Patatrack:

Accelerated Pixel Track reconstruction in CMS

Andrea Bocci[1], Vincenzo Innocente[1], Matti Kortelainen[2],
Felice Pantaleo[1], Roberto Ribatti[3],
Marco Rovere[1], Gimmy Tomaselli[3]

[1]CERN – Experimental Physics Department, [2]FNAL,
[3]Scuola Normale Superiore di Pisa

https://patatrack.web.cern.ch/patatrack

Spoiler Alert

- The Patatrack team has implemented a full pixel reconstruction of CMS running on GPUs
- The Patatrack Pixel Reconstruction is fully integrated in CMSSW
    - CMSSW has been improved to support heterogeneous workers
- The physics performance are improved wrt current CPU workflow
    - New algorithms and fitting techniques have been introduced
- The Patatrack Pixel Reco can achieve up to 1.5kHz on V100 and up to 800Hz on T4 with concurrent events in flight

CMS and LHC Upgrade Schedule

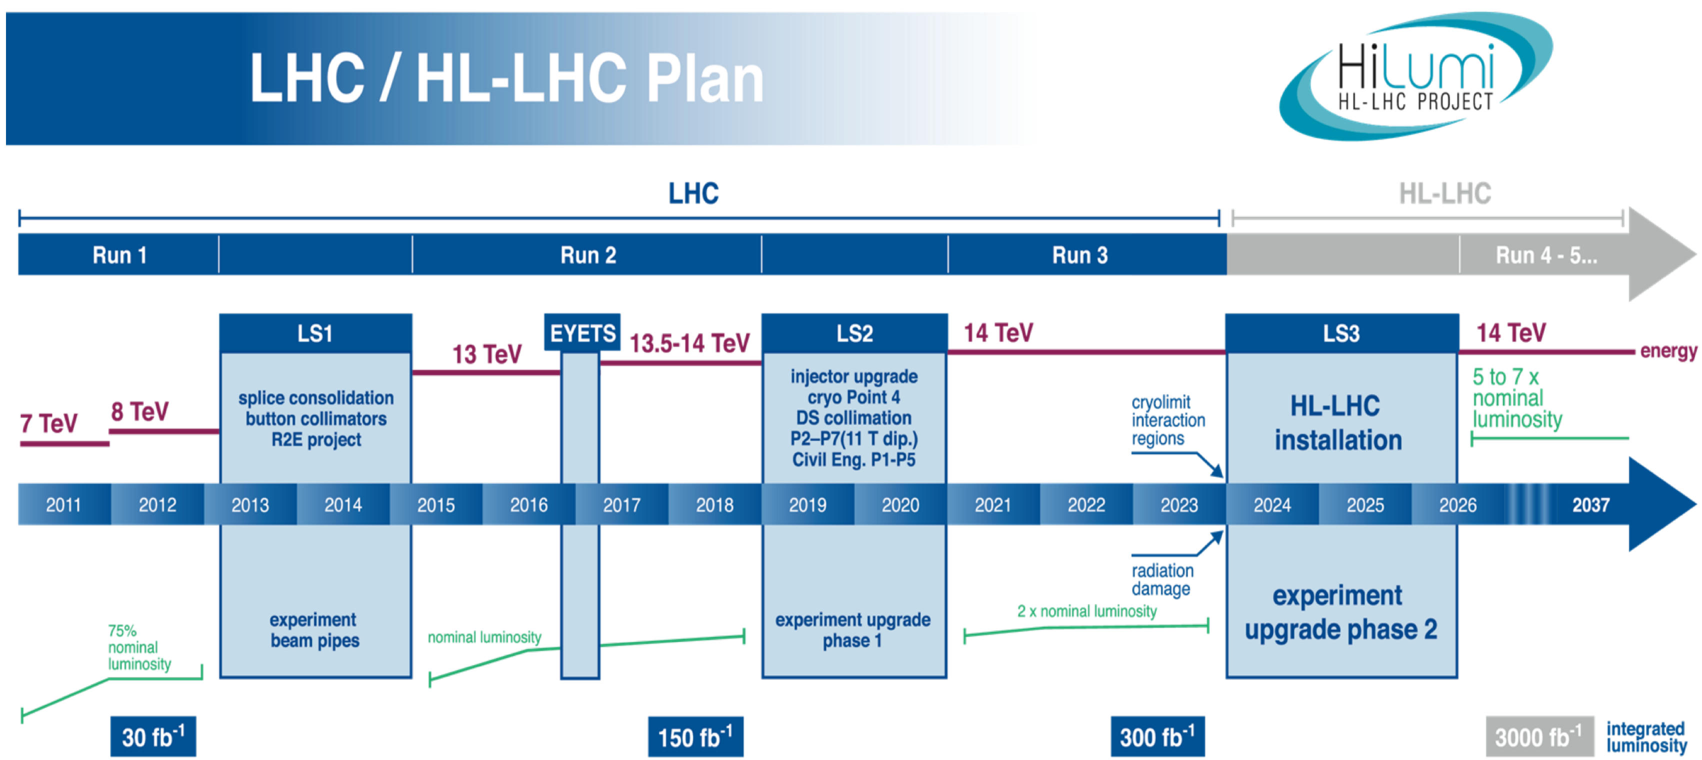

A reminder that…

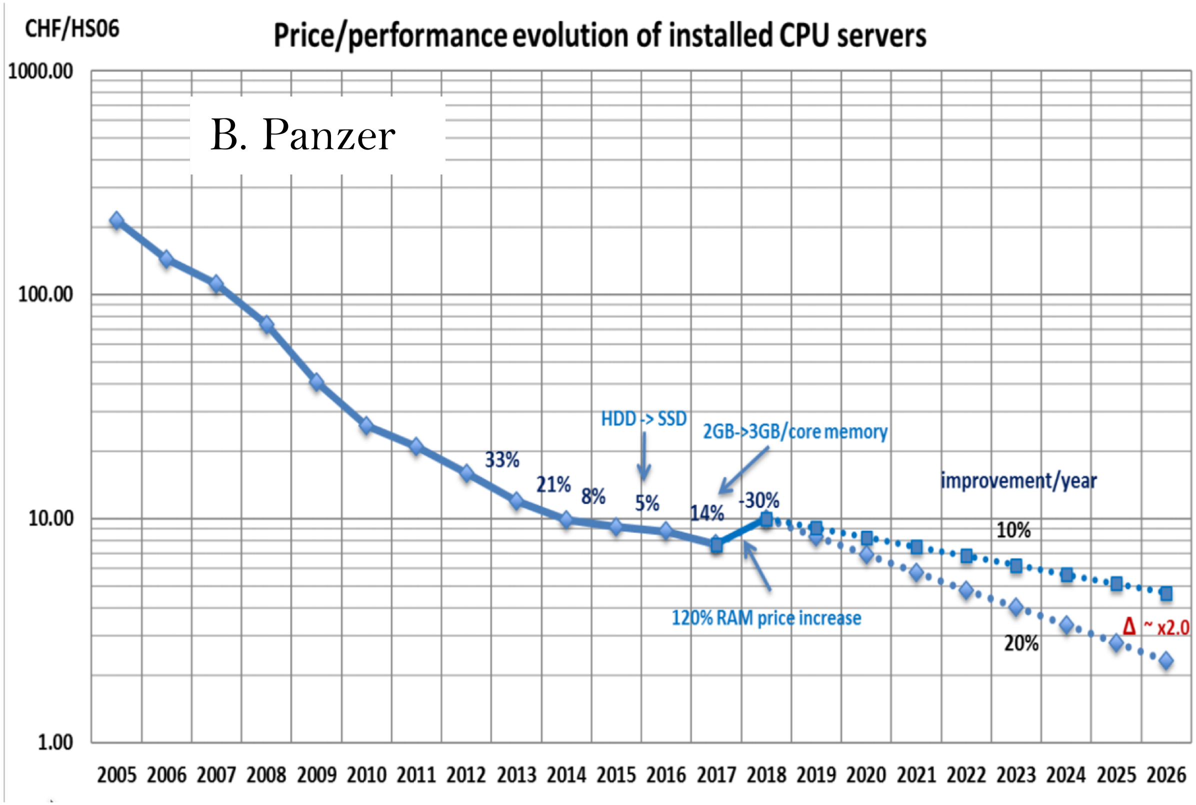

CPU evolution is not able to cope with the increasing demand of performance

…CERN is in a computing emergency

- Performance demand will increase substantially at HL-LHC
- 30x more CPU performance offline and online
- Traditional HW and SW will not suffice
- Major innovation required

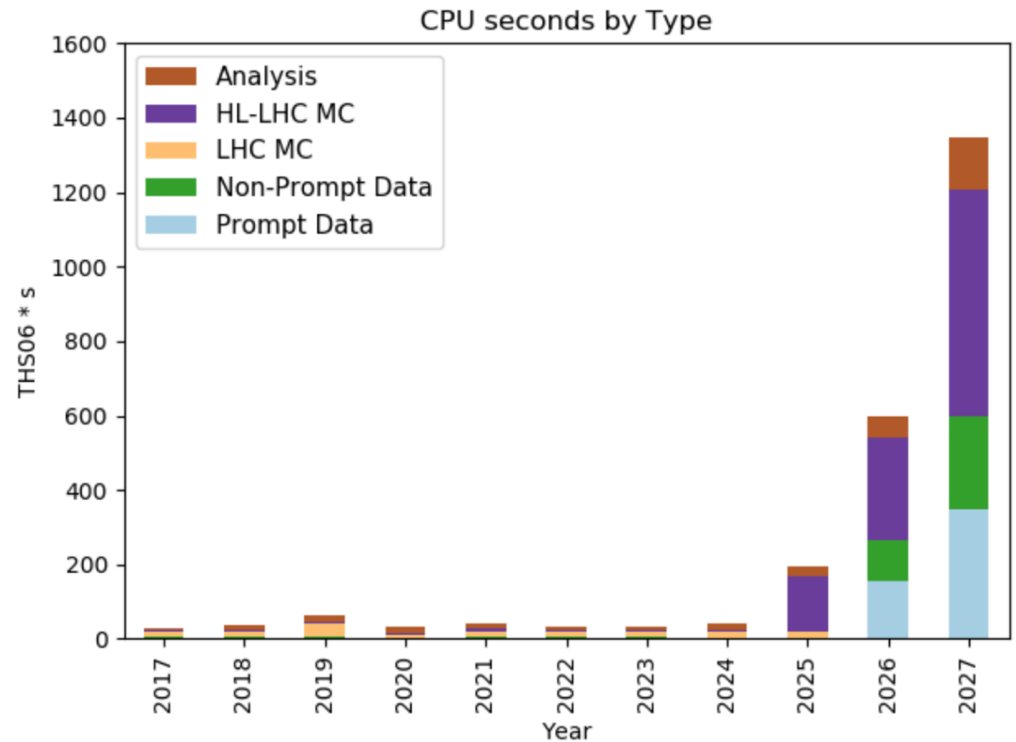

Patatrack

- Patatrack is a software R&D incubator
- Born in 2016 by a very small group of passionate people
- Interests: algorithms, HPC, heterogeneous computing, machine learning, software engineering
- Lay the foundations of the online/offline heterogeneous reconstruction starting from 2020s

Patatrack as an incubator

- Very few projects, although very promising, reach the production
- Many times getting from the PoC to a realistic demonstrator require skills outside the chosen technology
- Advices would be helpful along the full chain
- “Traditional” meetings are not helpful

In a computing emergency situation,

the failure of a group’s R&D to deliver the final PR becomes a failure of the entire collaboration

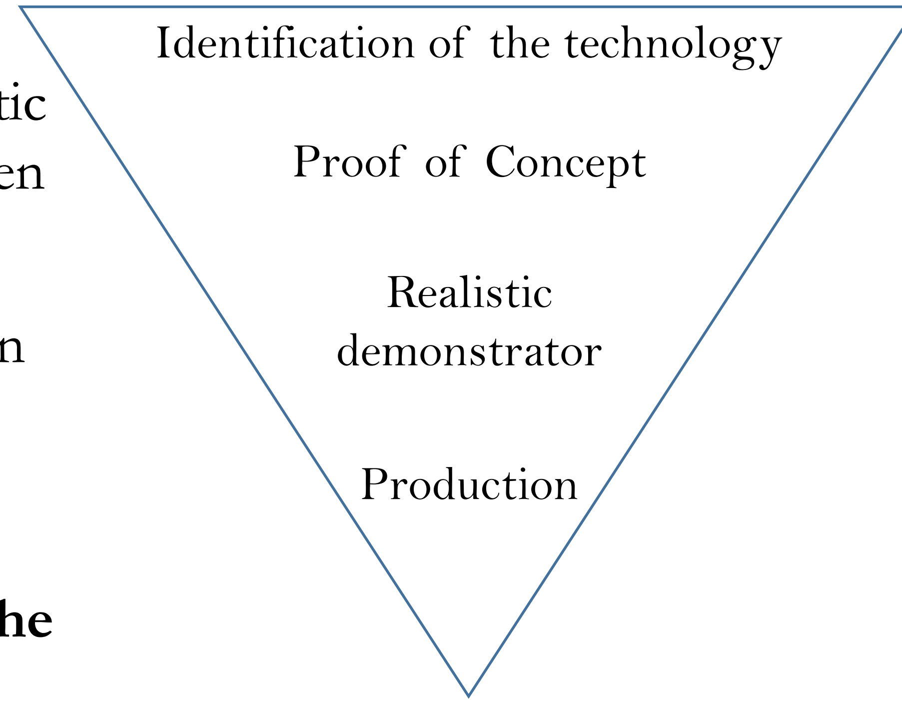

Incubate how?

