Task 3.3 report “Reduction of the RAW data size for HLT”

1st Next Generation Trigger technical workshop

Silvio Donato (INFN and University of Pisa) on behalf of the NGT Task 3.3 team

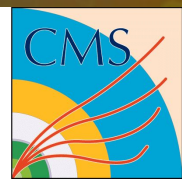

The team

- **Silvio Donato** (task leader, assistant prof. University of Pisa and INFN)
- **Simone Rossi Tisbeni** (NGT-CERN doctoral student, University of Bologna and INFN)

- **Synergy with "PRE: Partially Reconstructed Event" project PRIN2022BLJJLY:**
  - Vinaya Krishna (postdoc at University of Pisa)
  - Saswati Nandan (postdoc at INFN Pisa)
  - Silvio Donato (principal investigator)

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
INFN PISA

Outline

- **The team**
- **The goal**
- **The strategy**
- **The RAW event size in Run-3**
  - The strip RAW' compression in Run-3 pp-collisions
    - ongoing improvements
- **The event size in recent Phase-2 simulations**
  - Comparison with Run-3 objects
    - Inner tracker
    - Outer tracker
    - HGCal
  - Comparison with HLT-TDR numbers
- **Conclusions**

The goal

- **The Phase-2 event size of CMS is expected to be around 10 MB/event.**
  - **No way** to store **all events** accepted from the hardware trigger (750 kHz) in raw data format
    - 10 MB/event x 750 kHz = 7.5 TB/s, ie. about 10 EB/year!
- **The R^3 project aims to run an offline-like reconstruction on all events accepted by the hardware trigger (up to 750 kHz).**
- **The goal of Task 3.3 is to increase the maximum trigger available by reducing the size of events saved by CMS**

The strategy

- **NGT 3.3 will reduce the event size in three different ways and compression factors.**
  - **Lossy** compression, replacing raw data with **high-level** physics objects (eg. muons, electrons, **jets**, tracks?),
    - evolution of our current **scouting** data format,
    - very **strong compression** (around x100), **necessary to store all 750 kHz**,
      - limited possibility to re-reconstruct objects with newer algorithm or calibration.
  - **Lossy** compression, replacing raw data with **low-level** physics objects (eg. replacing tracker or calorimeter **raw** data with **reco's hit positions and energies**)
    - extension of **RAW'** currently used in HIon collisions in CMS,
    - **limited compression** (below x10),
      - will **allow** to **re-reconstruct** high-level objects with newer algorithm or calibration
  - **Lossless** compression:
    - testing new compression algorithms;
    - using physics objects (eg. tracks) to improve the raw data compression (see backup).
- **R^3 is key in all methods to achieve good compression with high quality physics objects**

The RAW event size in Run-3

Credits: V. Krishna

- **Usage of a specific Front-End Detector (FED) selector to measure the event size of each detector.**
- **Data from Muon dataset pp collisions pileup~60** (run=382216, fill=9804)
- **Total compressed event size: 1204 kB**
  - **Strips detector → 55%**
  - **Pixel detector → 35%**
  - ECAL preshower → 4%
  - All other detector < 10%

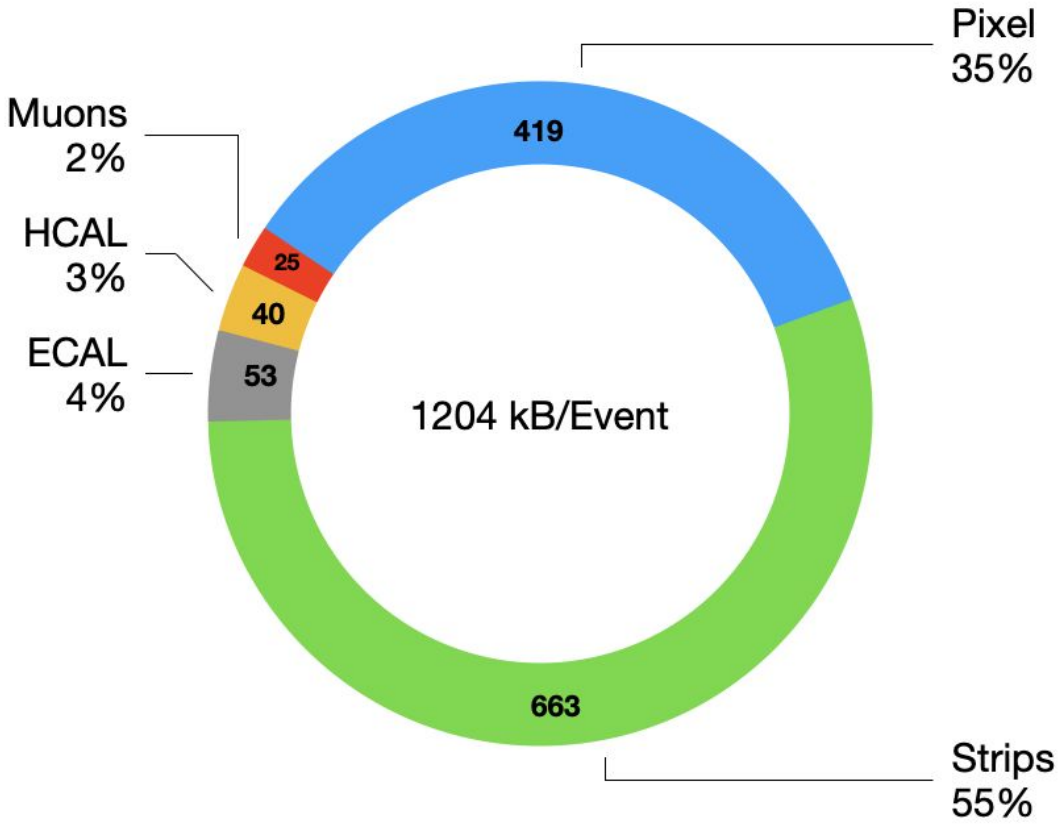

The RAW' idea: strip data compression

Credits: V. Krishna

- **The strip detector takes the largest fraction of the event size.**
- **The RAW' rationale is to replace**
  - the ADC amplitude numbers (8 bits x width)
- **with**
  - the average charge (8 bits) +
  - the length (8 bits) +
  - barycenter (16 bits)
- **RAW' compression used in PbPb collisions**
  - online since 2023
- **Very important for HIon collisions as DAQ bandwidth is the bottleneck of the trigger rate**

