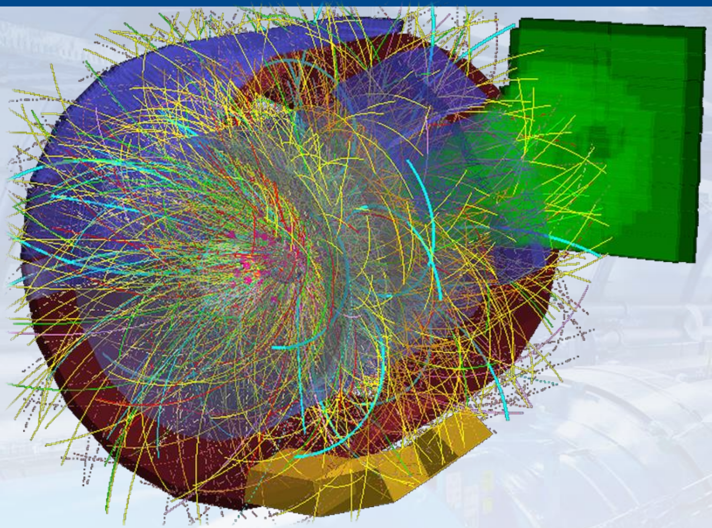
Proton-Lead

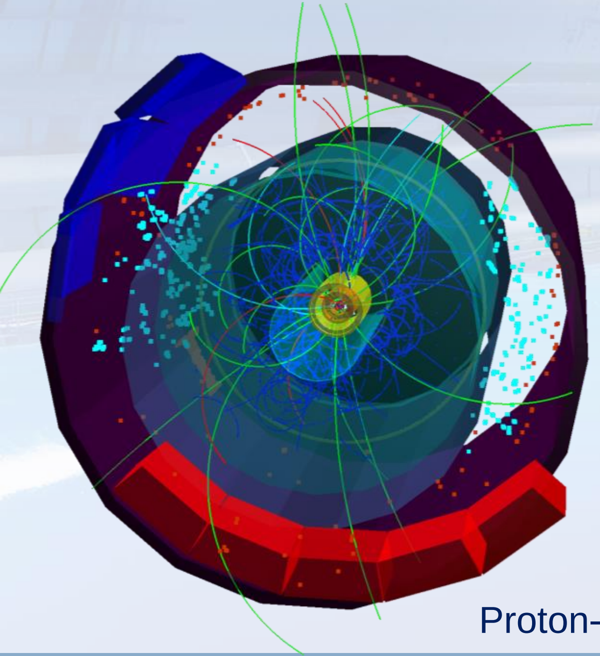
Proton-Proton

Run 2: O(1) kHz single events…

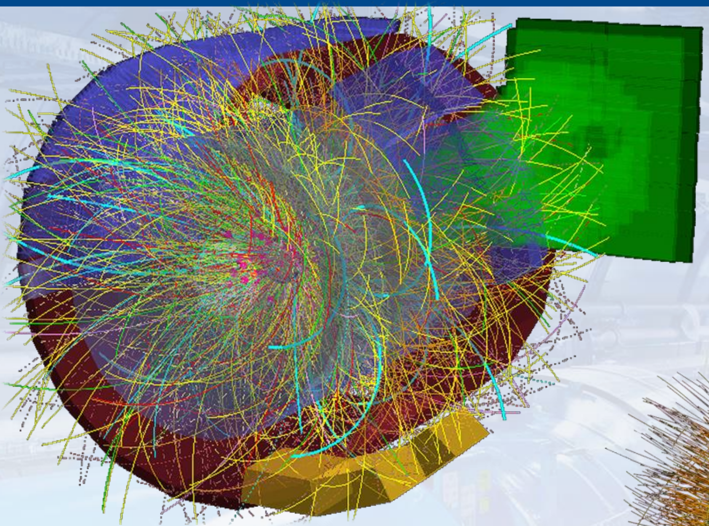

Proton-Lead

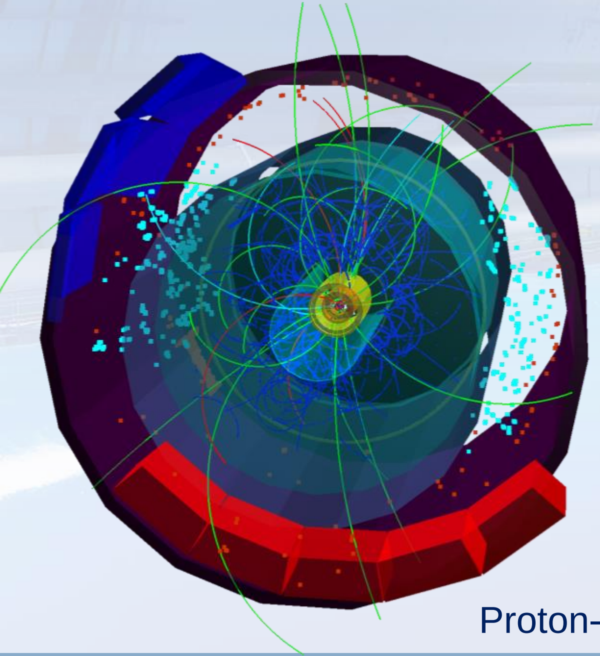

Proton-Proton

Lead-Lead

Run 2: O(1) kHz single events → Run 3: 50 kHz continuous data

- Overlapping events in TPC with realistic bunch structure @ 50 kHz Pb-Pb.
- Timeframe of 2 ms shown (will be 10 – 20 ms in production).
- Tracks of different collisions shown in different color.

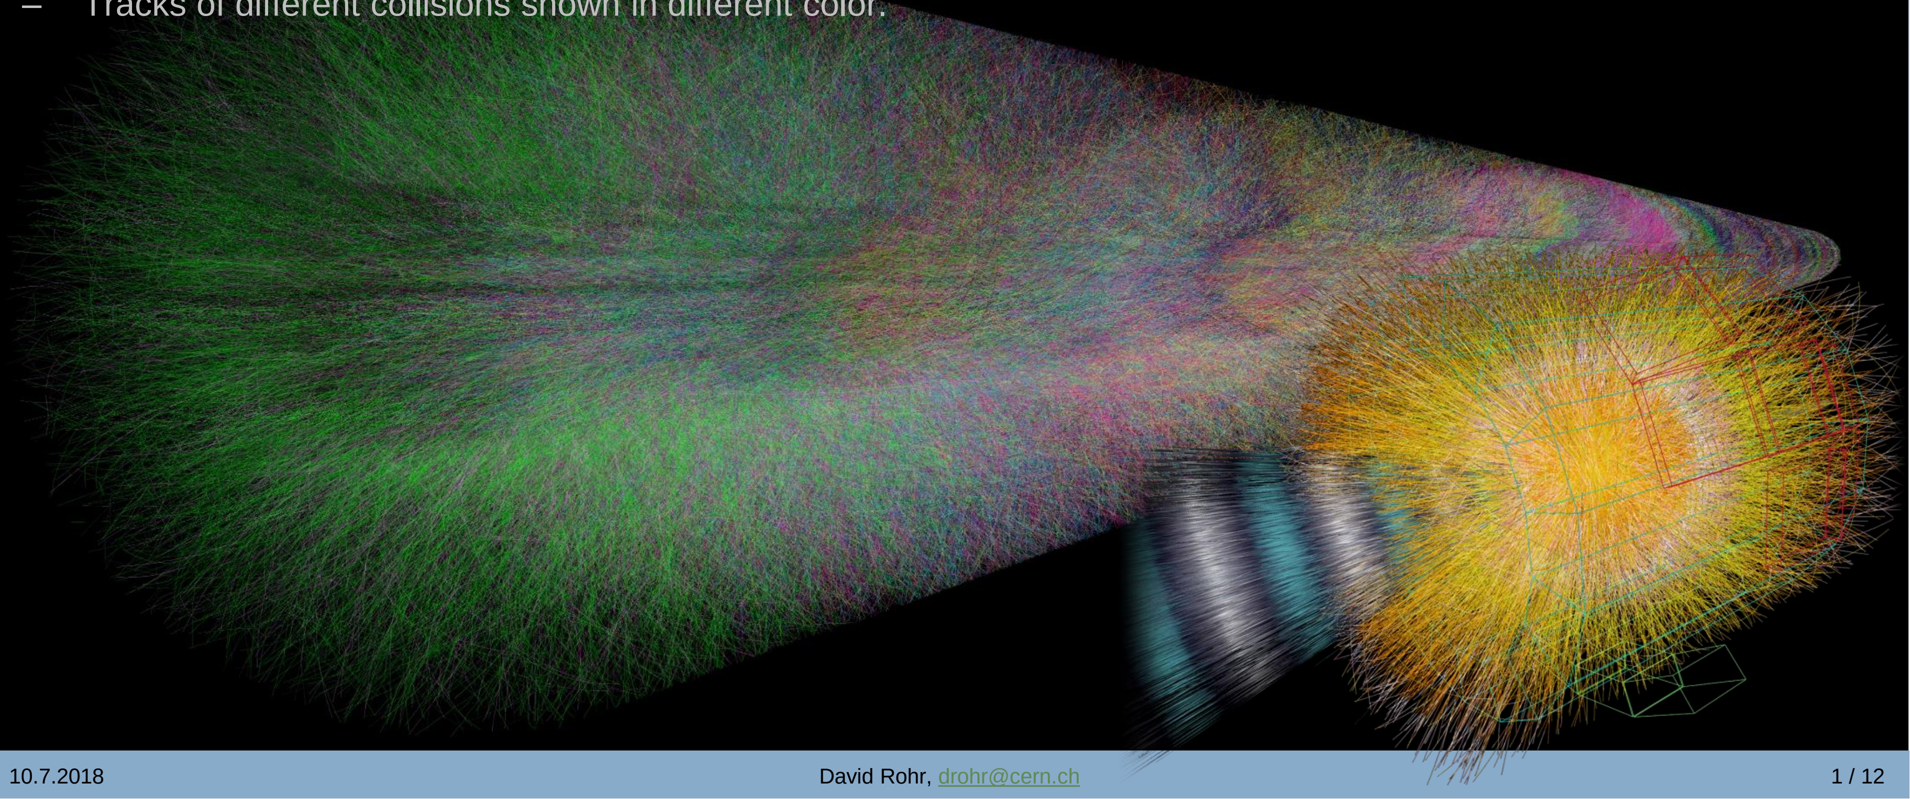

Run 2: O(1) kHz single events → Run 3: 50 kHz continuous data

- Overlapping events in TPC with realistic bunch structure @ 50 kHz Pb-Pb.
- Timeframe of 2 ms shown (will be 10 – 20 ms in production).
- Tracks of different collisions shown in different color.

- What are the challenges:
  - Reconstruct **50x** more events **online**.
  - Store **50x** more events (Needs TPC compression factor **20x** compared to Run 2 raw data size).
  - Reconstruct TPC data in continuous read out.
  - Cope with space charge distortions in the TPC.

- Scope of this talk:
  - Track reconstruction in the TPC (and related) for O².
    (ALICE Online Offline Computing Upgrade)

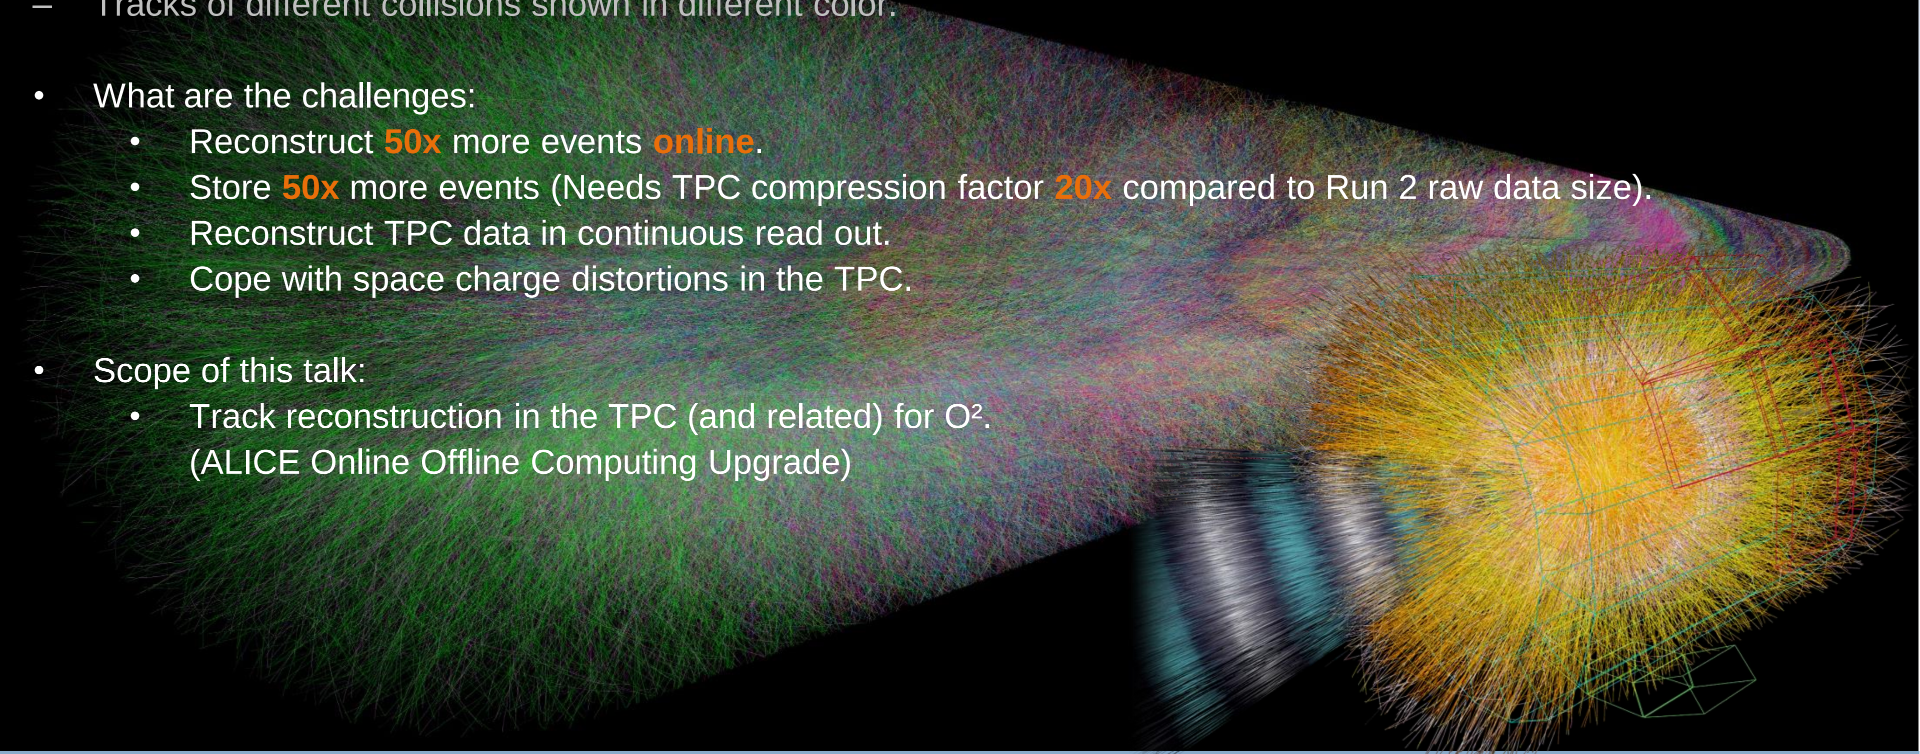

Tracking in ALICE in Run 3

- **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)