- By creating a safe place where developers can stay and focus and finally get some work done uninterrupted for few days
- Allowing people to learn from experts what they are interested in and apply it on a problem by joining new projects
- By seeding new ideas and creating new working groups that can survive after the hackathon week
- By promoting knowledge transfer hence speeding up development process
- By helping from Proof of Concept to Execution

Patatrack Hackathons

- 5 Hackathons in the past two years
- 20-30 participants per hackathon
- 28 institutes
- cooked 50kg of pasta

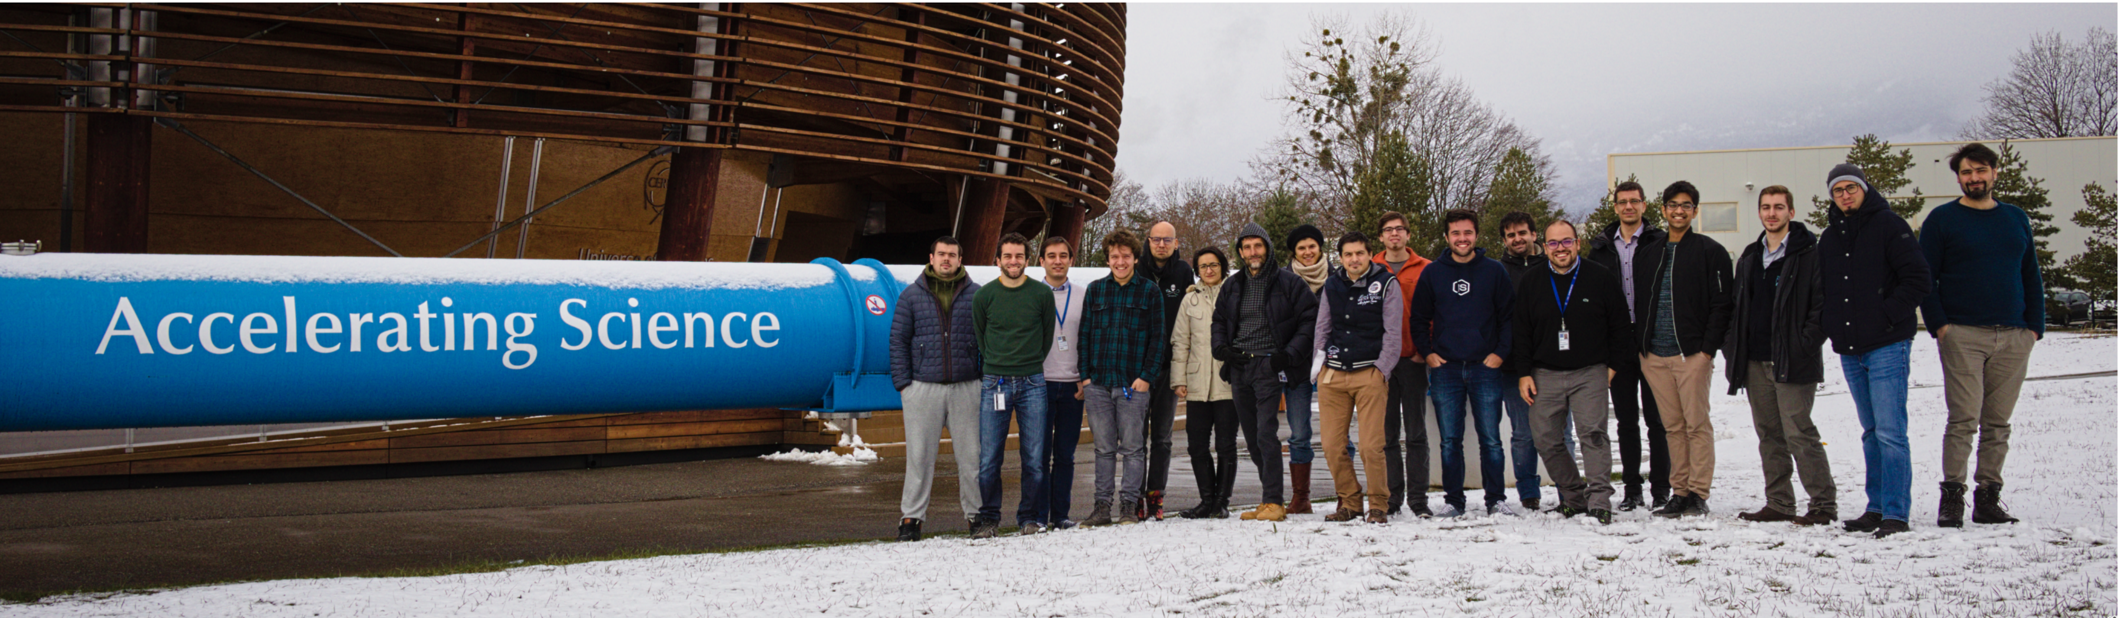

Patatrack Hackathons (ctd.)

- Usually 3-5 days at CERN Idea Square

Patatrack Hackathons (ctd.)

- Usually 3-5 days at CERN Idea Square
- Daily scrums used to receive feedback

Patatrack Hackathons (ctd.)

- Usually 3-5 days at CERN Idea Square
- Daily scrums used to receive feedback
- Experts available to work/design together

Patatrack Hackathons (ctd.)

- Usually 3-5 days at CERN Idea Square
- Daily scrums used to receive feedback
- Experts available to work/design together
- Creating a warm, familiar environment to foster networking, and team-working

Hackathons outcome

- 5 projects PoC -> ~Production
- 5 projects from idea to PoC
- 2 projects from idea to Realistic demonstrator
- 2 brand new working groups created
- Many new connections made
- Feedbacks averaged 4.8/5

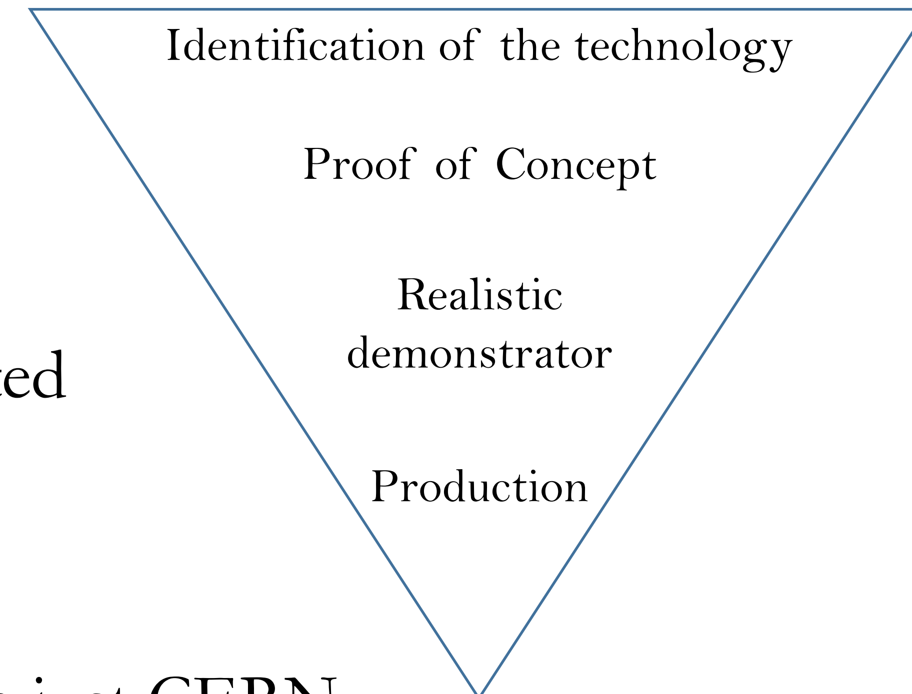

Not just CMS, non just CERN

Accelerating Pixel Tracks and Vertices during Run 3

- To ensure smooth operation of the heterogeneous HLT farm during run4
- Accelerated RAW data to Pixel Track and Vertices reconstruction by means of GPUs
- Complexity scales ~quadratically with respect to pile-up
- Today, at 50PU, Pixel Tracks are reconstructed only for ~10% of the events at the High-Level Trigger
- Profit from the upgrade of the Pixel to redesign the pixel track algorithm from scratch
- Integration in the CMS software framework