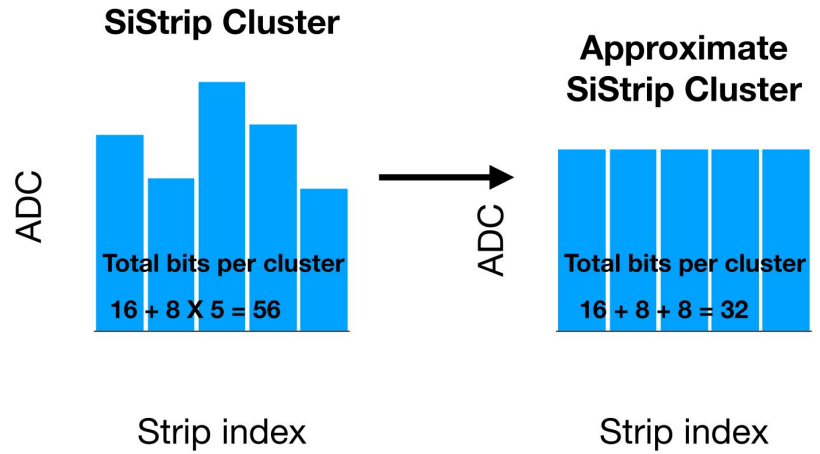

The RAW' in Run-3 pp collisions

Credits: V. Krishna

- **Here the measurement of RAW' in pp collisions:**
  - Strip detector from 663 kB (55%) to 425 kB (44%) [**-35.8%**]
  - Total event size from 1204 kB to 965 kB/event [**-20%**]
- **Strip detector is still the detector with the largest event size**
  - Focus in the **further compression of strip** in RAW'

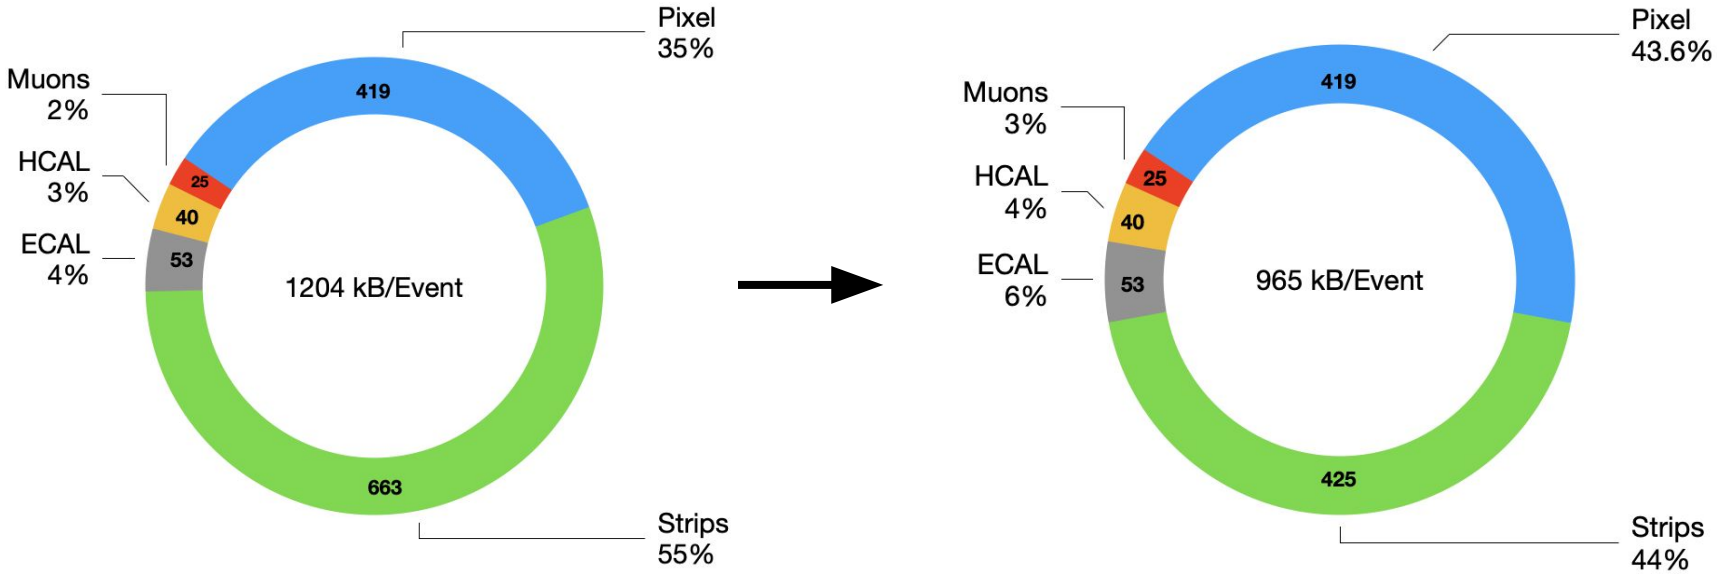

Credits: S. Nandan

Compressing further RAW' (Run-3)

- **Optimization of the number of bits used in**
  - barycenter
  - charge
  - width
- **Removal of strip clusters with low charge**
  - typically out-of-time pileup or noise
- **Example: 12-bit for barycenter + charge cut**
  - **~100% track efficiency**
    - **Track selection:** normalized $\chi^2 < 2$, $|p^T_{error} / p^T| < 0.10$, # of hits > 11
  - **- 17% in strip event size**

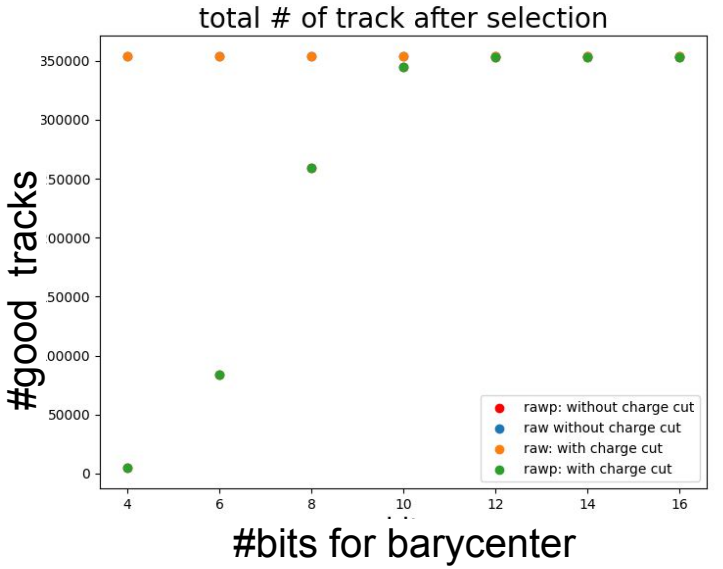

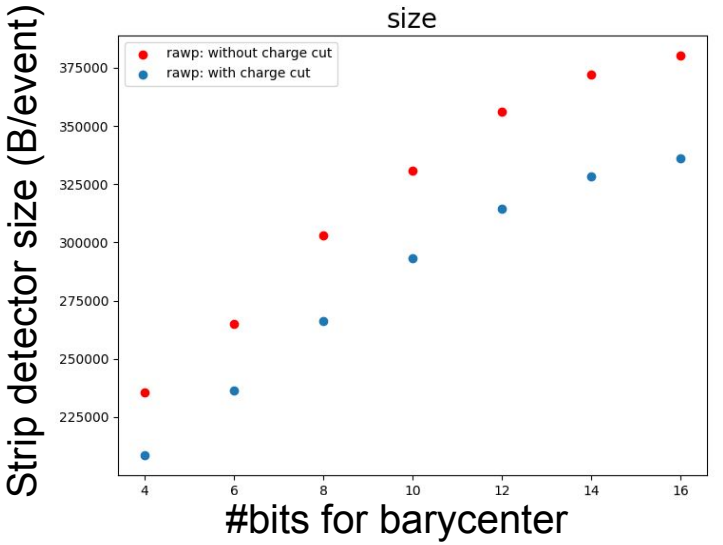

Credits: V. Krishna

Compressing further RAW' (Run-3)

