The ATLAS Trigger & Data AcQuisition (TDAQ) System

Catrin Bernius (SLAC)
on behalf of the ATLAS Collaboration

Mu2e-II Workshop
14. September 2020

LHC Conditions - Reminder

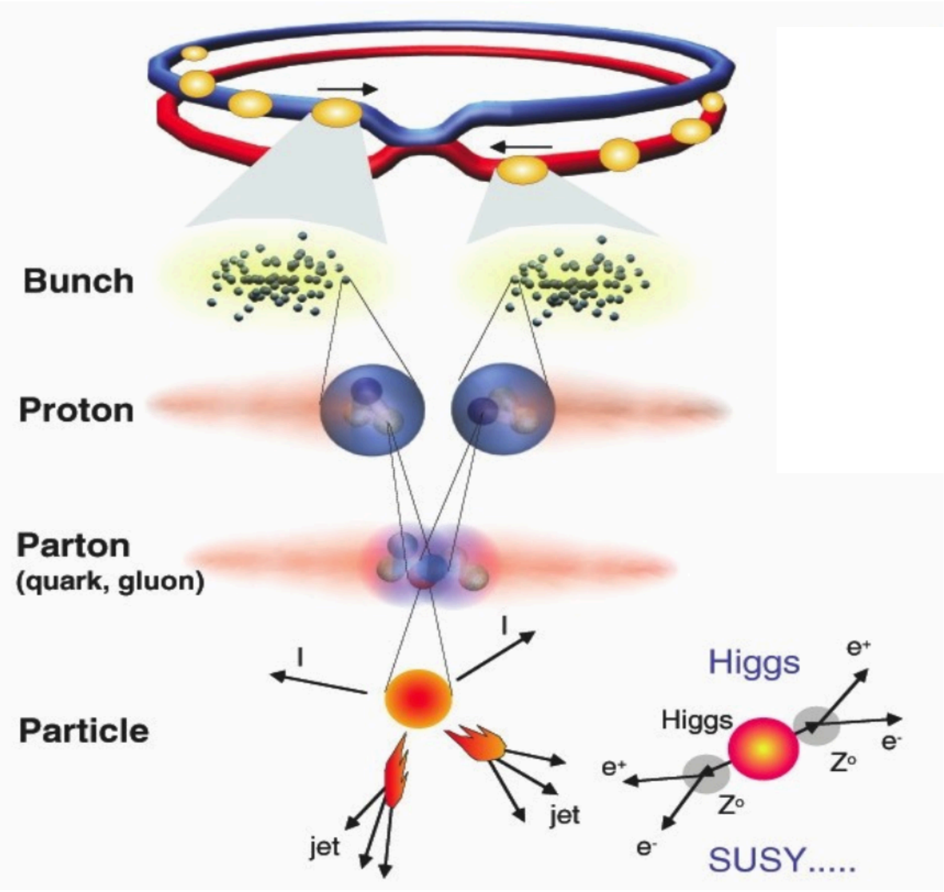

- Crossing of the two proton beams at an interaction point is called *Event*
  - Collisions in 4 points, surrounded by ATLAS, CMS, LHCb, ALICE
  - LHC beams are made up of proton bunches with 10^{11} protons/bunch
  - Colliding at 13 TeV (6.5 TeV per beam) and ~ 30 MHz in Run 2 (40 MHz if the machine is full with 1 bunch crossing every 25 ns)
  - In 2018: **up to 60 proton-proton collisions per bunch crossing**

Physics at the LHC

- Typical collision at LHC not necessarily what we call “interesting"
  - Need to reject orders of magnitude of soft QCD
- Interesting physics is 6-8 orders of magnitude rarer
  - Electro-weak (W/Z), top physics
  - At 13 TeV and 2e34 $cm^{-2}s^{-1}$ expecting ~ 600 Hz of W(→ leptons)
- LHC was built to explore and search for even more rare physics
  - Higgs produced in about 1 out of 10^9 collisions with the detection rate being even lower
  - e.g. at 13 TeV and 2e34 $cm^{-2}s^{-1}$ expect ~ 0.01 Hz of ttH production

Saving all events at LHC is not useful (even if we could)!