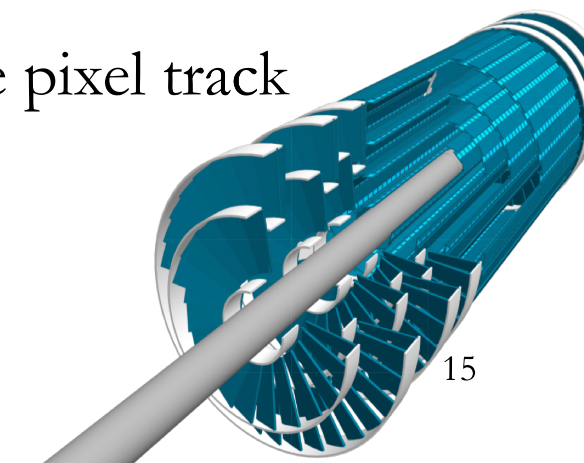

Terminology

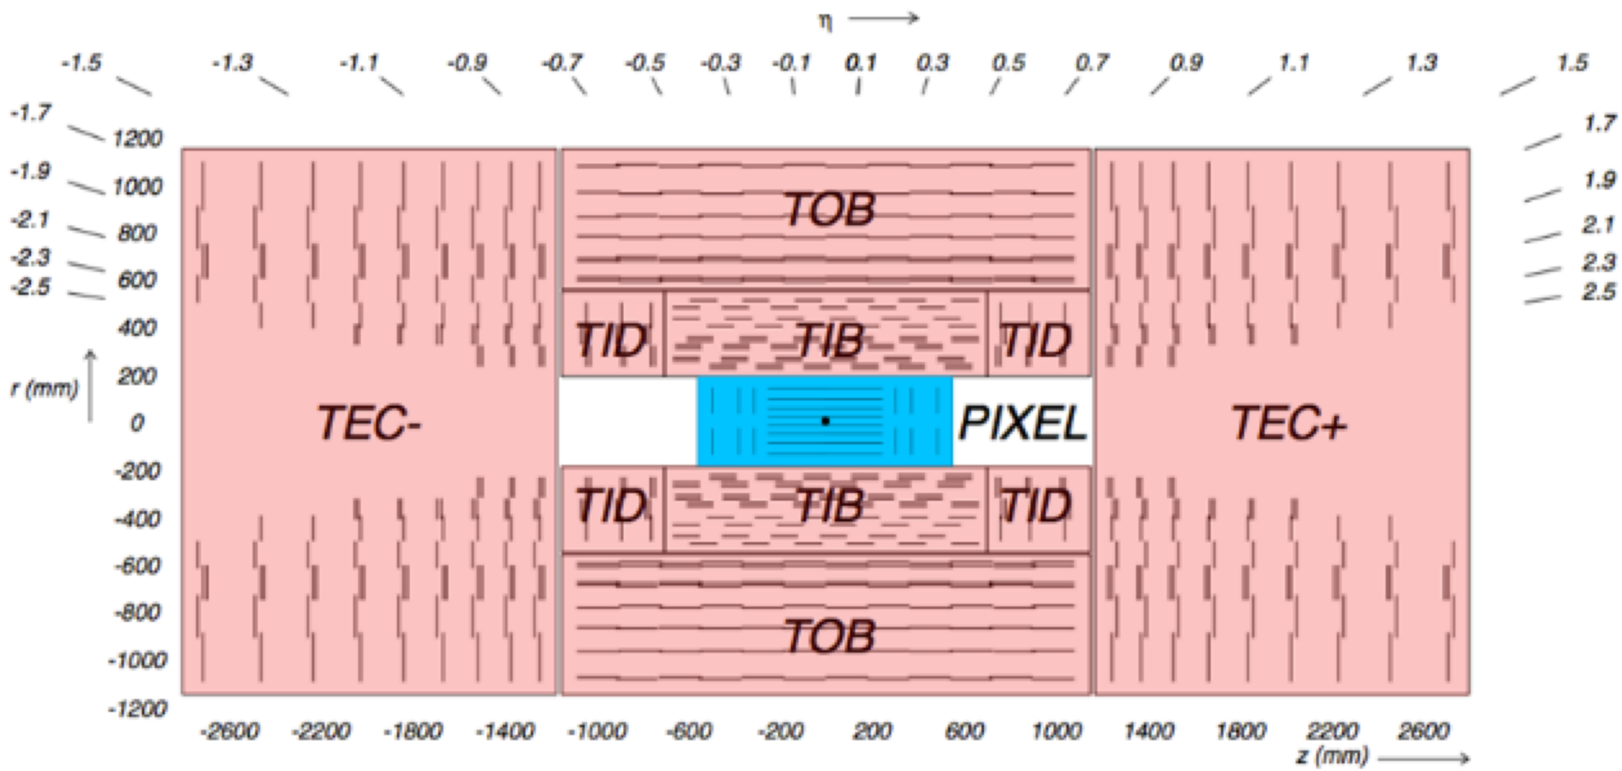

Online:

- Pixel-only tracks used for fast tracking and vertexing

Offline:

- Pixel tracks are used as seeds for the Kalman filter in the strip detector

Patatrack Pixel Reconstruction Workflow

- Full Pixel Track reconstruction in CMSSW
  - from Raw data decoding to Primary Vertices determination
- Raw data for each event is transferred to the GPU initially (~250kB/event)
- At each step data can be transferred to CPU and used to populate “legacy” event data
- The standard validation is fully supported
- Integer results are identical
- Small differences in the results of floating point can be explained by differences in re-association

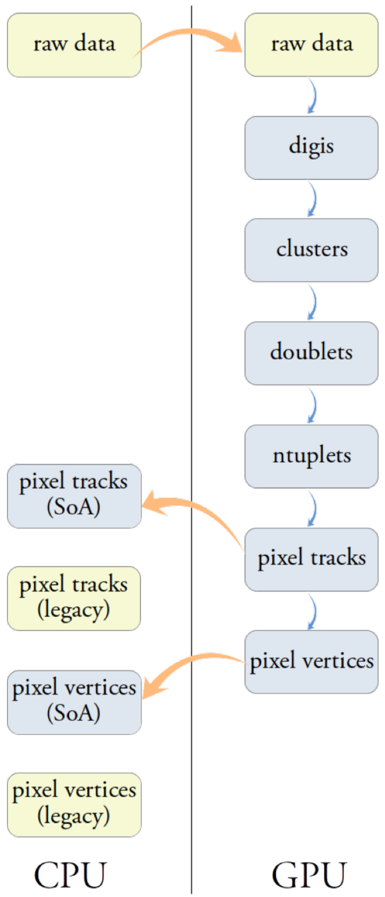

Doublets

- The local reconstruction produces hits
- Doublets are created opening a window depending on the tracking region/beamspot and layer-pair
  - The cluster size along the beamline can be required to exceed a minimum value for barrel hits connecting to an endcap layer
- Hits within the bins are connected to form doublets if they pass further "alignment cuts" based on their actual position
- In the barrel the compatibility of the cluster size along the beamline between the two hits can be required
- The cuts above reduce the number of doublets by an order of magnitude and the combinatorics by a factor 50
  - "Purity" @PU50 ~5%

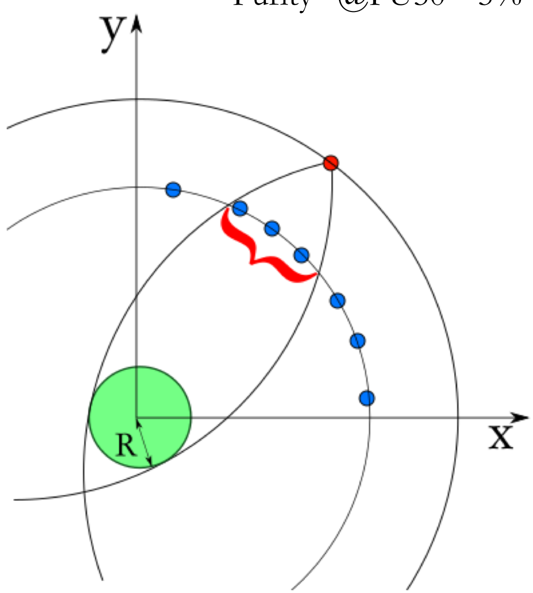

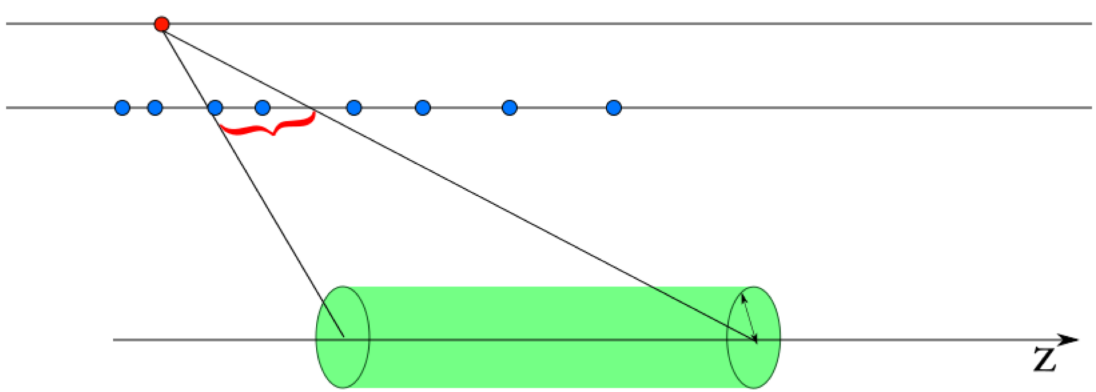

Cellular Automaton-based Hit Chain-Maker

The CA is a track seeding algorithm designed for parallel architectures

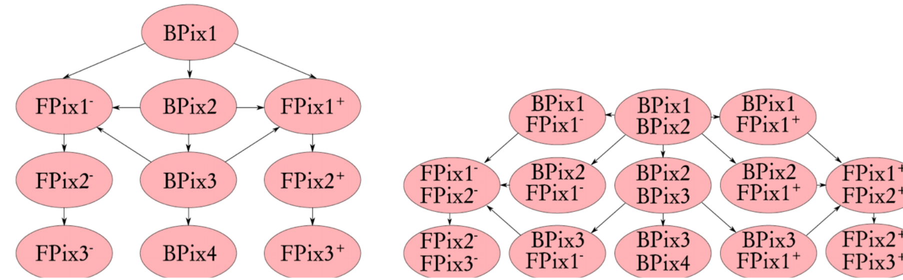

It requires a list of layers and their pairings

- A graph of all the possible connections between layers is created
- Doublets aka Cells are created for each pair of layers, in parallel at the same time
- Fast computation of the compatibility between two connected cells, in parallel
- No knowledge of the world outside adjacent neighboring cells required, making it easy to parallelize

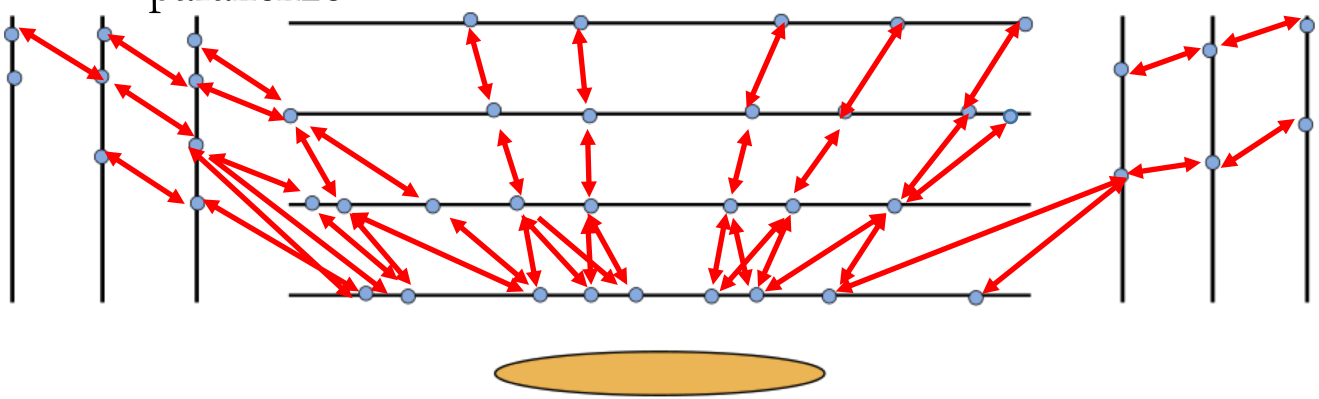

- Better efficiency and fake rejection wrt previous algo
- Since 2017 data-taking has become the default track seeding algorithm for all the pixel-seeded online and offline iterations

- In the following, at least four hits are required, but triplets can be kept to recover efficiency where geometric acceptance lacks one hit

CA compatibility cuts

- The compatibility between two cells is checked only if they share one hit
  - AB and BC share hit B
- In the R-z plane a requirement is alignment of the two cells
- In the cross plane the compatibility with the beamspot region

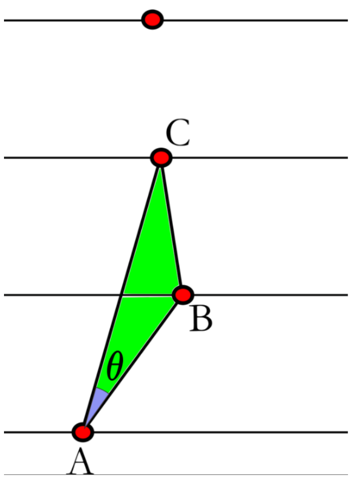

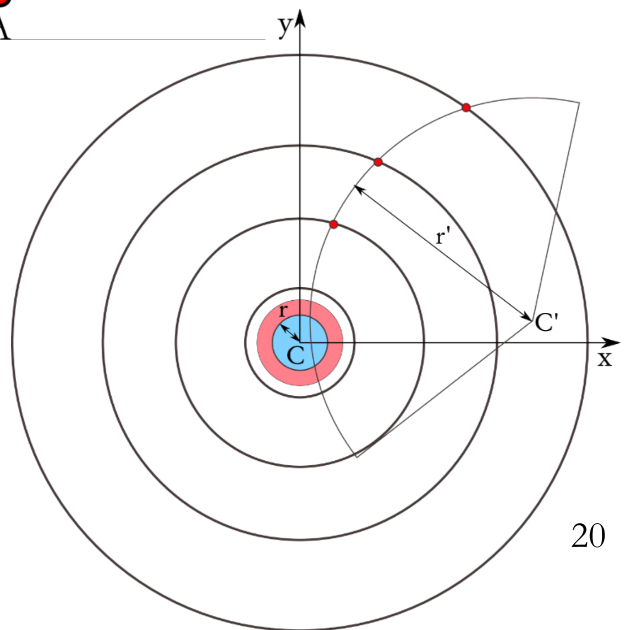

Fishbone

- After using the CA for producing N-tuplets, “fishbone” seeds can be produced to account for module/layer overlaps
- Only highest grade n-tuplet is fitted and duplicate doublets are filtered out