- **Only ~10% on strip clusters are associated with any reconstructed track!**
  - 90% of clusters have no associated tracks
- **The source of 90% of clusters is under study, some origins are:**
  - clusters not passing quality criteria
  - noise and out-of-time pileup
  - track inefficiency → no associated clusters
    - nuclear interaction
    - esp. for very soft tracks (loopers)
  - *caveat: clusters excluded in tracking appears as "not associated" by design (eg. low charge clusters)*
- **Investigating possible ways to remove "useless" clusters to reduce strip event size.**

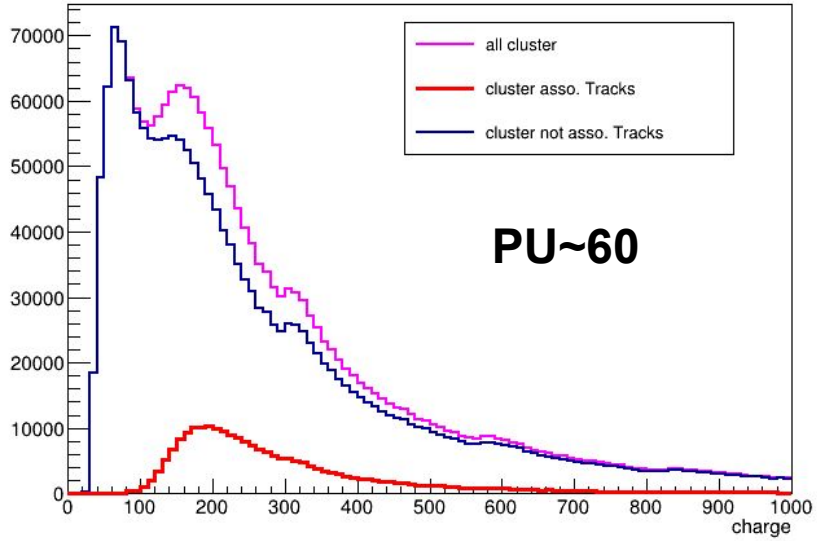

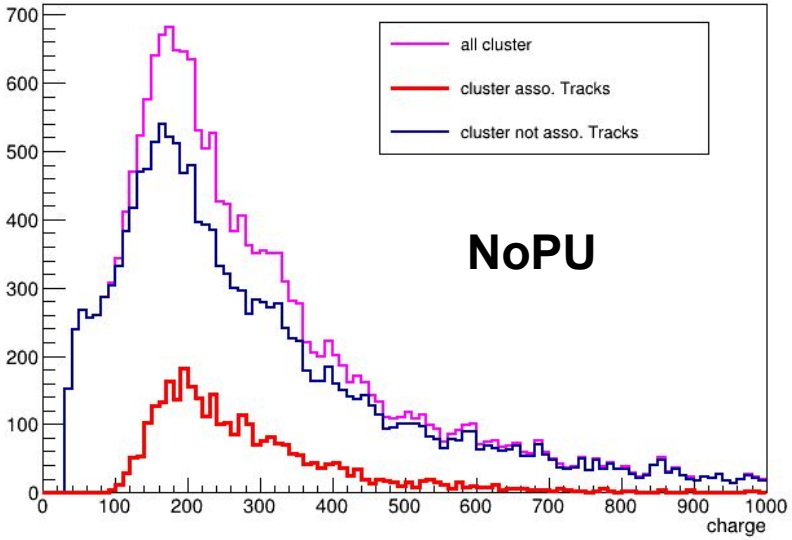

Slide source

The Phase-2 upgrade

- **New detectors**
  - **H**igh-**G**ranularity **C**alorimeter
  - **M**IP **T**iming **D**etector
  - More sophisticated L1 trigger
- **Larger not only because of the larger pileup but also for new detector with higher complexity**

