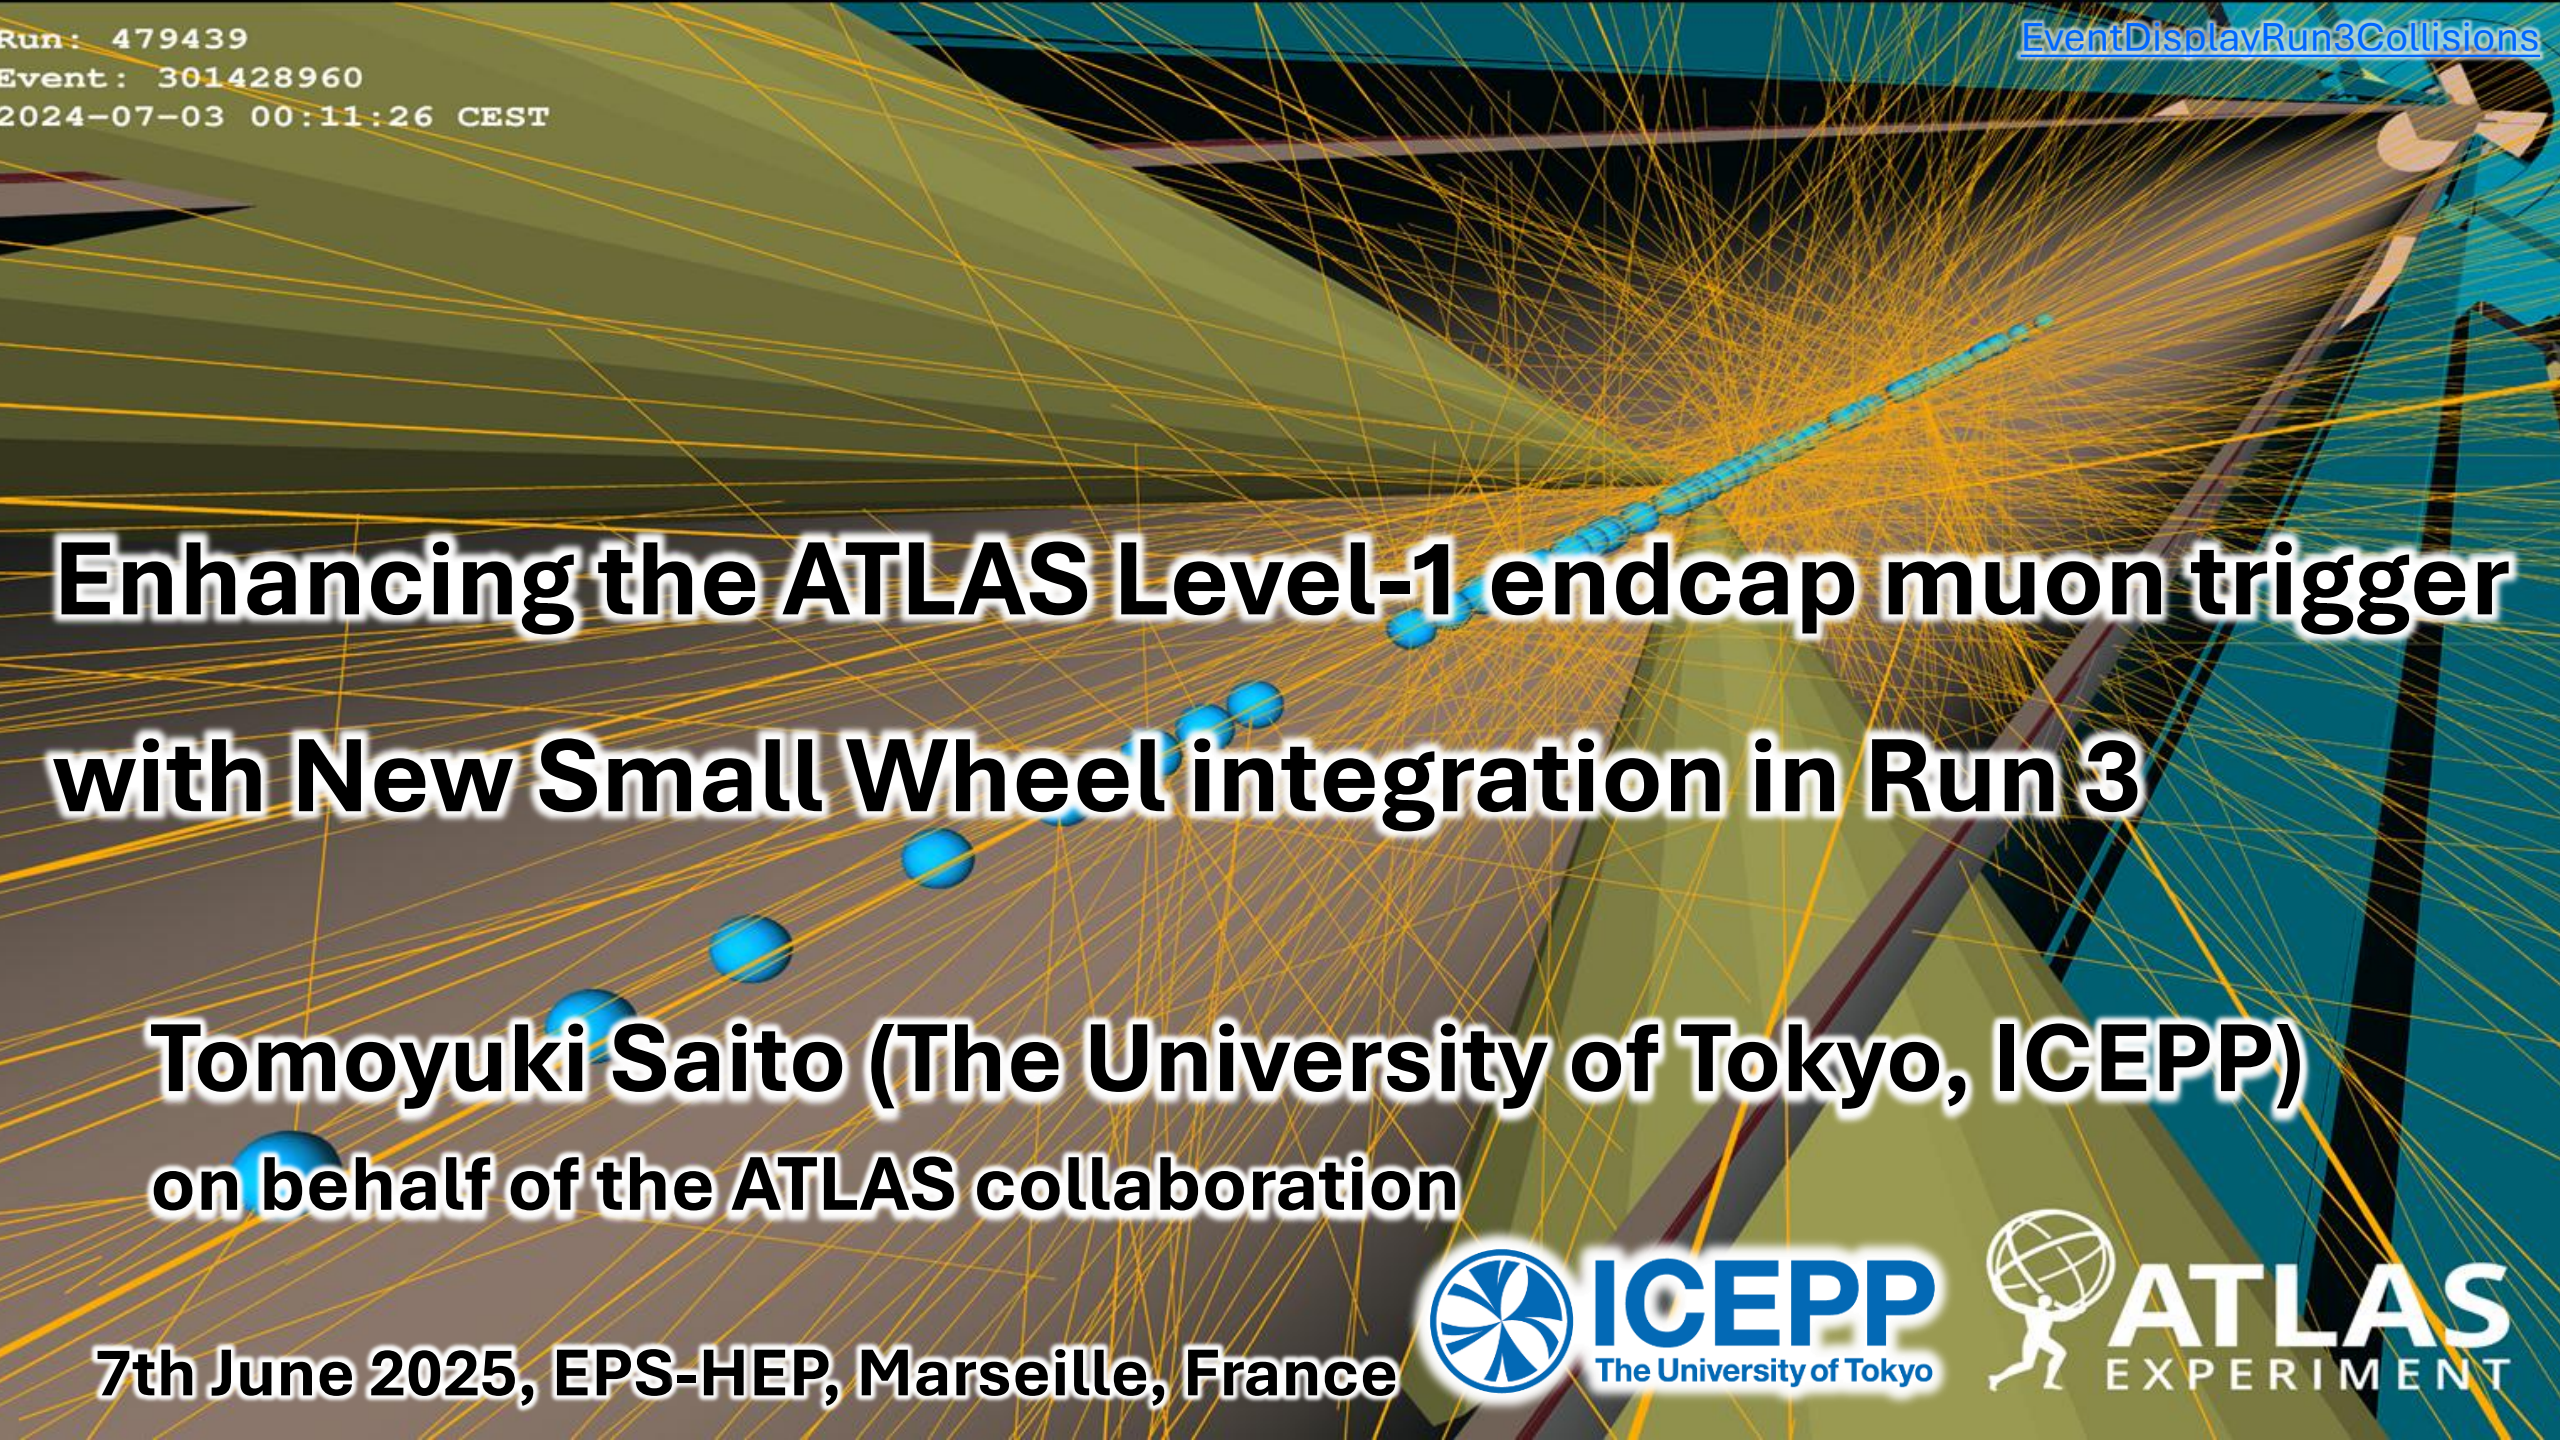

Enhancing the ATLAS Level-1 endcap muon trigger with New Small Wheel integration in Run 3

Tomoyuki Saito (The University of Tokyo, ICEPP)

on behalf of the ATLAS collaboration

7th June 2025, EPS-HEP, Marseille, France

ICEPP The University of Tokyo

ATLAS EXPERIMENT

LHC/ATLAS operation and Level-1 trigger rate in 2024

The LHC delivered the largest amount of data to the ATLAS experiment in 2024.

- Remarkable achievement, but also a major challenge for the ATLAS data-taking.
- The Level-1 trigger rate reached the ATLAS's limit, causing deadtime during data-taking.
  - Operating at L1 trigger rate~95 kHz with ~6% deadtime at a pileup of ~60 at the beginning of 2024.