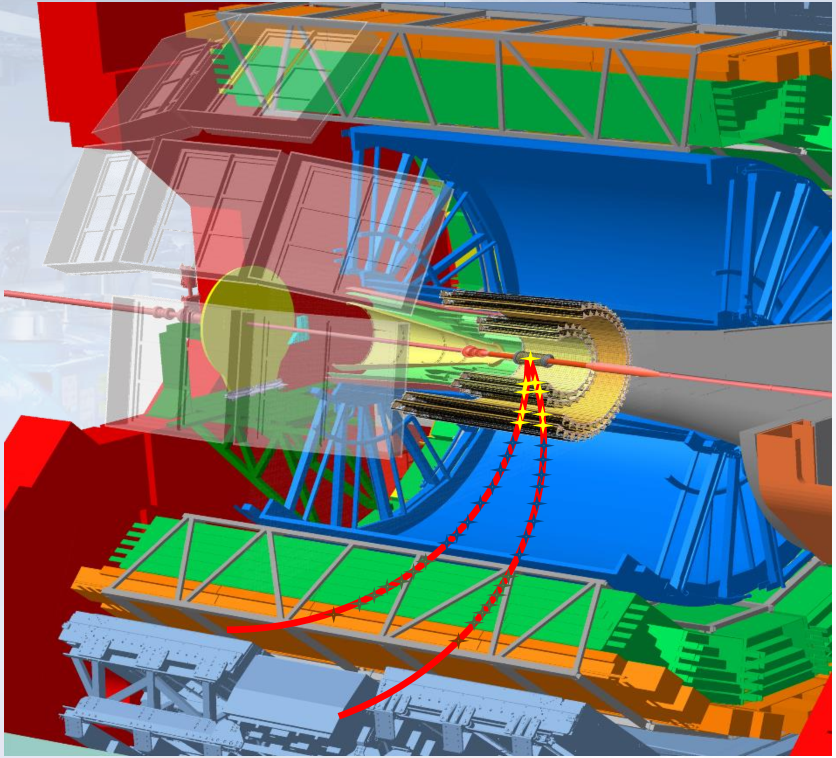

Tracking in ALICE in Run 3

- **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)

- **Two reconstruction phases in Run 3:**
  - Synchronous reconstruction (during data taking):
    - **Calibration**
    - **Data compression**
  - Asynchronous reconstruction (when no beam):
    - **Full reconstruction with final calibration**

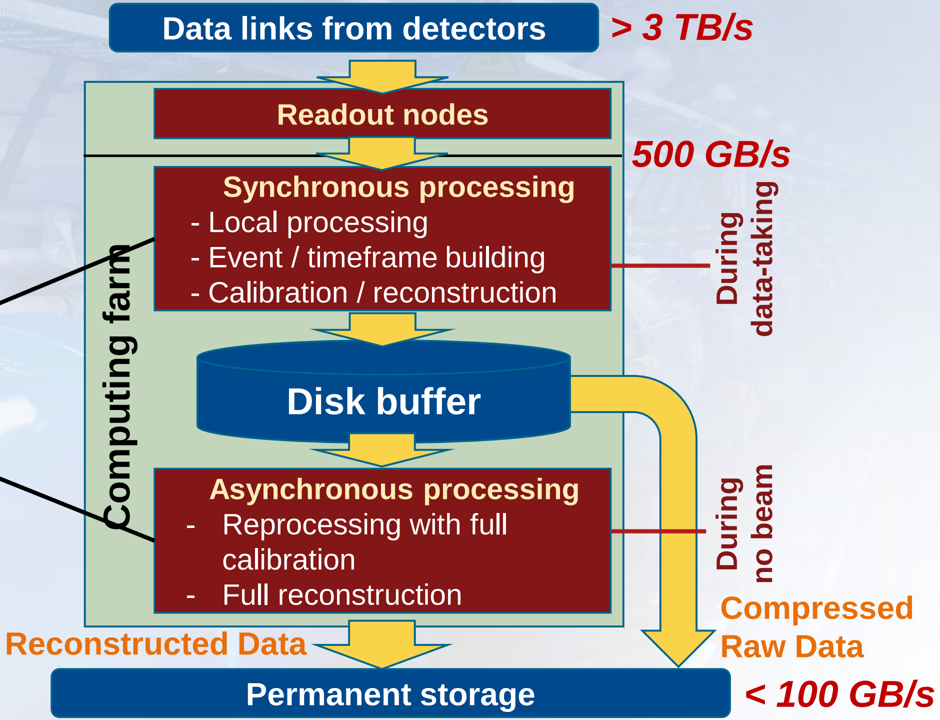

Tracking in ALICE in Run 3

- **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)

- **Two reconstruction phases in Run 3:**
  - Synchronous reconstruction (during data taking):
    - **Calibration**
    - **Data compression**
  - Asynchronous reconstruction (when no beam):
    - **Full reconstruction with final calibration**

Partial ITS + TPC + TRD tracking
- reduced statistics sufficient
(calibration based on matching of TPC / ITS / TRD tracks and TPC residuals v.s. TRD-ITS refit: see arXiv:1709.00618)

Full TPC tracking
- cluster to track residuals → better entropy coding
(needs track refit in distorted coordinates: see arXiv:1709.00618)
- removal of tracks not used for physics

Second tracking pass with final calibration

Tracking in ALICE in Run 3

- **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)

- **Two reconstruction phases in Run 3:**
  - Synchronous reconstruction (during data taking):
    - **Calibration**
    - **Data compression**
  - Asynchronous reconstruction (when no beam):
    - **Full reconstruction with final calibration**

Partial ITS + TPC + TRD tracking
- reduced statistics sufficient
(calibration based on matching of TPC / ITS / TRD tracks and TPC residuals v.s. TRD-ITS refit: see arXiv:1709.00618)

Full TPC tracking
- cluster to track residuals → better entropy coding
(needs track refit in distorted coordinates: see arXiv:1709.00618)
- removal of tracks not used for physics

Second tracking pass with final calibration

- **This means:**
  - **Full TPC online tracking @50 kHz Pb-Pb.**
  - **Reduced ITS + TRD online tracking, full tracking in phase 2.**
  - → **TPC Defines peak compute load, ITS + TRD must be fast enough at reduced statistics.**