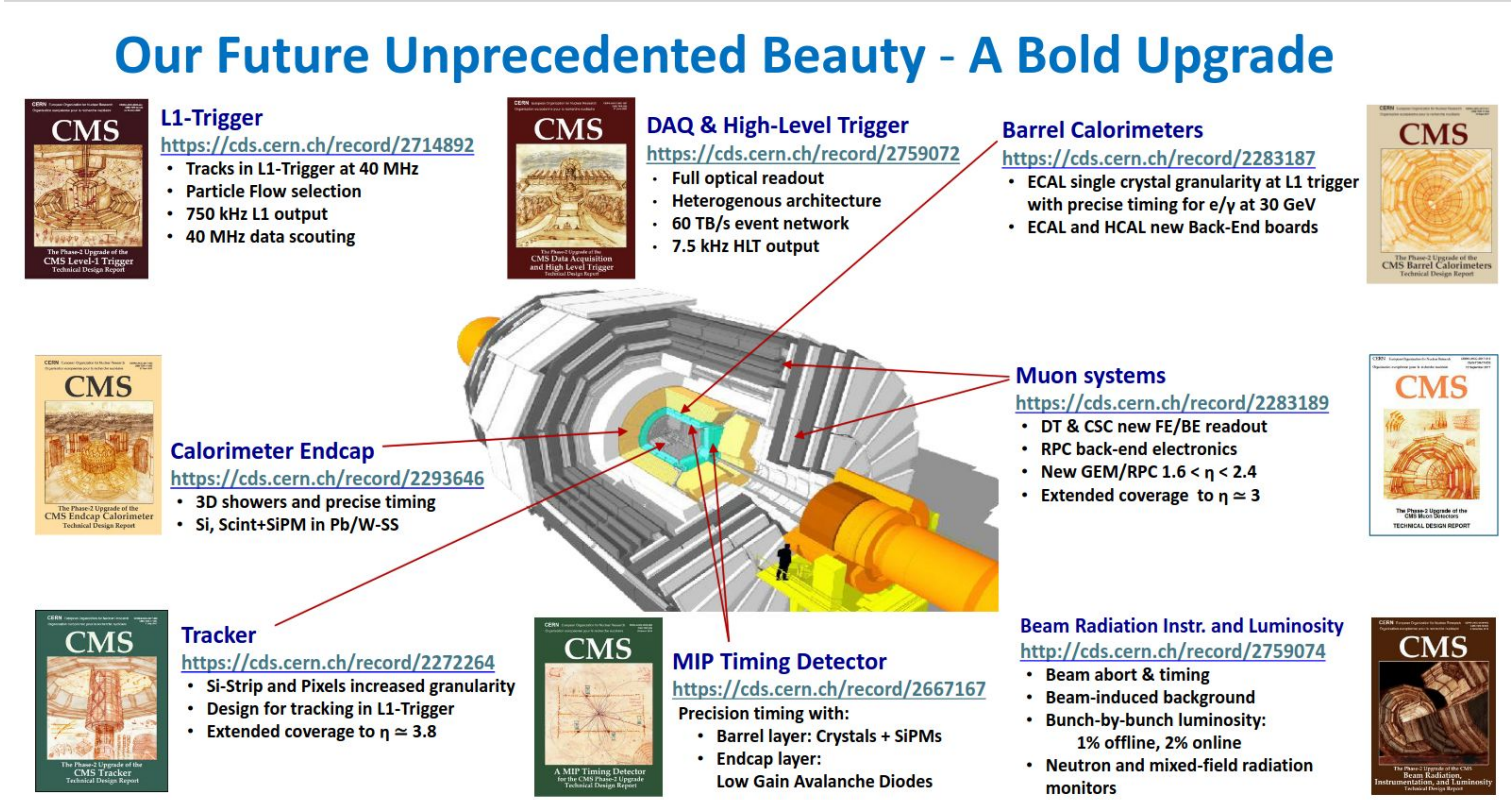

The RAW event size in Phase-2

Credits: S. Rossi Tisbeni

- **The RAW event content is not yet available for Phase-2**
  - “raw” in current phase-2 simulations with pileup 200 is only 300 kB/event instead of ~10MB → very few detector are included
- **The simulated digitalization of the detector (simDigi) have been used to estimate the raw detector size.**
- **Run-3 simulations have been used as a test of the goodness of the estimate**
- **The low-level reconstructed objects (detector local reconstruction) have been used as an estimate of a possible compressing in the RAW’ fashion.**
- **In the next slides:**
  - Run-3: TTbar sample with pileup 60 (2024 conditions)
  - Phase-2: TTbar sample with pileup 200 (2027 conditions)

The pixel detector (Run-3)

Credits: S. Rossi Tisbeni

- **RAW detector size [per event compressed]: 286 kB**
- **All ‘simSiPixelDigis’ objects: 736 kB**
  - excluding PixelDigiSimLinkedmDetSetVector: **279 kB**
- **All ‘siPixelDigis’ objects: 281 kB**
- **Low-level reconstructed objects:**
  - siPixelClusters: 286 kB
    - siPixelClusters contains all pixels ADC counts
  - siPixelRecHits: 48 kB
    - **potential large compression saving recHits instead of raw (-83%)**

The inner tracker (Phase-2)

Credits: S. Rossi Tisbeni

- **No inner tracker RAW in Phase-2 yet**
- **simSiPixelDigis: 6920 kB**
  - excluding PixelDigiSimLinkedmDetSetVector: 1357 kB
- **No siPixelDigis yet**
- **Low-level reconstructed objects:**
  - siPixelRecHits: 142 kB
    - **potential large compression saving recHits instead of raw (-89%)**

The strip detector (Run-3)

Credits: S. Rossi Tisbeni

- **RAW detector size [per event compressed]: 731 kB**
- **All ‘simSiStripDigis’ objects: 4582 kB**
  - excluding StripDigiSimLinkedmDetSetVector: **849 kB**
- **All siStripDigis objects: 850 kB**
- **Low-level reconstructed objects:**
  - siStripClusters: 727 kB
    - siStripClusters contains all strips ADC counts
  - all siStripMatchedRecHits objects: 366 kB
    - potential compression saving recHits instead of raw (-50%) similar to our current RAW’ compression

The outer tracker (Phase-2)

Credits: S. Rossi Tisbeni

- **No outer tracker RAW in Phase-2 yet**
- **This collection is currently used in the inner tracker simulation:**
  - mix, Tracker: 699 kB (pileup mixing module)
- **No siStripDigis yet**
- **Low-level reconstructed objects:**
  - siPhase2Clusters: 587 kB
- **Note: in Phase-2 the outer tracker will be fully digitalized:**
  - only 0 or 1 per strip, no charge deposit, barycenter is semi-integer number
  - this explains the reason of
    - siPhase2Clusters size is very close to siPhase2Clusters (size ratio 84%)
    - **siPhase2Clusters size at PU200 (587 kB) is smaller than Run-3 siStripClusters at PU60 (727 kB)!**

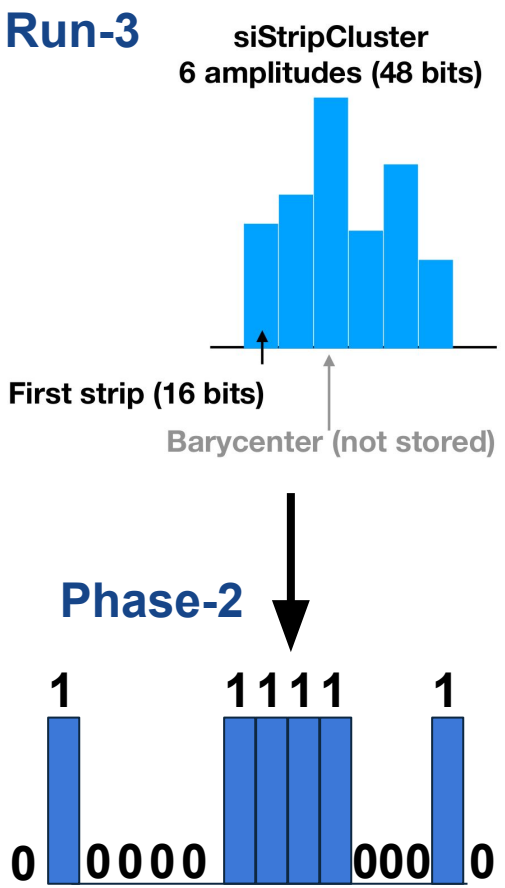

The ECAL (Run-3 and Phase-2)

Credits: S. Rossi Tisbeni

- **ECAL: link**
  - RAW: 37 kB (only Run-3)
  - ‘simEcalDigis’: 36 kB (Run-3) vs 65 kB (Phase-2)
  - ‘ecalDigis’ (no trigger primitives): 49 kB (Run-3) vs 70 kB (Phase-2)
- **Larger ECAL size in Phase-2 because of timing?**
- **Low-level reconstructed objects:**
  - ‘ecalMultiFitUncalibRecHit’ 110 kB (Run-3) vs 275 kB (Phase-2)
  - ‘ecalRecHit’ 38 kB (Run-3) vs 81 kB (Phase-2)
    - ‘ecalDetailedTimeRecHit’ 89 kB (only Phase-2)
    - ‘reducedEcalRecHits’ 43.6 kB (Run-3) vs 30.9 kB (Phase-2)
      - ‘reducedEcalRecHitsEB’ 17.5 kB (Run-3) vs 30.7 kB (Phase-2)
  - ‘particleFlowClusterECAL’ 16 kB (Run-3) vs 22 kB (Phase-2)

The HCAL (Run-3 and Phase-2)