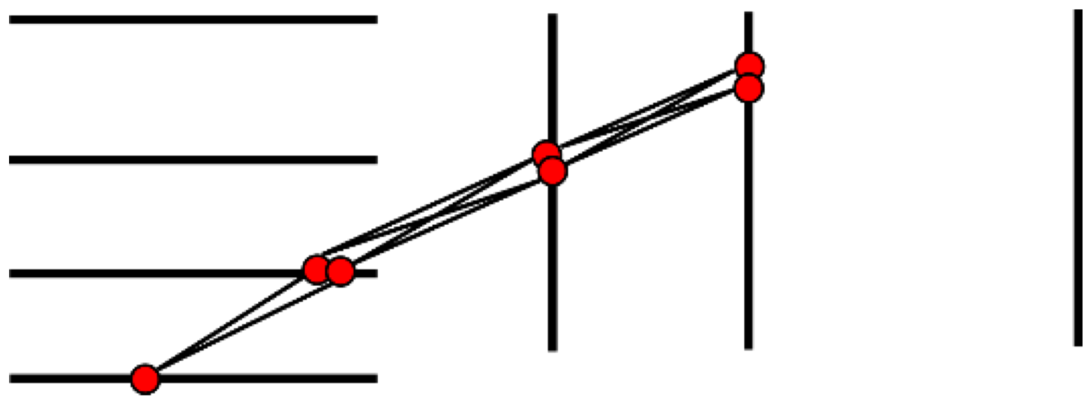

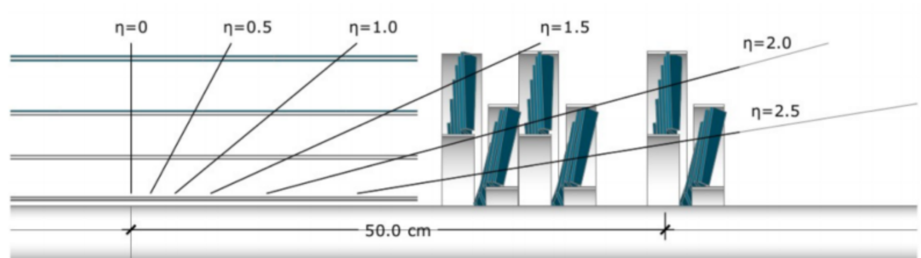

Fits

Pixel track "fit" at the HLT is still using 3 points for quadruplets and errors on parameters are loaded from a look-up table[eta][pT]

The Patatrack Pixel reconstruction includes two Multiple Scattering-aware fits:

- Riemann Fit
- Broken Line Fit

They allow to better exploit information coming from our 4-layer pixel detector and improve parameter resolutions and fake rejection

Fits - Implementation

Both the Riemann and the Broken Line fits have been implemented using Eigen

Eigen is a C++ template library for linear algebra, matrix and vector operations

This allows perfect code portability between CPU and GPU implementation and bitwise-matching of the results

Fits - Algorithm

Fitting procedure:

- Fast circle fit: estimate of p for MS, estimate of the radius/center
- Circle fit: d0, pT, phi
- Line fit: d_z, cot(theta)
- Return line and circle chi2

Riemann Fit:

- MS included in the covariance matrix

Broken Line Fit

- Fit of the broken line includes MS kinks in the design

Final Cleaning

- Among Tracks with at least one shared doublet only one with best chi2 retained
- Tracks "rejected" if fails chi2, TIP, ZIP or pt cut
- p_T>0.3 GeV, $|d_0|$<0.5 cm, $|z_0|$<12 cm

Performance Definitions

Physics performance:

- 20000 MC ttbar events <PU> = 50, design conditions, 25ns, sqrt(s)=13TeV
- Matching of reconstructed tracks with simulated ones requires that all hits of the reconstructed track come from the same simulated track
- Efficiency: number of matched reconstructed tracks divided by number of simulated tracks
- Fake Rate: number of non-matched reconstructed tracks divided by number of reconstructed tracks
- Efficiency is computed only with respect to the hard scatter.
- Efficiency has the following implicit cut: $|d_0| < 3.5$ cm additionally to the cuts quoted in the plots
- Duplicate is a reconstructed track matching to a simulated track that itself is matched to >= 2 tracks

Physics Performance - Efficiency

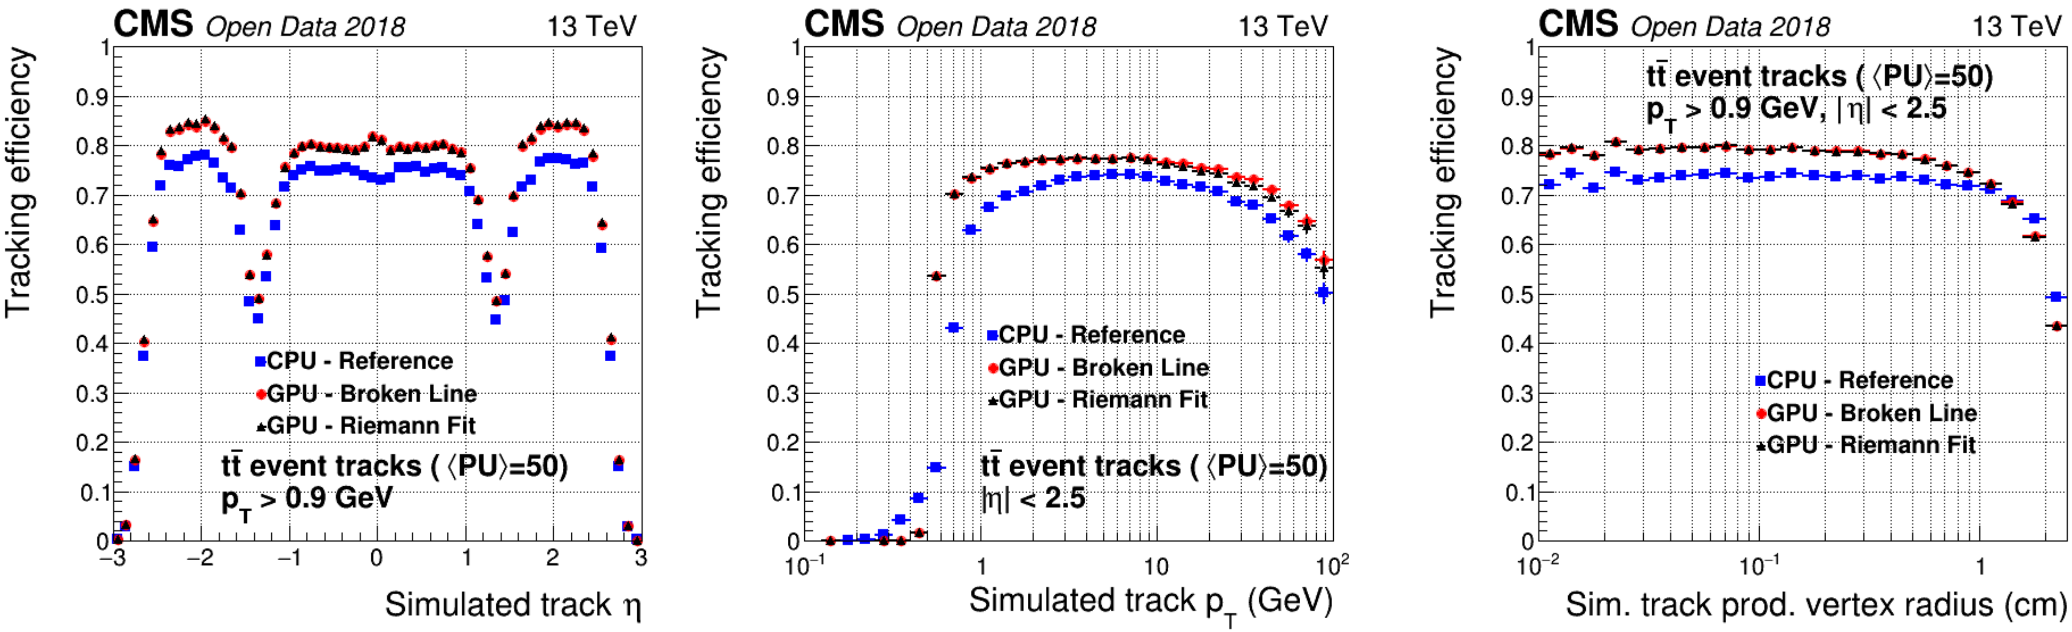

Track reconstruction efficiency as a function of simulated track η, p_T, and production vertex radius.