The tracking challenge

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

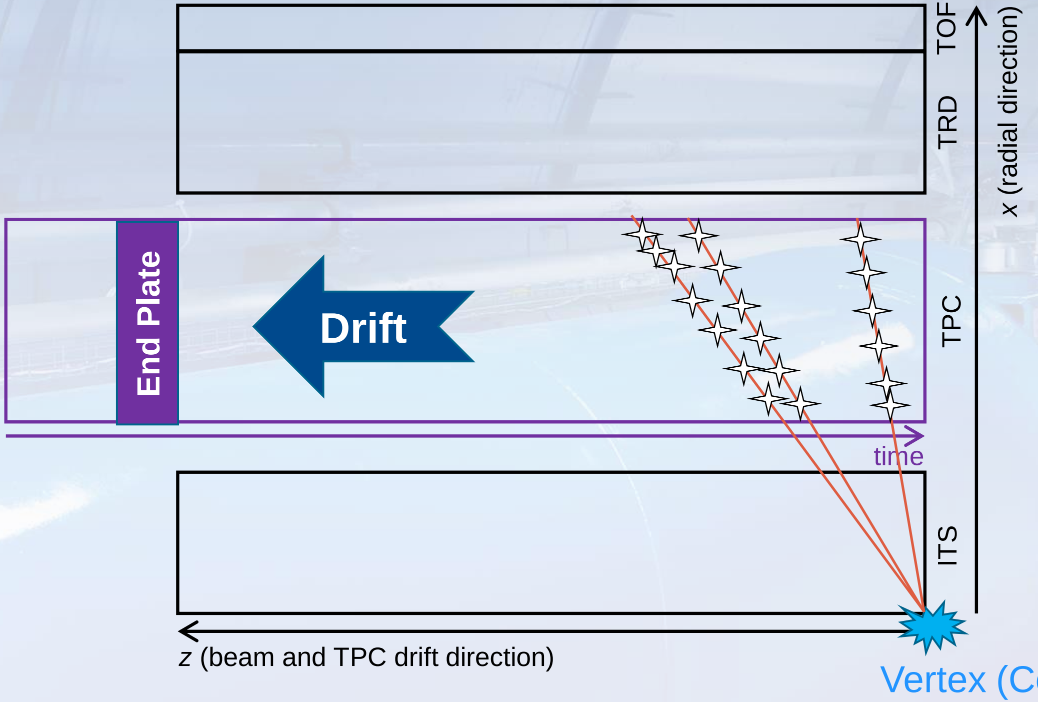

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**

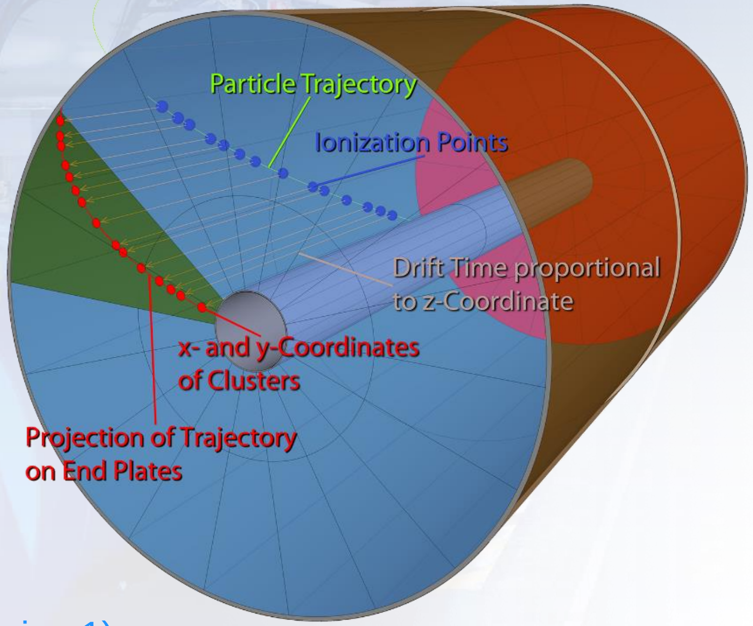

The tracking challenge

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

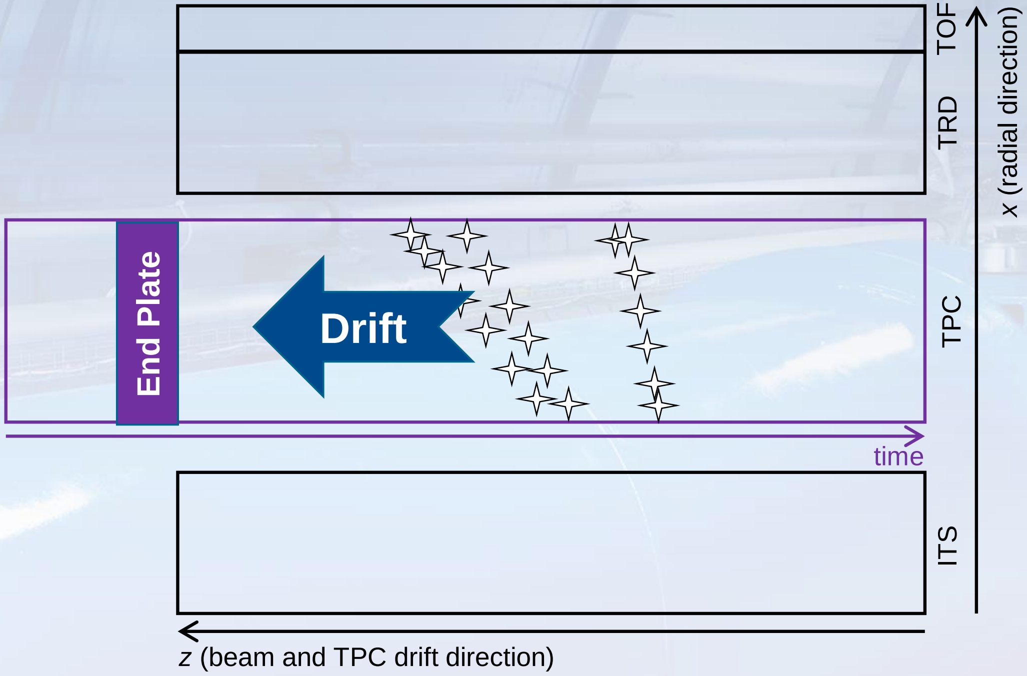

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**

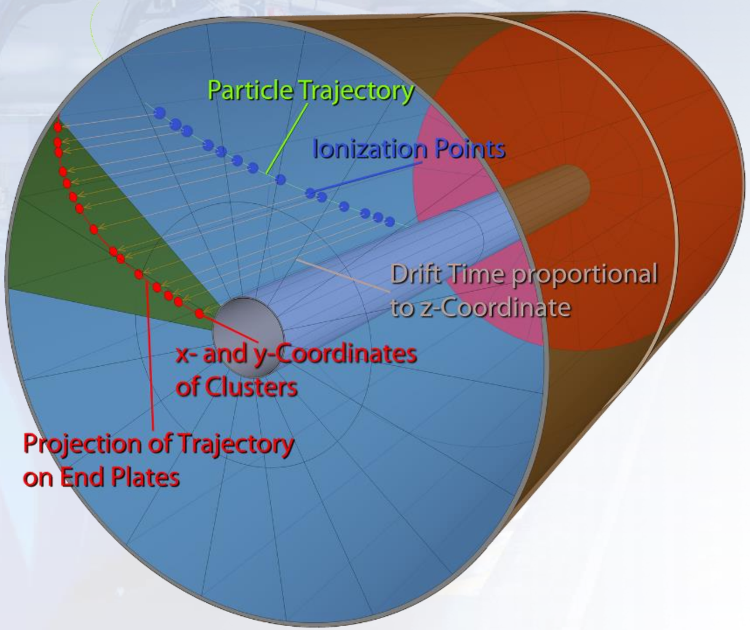

The tracking challenge

- **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.**

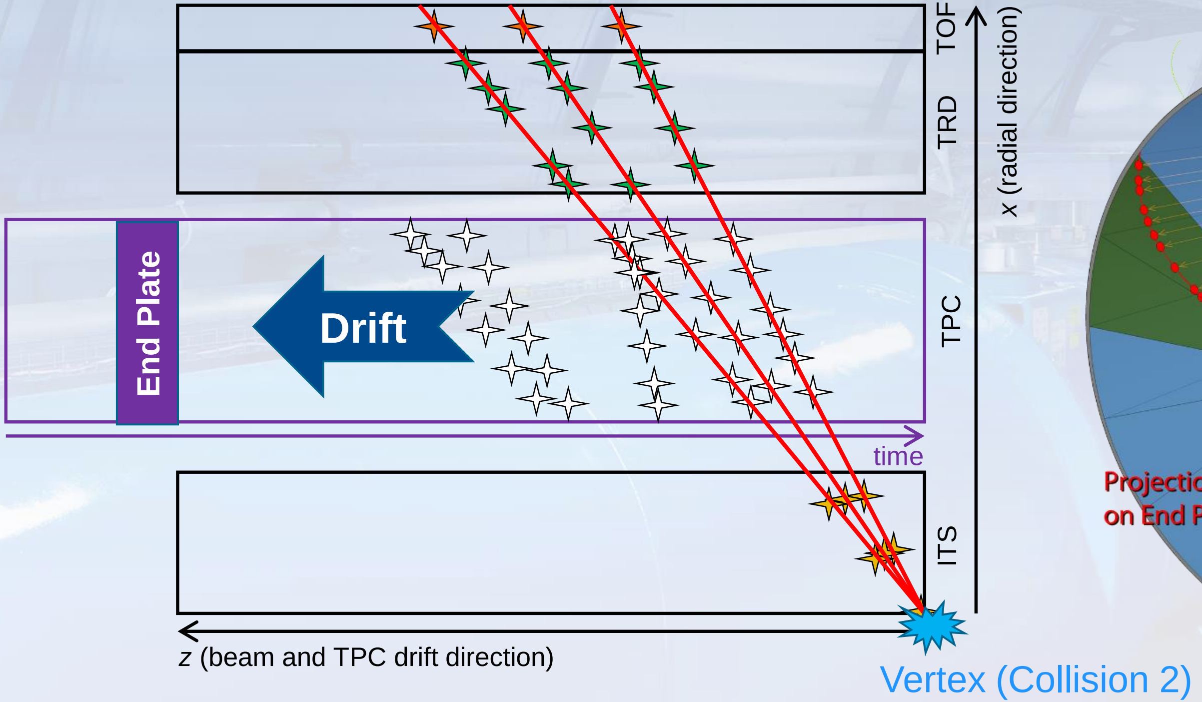

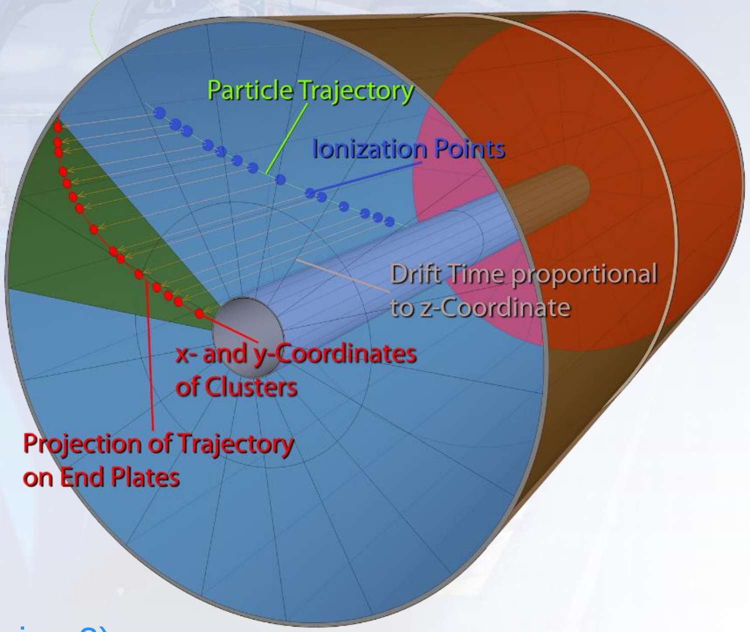

The tracking challenge

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

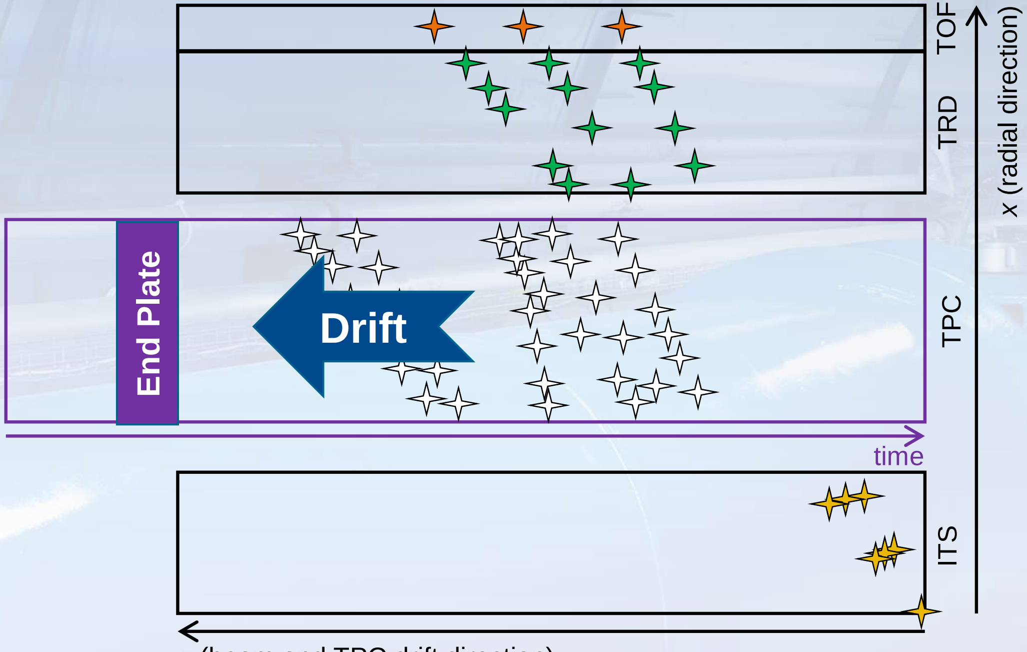

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**

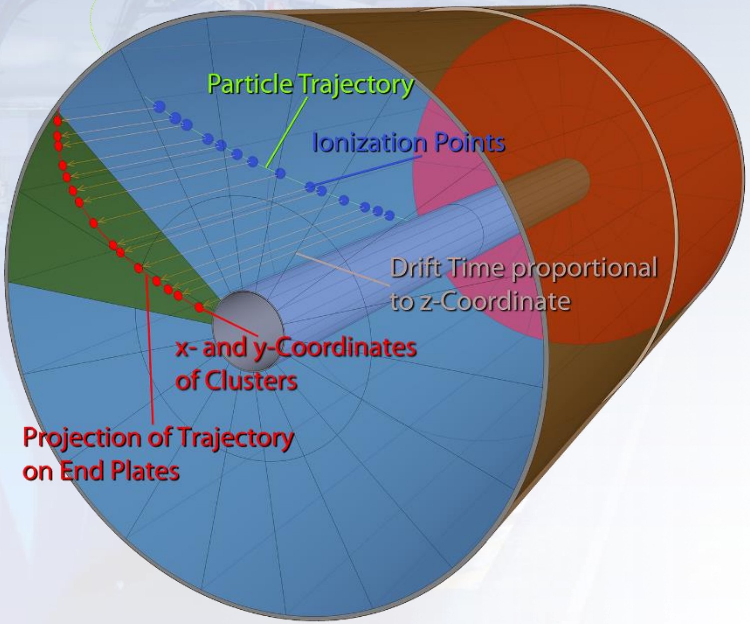

The tracking challenge

- **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.**

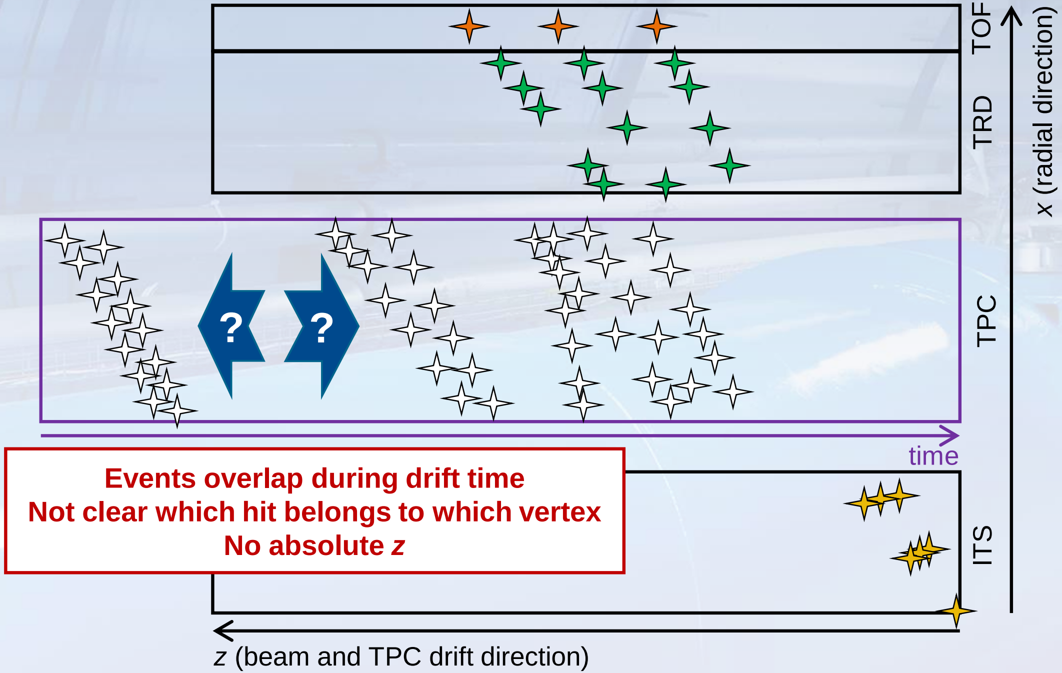

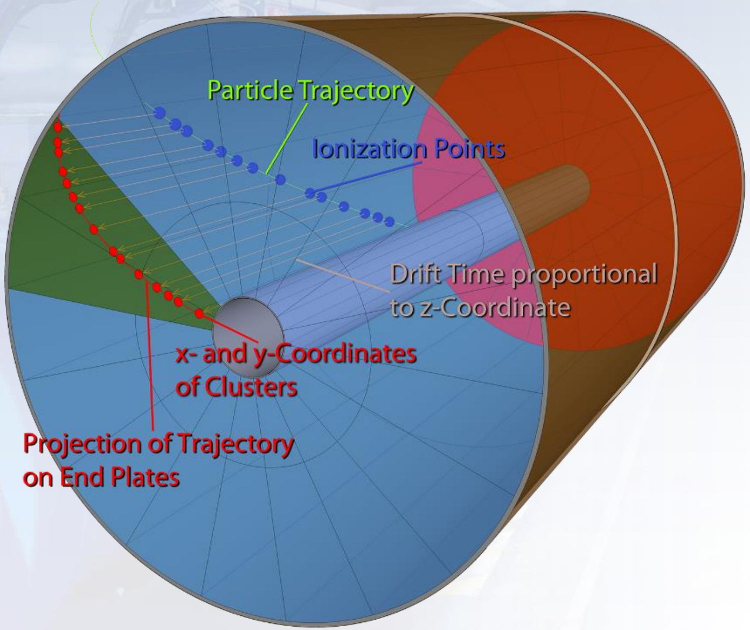

The tracking challenge

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

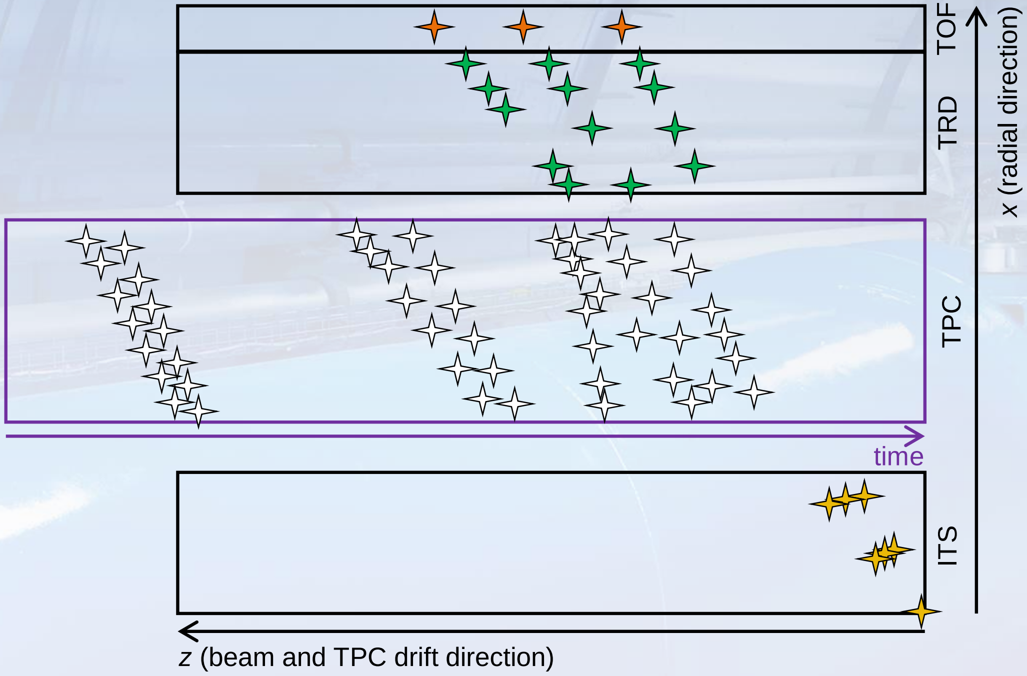

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**

The tracking challenge

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

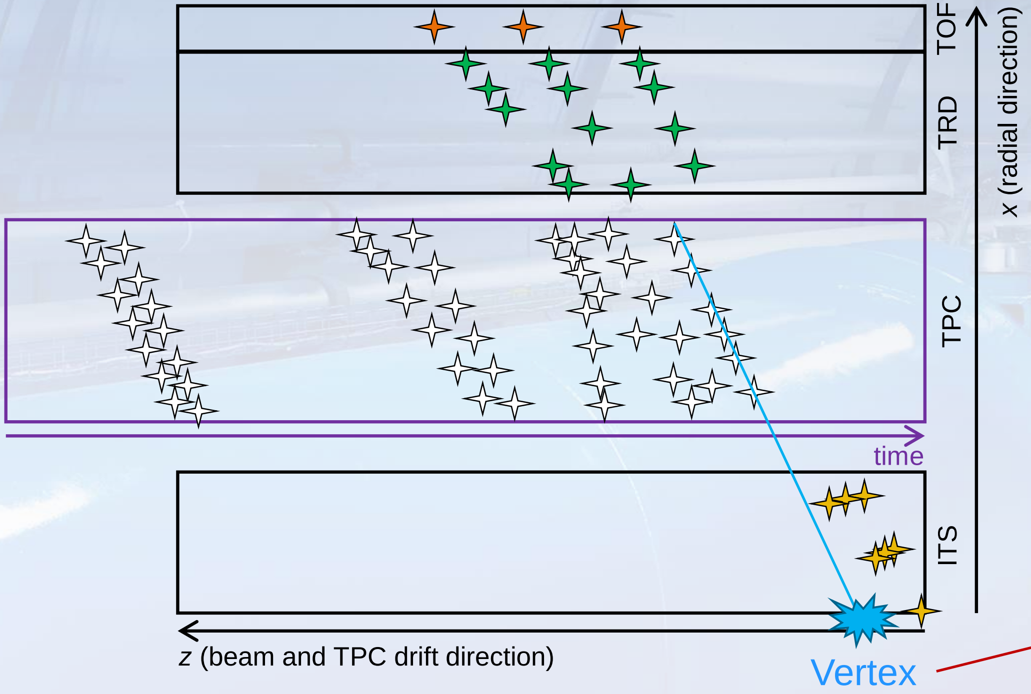

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**

$z \sim t - t_{Vertex}$

→ 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.