- The high Level-1 trigger rate limited the ability of ATLAS to operate at even higher pileup.

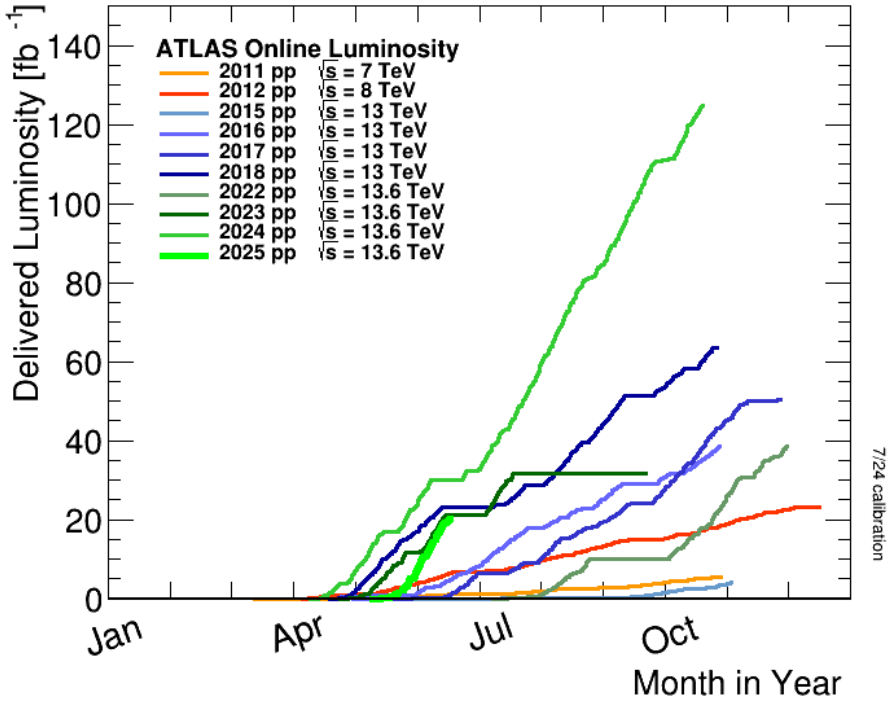

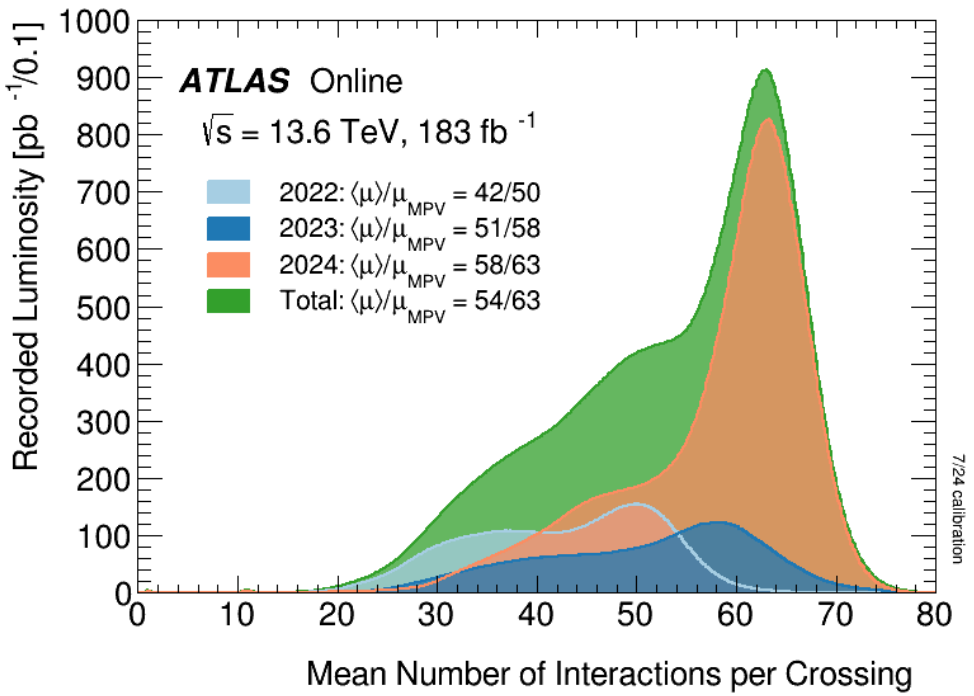

Pushing the system to its limits to fully exploit high-intensity beams delivered by the LHC.

LuminosityPublicResultsRun3

ATLAS challenge on high Level-1 trigger rate

High level-1 trigger rate limited ATLAS data taking efficiency and limited the achievable pileup.

Level-1 trigger rate and high-μ test in 2024

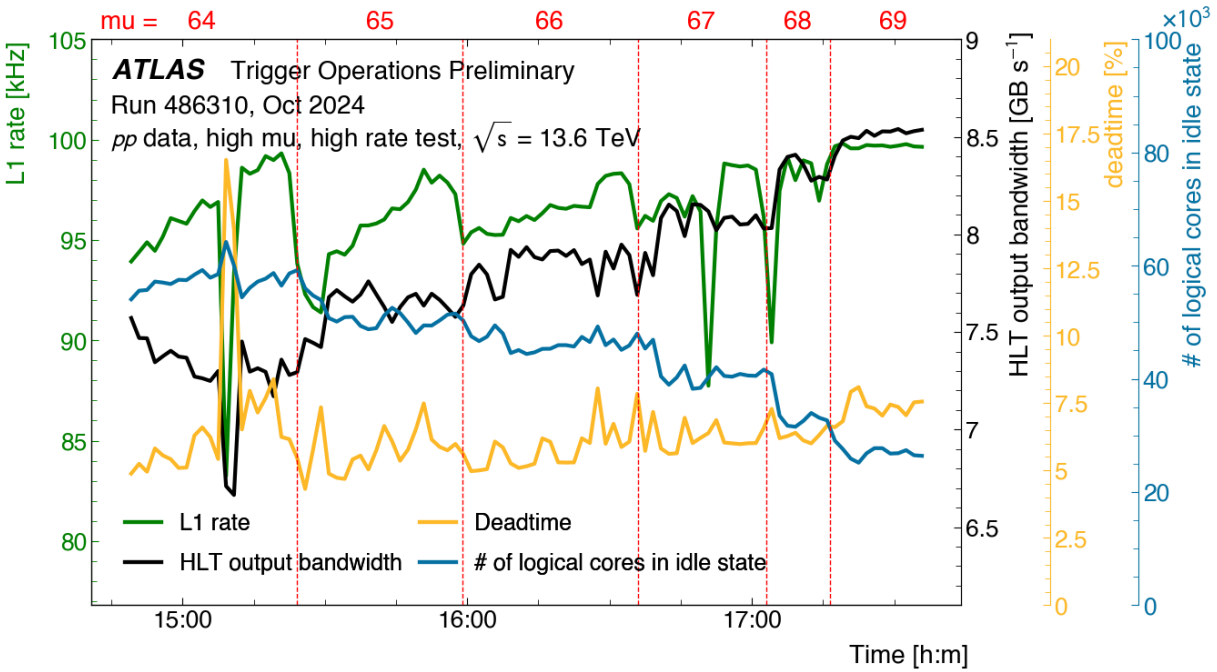

The deadtime increased with the L1 trigger rate

TriggerOperationPublicResults.