Physics performance - Duplicates

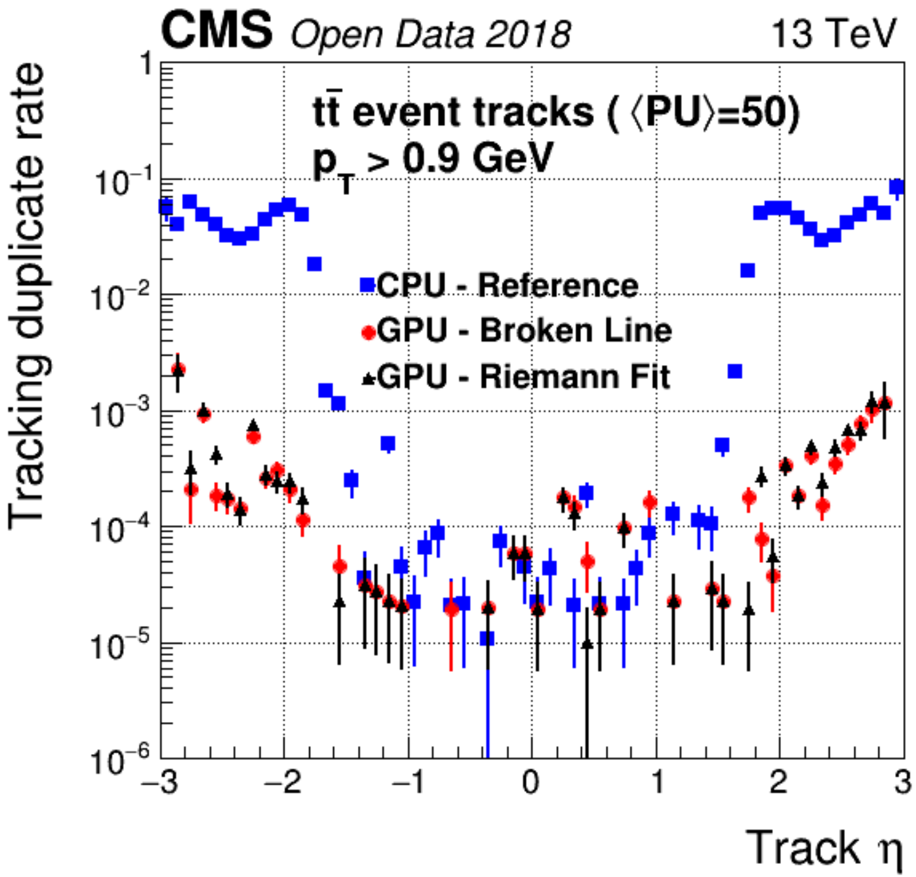

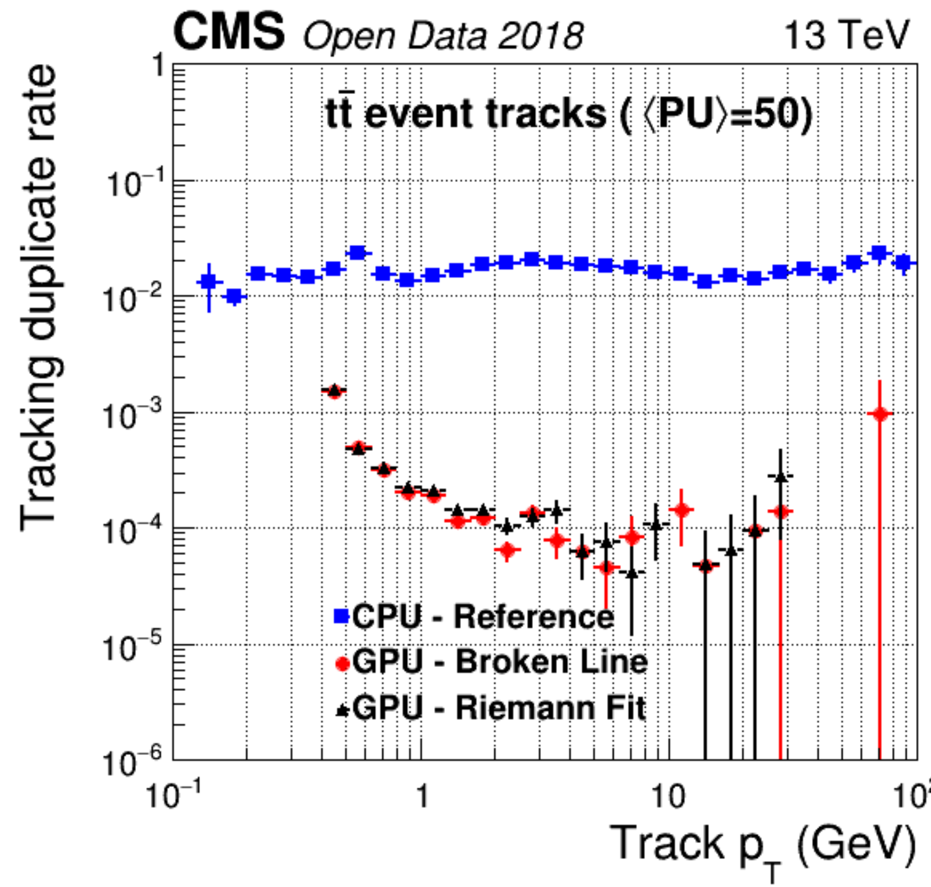

Track reconstruction duplicate rate as a function of reconstructed tracks η, p_T

Physics performance – Fakes

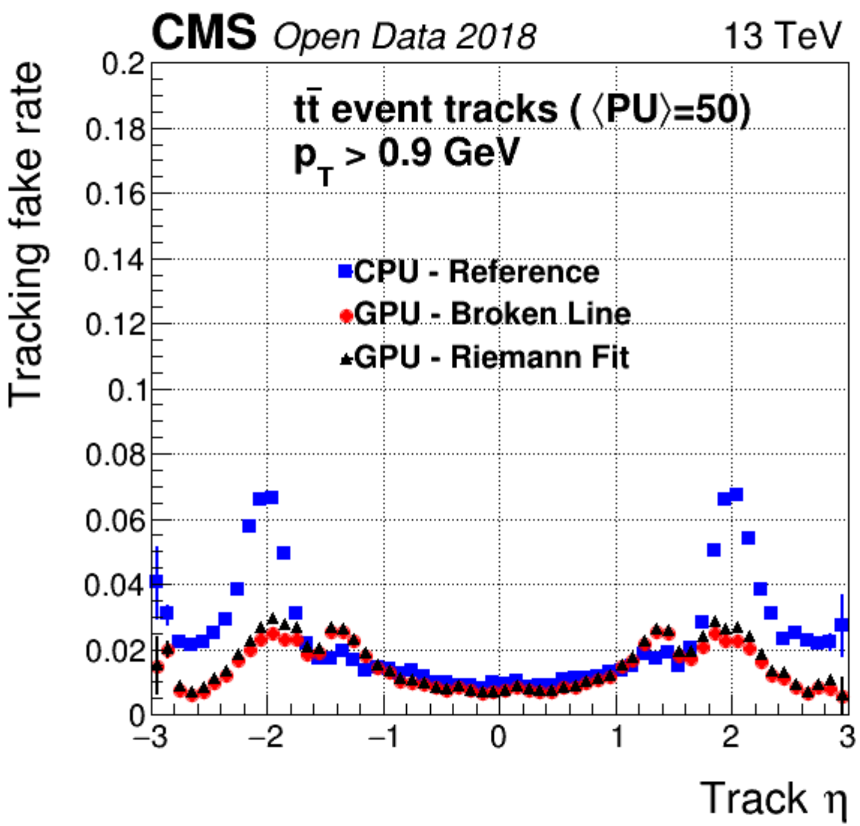

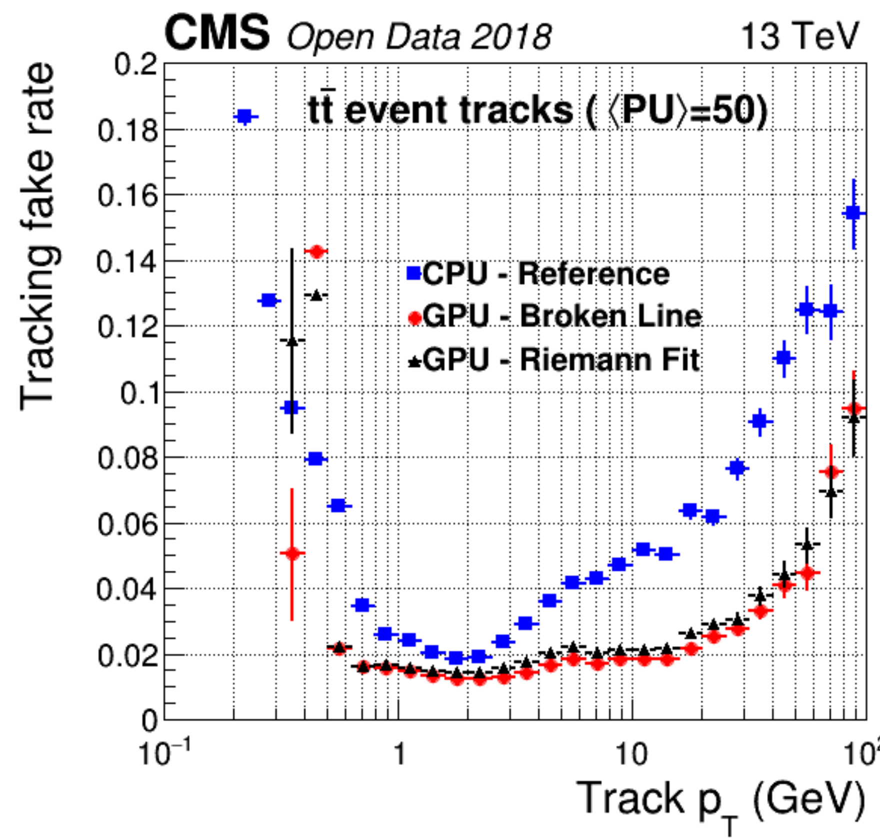

Track reconstruction fake rate as a function of reconstructed tracks η, p_T

Physics Performance – Fit Pulls

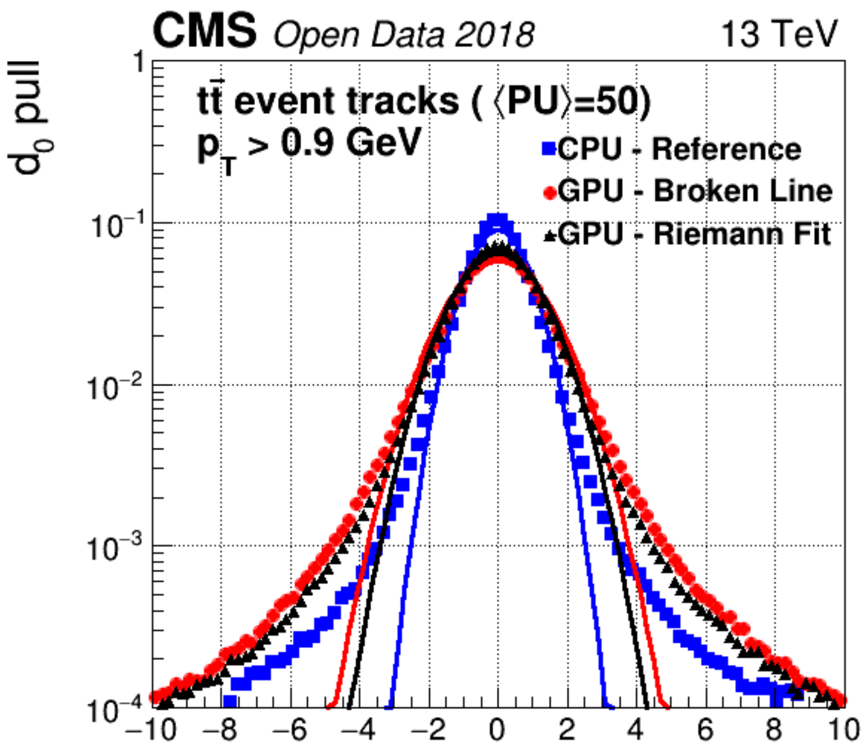

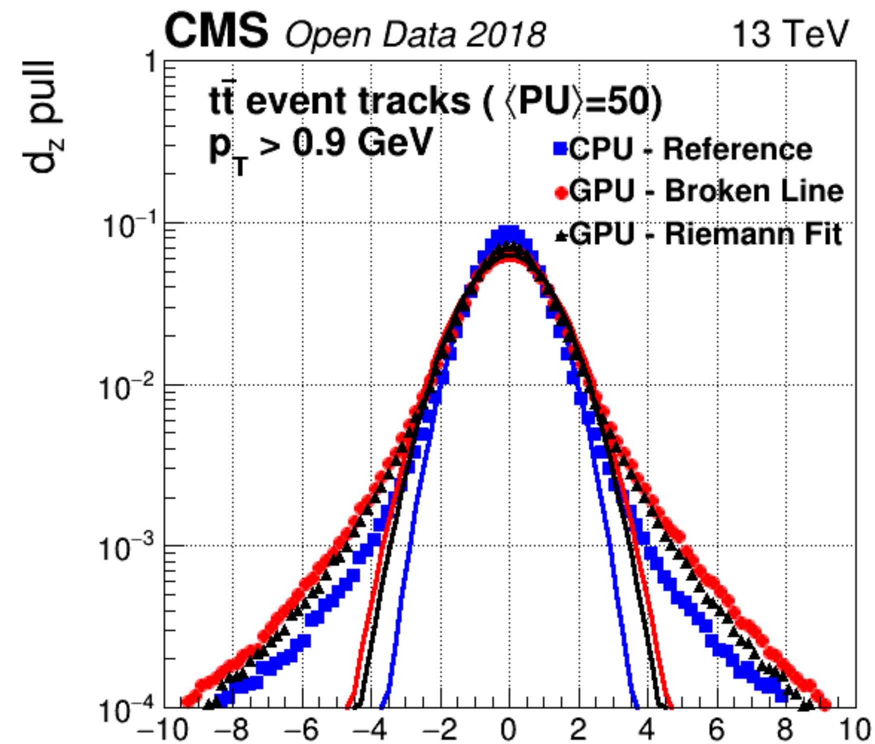

	σ - Reference	σ - Broken Line	σ - Riemann Fit
d_0	0.84	1.32	1.18
d_z	0.97	1.28	1.20