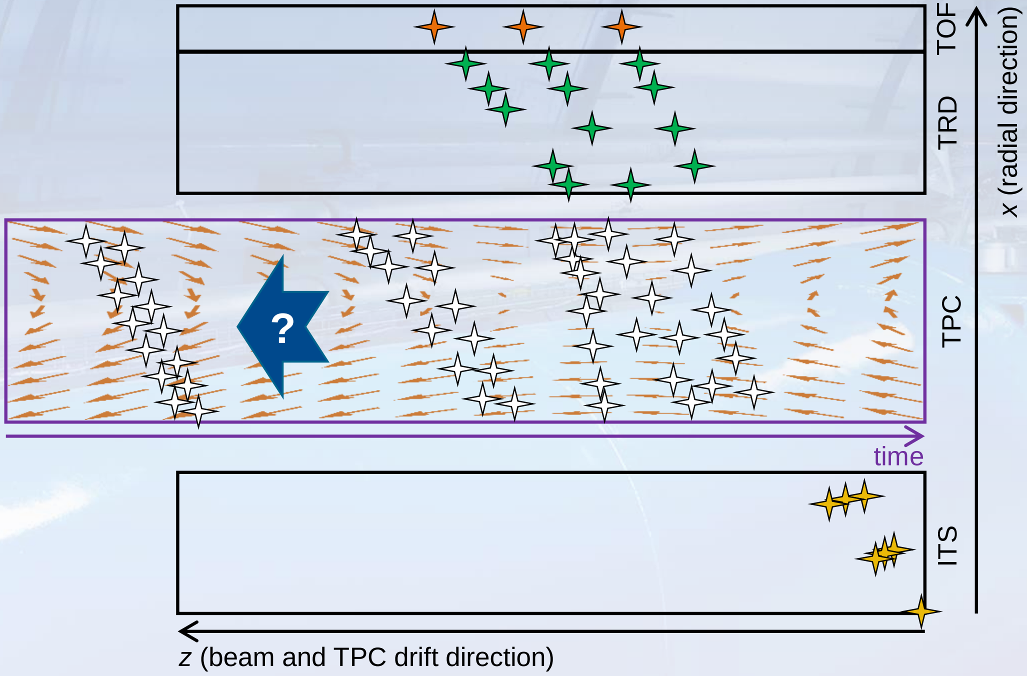

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

The tracking challenge

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

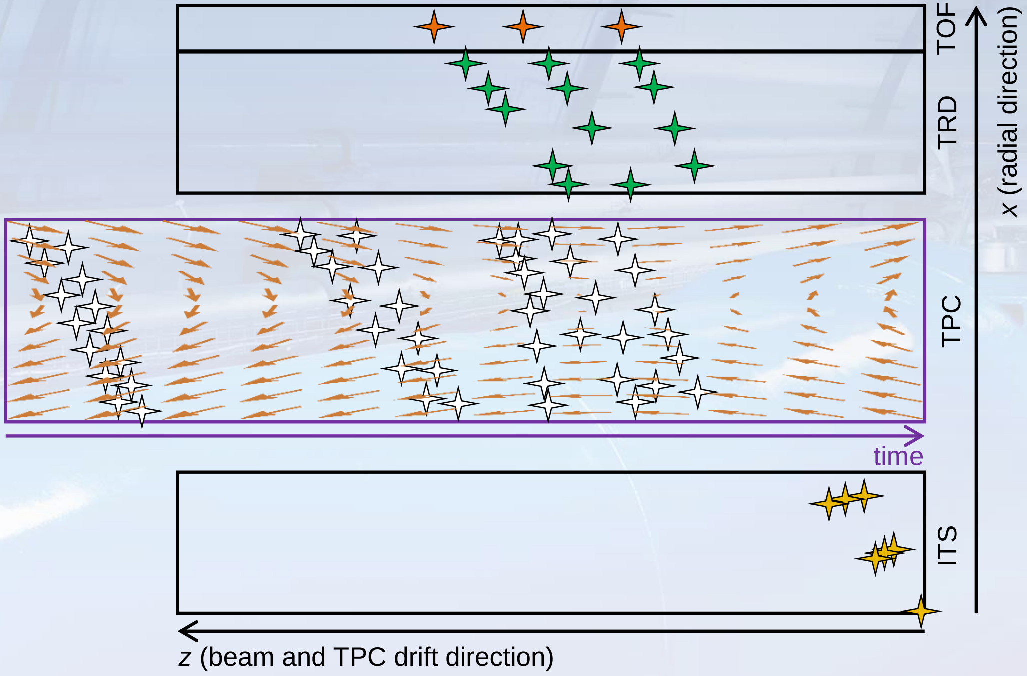

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

The tracking challenge – How the tracking will work

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

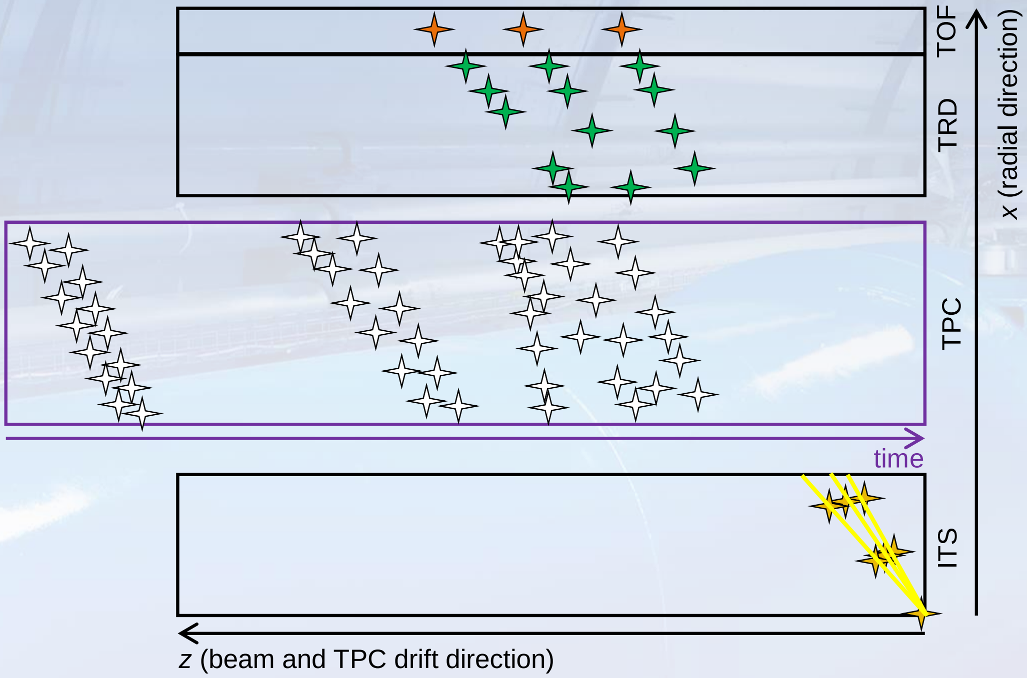

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- Standalone ITS tracking.

The tracking challenge – How the tracking will work

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

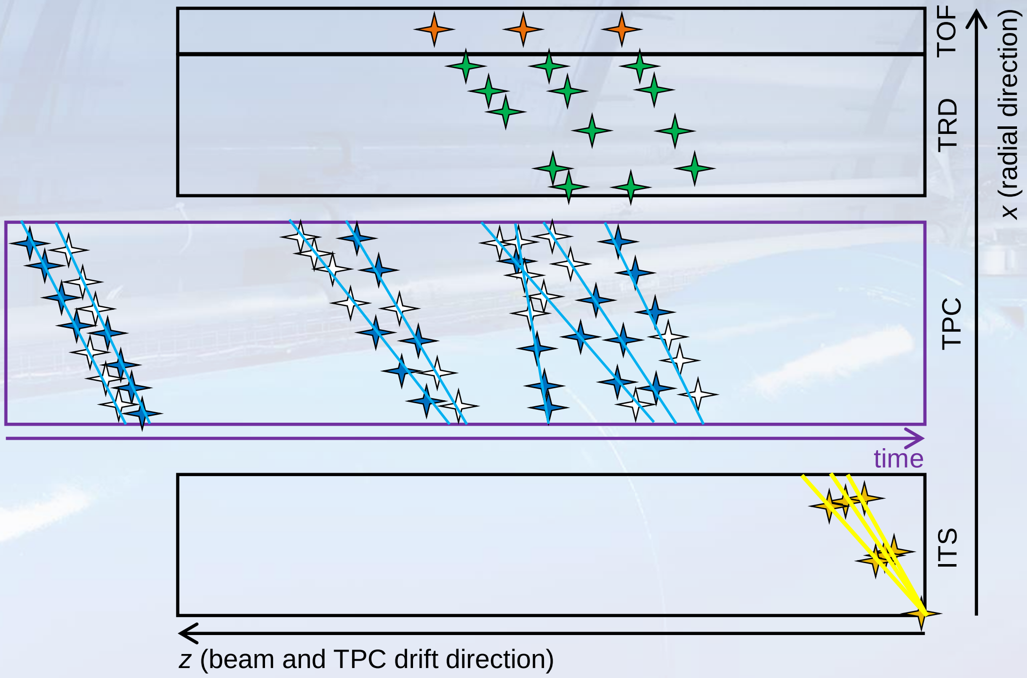

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

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

Precise tracking needs *z* for:
- Cluster error parameterization
- Inhomogeneous B-field
- Distortion correction

Effects smooth →
irrelevant for initial trackletting

The tracking challenge – How the tracking will work

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

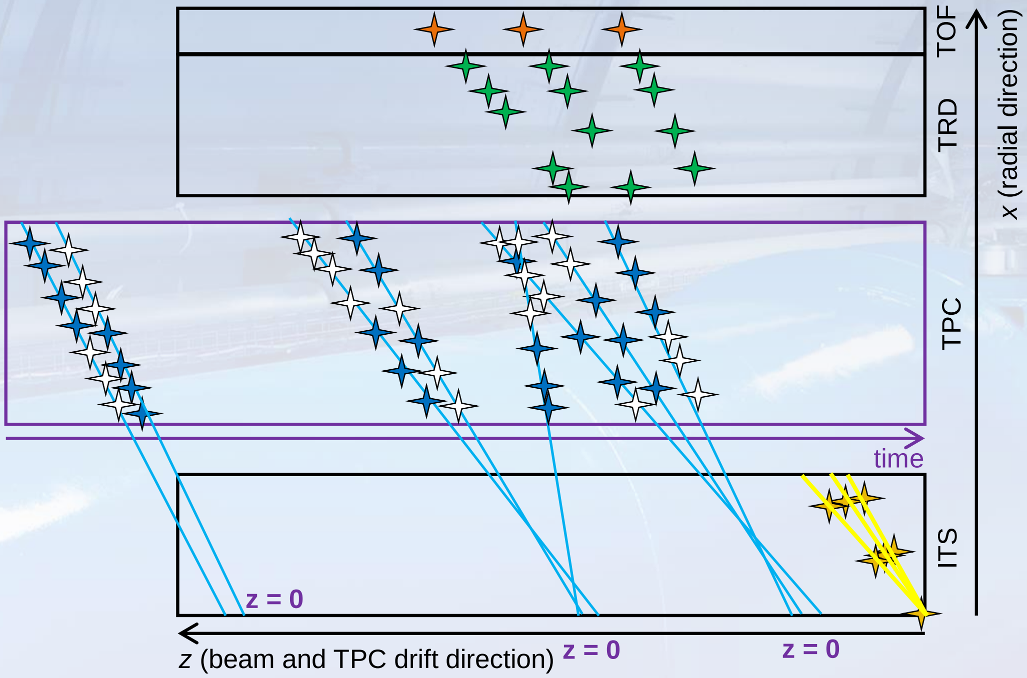

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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.

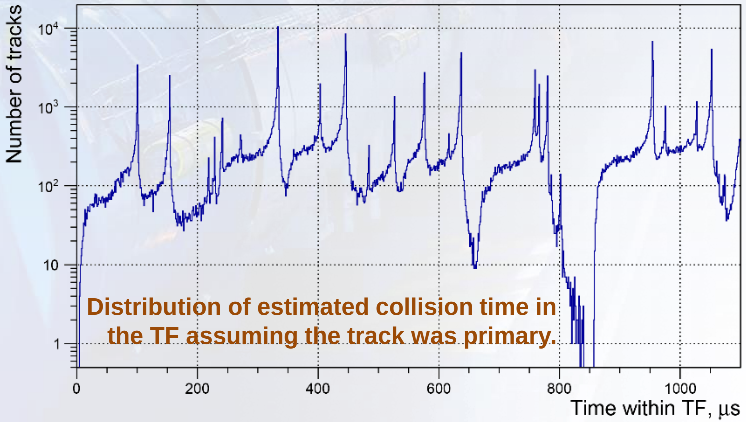

The tracking challenge – How the tracking will work

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

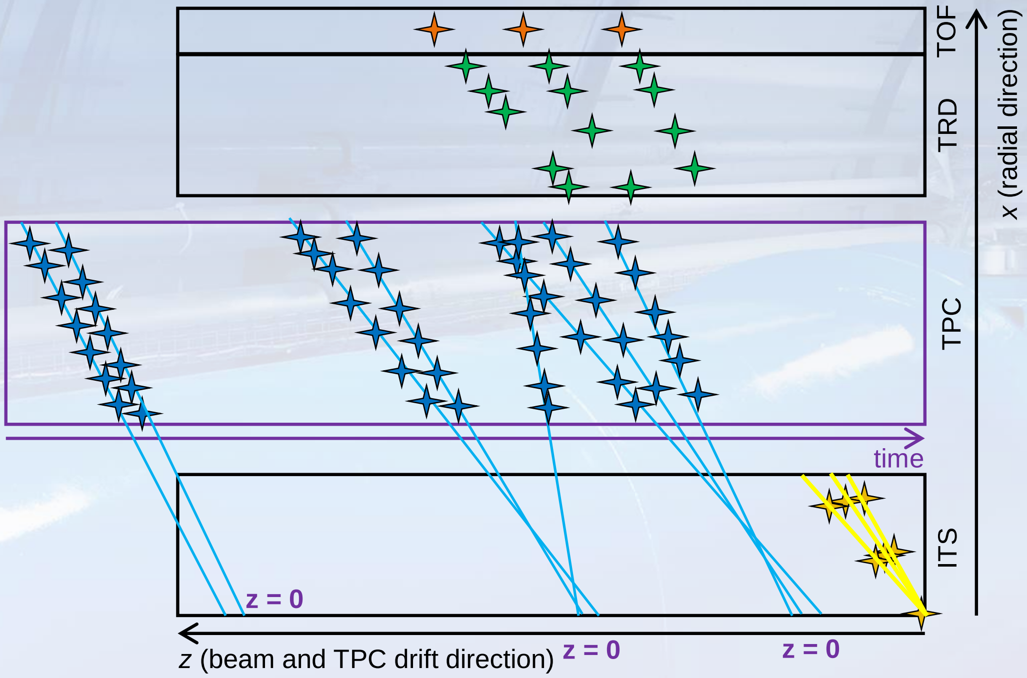

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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 – How the tracking will work

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

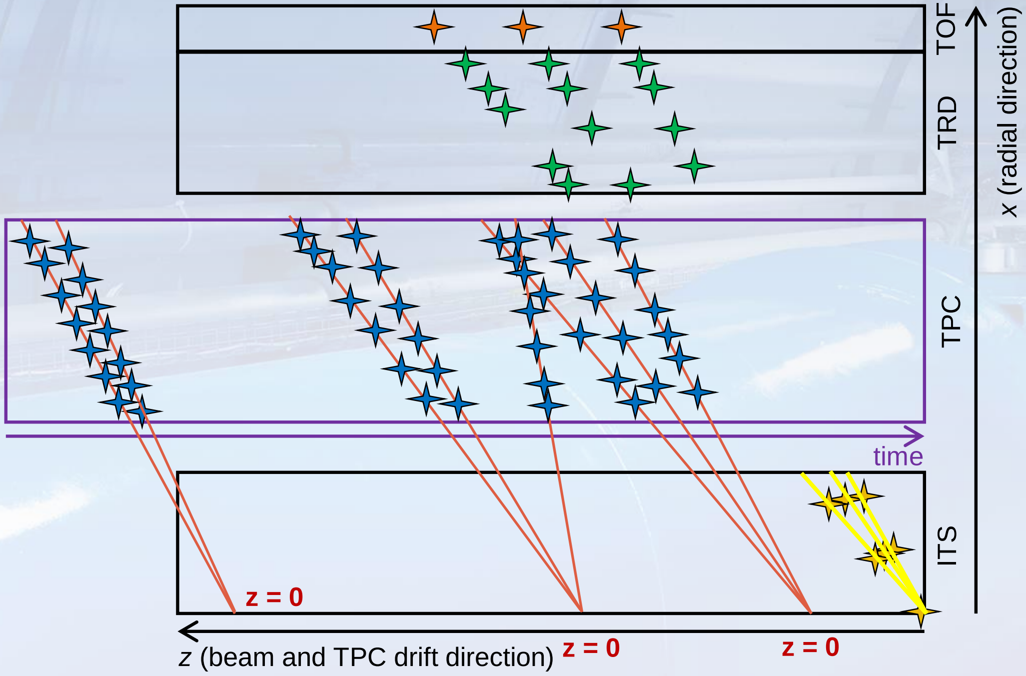

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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