Credits: S. Rossi Tisbeni

- **HCAL: link**
  - RAW: 142 kB (only Run-3, extrapolated(*))
  - 'simHcalDigis': 134 kB (Run-3) vs 80 kB (Phase-2)
  - 'hcalDigis' (no trigger primitives): 163 kB (Run-3) vs 96 kB (Phase-2)
- **Smaller HCAL size in Phase-2 because of HGCal (no endcap)**
- **Low-level reconstructed objects:**
  - 'reducedHcalRecHits' 12.0 kB (Run-3) vs 10.1 kB (Phase-2)
  - 'particleFlowClusterHCAL' 187 kB (Run-3) vs 3.3 kB (Phase-2)

(*) FED selector crashes when using HCAL FED. HCAL RAW size computed as total RAW minus all detectors excepted HCAL

The HGCal (Phase-2)

Credits: S. Rossi Tisbeni

- **No HGCal RAW in Phase-2 yet**
- **All ‘simHGCalUnsuppressedDigis’ objects: 2685 kB**
- **‘hgcalDigis’: 2685 kB**
- **Low-level reconstructed objects:**
  - all ‘HGCalUncalibRecHit’ —> 1114 kB
  - all ‘HGCalRecHit’ —> 1626 kB
  - ‘particleFlowClusterHGCalFromSimCl’ objects: 3162 kB
  - ‘particleFlowClusterHGCal’ objects: 79 kB

Event size summary (Phase-2)

Credits: S. Rossi Tisbeni

- **HGCal is the largest detector (2.6 MB/event)**
  - saving uncalibrated recHits 1.1 MB/event (-59%)
- **The inner tracker is the second largest detector (1.4 MB/event)**
  - saving pixel recHits could reduce the size down to 142 kB/event (-89%)
    - a similar reduction is found using Run-3 data
- **The size of the outer tracker with PU 200 is smaller (587 kB) than the strip detector with PU 60 because of the digitalization of the strip detector**
  - almost no compression using a RAW' approach
    - might be still possible to apply a cleaning of "useless" strip clusters to reduce the size
- **Caveat: the definition of new data formats (eg. 'HGCalUncalibRecHit') are still in progress and their size might change in the near future.**

Event size summary from HLT TDR

- **Uncompressed detector size from Phase-2 HLT TDR (2021)**
  - Some changes wrt recent Phase-2 simulation
    - overall picture is matching
    - part of differences because of compression
      - eg. ECAL compression 300 kB → 70 kB
      - still some difference with latest number
        - eg. ECAL 300 kB vs 600 kB
    - some important detectors to be computed
      - MTD, CSC, DT, L1T+HGCal TPG

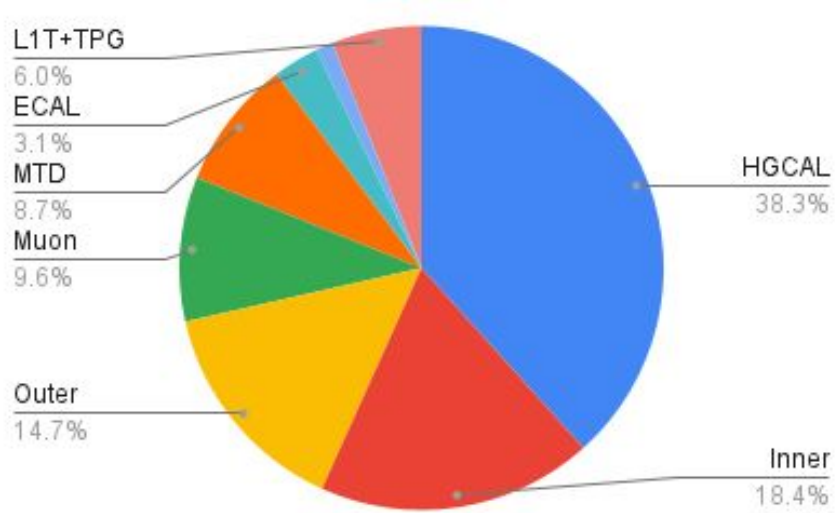

Subsystem	Size (MB/event)	
Inner Tracker	**1.44**	**1.44**
Outer Tracker - PS	**0.72**	**1.15**
Outer Tracker - 2S	**0.43**	
MIP Timing Det. - BTL	0.24	0.68
MIP Timing Det. - ETL	0.44	
ECAL Barrel	0.60	0.60
HCAL Barrel	0.24	0.33
HCAL HO	0.03	
HCAL HF	0.06	
HGCAL	**3.00**	**3.00**
Muon DT	0.15	0.76
Muon CSC	0.47	
Muon GEM - GE1/1	0.00	
Muon GEM - GE2/1	0.00	
Muon GEM - ME0	0.12	
Muon RPC	0.01	
Total	**8.42**	**7.96**

Conclusions

- **The replacement of raw data with low-level reconstructed objects is already used in PbPb collisions for the strip detector (RAW').**
  - The usage of strip RAW' compression to Run-3 pp collisions data would decrease the total event size of -30%
    - studies are ongoing to enhance the compression factor
- **The Phase-2 event size has been studied using recent simulations.**
- **A possible compression of biggest detectors using low-level variables gives**
  - HGCal → uncalibrated recHits is 59% of smaller than raw data
  - Inner tracker → siPixelRecHits would reduce the data size of 89%
  - Outer tracker → the event size is already small because of the digitizer (600 kB/event)
    - further reduction might come from the cleaning of fake tracker clusters

BackUp

Lossless compression using tracks

J.Phys.: Conf. Ser. 396 012043

- ALICE Time Projection Chamber (TPC) data is compressed online using track information
- The compression is lossless: clusters variable are replaced with residual variable wrt ass. tracks
- Replacing flat distribution with “peaked” distribution allows for a stronger Huffman compression
- A similar idea might be applied to CMS outer tracker, especially for ‘looper’ tracks
- Reference: J.Phys.: Conf. Ser. 396 012043