Dominant deadtime: Protective measures for detector readout from the high Level-1 rate (complex deadtime)

- Each detector stores data in its frontend pipeline buffer until the arrival of a Level-1 accept signal.
- At high L1 rate, the buffers are more likely to be full, resulting in deadtime due to the limited buffer size.

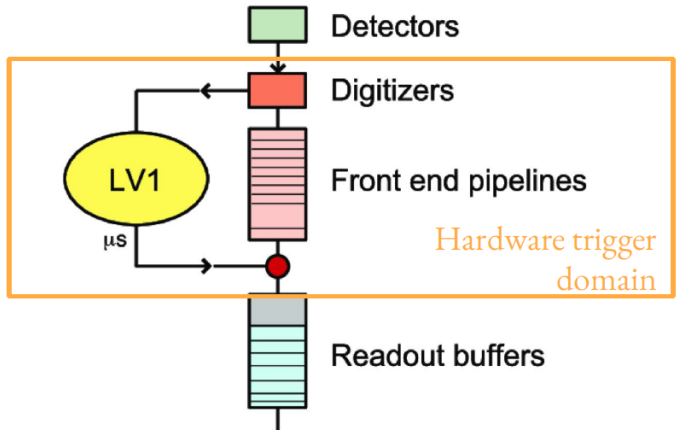

- The fraction of the complex deadtime, which keeps the buffer from overflow, has a non-linear effect for the L1 rate

The only effective mitigation is to suppress the Level-1 rate during Run 3.

ATLAS Level-1 trigger system

The Level-1 trigger system, built using custom electronics, reduces the trigger rate from 40 MHz to ~100 kHz with a maximum latency of 2.5 μs.

- It processes data from the calorimeters in L1Calo and the muon detectors in L1Muon.

Level-1 trigger rate

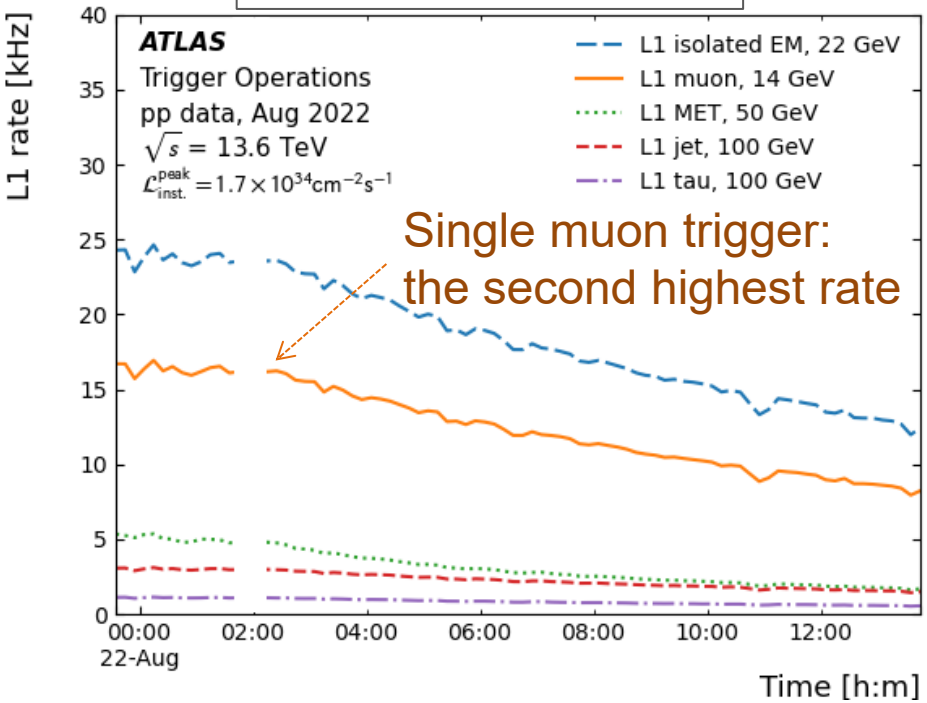

An increase in the L1 rates can lead to an unacceptable rise in ATLAS deadtime.

The suppression of the L1Muon rate can significantly improve ATLAS data taking.

TriggerOperationPublicResults.

ATLAS Trigger DAQ in Run3

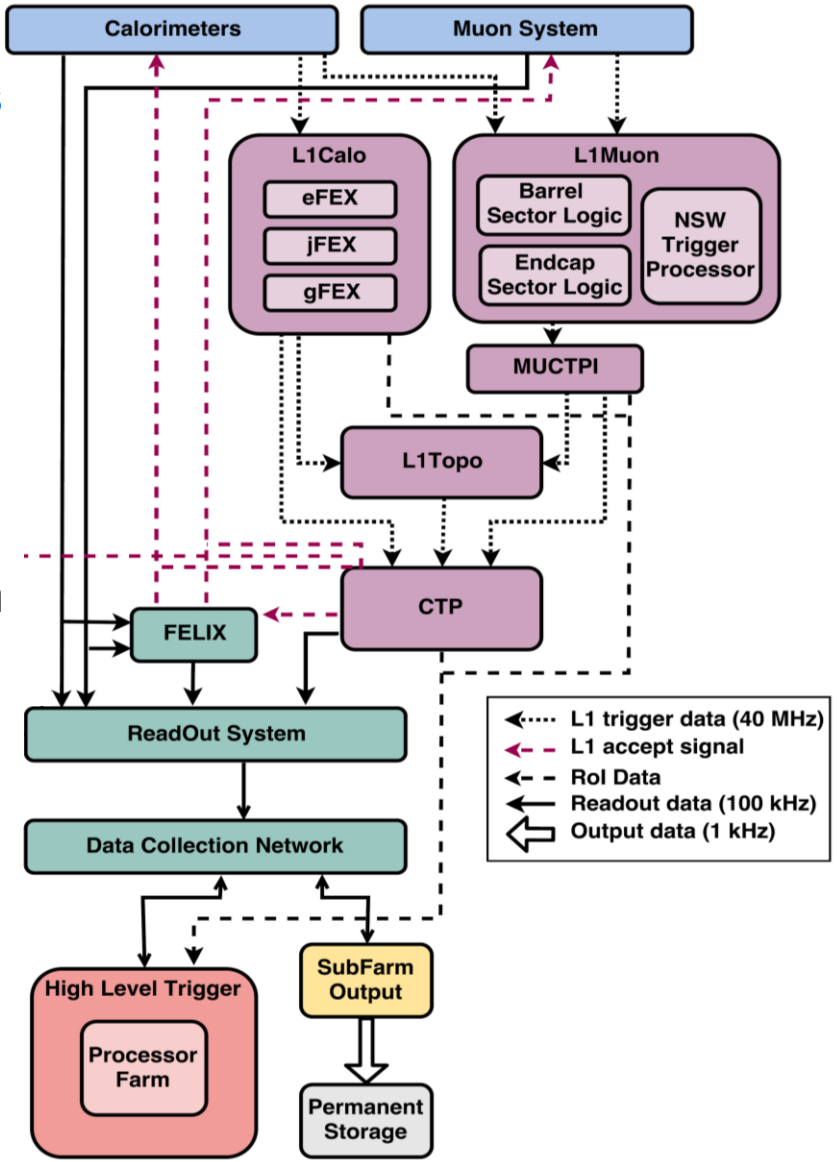

Level-1 Muon endcap trigger

The Level-1 muon endcap trigger (1.05<|η|<2.4) based on the Thin Gap Chamber (TGC) at the middle muon station (Big Wheel).

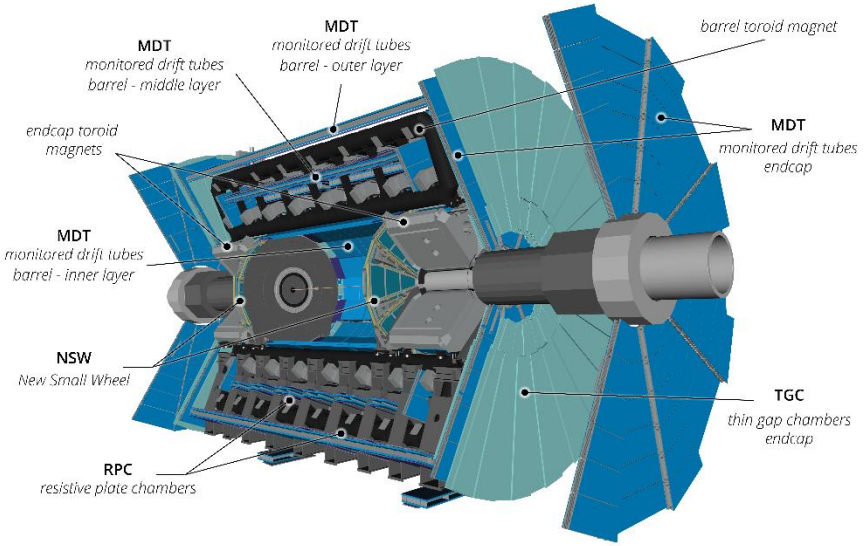

- The TGC detector consists of 3 stations
- Frontend electronics on detector
  - 2/3 layer coincidence in TGC1 and 3/4 layer coincidence in TGC2-3 independently for wire and strip data (R/ϕ)
- Backend electronics (Sector Logic)
  - R-ϕ coincidence to identify muon track candidates
- Point angle measurement
  - Intrinsic backgrounds of charged hadron particles from hadronic interactions inside the toroidal magnet structures.