David Rohr, drohr@cern.ch

The tracking challenge – How the tracking will work

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

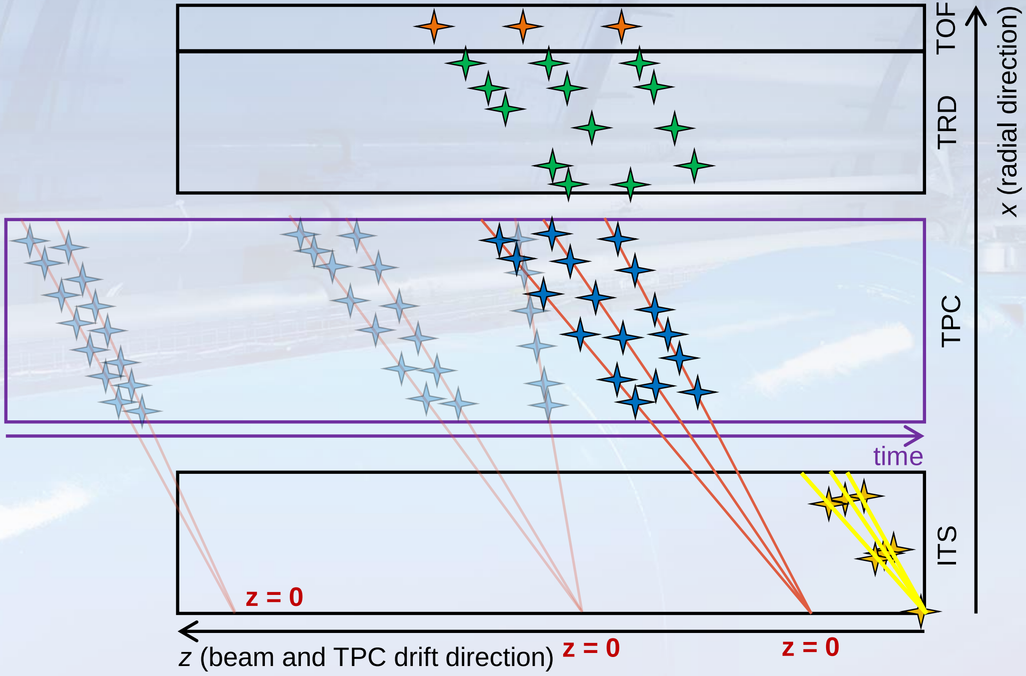

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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 – How the tracking will work

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

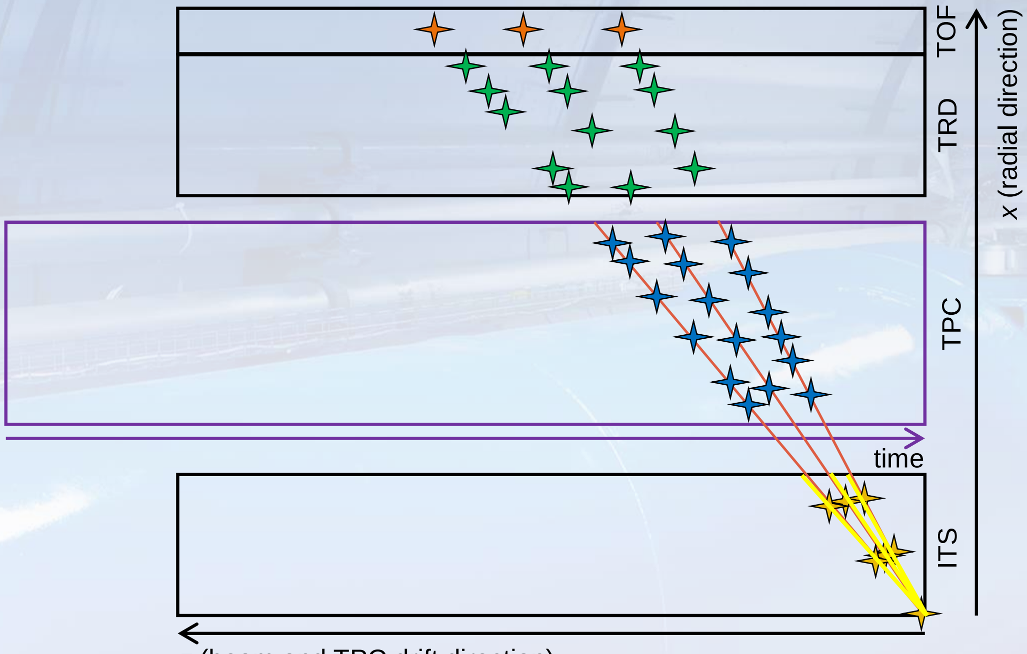

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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 – How the tracking will work

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

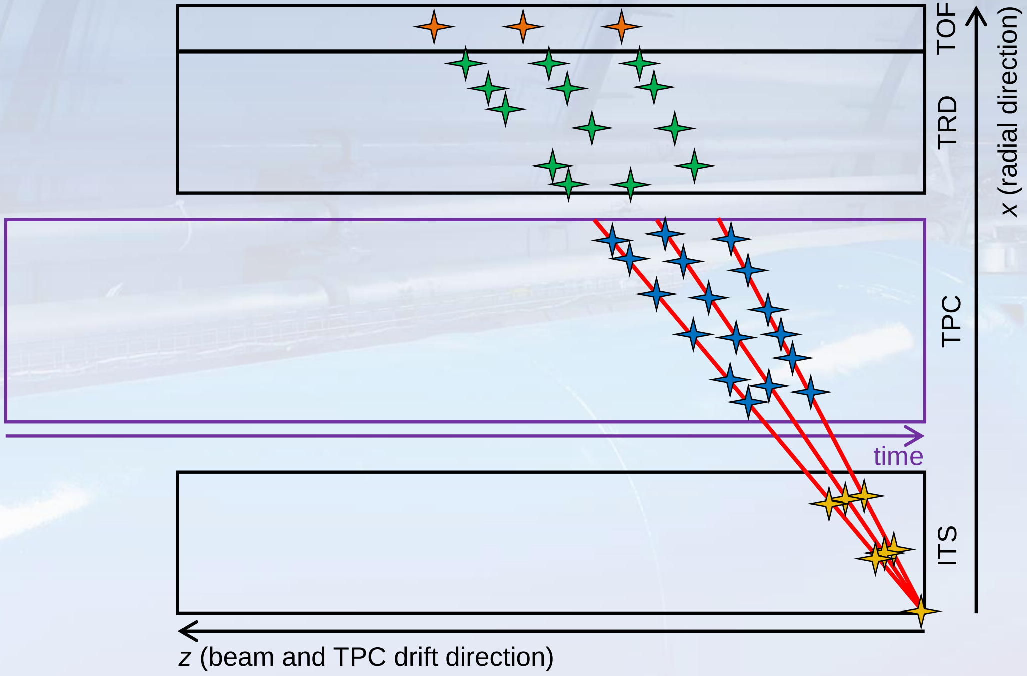

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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 – How the tracking will work

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

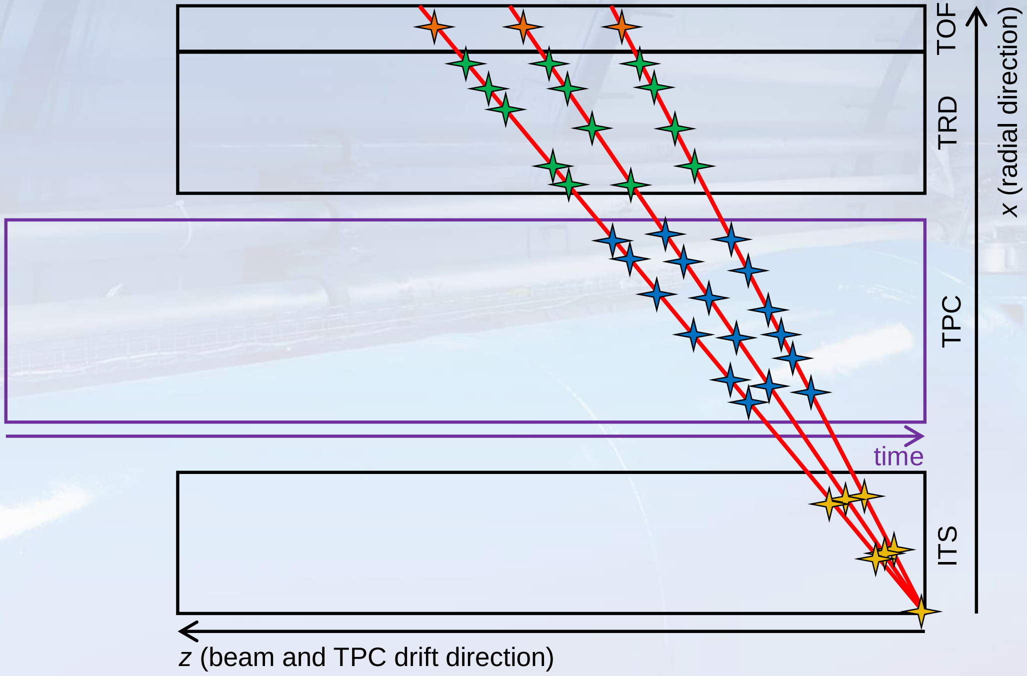

- **Problem: TPC clusters have no defined *z*-position but only a time. They can be shifted in *z* arbitrarily.**
- **GEM amplification produces ions that deflect the electrons during the drift. The correction of these space-charge distortions requires the absolute *z* position.**

- 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 – How the tracking will work

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

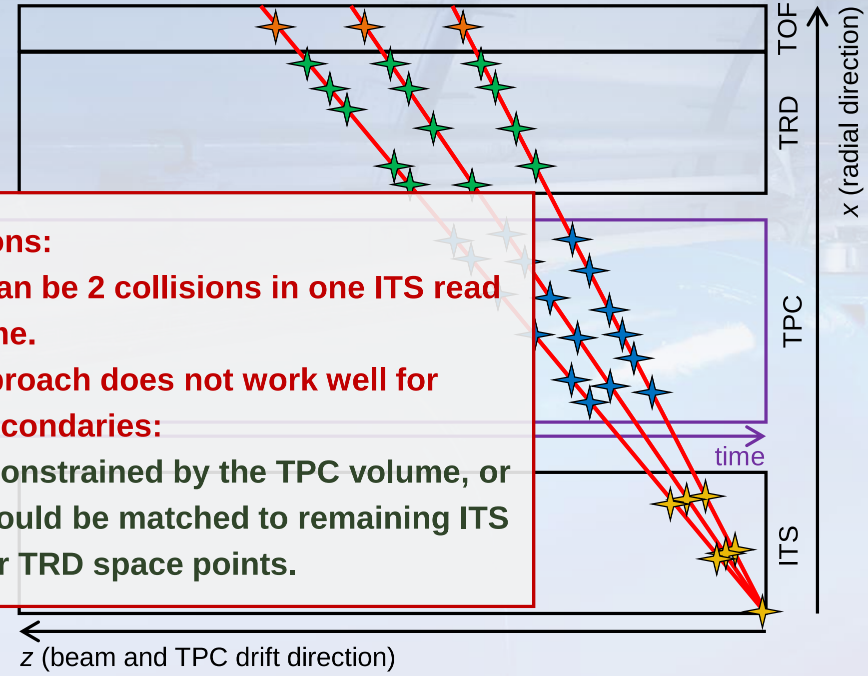

Complications:

- **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 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 amplification 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.

ALICE TPC Tracking Status

- **ALICE TPC tracking for O² developed.**
  - Derived from Run 2 HLT tracking.
    - Can track 40.000.000 tracks / second in the ALICE HLT.

  (See https://indico.physics.lbl.gov/indico/event/149/contributions/222/attachments/216/230/berkeley-tracker.pdf)

  - Based on Cellular Automaton and Kalman Filter.

  - Uses GPU-acceleration to meet compute constraints.
    - Generic source code that run on CPU (with OpenMP) and GPU (CUDA and OpenCL)
    - Identical results from GPU and CPU version.

  - Adapted for ALICE O² software, available in standard software installation since 2018.
    - Improved efficiency and resolution compared to Run 2 in order to match offline quality.
    - Added low-p_T tracking to enable cluster rejection needed for Run 3 data reduction.

Tracking efficiency (Run 2, O²/HLT v.s. Offline – Pb-Pb)

- New HLT / O² tracking shows comparable efficiency to Run 2 offline tracking.
- In certain situations the new tracking is already superior thanks to tuning for Run 3 conditions.

- 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.
- Same calibration for offline / HLT.
- Same cluster error parameterization.

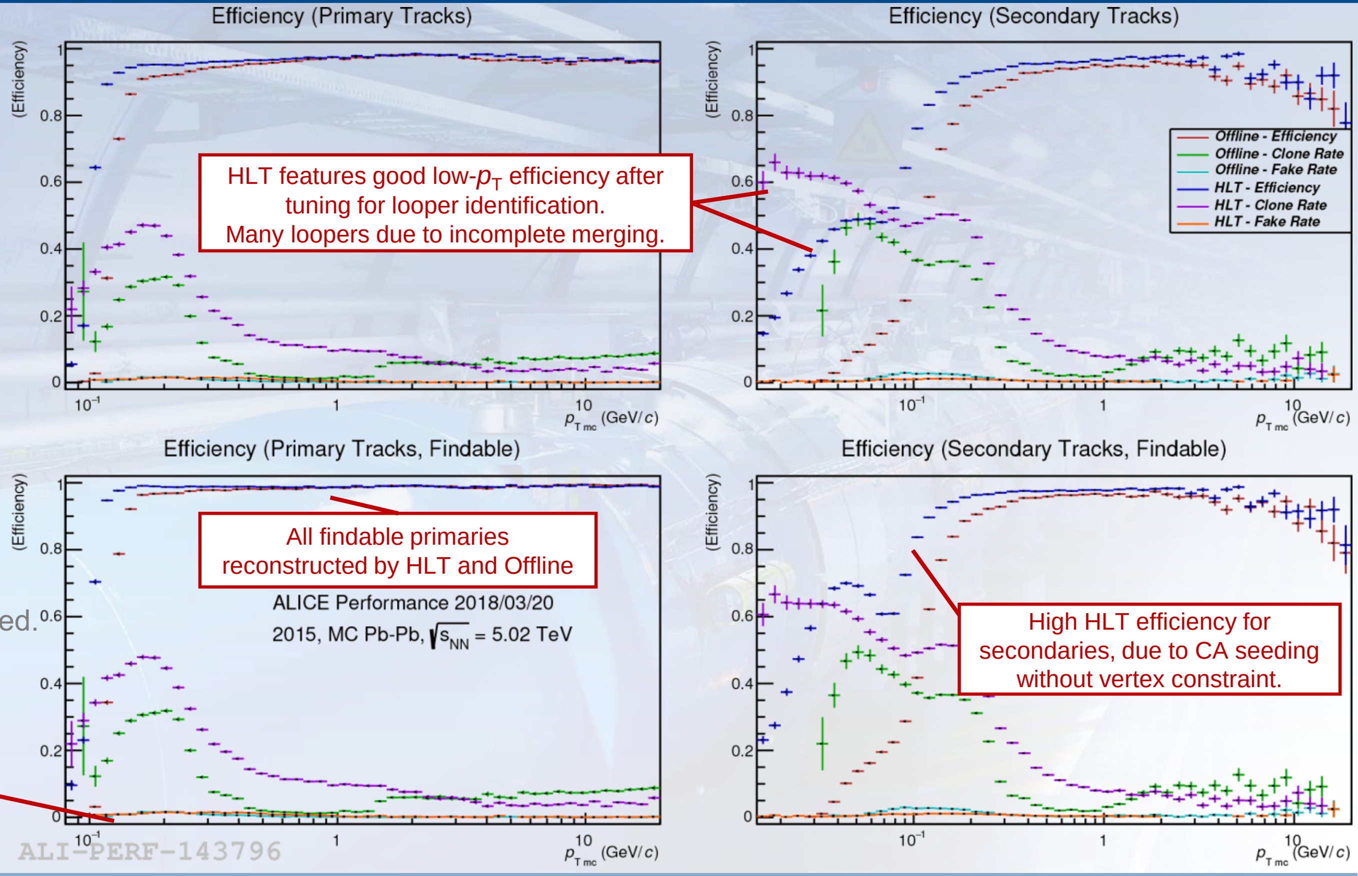

Practically zero fake rate for both trackers

(See backup for proton-proton plot)

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.