TPC clusters and track association

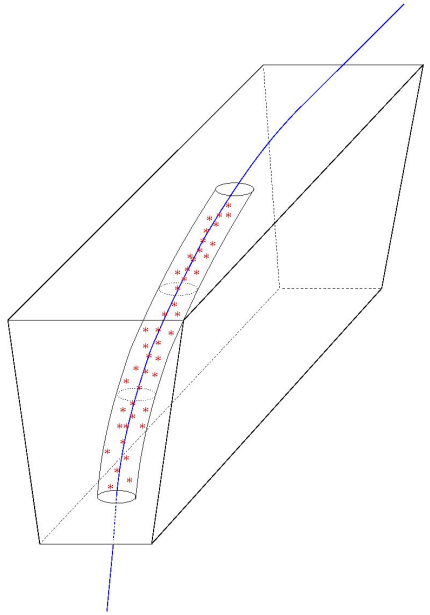

Time distribution

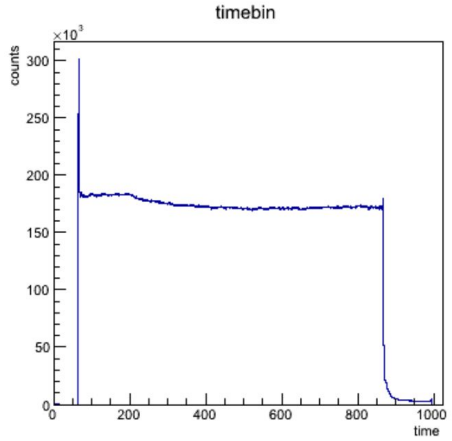

Time **residual** cluster distribution wrt tracks

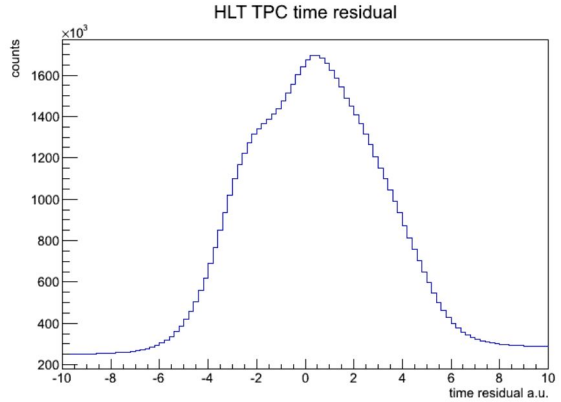

Cluster charge Fix

Credits: V. Krishna

The PR for approximated SiStripCluster made by HION group.

We redefined the average charge, which returns to correct round off value.

Previously, Avg. charge = $\dfrac{\texttt{cluster.charge()}}{\texttt{cluster.size()}}$ New, Avg. charge = $\dfrac{\texttt{cluster.charge()} + \texttt{cluster.size()}/2}{\texttt{cluster.size()}}$

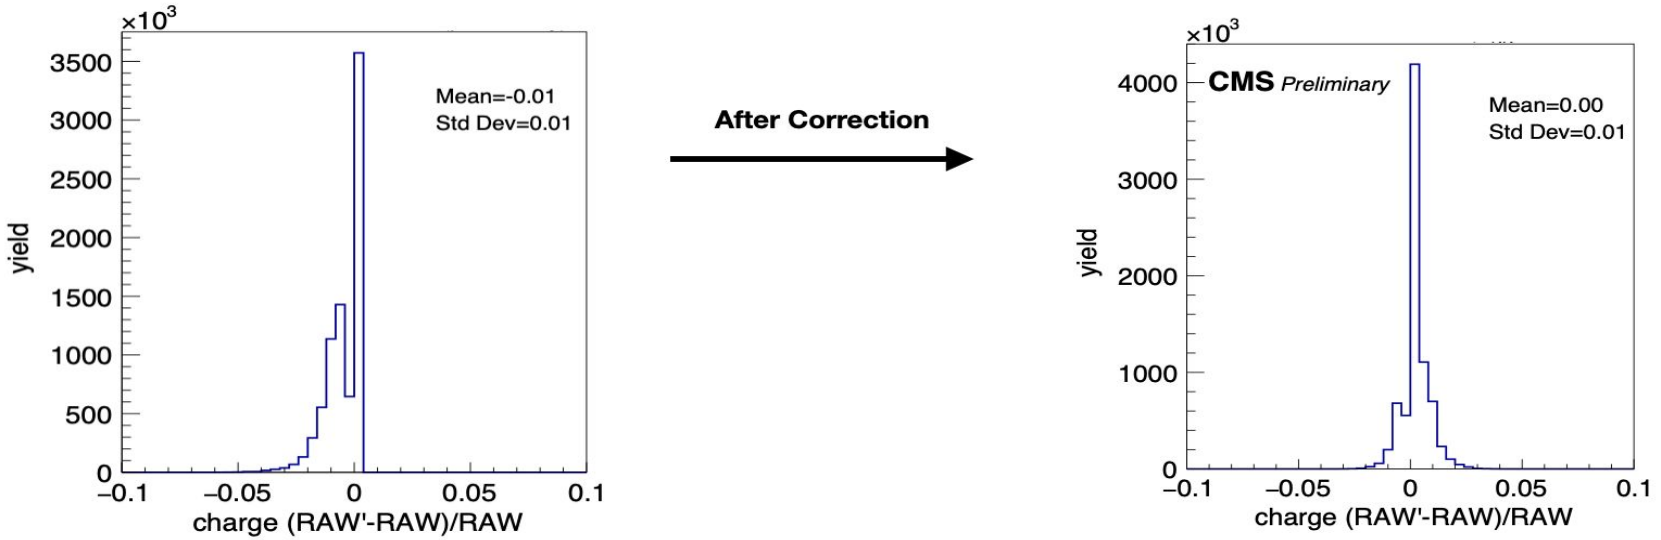

Scouting - Event Size (2023)

Objects	Size (kB/ev)
PF candidates	4.1
Electrons/Photons	1.4
Pixel tracks	1.0
Primary vertices	0.3
Muons	0.1
Jets	0.1
L1 trigger data	0.7
HLT trigger bits	0.0
TOTAL	**7.7**

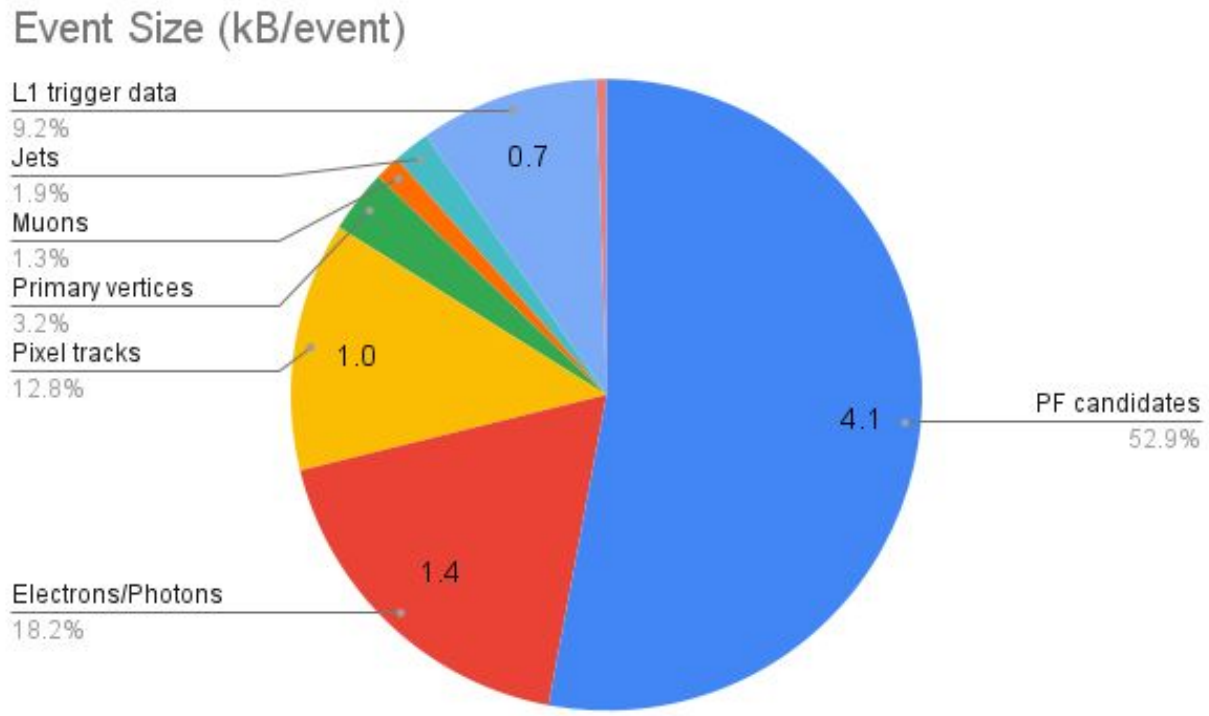

pileup=60, lumi= fill=9049, run=370355, ls=73

- ~135 times smaller than the full RAW event content (1.06 MB/event)
- **20 kHz** @ 7.7 kB/event are “equivalent” to **~150 Hz** @ 1.06 MB/event (full RAW)