Physics Performance – Fit Pulls

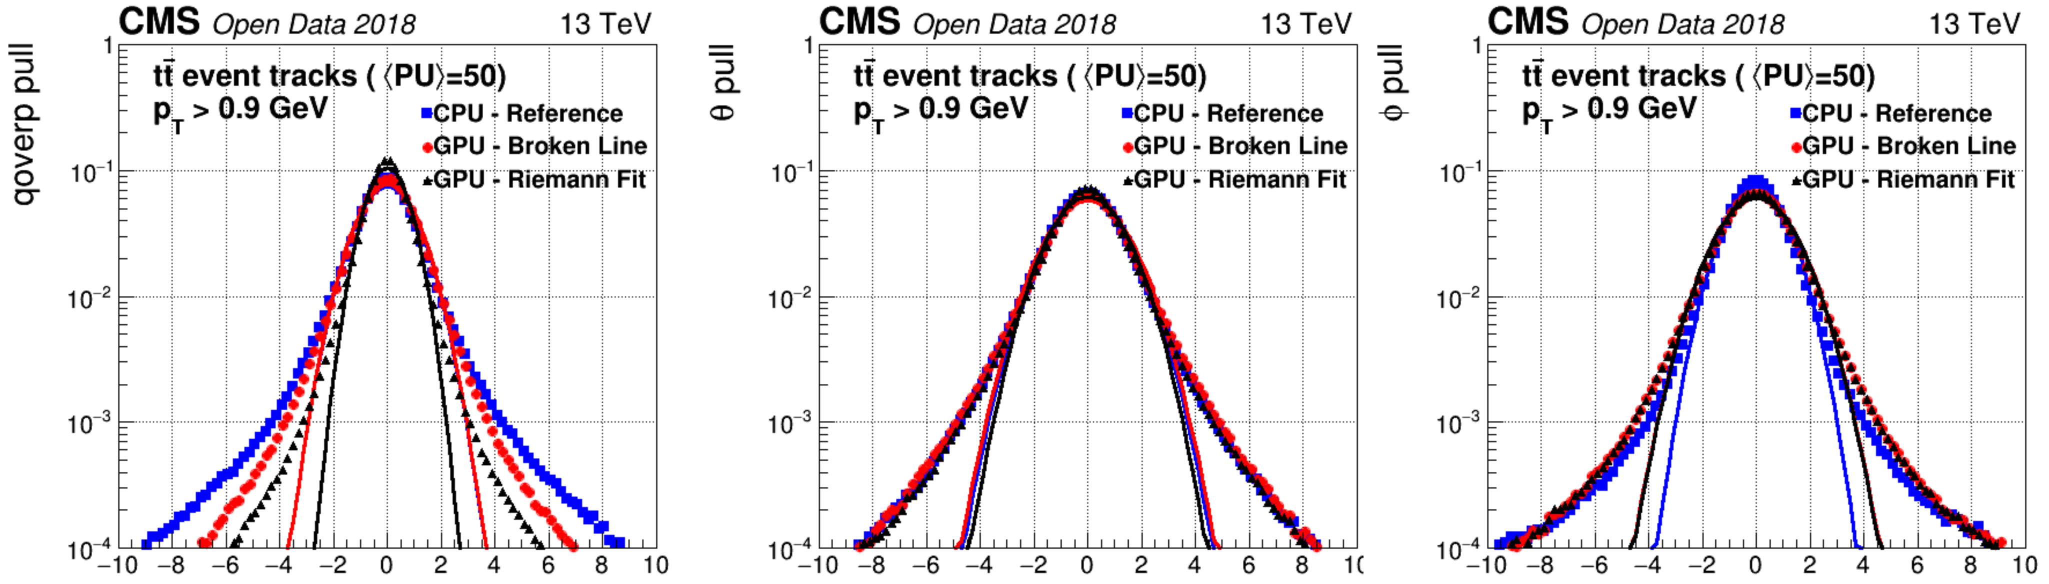

	σ - Reference	σ - Broken Line	σ - Riemann Fit
qoverp	0.99	0.99	0.72
θ	1.29	1.33	1.22
φ	1.02	1.28	1.27

Physics Performance - Resolutions

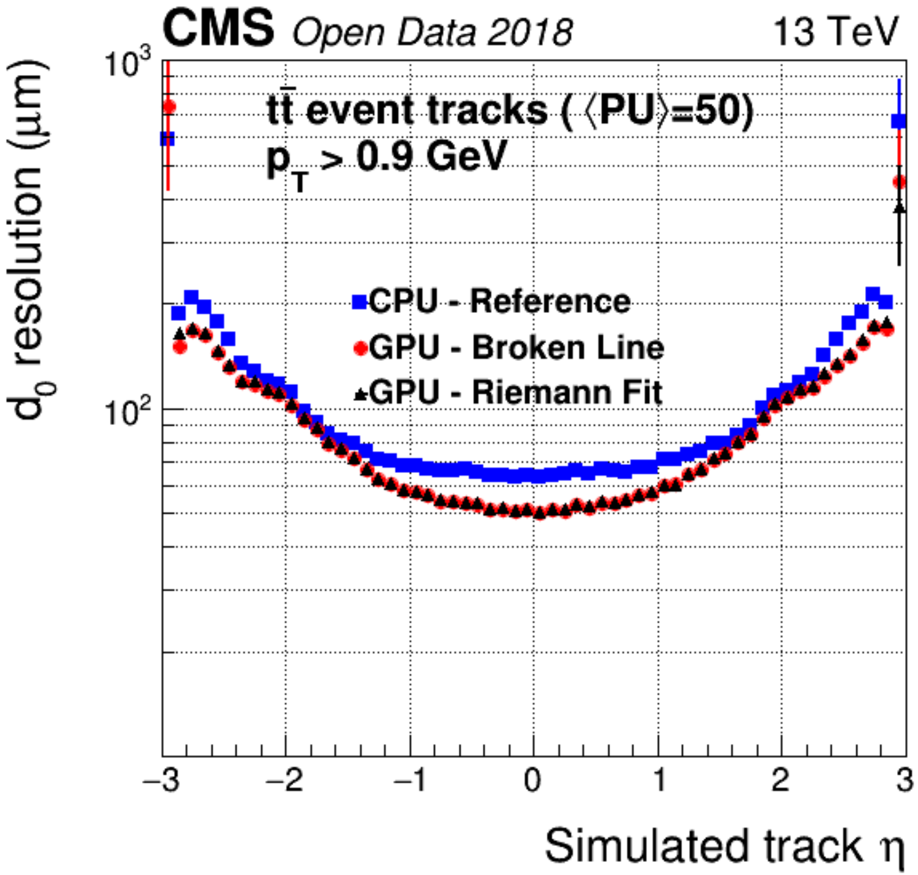

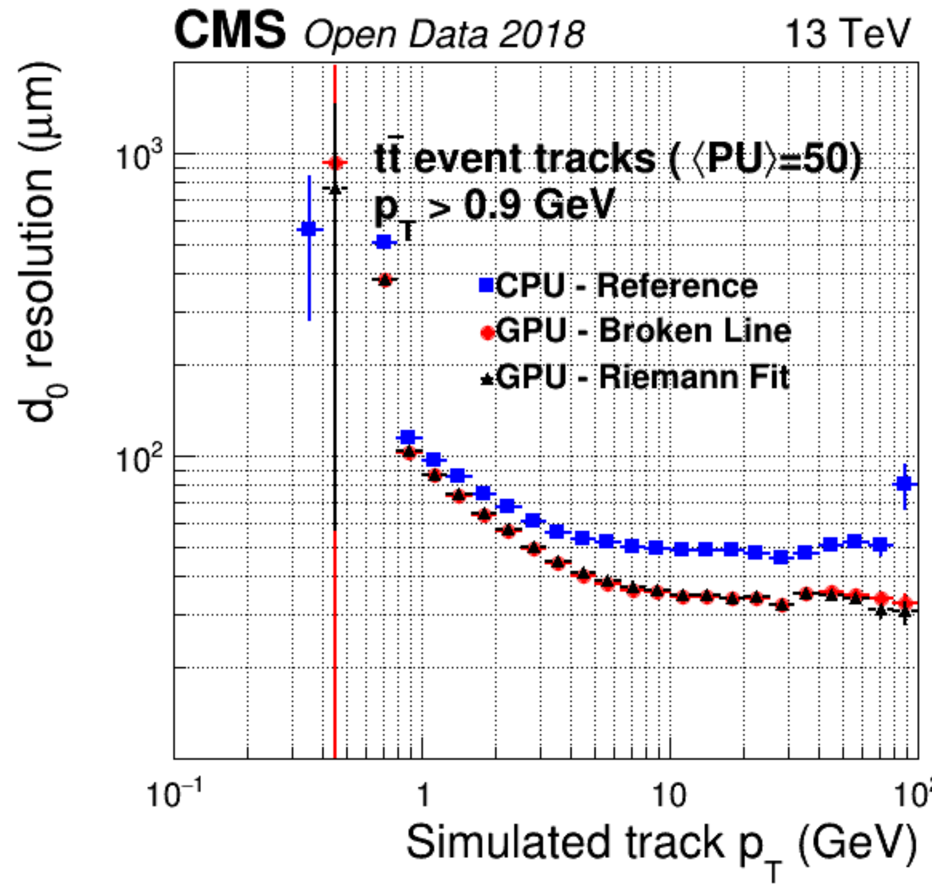

Track resolution of the transverse impact parameter as a function of simulated track η and p_T

Physics Performance - Resolutions

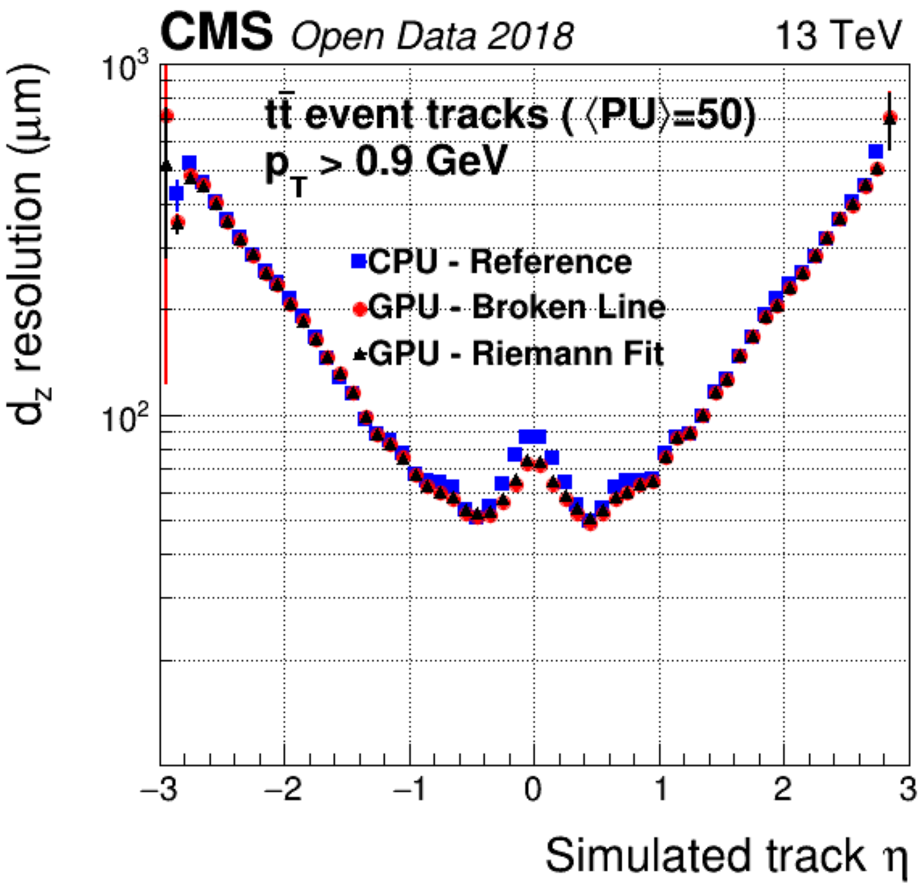

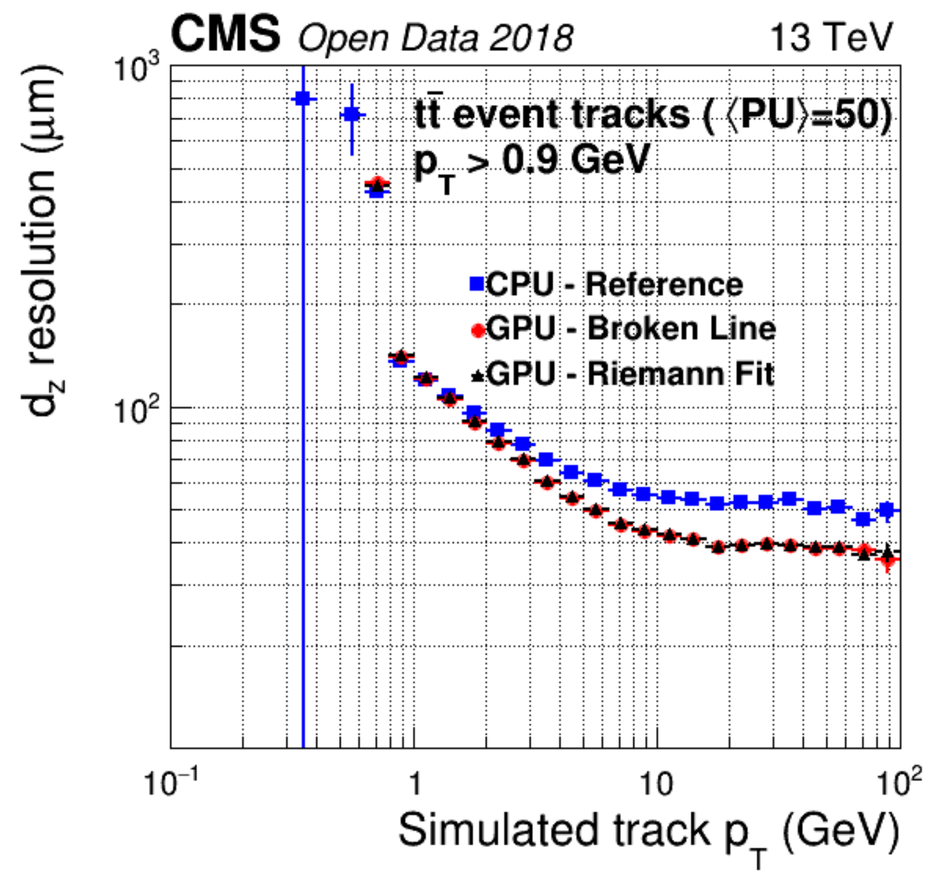

Track resolution of the longitudinal impact parameter as a function of simulated track η and p_T