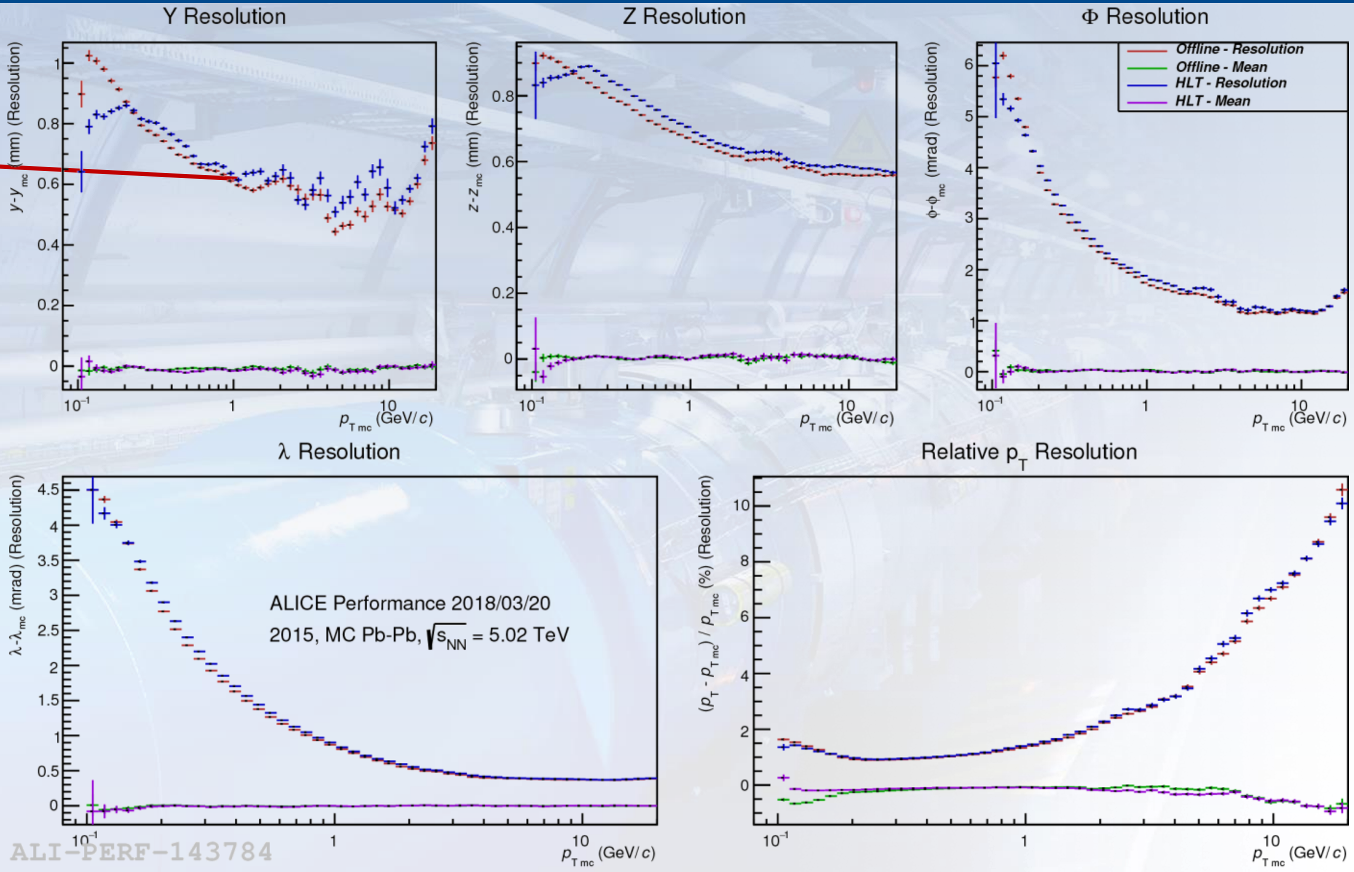

(See poster of S. Gorbunov: #423 !)

(See backup for without distortions.)

Track fit / finding stability

- Various benchmarks ensure same results in O2 scenario as during Run 2:
  - *z*-independent tracking: Take “normal” Run 2 event, forget about absolute z an process like O2 time frame:
    - Identical efficiency, negligible resolution decrease for secondaries.

(See backup for respective figures)

Track fit / finding stability

- Various benchmarks ensure same results in O2 scenario as during Run 2:
  - *z*-independent tracking: Take “normal” Run 2 event, forget about absolute z an process like O2 time frame:
    - Identical efficiency, negligible resolution decrease for secondaries.
  - In bunch pile up: μ = 100 to 1000 in proton-proton (μ = simultaneous collisions per bunch crossing):
    - No change up to μ = 300, minor efficiency decrease for secondaries above μ = 300, tracking still working at μ = 1000 at reduced efficiency.
    - No effect on resolution.

(See backup for respective figures)

Track fit / finding stability

- Various benchmarks ensure same results in O2 scenario as during Run 2:
  - *z*-independent tracking: Take “normal” Run 2 event, forget about absolute z an process like O2 time frame:
    - Identical efficiency, negligible resolution decrease for secondaries.
  - In bunch pile up: μ = 100 to 1000 in proton-proton (μ = simultaneous collisions per bunch crossing):
    - No change up to μ = 300, minor efficiency decrease for secondaries above μ = 300, tracking still working at μ = 1000 at reduced efficiency.
    - No effect on resolution.
  - Length of time frame (100 μs to 20 ms):
    - No difference (with fix for limited single precision float accuracy).

(See backup for respective figures)

Track fit / finding stability

- Various benchmarks ensure same results in O2 scenario as during Run 2:
  - *z*-independent tracking: Take “normal” Run 2 event, forget about absolute z an process like O2 time frame:
    - Identical efficiency, negligible resolution decrease for secondaries.
  - In bunch pile up in pp: μ = 100 to 1000 in proton-proton (μ = simultaneous collisions per bunch crossing):
    - No change up to μ = 300, minor efficiency decrease for secondaries above μ = 300, tracking still working at μ = 1000 at reduced efficiency.
    - No effect on resolution.
  - Length of time frame (100 μs to 20 ms):
    - No difference (with fix for limited single precision float accuracy).
  - **TPC occupancy: Single event to 50 kHz time frame Pb-Pb**:
    (see figure on the right)
    - Resolution identical.
    - Small efficiency decrease below 150 MeV/*c*.
    - Clone rate of short low-p_T tracks increases with occupancy.

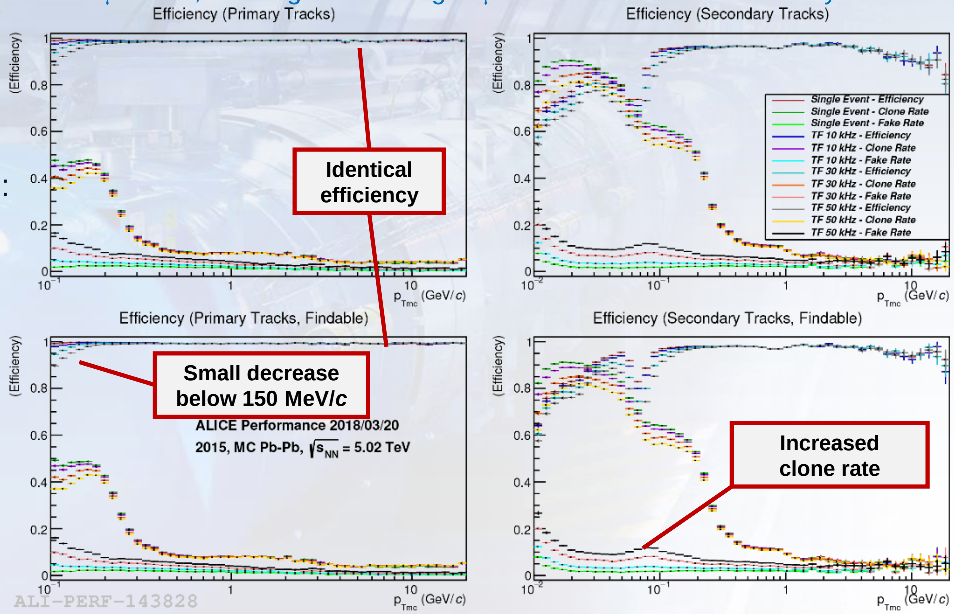

(See backup for respective figures)

Track fit / finding stability

- Various benchmarks ensure same results in O2 scenario as during Run 2:
  - *z*-independent tracking: Take “normal” Run 2 event, forget about absolute z an process like O2 time frame:
    - Identical efficiency, negligible resolution decrease for secondaries.
  - In bunch pile up: μ = 100 to 1000 in proton-proton (μ = simultaneous collisions per bunch crossing):
    - No change up to μ = 300, minor efficiency decrease for secondaries above μ = 300, tracking still working at μ = 1000 at reduced efficiency.
    - No effect on resolution.
  - Length of time frame (100 μs to 20 ms):
    - No difference (with fix for limited single precision float accuracy).
  - **TPC occupancy: Single event to 50 kHz time frame Pb-Pb**: *(see figure on the right)*
    - Resolution identical.
    - Small efficiency decrease below 150 MeV/*c*.
    - Clone rate of short low-p_T tracks increases with occupancy.

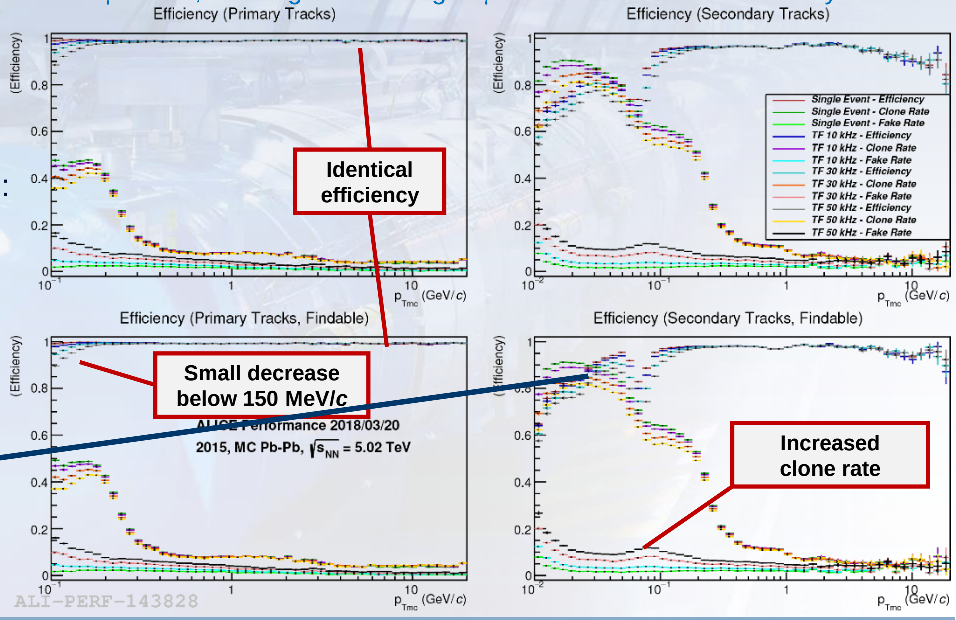

Good efficiency for low-p_T looping secondaries.

- Absolutely crucial for rejecting tracks not used for physics.
- High clone rate due to incomplete implementation of low-p_T merging.

(See backup for respective figures)

Tracking time

- Speed-up normalized to single CPU core.
    - Red curve: exactly the speed-up.
    - Other curves: corrected for required CPU resources.
        - How many cores does the GPU replace.

- Significant gain with newer GPU (blue v.s. green).

- Compared to Run 2 offline, One GPU replaces **> 800 CPU cores** (**blue * red**). (**at same efficiency / resolution**)

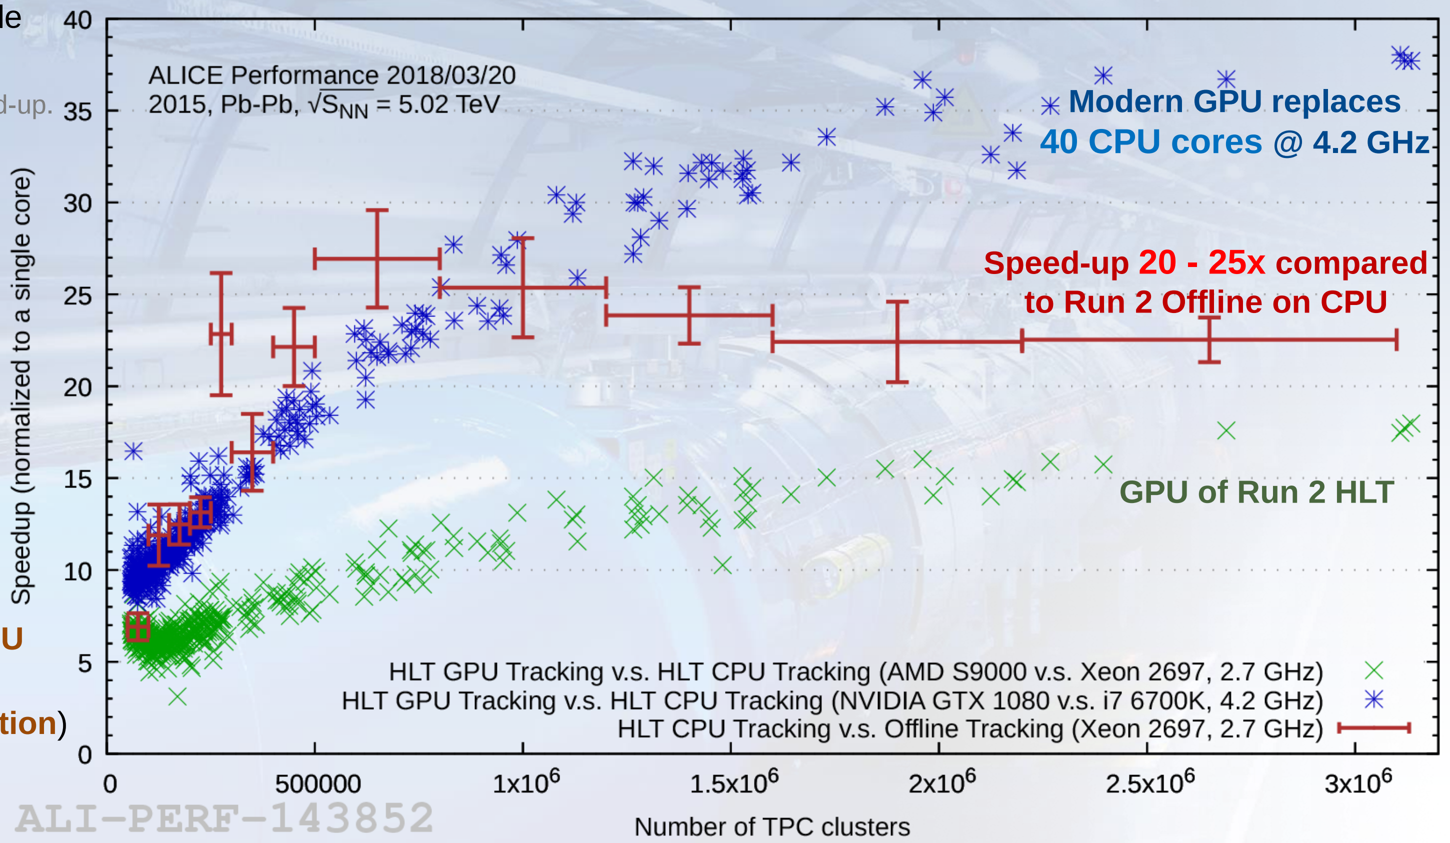

TPC Data Compression

- TPC Data compression involves 3 steps:
  1. Entropy reduction (Track model, logarithmic precision, etc.)
  2. Entropy encoding (Huffman, Arithmetic, ANS)
  3. Removal of tracks not used for physics.

- Steps 1 + 2 implemented for Run 2.
  - Current compression factor **8.3x**.
  - Prototype for Run 3 achieves factor **9.1x**.

- **Missing factor ~2x** to reach total reduction factor **20x** for Run 3.