The milestones

- **The first year milestones:**
  - Production of a report illustrating the impact in terms of RAW data size coming from each detector, and suggesting the major area of interest/intervention.
  - **Report on the impact of RAW data compression and of their replacement with low-level reconstructed quantities (RAW') (contractual milestone)**

The milestones

- **The contractual milestones**
  - **1 year**: **Report on the impact of RAW data compression and of their replacement with low-level reconstructed quantities (RAW')**
  - **3 year**: A comparative analysis of the different data reduction approaches and their impact on the accuracy of the physics reconstruction is published.
- **The detailed milestones**
  - **12 m** (was 6 m): Production of a report illustrating the impact in terms of RAW data size coming from each detector, and suggesting the major area of interest/intervention.
  - **18 m**: Assessment of two different approaches to data compression: lossless compression on accelerators, and replacement of part of the RAW data with low-level reconstructed quantities (RAW').
  - **30 m**: Implementation of the most promising solutions replacing part of the RAW data with low-level reconstructed quantities in the CMS reconstruction software.
  - **42 m**: Assessment of the trade-off between the data reduction and the impact on the high-level physics reconstruction, as part of the RAW data are replaced by high-level reconstruction objects.
  - **54 m**: Implementation of the most promising solutions replacing part of the RAW data with high-level reconstructed quantities in the CMS reconstruction software.
  - **60 m**: Large scale deployment and testing of the whole infrastructure in view of the readiness for the 2029 HL-LHC data-taking.