A significant challenge is the high background rate, causing extra rate to the ATLAS Level-1 trigger rate.

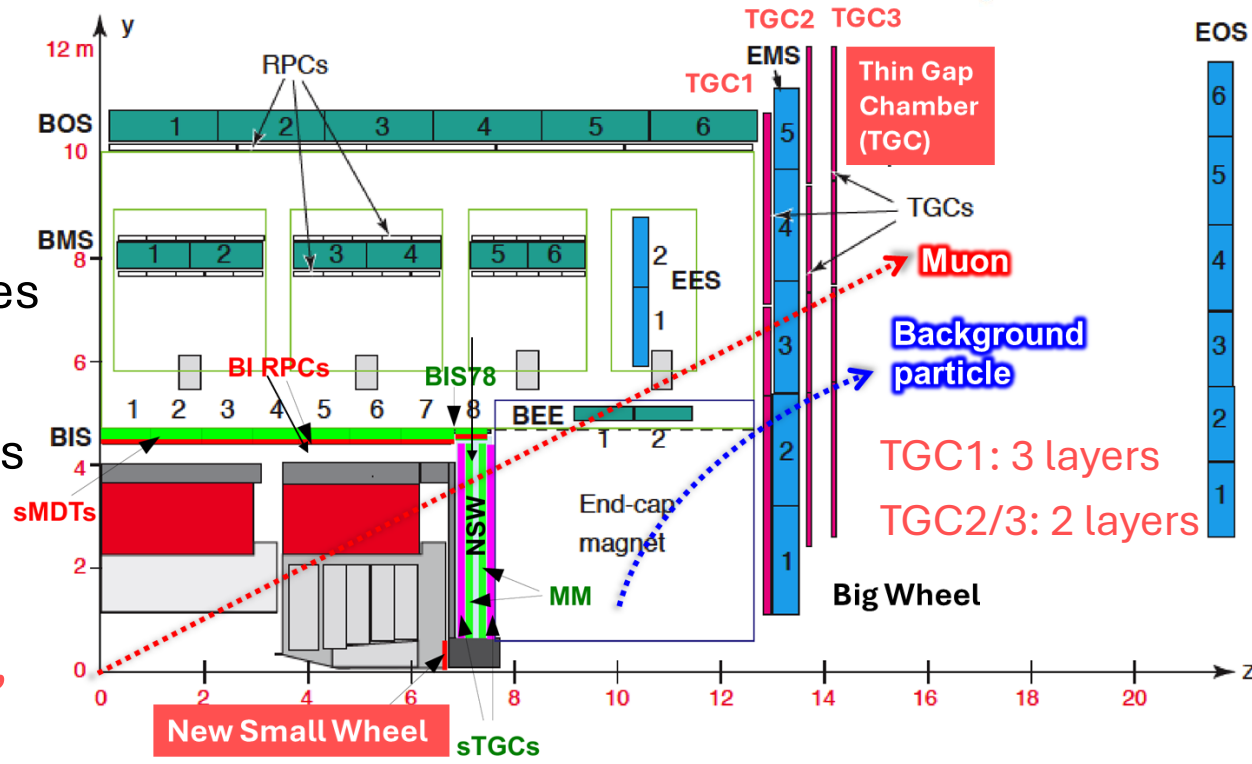

High background rate in Level-1 Muon endcap trigger

To mitigate the high background rate, the L1Muon endcap system needs detector information radially inside the toroidal magnet as well as the TGC.

- The New Small Wheel (NSW) detectors installed at the inner muon station before Run3
- The NSW integration into the L1 Muon was implemented to reduce the fake trigger rate and suppress the Level-1 trigger rate.

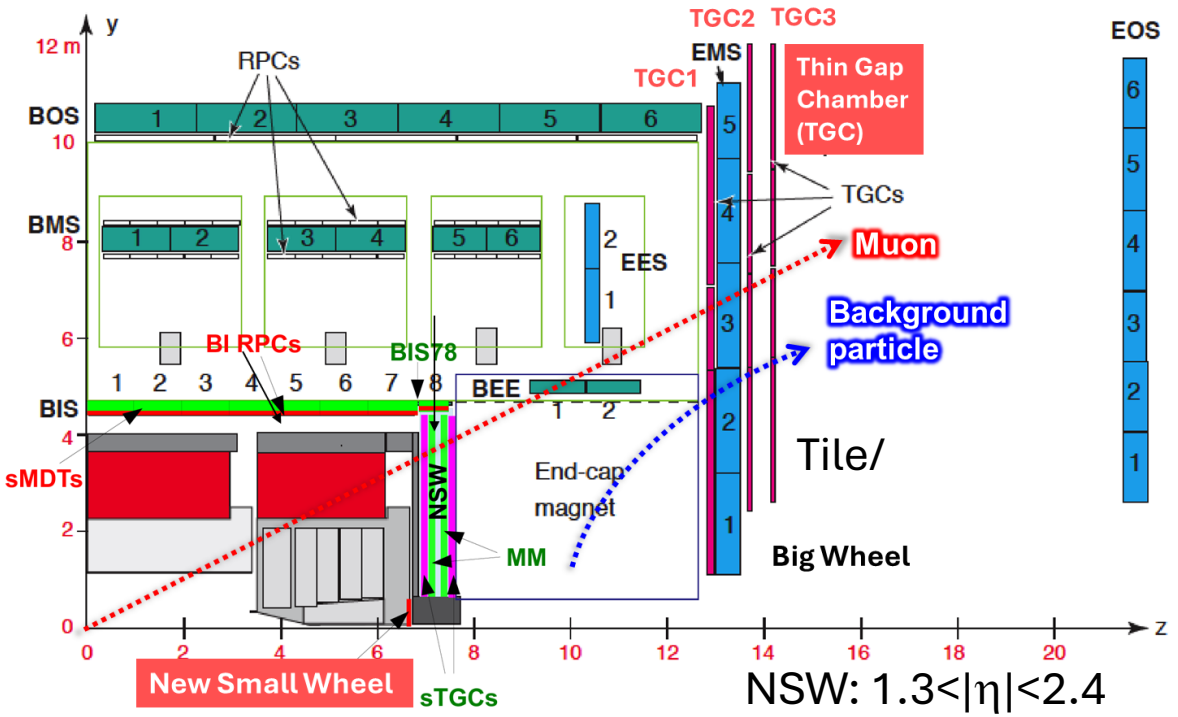

NSW: 1.3<|η|<2.4
Tile/TGC-EIL4: 1.05<|η|<1.3

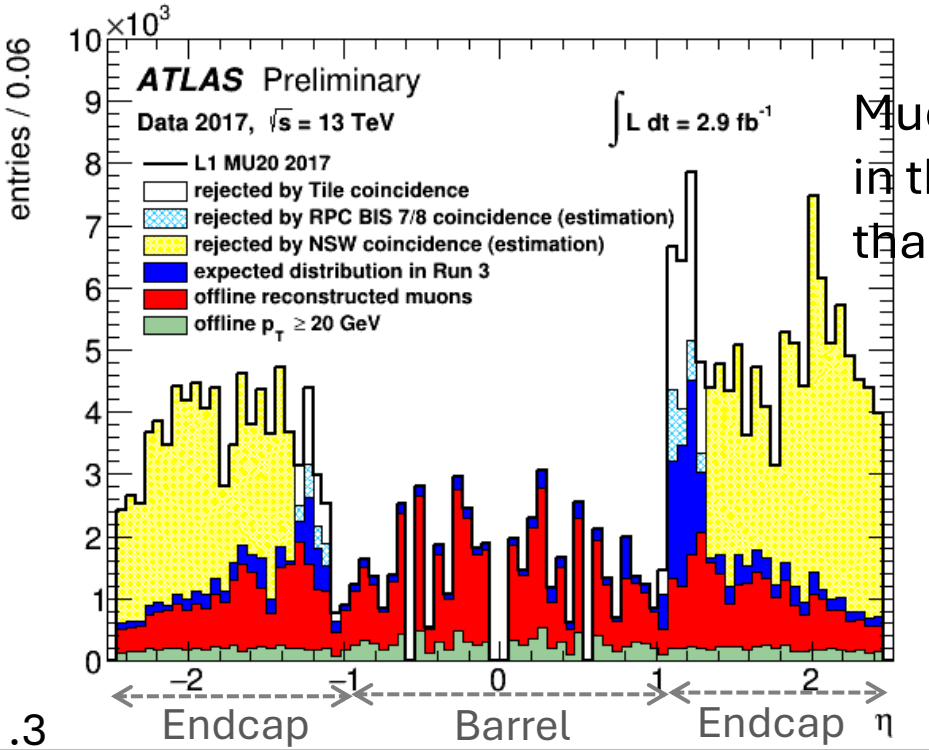

Much higher rate in the endcap than in the barrel.

[L1MuonTrigger PublicResults]

New Small Wheel for Level-1 Muon Trigger

The NSW is composed of two technologies, small-strip TGC (sTGC) with 8 layers and Micromegas (MM) with 8 layers, providing typically a single-hit spacial resolution of ~100 μm.

- Track segments independently reconstructed in the sTGC and the MM by layer coincidences, merged at the NSW Trigger Processor (NSW-TP), which sends the results to the Sector Logic.
  - Currently, sTGC pad and MM signals are integrated (sTGC strip integration will come).
- The NSW-TP consists of an ATCA board, 2 mezzanine cards with 2 FPGAs, each with optic links.
- 36 NSW-TP boards for the entire NSW region.