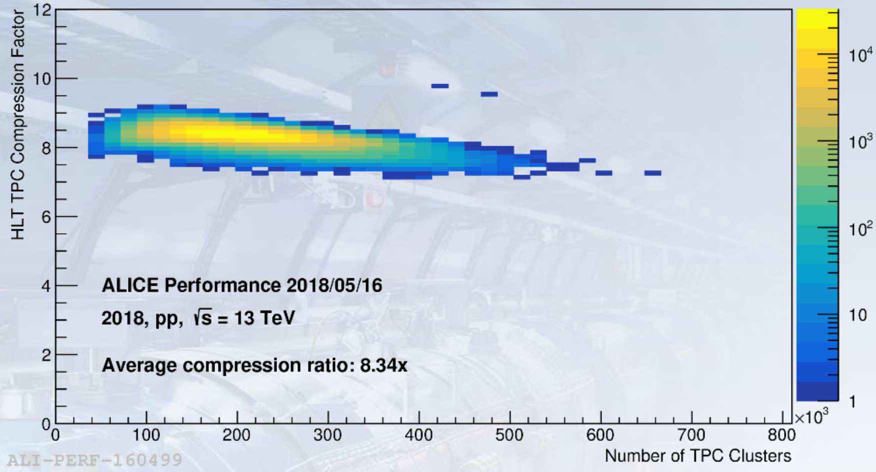

TPC Data Compression

- TPC Data compression involves 3 steps:
  1. Entropy reduction (Track model, logarithmic precision, etc.)
  2. Entropy encoding (Huffman, Arithmetic, ANS)
  3. Removal of tracks not used for physics.

- Steps 1 + 2 implemented for Run 2.
  - Current compression factor **8.3x**.
  - Prototype for Run 3 achieves factor **9.1x**.

- **Missing factor ~2x** to reach total reduction factor **20x** for Run 3.
  - Remove non-physics tracks < 50 MeV/*c*.
  - Remove additional legs of looping tracks.
  - Remove track segments with high inclination angle.

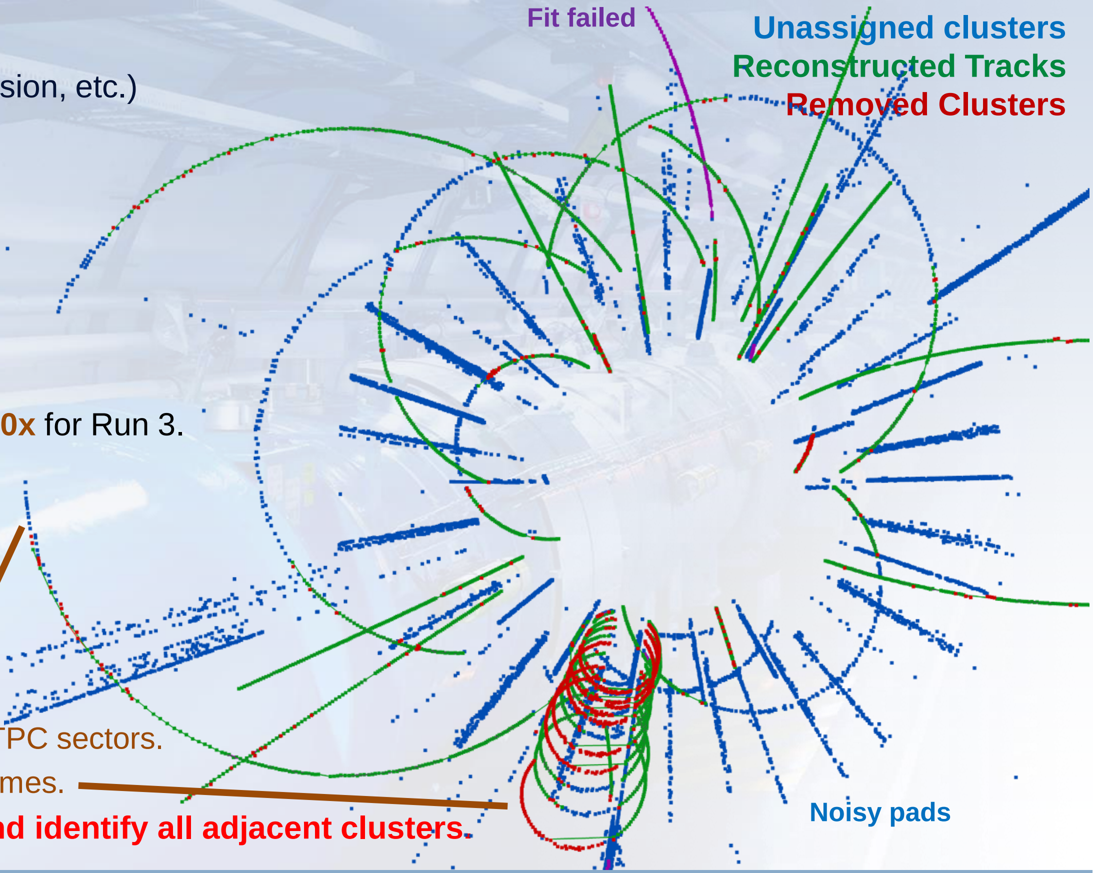

- **Low-p_T merging still incomplete:**
  - Long arcs with high inclination angle over multiple TPC sectors.
  - Only one side of helix crossing sector boundary n times.
- **Current Task: Fix merging, extrapolate tracks and identify all adjacent clusters.**

Cluster removal for O²

- Cluster attachment v.s. p_T (integrated):
  - **Cyan: all clusters**
  - **Green: clusters attached to wrong track.**
  - **Red: clusters attached to the correct track - used in track fit.**
  - **Blue: correctly attached and adjacent clusters.**
  - **Purple: All clusters (if attached or not) of a reconstructed track.**

- Clusters below 10 MeV/*c* not accessible by tracking.
  - Very low-p_T looping tracks.
  - Charge clouds by low-p_T protons.
  - Should be identified by different algorithms.
- Majority of hits below 200 MeV/*c belong to additional legs of looping tracks.*

→ *Potential to remove ~50% of clusters in total.*

→ *Can gain missing 2x compression factor.*

Attachment of adjacent clusters – work in progress.

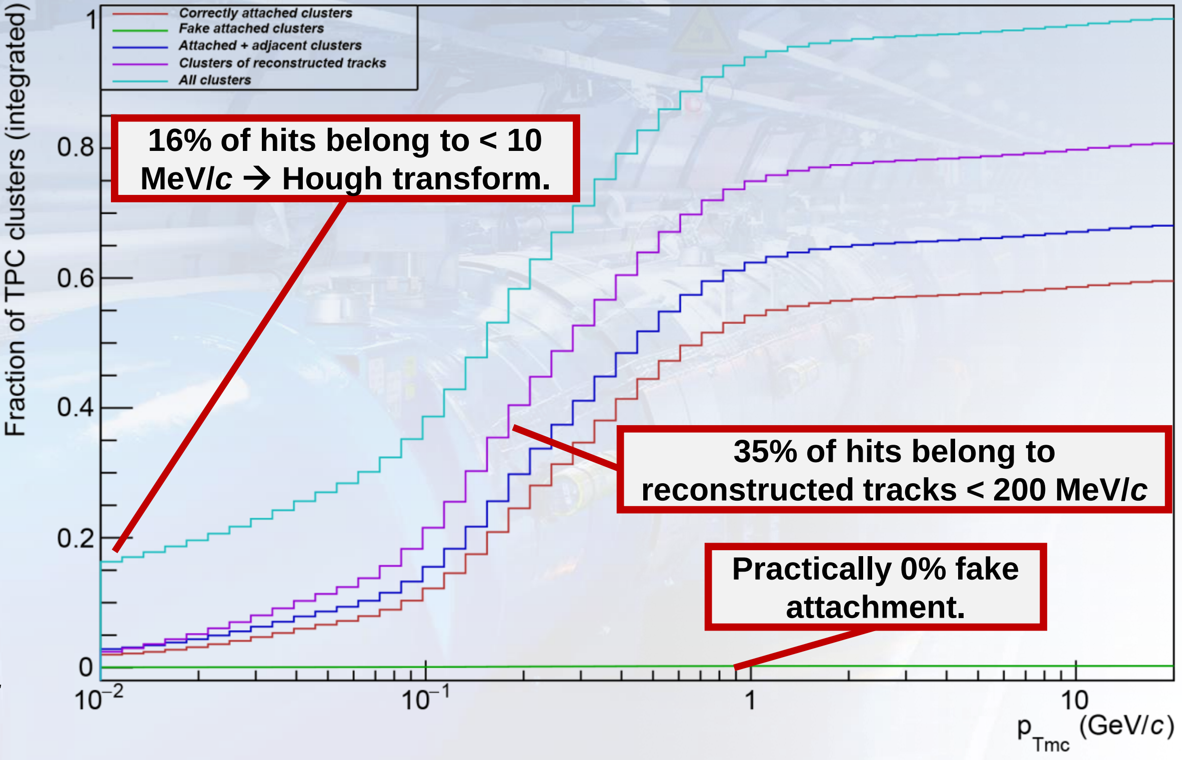

Global Tracking (TPC + ITS + TRD)

(R. Shahoyan, M. Puccio, O. Schmidt)

- **TPC-TRD Prolongation TRD developed within HLT framework.**
  - Good efficiency so far, comparable to offline.
  - Online version uses only TRD tracklets.
  - Decrease for low-p_T due to absence of TRD hits in Run 3.
  - Reduced purity in Pb-Pb due to large amount of TRD fake tracklets.

- **Status of ITS tracking**:
  - GPU-accelerated ITS standalone tracking under development. (first version available)
  - TPC to ITS track matching available (comparable purity as in Run 2).
  - See poster of M. Concas: #323 !

- **Next steps**:
  - Work on combined TPC + ITS + TRD tracking and fit on GPU without intermediate data transfer.
  - Test TPC calibration procedure using TPC + ITS + TRD tracking.

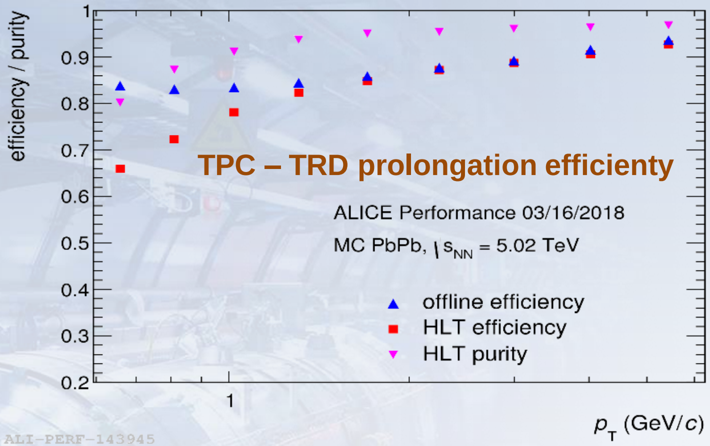

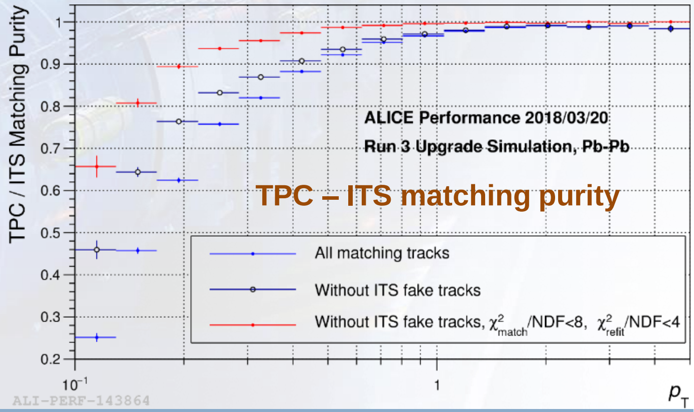

Summary

- **ALICE TPC Track reconstruction for Run 3 derived from Run 2 Online Tracking.**
  - Cellular Automaton + Kalman Filter.
  - Runs on GPUs (Common source code for CPU / GPU with OpenMP / CUDA / OpenCL).
- **Enormous speed-up compared to Run 2 offline.**
  - 20x – 25x speed-up on single CPU core.
  - GTX 1080 GPU replaces ~800 CPU cores (running Run 2 offline code).
  - Processing of 23 ms time frame needs ~20 seconds on one EPN. (Compute farm has ~1500 EPNs).
- **Tracking independent from absolute *z*-position (needed to process time frames).**
  - Same efficiency and resolution as Run 2 offline (some decline for deep secondaries).
  - Small decline in efficiency for short low-pT secondaries with 50 kHz time frames as compared to single events.
    - Unavoidable due to higher occupancy.
- **Need TPC data compression factor 20x (compared to Run 2 raw data size).**
  - Factor 8.3 in Run 2, Run 3 prototype achieves 9.1.
  - Potential to gain missing factor 2 by removing clusters not used for physics.
    - Removal of clusters of low-p_T tracks down to 10 MeV/*c* already working in tracking.

BACKUP

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)

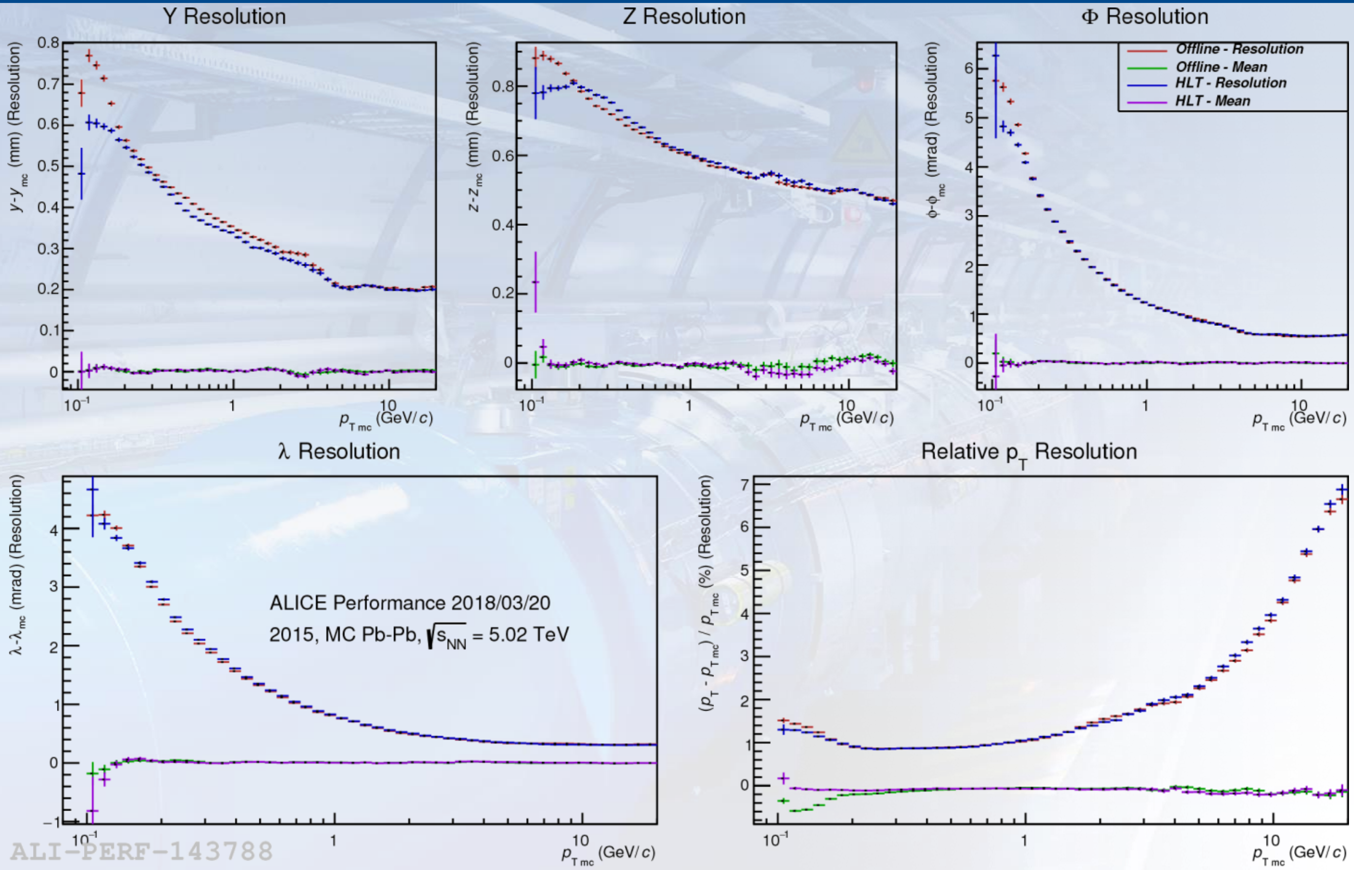

Tracking efficiency / resolution (Run 2, HLT v.s. Offline – pp)