Physics Performance - Resolutions

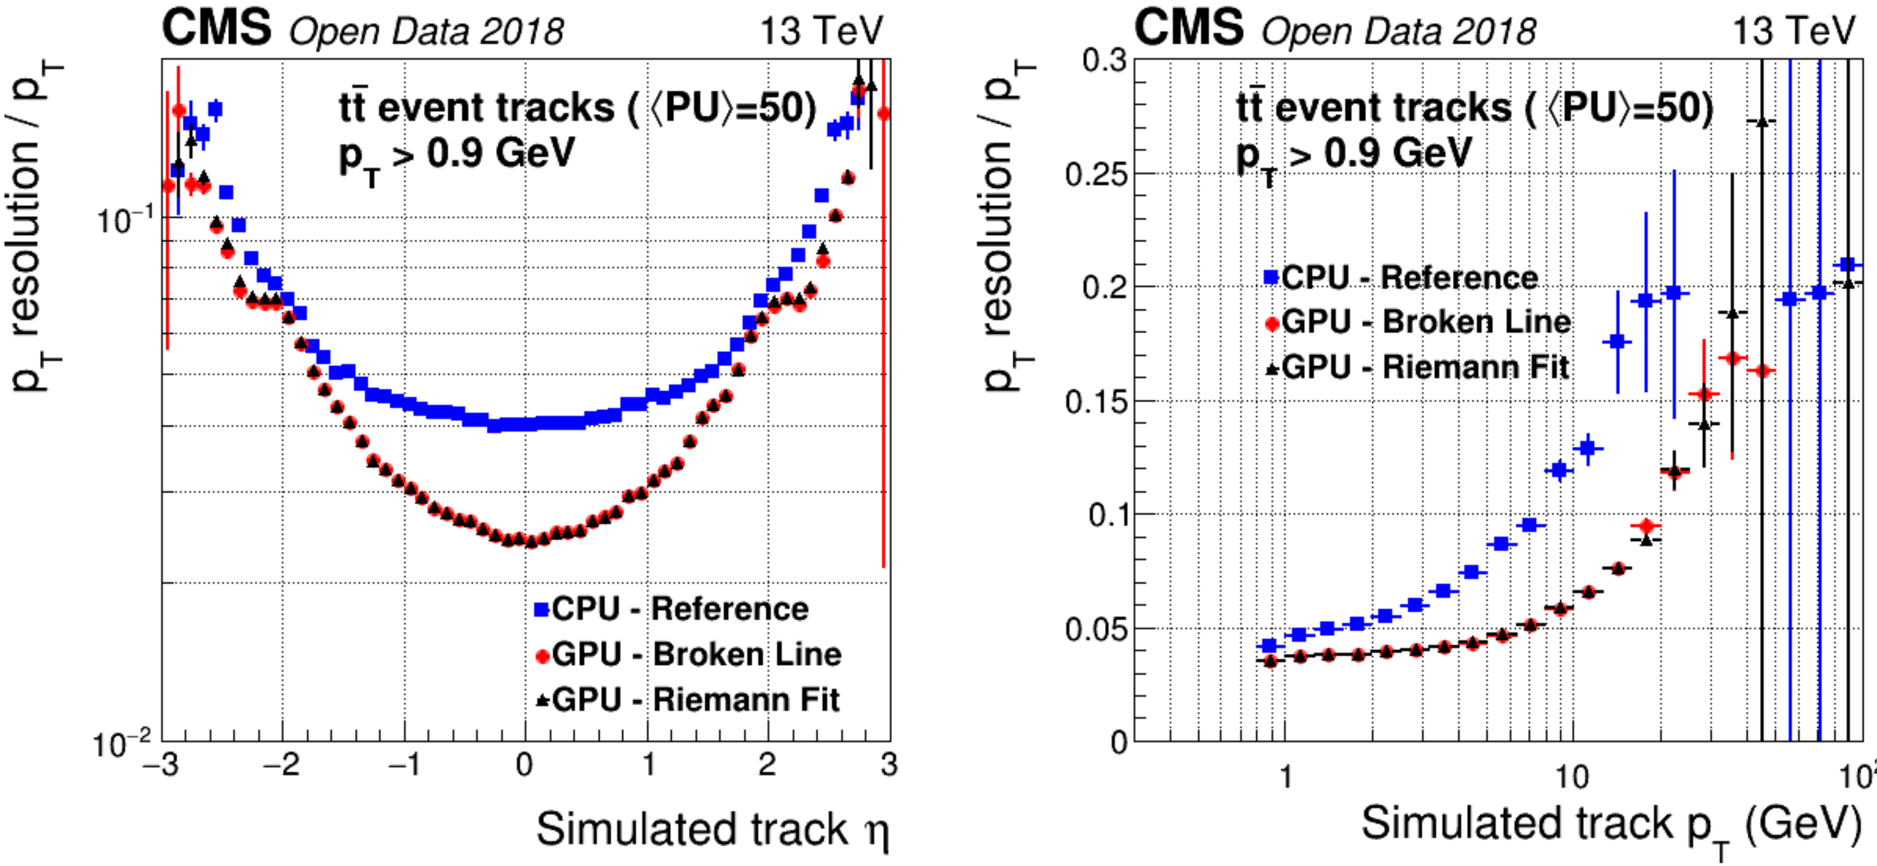

Track reconstruction resolution of p_T as a function of simulated track η and p_T

Computational Performance

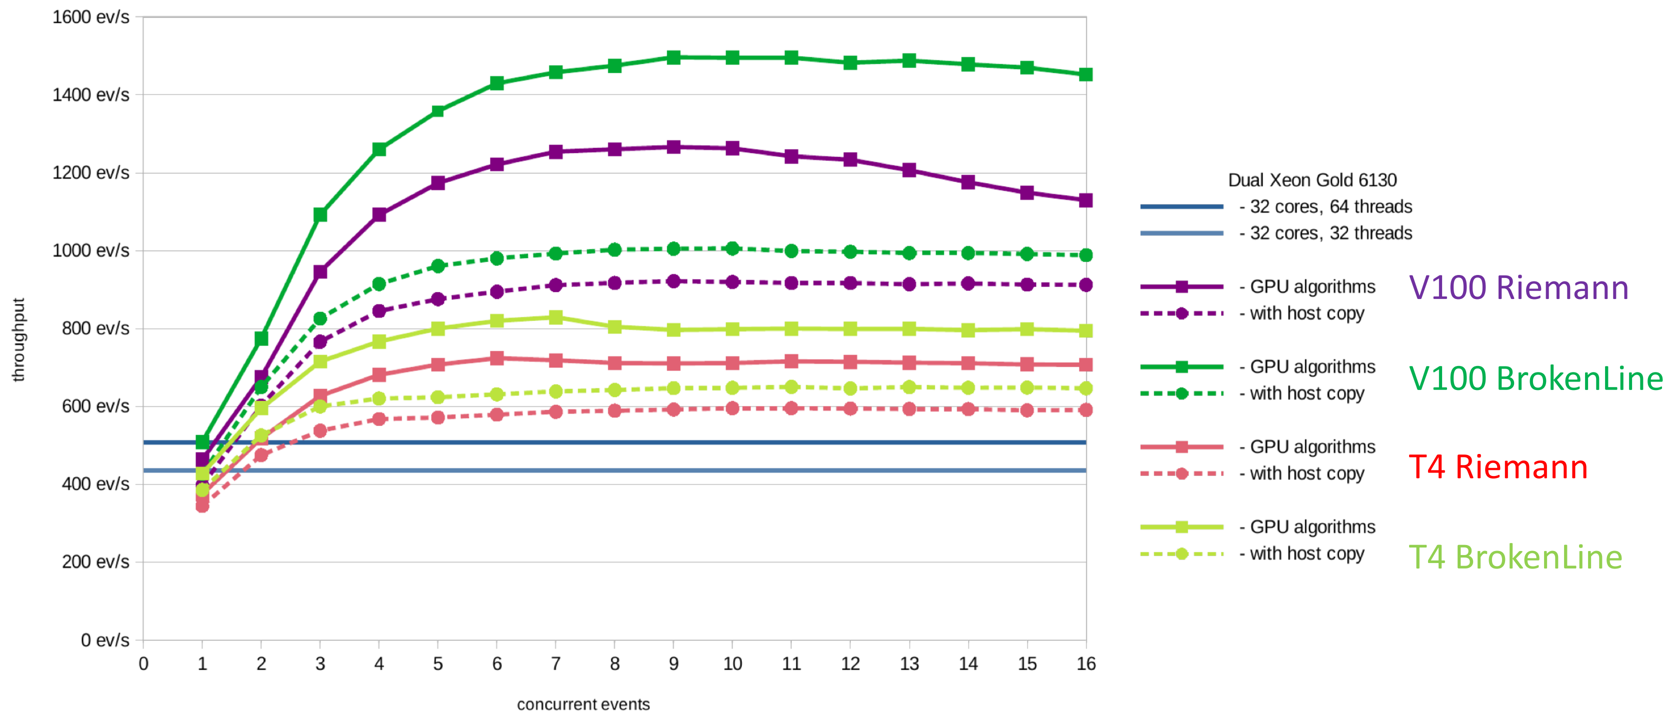

Pixel reconstruction consumers can either work directly on the GPU or ask for a copy of the tracks and vertices on the host