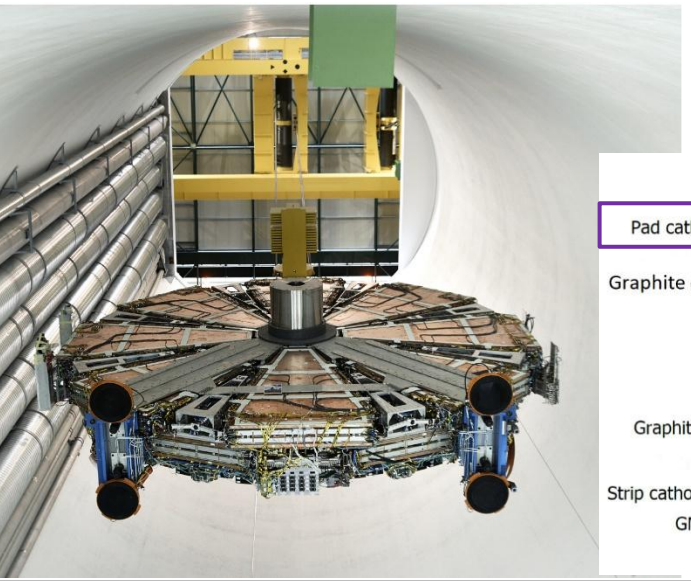

small-strip TGC (sTGC)

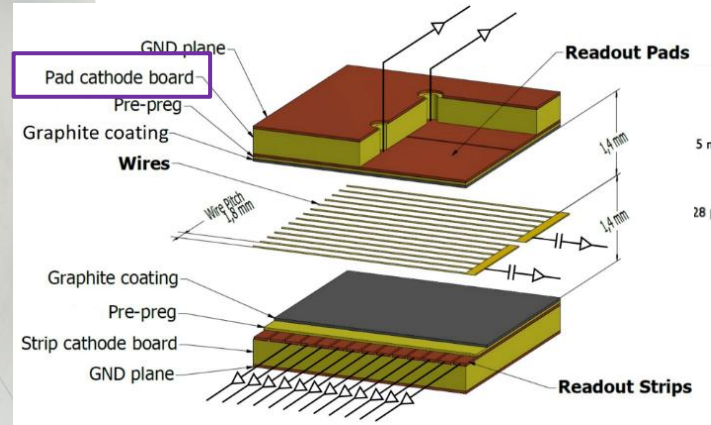

Micromegas (MM)

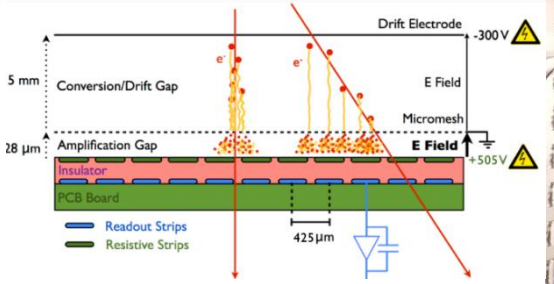

NSW Trigger Processor (NSW-TP)

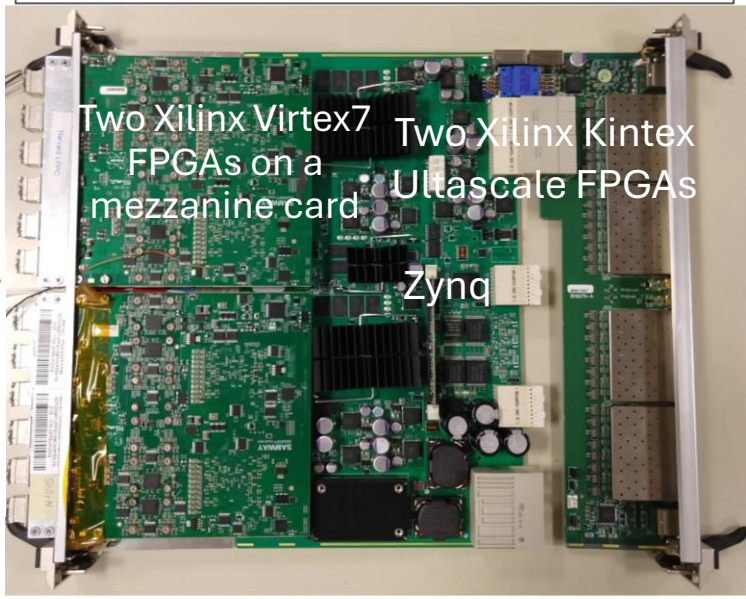

Level-1 muon trigger processor (Sector Logic)

The Sector Logic (SL) is the central component in the L1Muon trigger decision.

- Upgraded for Run3 to provide increased FPGA resources and additional high-speed transceivers for the NSW integration.
- Main processor: Xilinx Kintex-7 FPGA
  - A sufficient number of multi-gigabit transceiver (GTX) available and connected to the front panel, allowing the board to receive data from the NSW and transmit the final processing results to the Muon CTP Interface (6.4 Gbps).
  - The large number of Block RAMS essential for implementing low-latency trigger processing.
- 14 G-link receivers to receive the TGC wire and strip data.
- The Level-1 Muon endcap system consists of 72 SL boards.

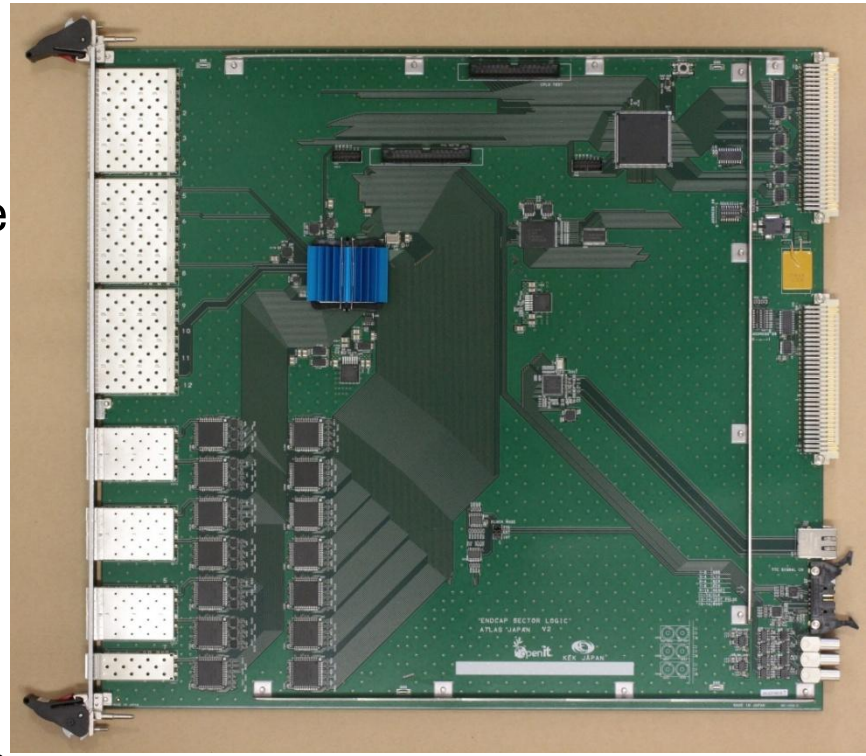

Muon identification algorithm in Sector Logic

The trigger algorithm identifies muon candidates in two stages.

Stage-1: Muon candidates are identified based on TGC wire/strip(R/ϕ) data (TGC-only trigger)

- 7-layer R/ϕ coincidence for TGC wire and strip and p_T determination.

Stage-2: Fake candidates are rejected by incorporating the NSW information into the TGC-only trigger candidates (TGC+NSW trigger)

➜ Significant reduction of the fake trigger rate

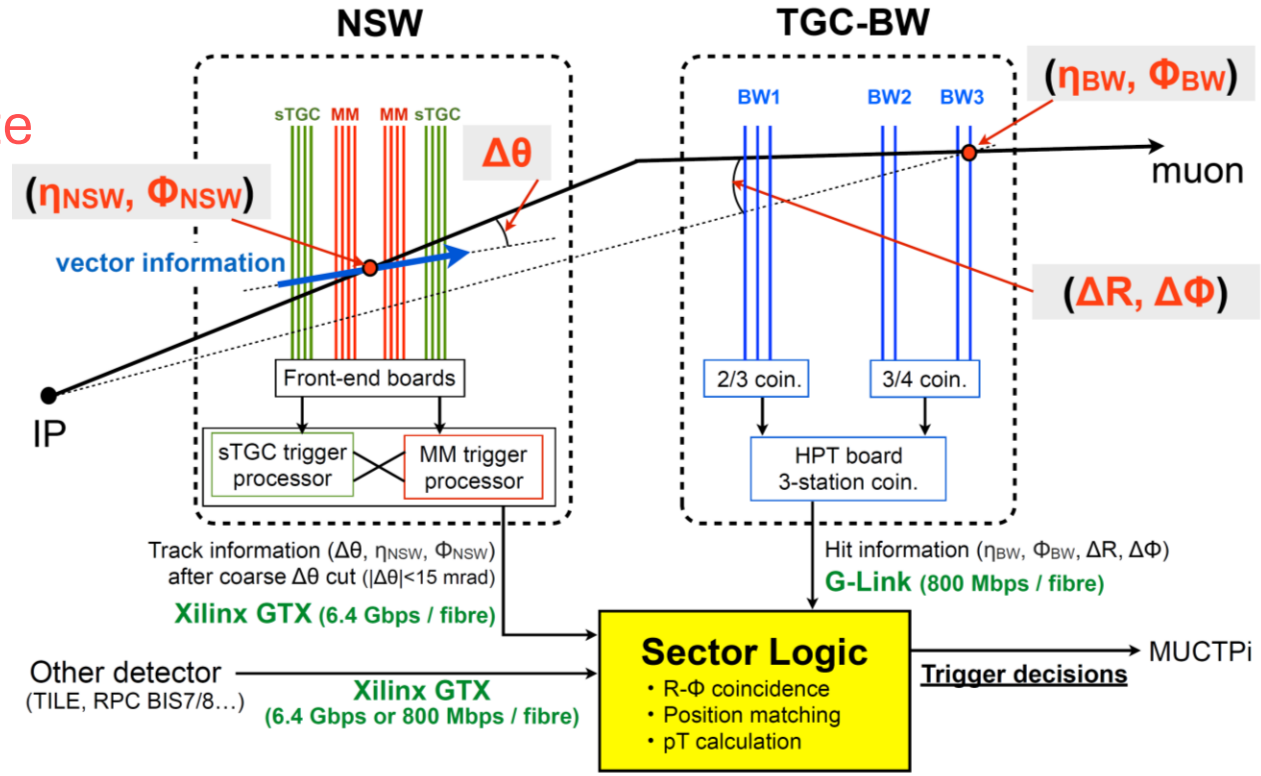

Implementation point:

- The level-1 trigger latency is severely limited.
- The algorithm is mainly processed using the Look-Up-Tables (LUTs), where pre-processed results are stored for rapid retrieval during processing.