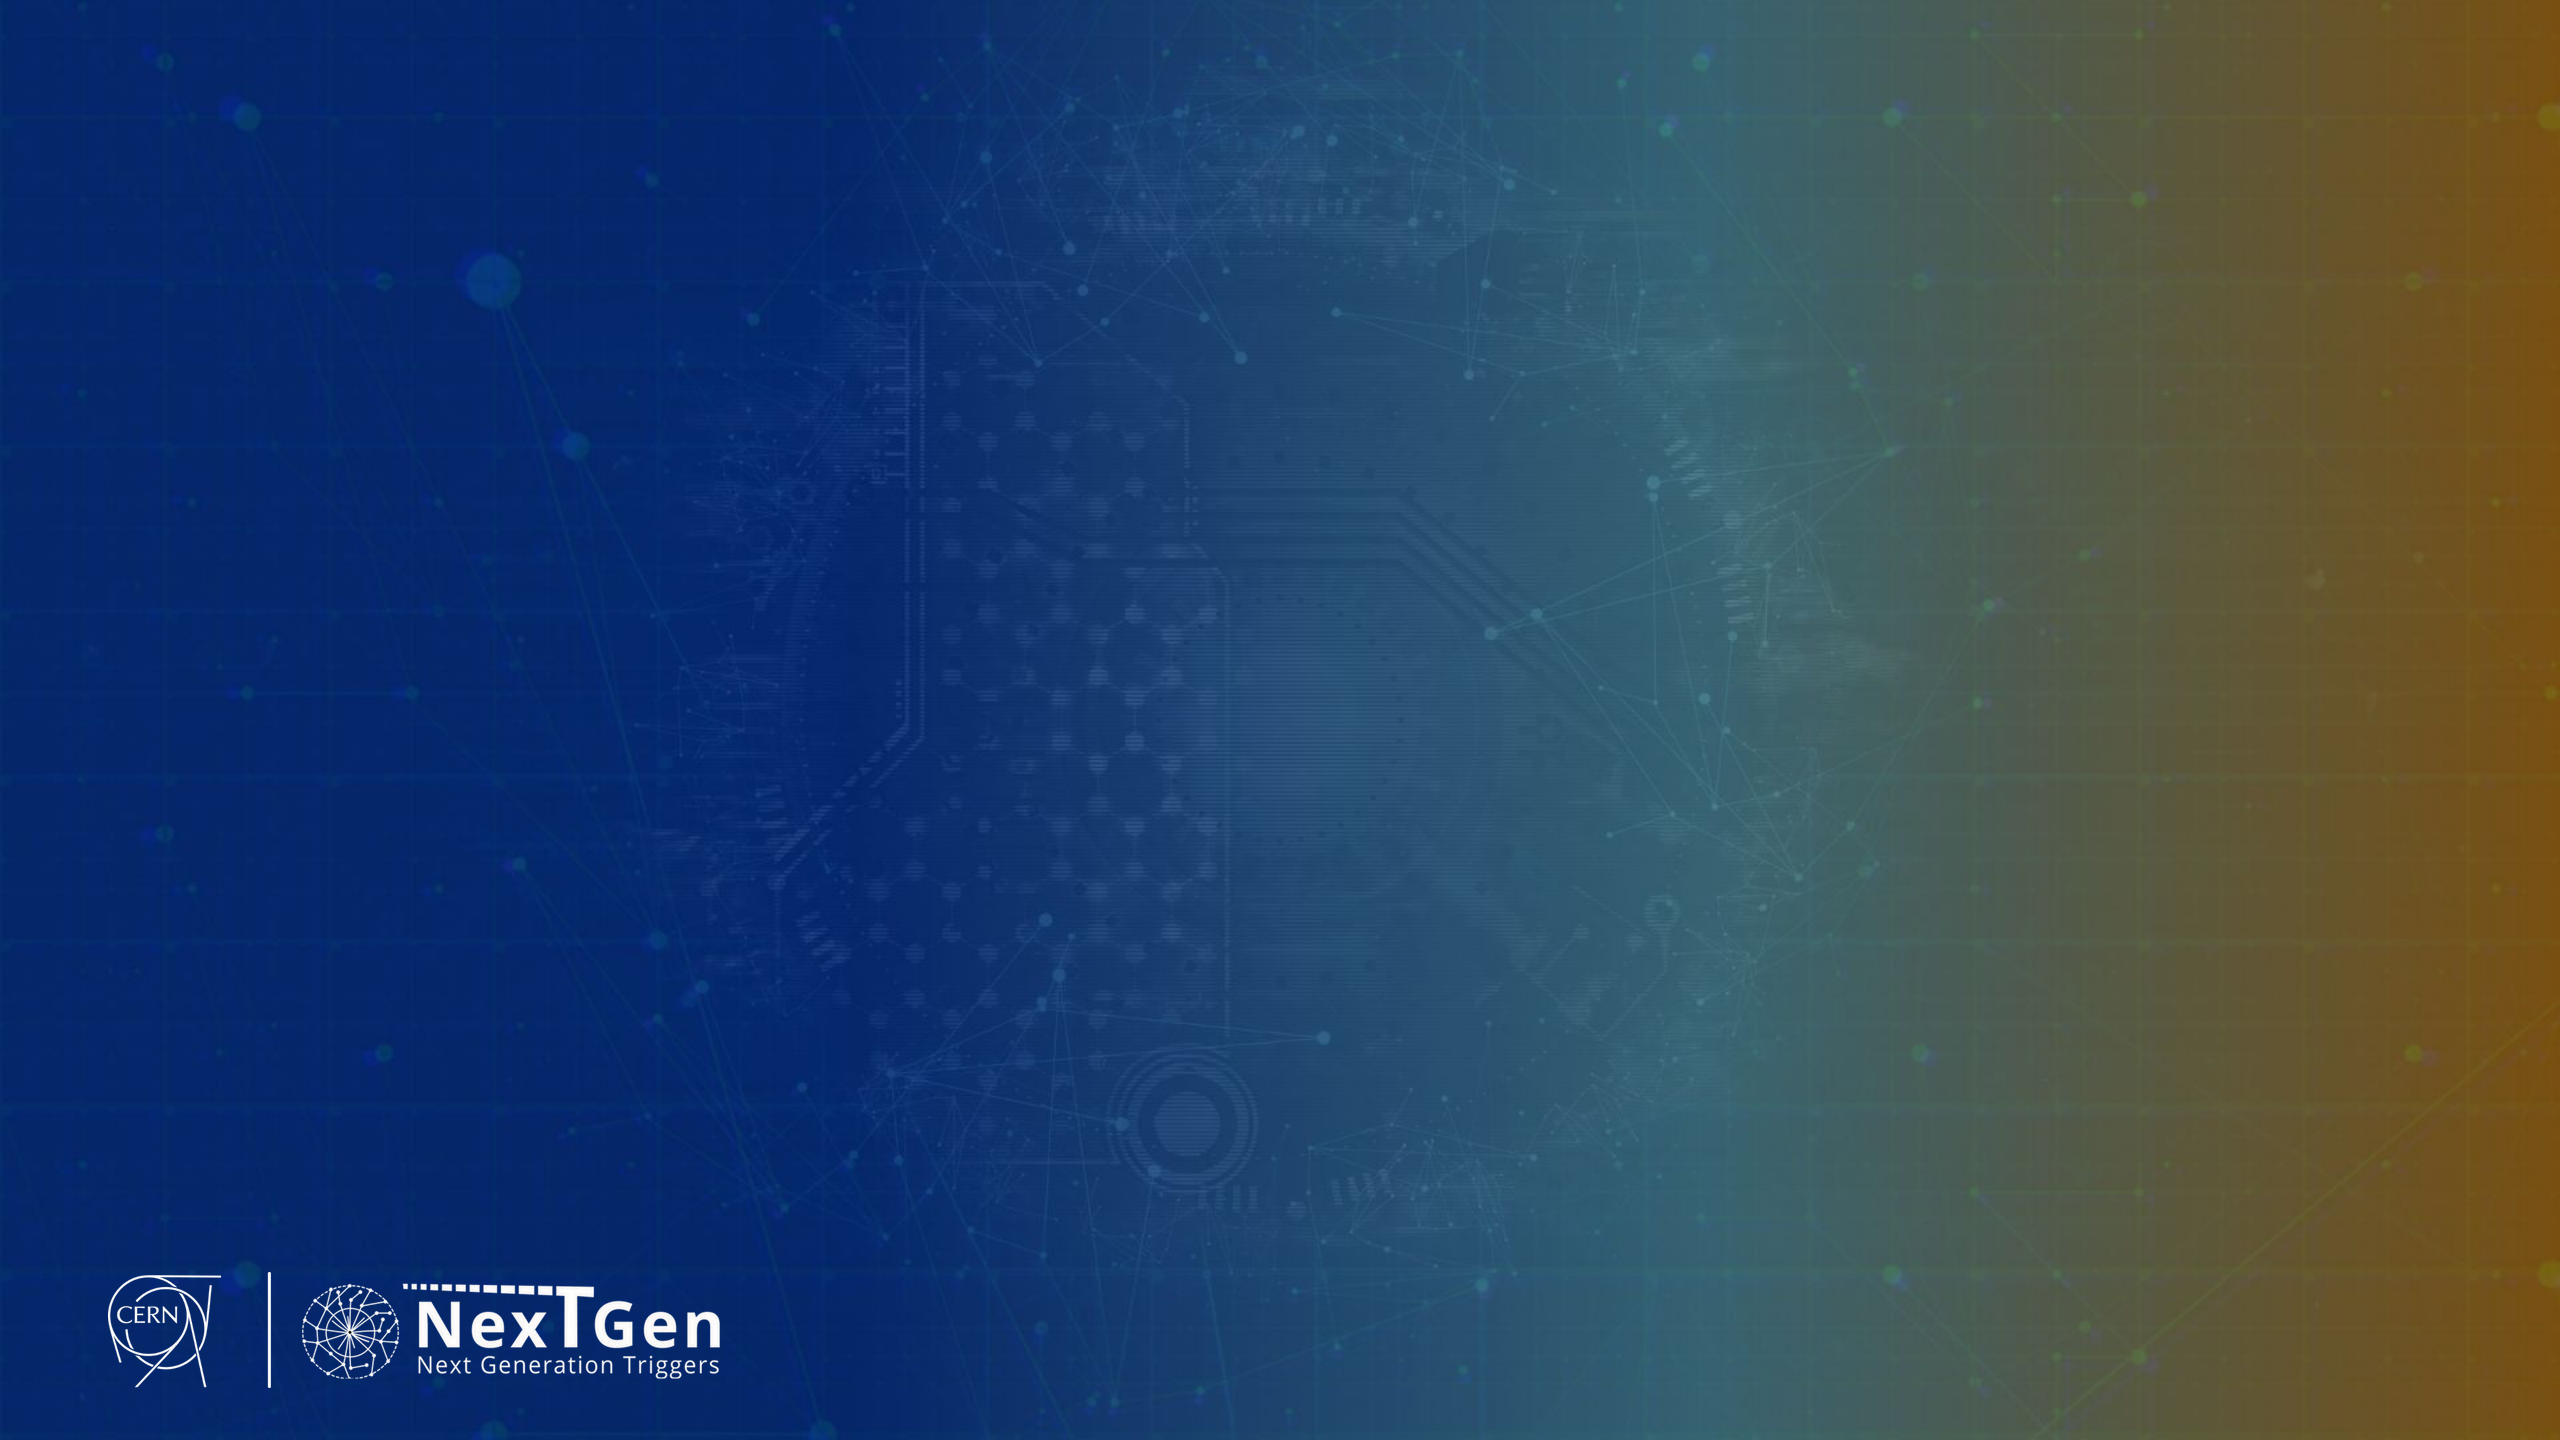
CERN
NexTGen
Next Generation Triggers