Computational Performance

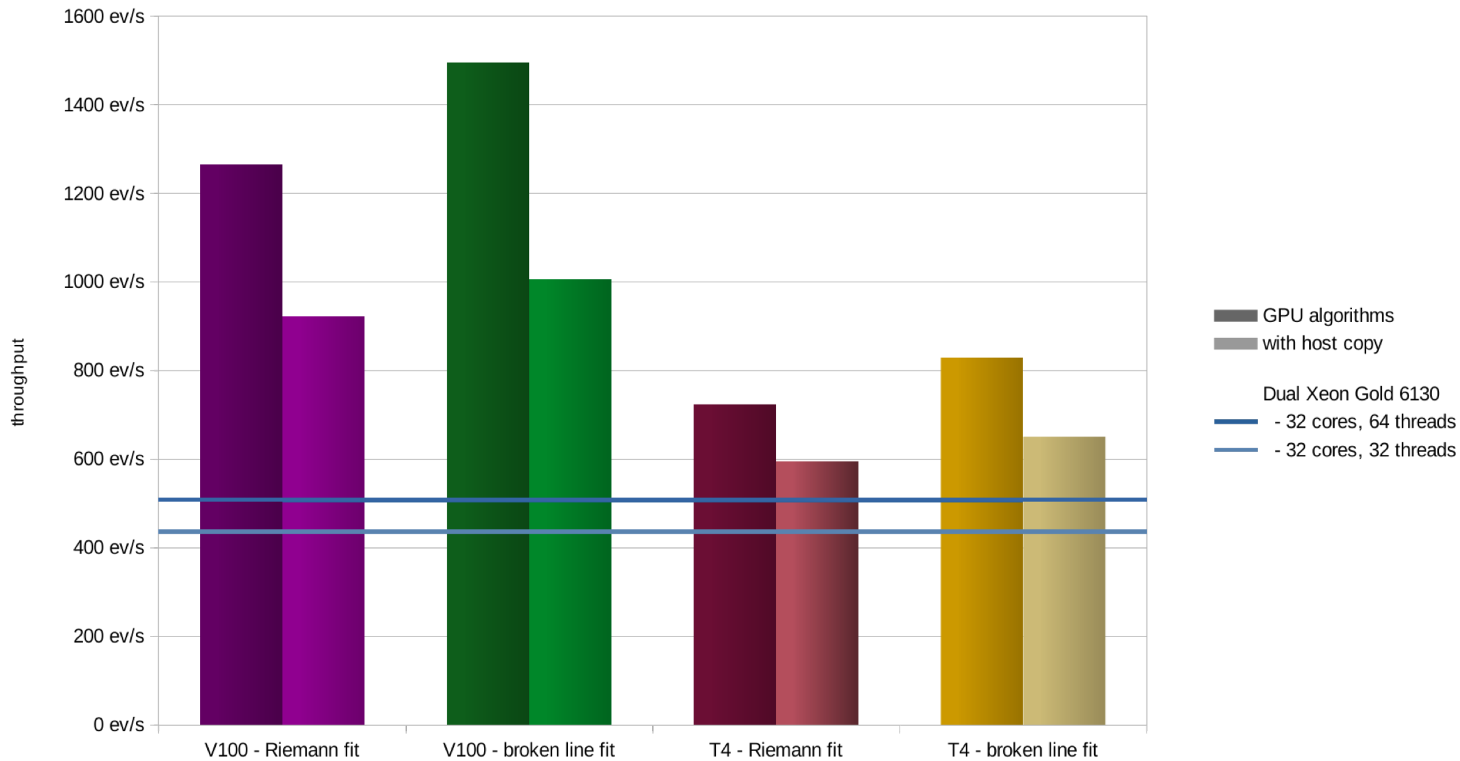

Pixel reconstruction consumers can either work directly on the GPU or ask for a copy of the tracks and vertices on the host

Back on the envelop calculation

- If one node costs 7x 🍌
- The HLT farm would cost 7000 x 🍌
- HLT farm not running Pixel Tracking (on 10% of the events) would cost 700 x🍌 less
- A NVIDIA T4 costs 2x🍌
- To run Patatrack Pixel Reconstruction on 100kHz we need 154 T4, that would cost 308x 🍌
- We can eat the other 392x 🍌 and enjoy a way better Pixel Reconstruction

Conclusion

- A GPU-based full reconstruction of the Pixel detector from RAW data decoding to Pixel Tracks and Vertices determination has been implemented
- This reconstruction is fully integrated in the CMS Software
  - Conversion to the legacy data formats and the standard validation can be run on demand
- Can achieve better physics performance, faster computational performance at a lower cost with respect to the baseline solution
- The focus during LS2 will be to maximize code sharing to have the very same workflow running on GPUs and CPUs
  - Already achieved for many critical algorithms

Backup

Physics performance – B hadrons

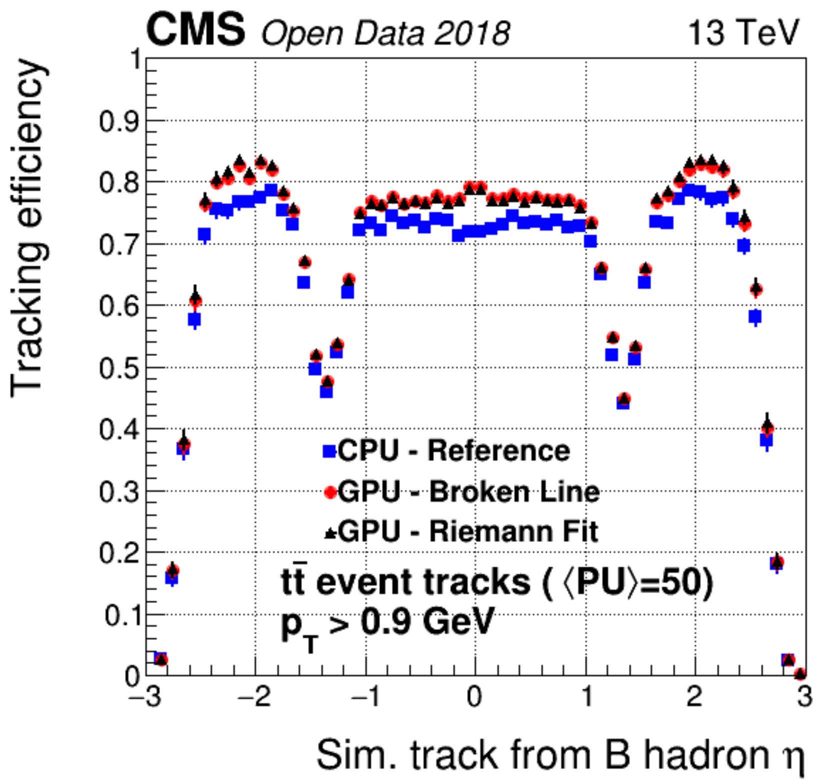

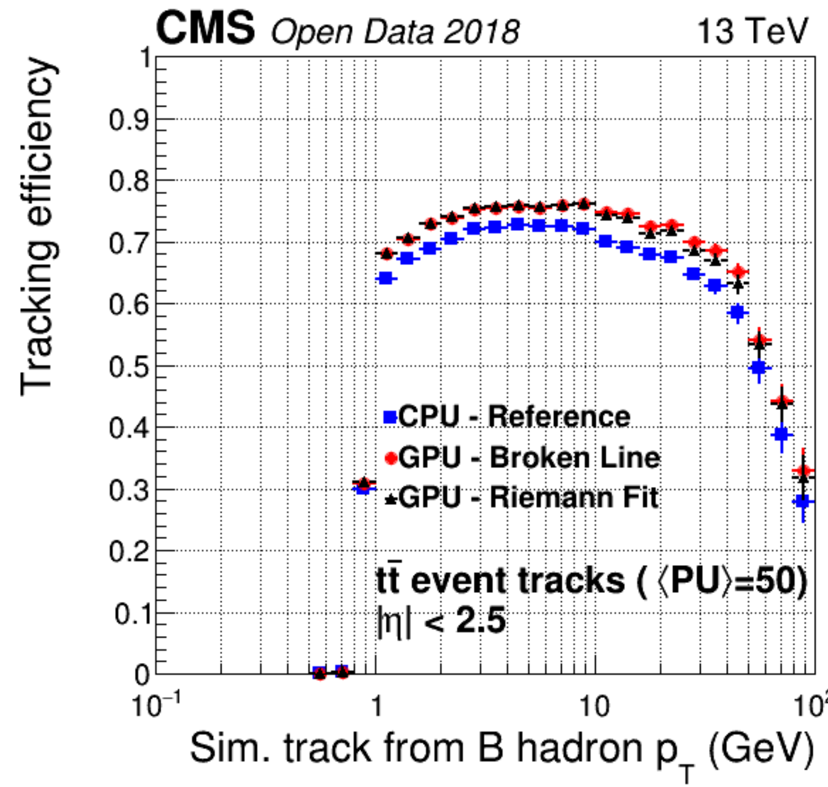

Track reconstruction efficiency as a function of simulated track from B hadron η, p_T

Triplet in the hole

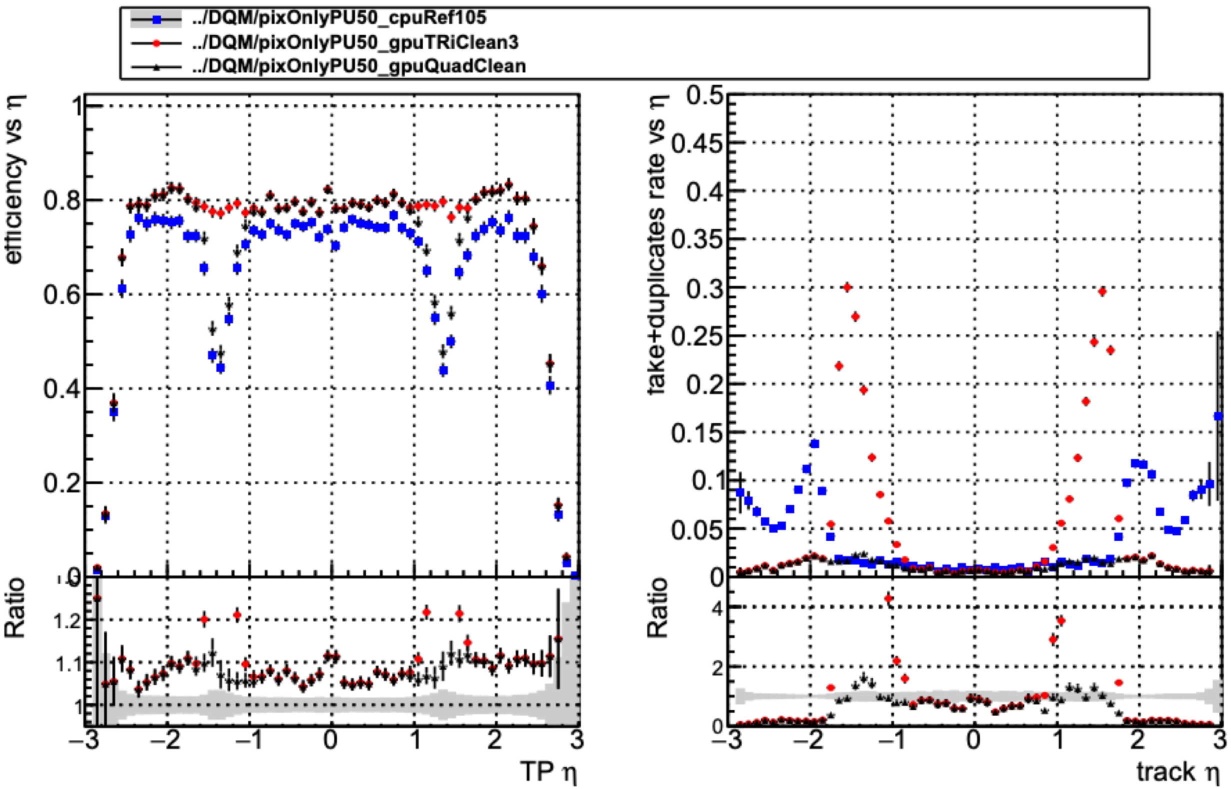

Tracks associated to PV

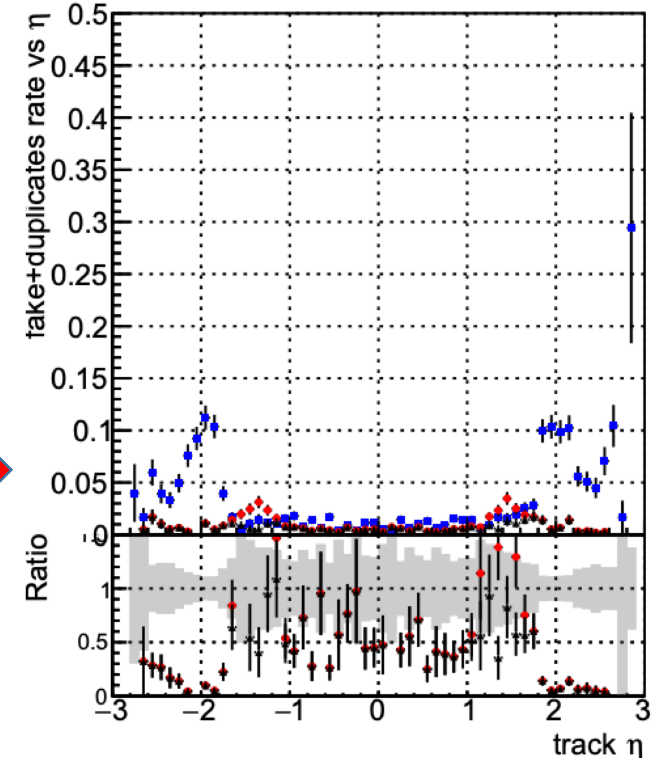

Future directions

The whole CA is centered around doublets + links

- They can be "augmented" with any (meta)-information
- The same set can be easily reused in "iterations"
  - Just mask the unwanted doublets (even the links in principle)