NSW integration process in 2024

NSW integration into the Level-1 muon was carried out in stages to ensure no data loss during the integration process.

- The TGC+NSW trigger or TGC-only trigger can be switched independently for each small unit.
  - allowing us to activate the TGC+NSW trigger selectively for well-performing NSW regions.
- Thorough check on the efficiency in each region of NSW
  - Masks implemented to the regions with low efficiency
  - Fine-tuning of signal timing alignment to further improve the efficiency

Finally, the full Integration of all the NSW sectors completed in May 2024

TGC+NSW trigger efficiency

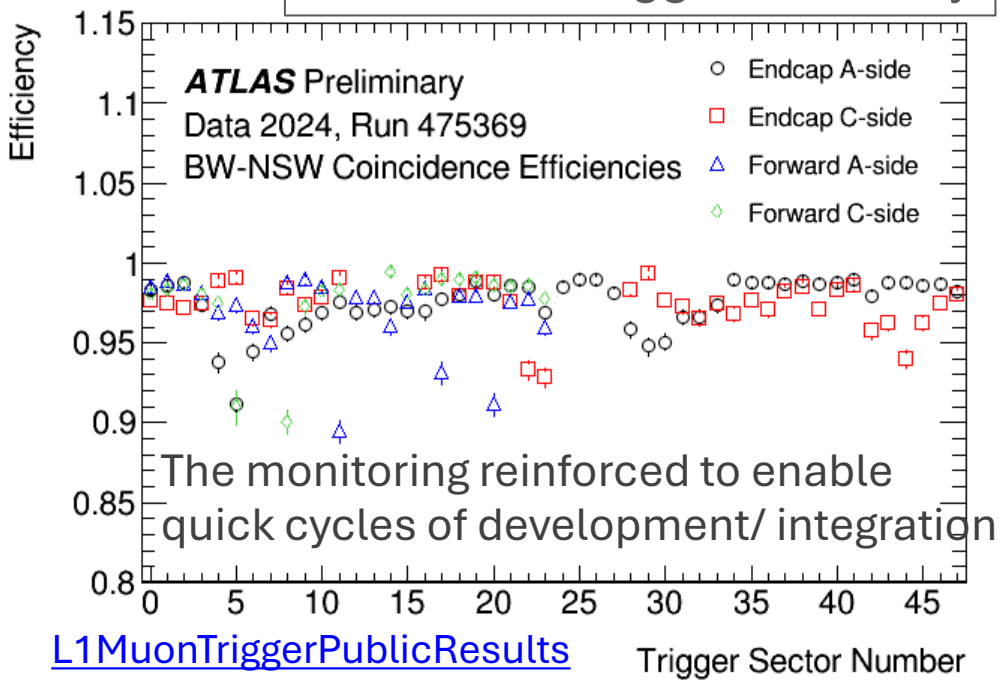

The monitoring reinforced to enable quick cycles of development/ integration

L1MuonTriggerPublicResults

Date in 2024	Integrated region	
	sTGC pad	MM
16 April	67%	-
7 May	87%	3%
24 May	92%	12%
28 May	100%	100%

Level-1 muon rate reduction by NSW integration

Level-1 muon rate transition in 2024

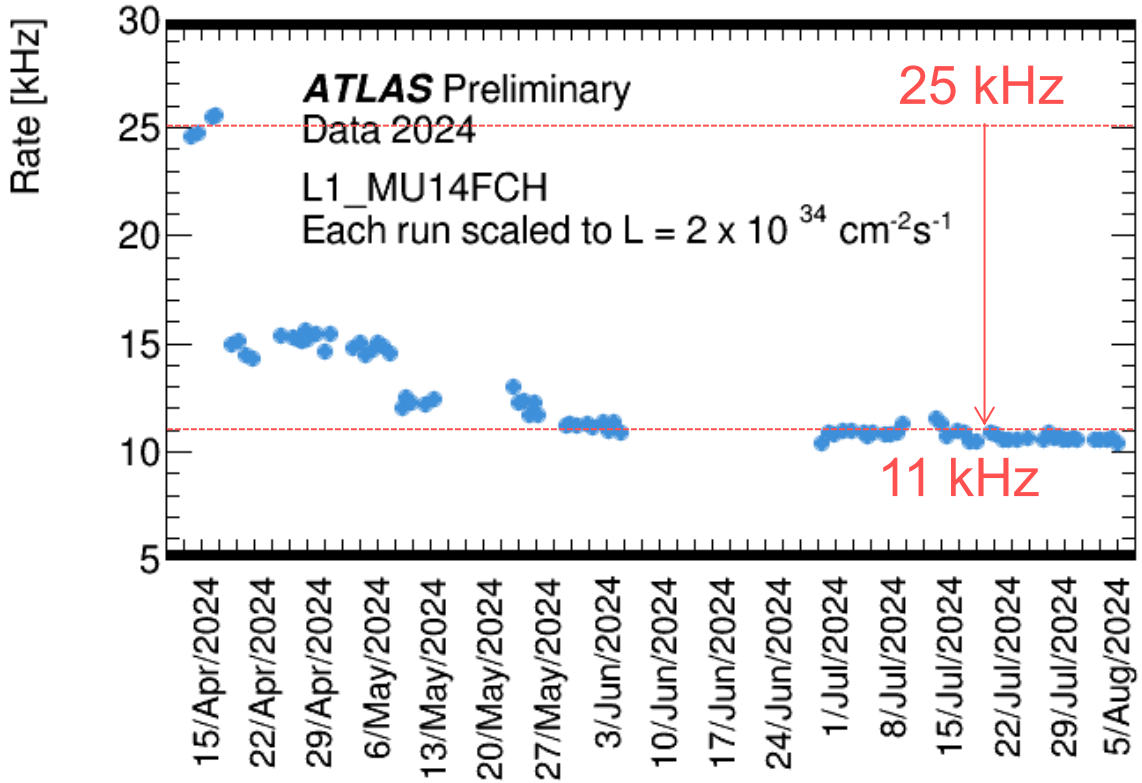

- Rate reduction at stages in the integration process
- 15% reduction of the overall Level-1 trigger rate

L1MuonTriggerPublicResults

Level-1 muon rate reduction for each η region

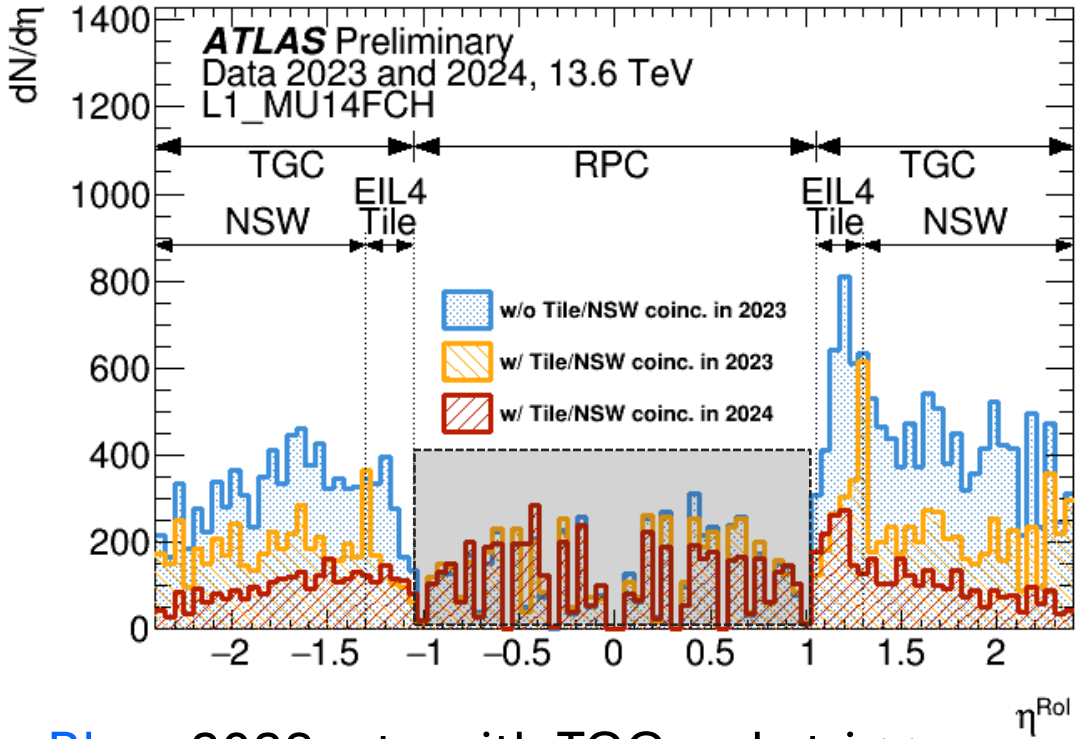

- Blue: 2023 rate with TGC-only trigger.
- Red: 2024 rate after the TGC-NSW integration
- Significant rate reduction across the high-rate endcap regions → Successful integration!