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

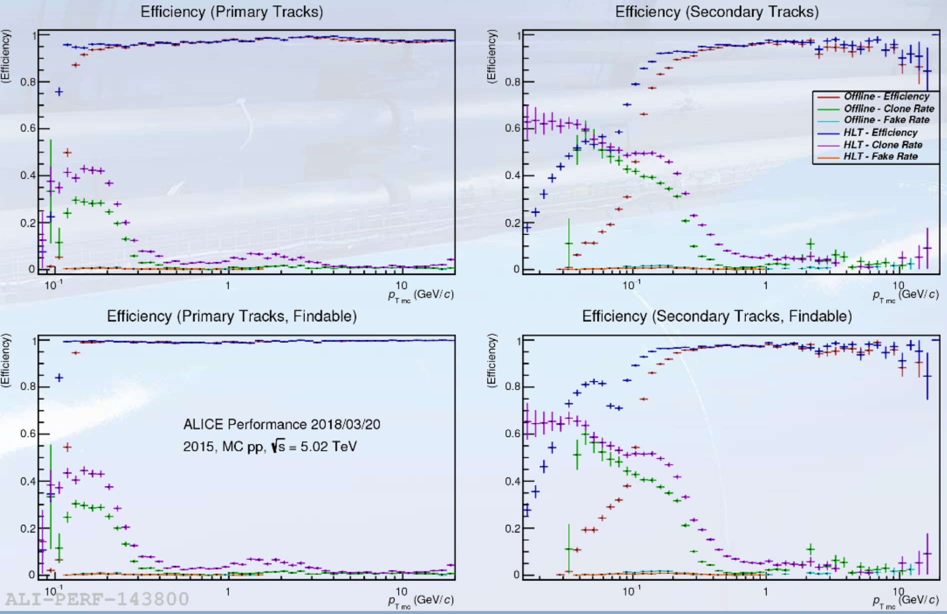

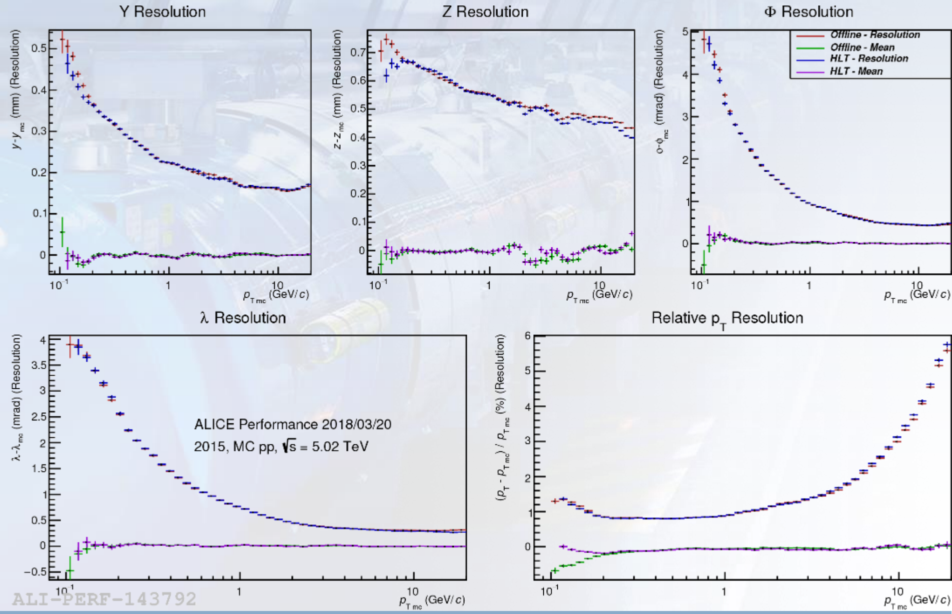

Tracking time frames at different interaction rates

- 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.

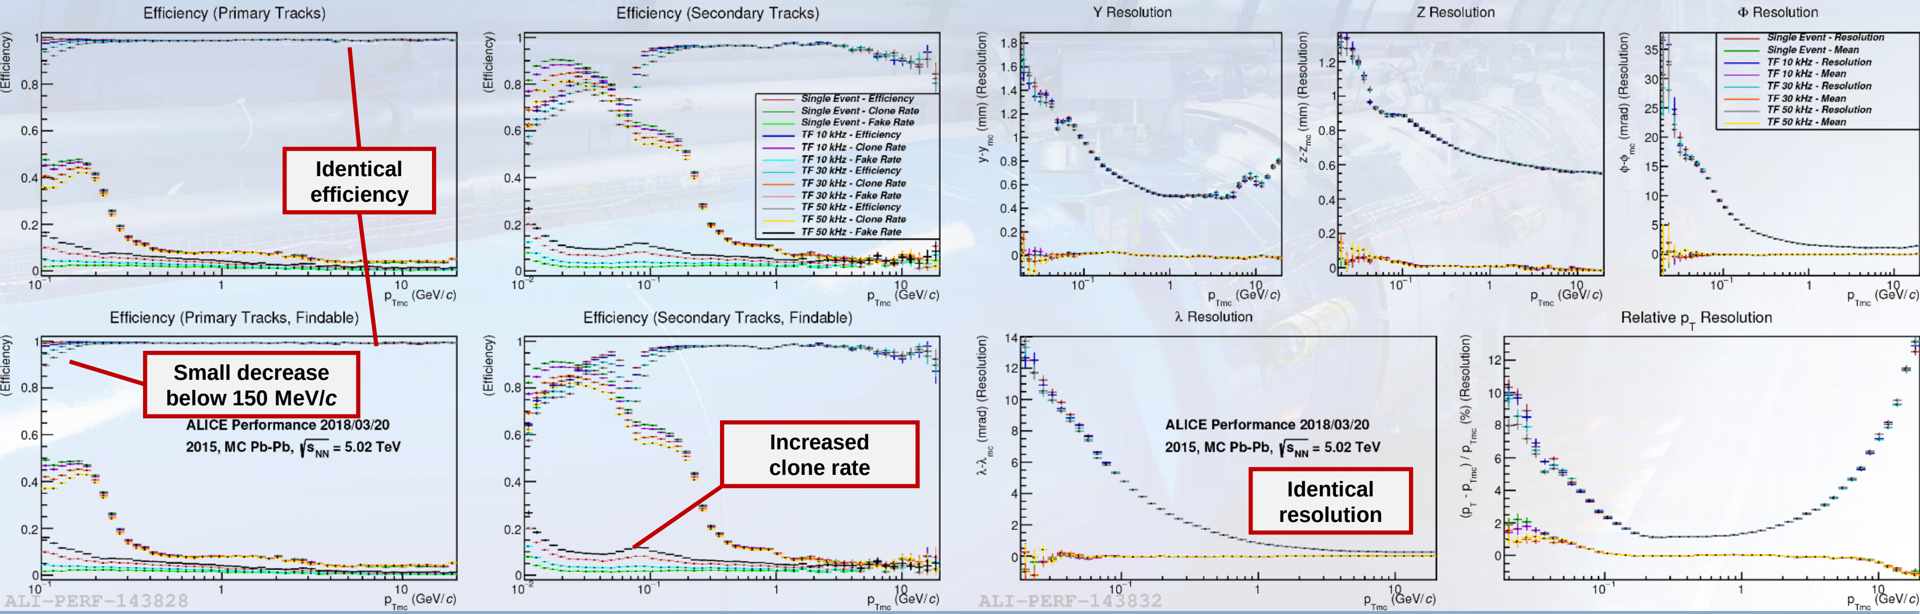

Multiplicity / event pile-up (pp)

- 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.

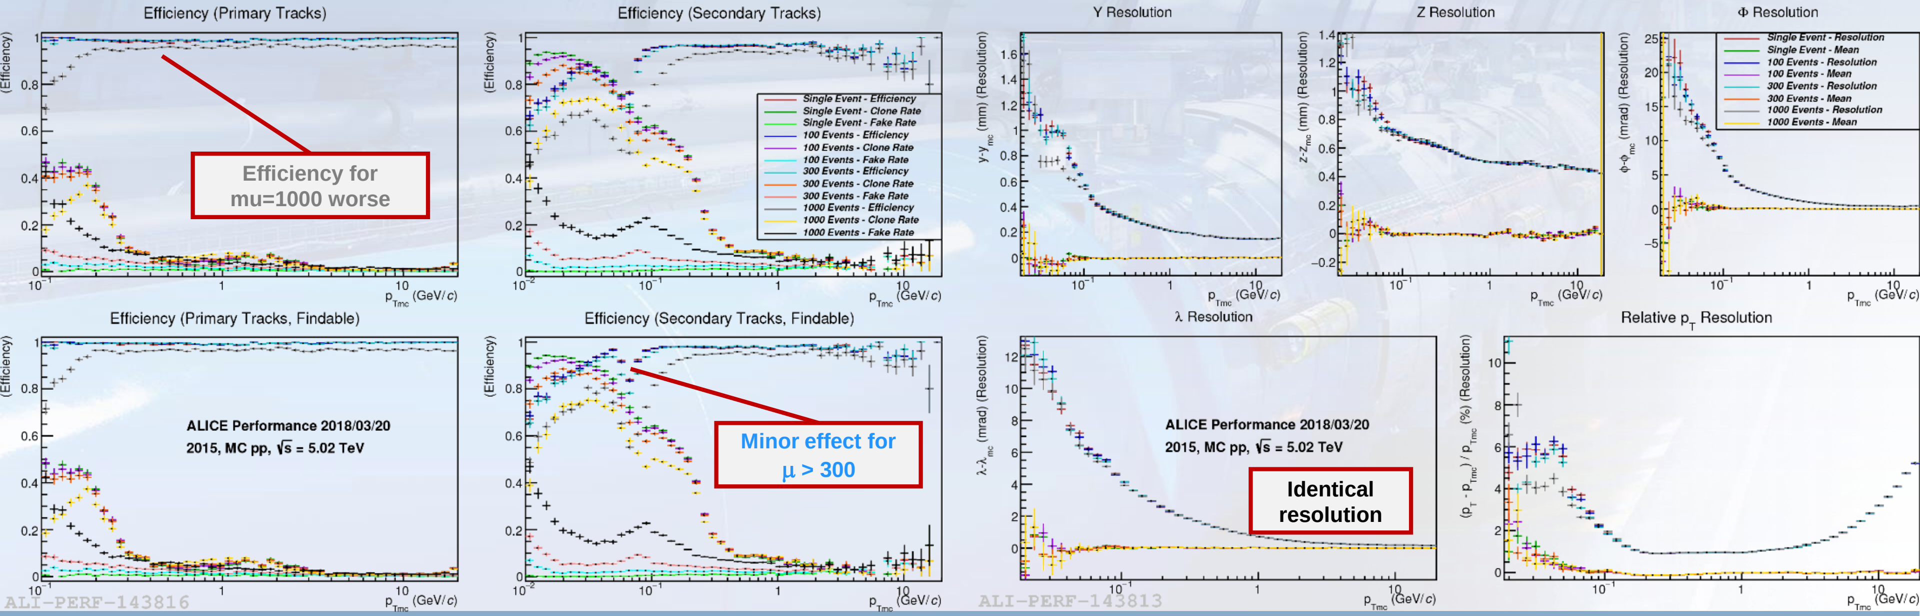

Normal tracking / z-independent tracking

- 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.**
  - Due to slightly larger errors, more tracks are merged.

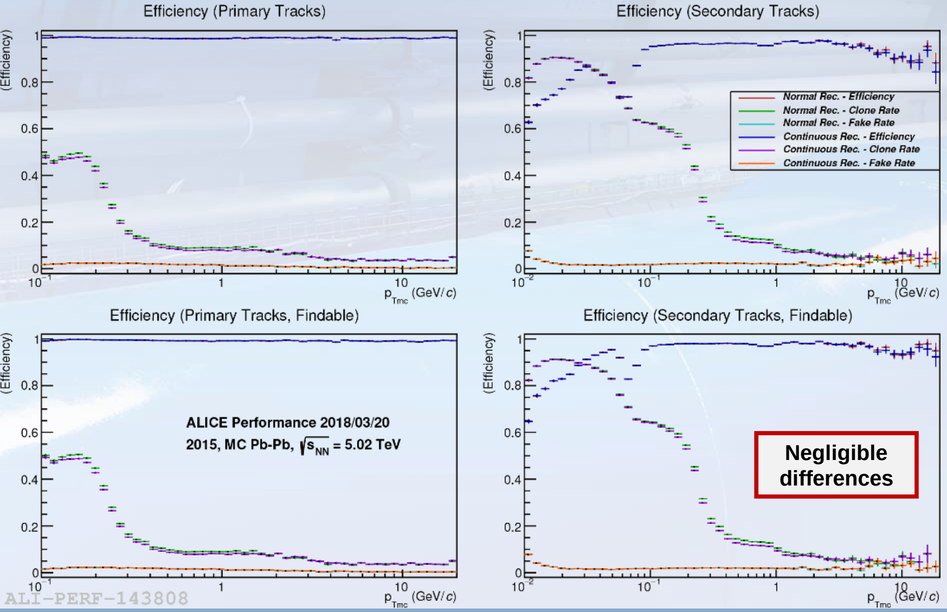

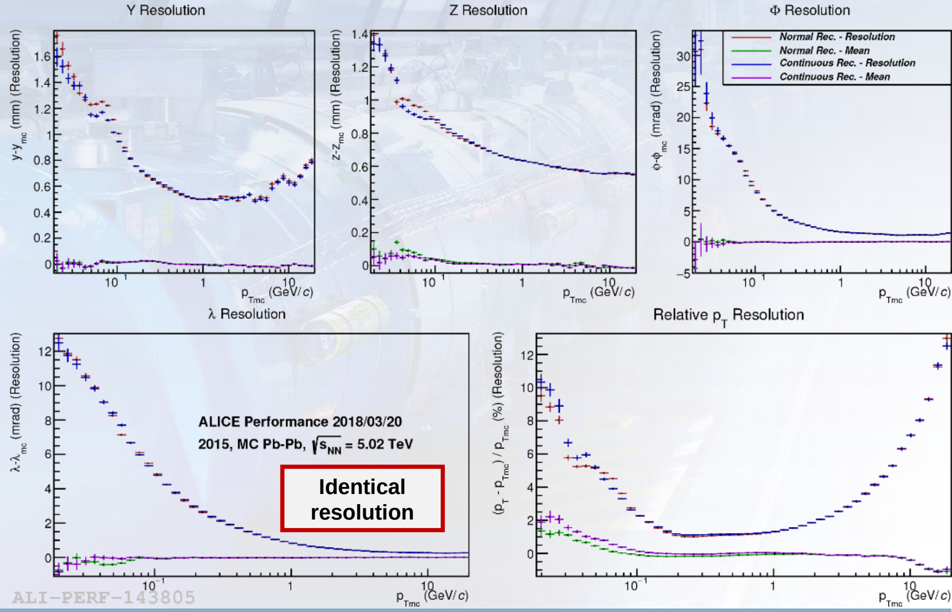

Length of time frame

- **Identical result independent of length of time frame.**
- Before, efficiency / resolution decreased with long time frames.
  - Completely fixed.
  - Floating point problems avoided by z-independent tracking (track fit happens in |z| < 250 cm).
  - Fixed precision for storing clusters (16 bits as used in the HLT insufficient for full TF).
  - Some other minor problems solved.