Improvements of ATLAS by L1 trigger rate reduction

ATLAS pileup in luminosity levelling

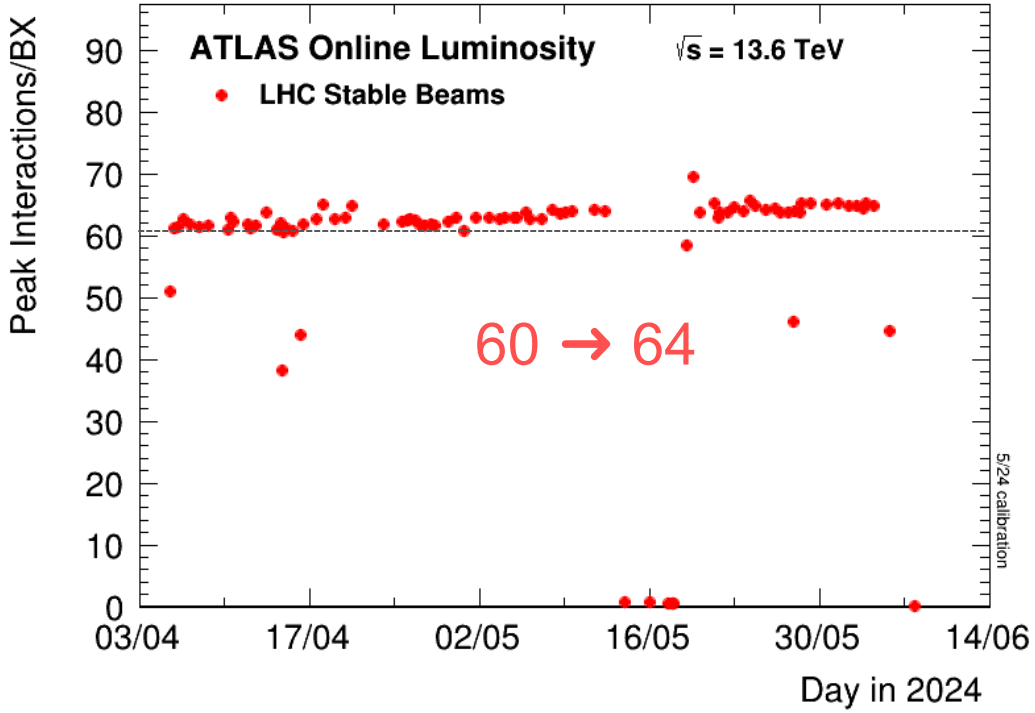

The pileup during luminosity levelling increased from 60 to 64 due to the Level-1 trigger rate reduction.

LuminosityPublicResults

ATLAS data taking efficiency

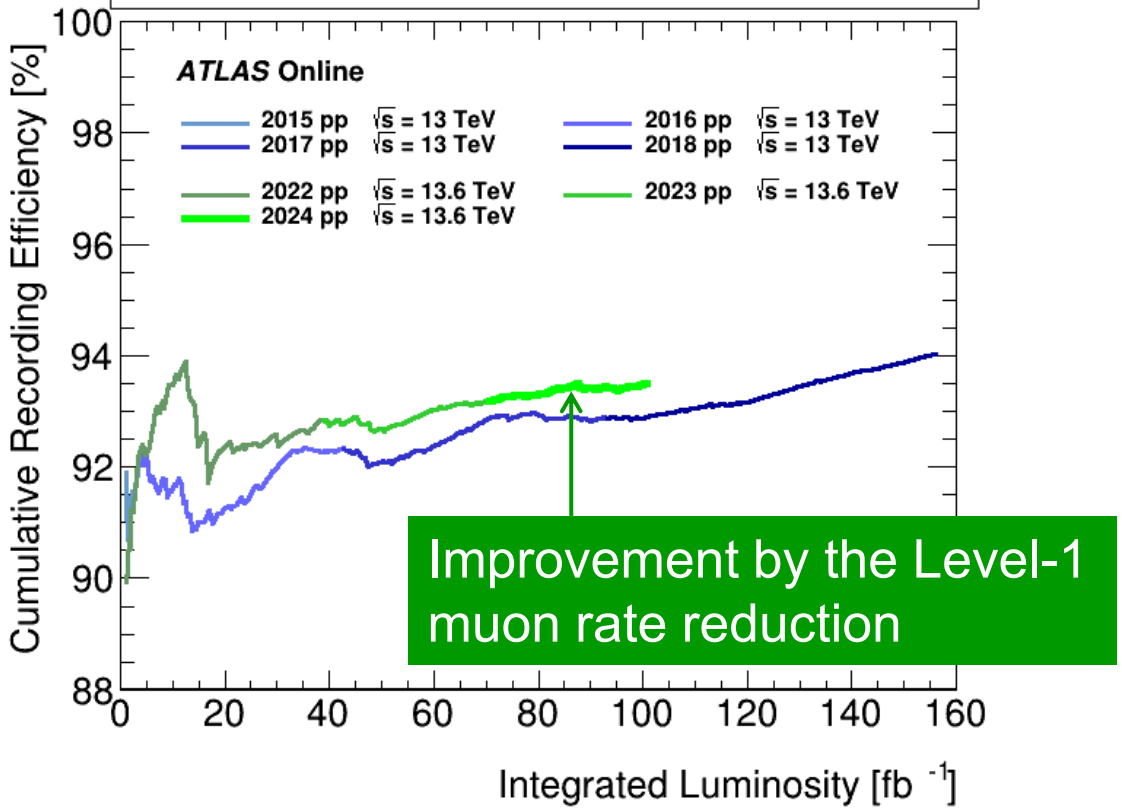

Higher efficiency in 2024 despite operating at higher pileup compared to 2023.

Improvements of TGC+NSW trigger in 2025

There is potential to further improve the efficiency of TGC+NSW trigger in 2025.

- ~4% lower efficiency of TGC+NSW trigger (red), compared to the TGC-only trigger (blue).
- The NSW-MM can mitigate the inefficiency of NSW-sTGC.
- Additionally, there is still room to optimise the window size of the TGC-NSW coincidence requirement in the SL.

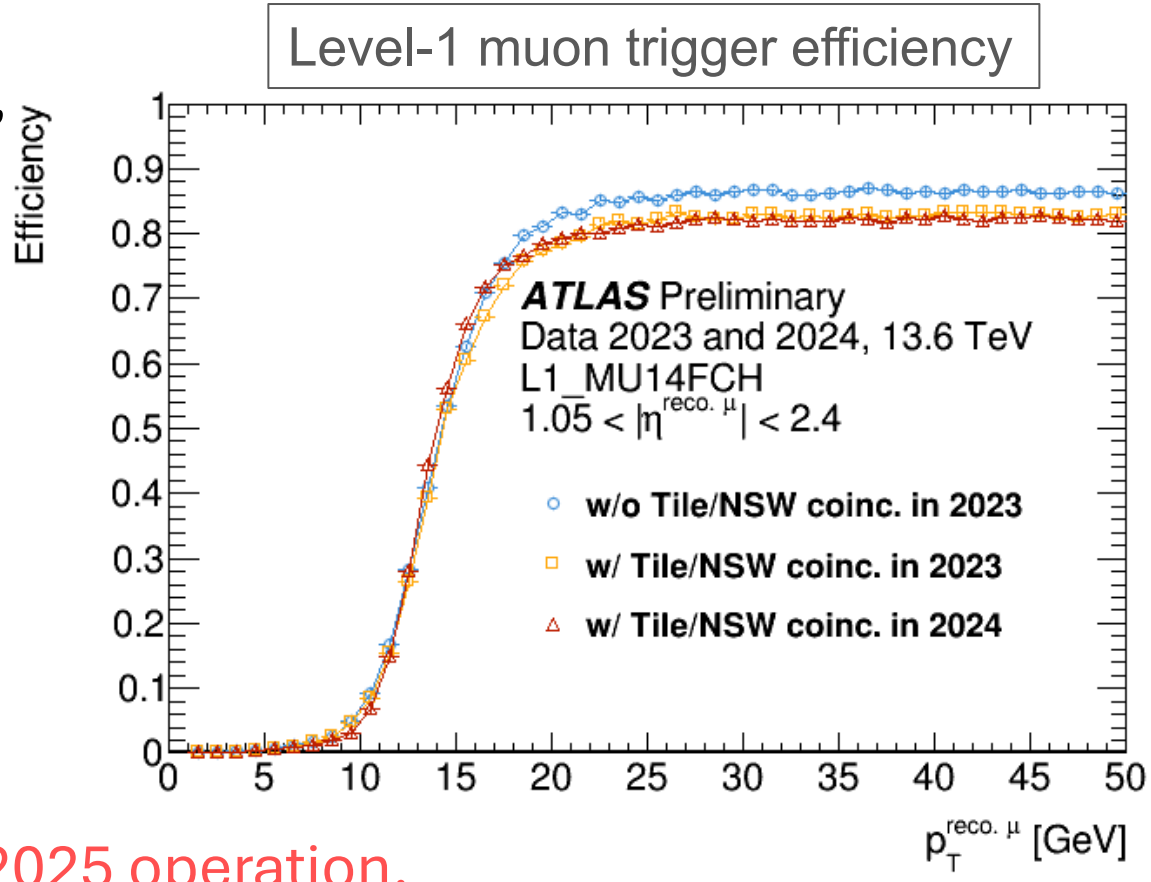

The efficiency is getting improved in the ongoing 2025 operation, getting close to the TGC-only trigger.

L1MuonTriggerPublicResults

Summary

The challenge in the ATLAS Level-1 muon endcap trigger is the high background rate, which contributes significantly to the overall Level-1 trigger rate in the ATLAS experiment.

- The L1 muon endcap trigger utilises the TGC signals in the middle muon stations, but still suffers from a high fake trigger rate due to the background particles.
- We achieved significant improvements in muon-background reduction by integrating the NSW detectors into the L1 trigger.
- The NSW detectors were fully integrated into the L1Muon trigger in 2024.
- This advancement reduced the Level-1 muon trigger rate by ~14 kHz, which improves the overall data-taking efficiency of the ATLAS and allows the ATLAS to run at high pileup (64).

ATLAS Outstanding Achievement Awards 2025

Backup

ATLAS Complex deadtime

Complex deadtime arises due to limitations in the ability of the detector front-end electronics to process and buffer multiple events that arrive too close in time. It is managed by a set of rules in the Central Trigger Processor (CTP), which enforces certain conditions to protect the system from being overloaded. These rules are based on:

- The **rate** of Level-1 Accept (L1A) triggers in short time intervals
- The **spacing** between L1A triggers (e.g., a minimum number of bunch crossings between two triggers)
- The **buffering capacity** of sub-detectors and readout electronics

It is defined by X (in units of L1A) being the size of the bucket (i.e. the front-end buffer) and R (in BC) being the time it takes to leak 1 L1A. With these numbers the trigger rate, on average, is limited to X triggers in a time period of X × R bunch crossings. T

Background particles

The majority of the background rate originates from beam-induced backgrounds that produce relatively low-energy protons.

- These protons, in turn, generate both random and correlated hits.
- These are more positively charged background particles than negatively charged ones.
- The direction of the toroidal magnetic field is the same on both the A- and C-sides.
- Due to the curvature induced by the toroidal field, positively charged background particles are more likely to pass through the TGC on the A-side.
- In contrast, this is not the case at C-side, where positively charged particles are less likely to reach the TGC.

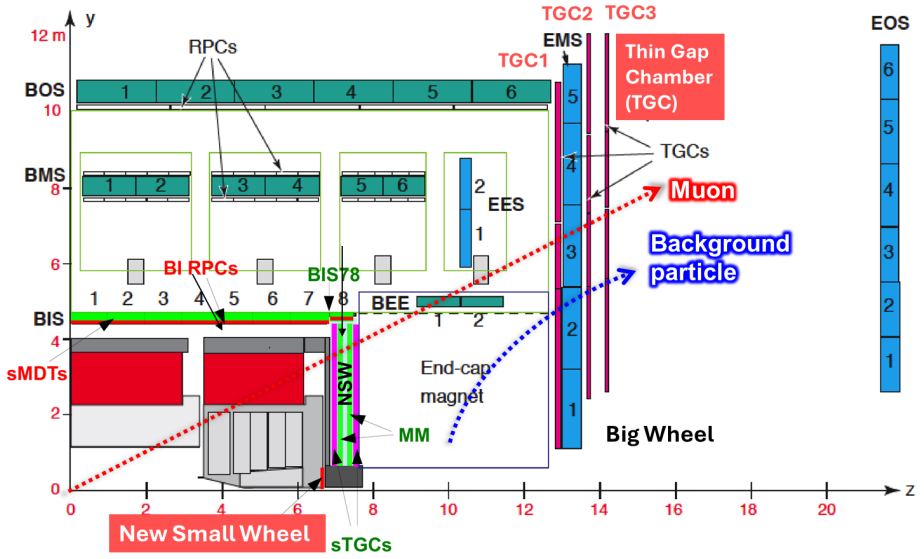

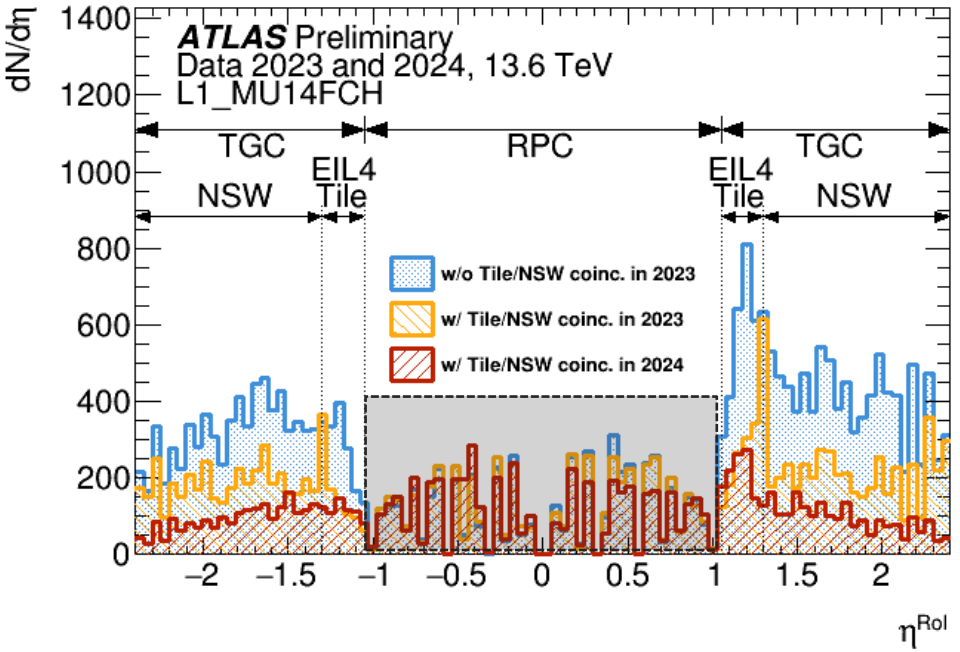

NSW trigger path

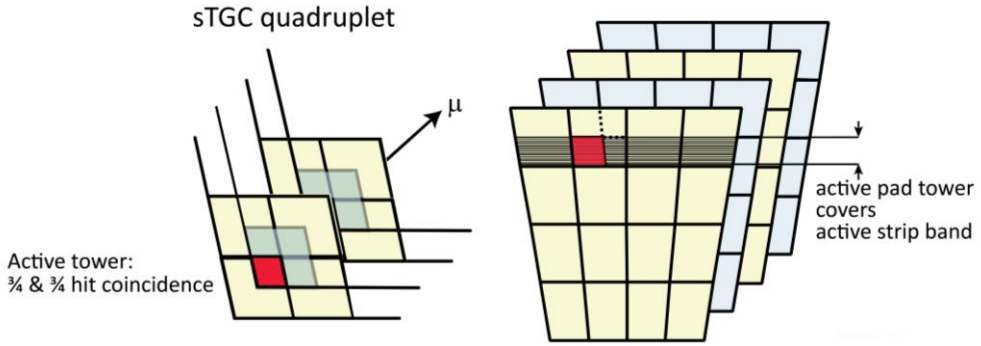

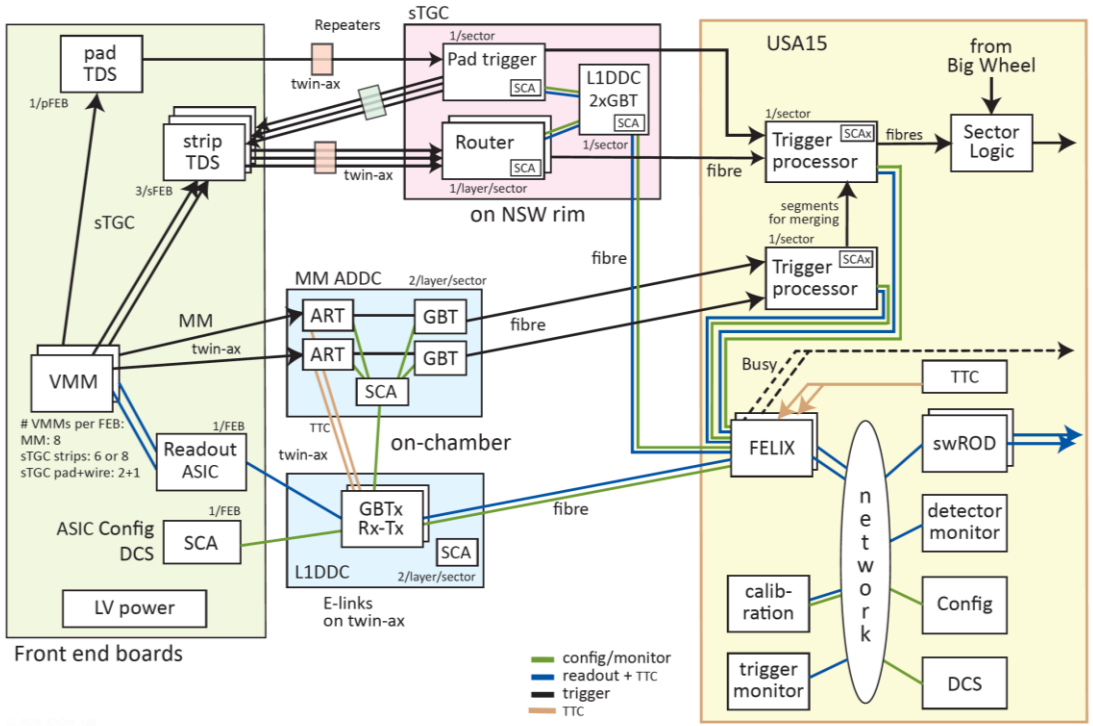

The sTGC and MM Trigger Processors separately find track segments from hits in their respective detectors.

sTGC:

- Currently, the pre-trigger of pad signal formed at the Pad trigger board is used.
- The pads in half of the layers are shifted by half a pad in both directions to increase the resolution.
- Eight-layer towers pointing to the IP are defined by the overlapping physical pads.

MM:

- The concept utilizes the fine pitch of the MM detectors and the spread of ionization charge for particles crossing the detector at an angle.
- The VMM sends out the address of the channel that presents the earliest signal in every BC. For the fine ~0.45mm strip pitch, this is a good a approximation of the coordinate perpendicular to the strip direction.