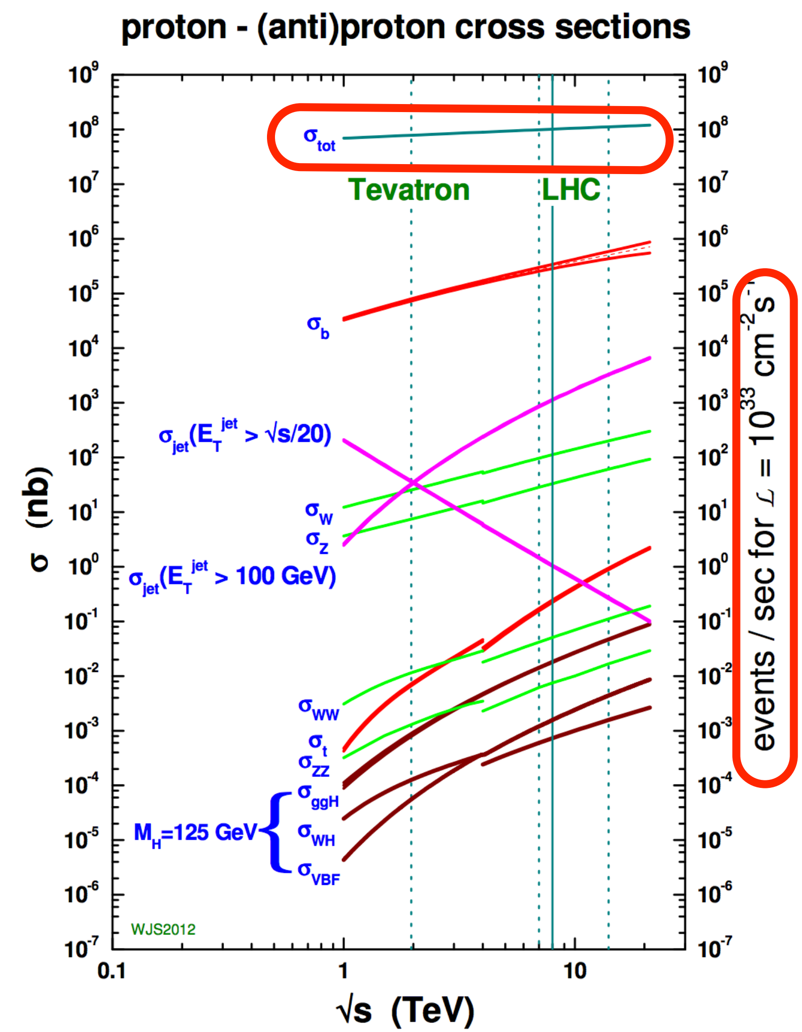

James Stirling

Trigger & Data Acquisition (TDAQ) System

- Recording the interesting physics is a challenge
  - **ATLAS detector is BIG**
    - ~100 million channels
    - Up to 2.2 MB of RAW data per event (dependent on running and recording conditions)
  - **Rate of delivered collisions is high**
    - In 2018 ~30 MHz measurement rate

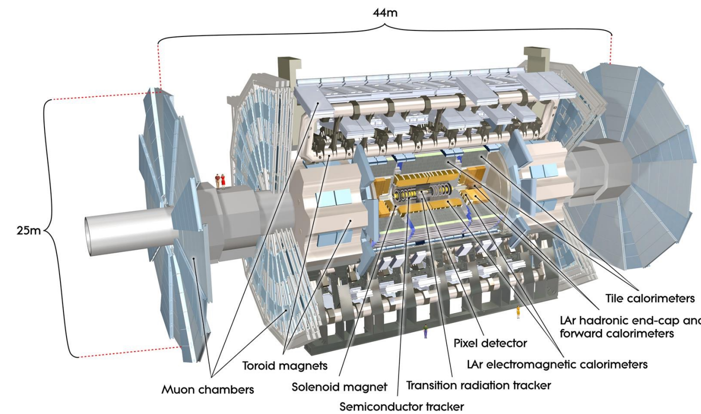

- **Data Acquisition (DAQ)** is responsible for
  - collecting data from detector systems (detector read-out),
  - digital conversion and
  - recording them to mass storage for offline analysis (data flow)
- **Trigger** is responsible for **real-time (online) selection** of the subset of events to be recorded

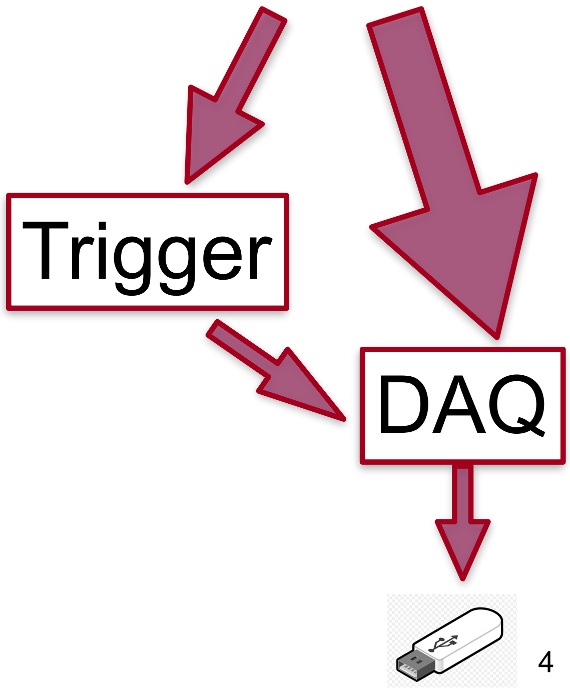

The ATLAS Trigger System

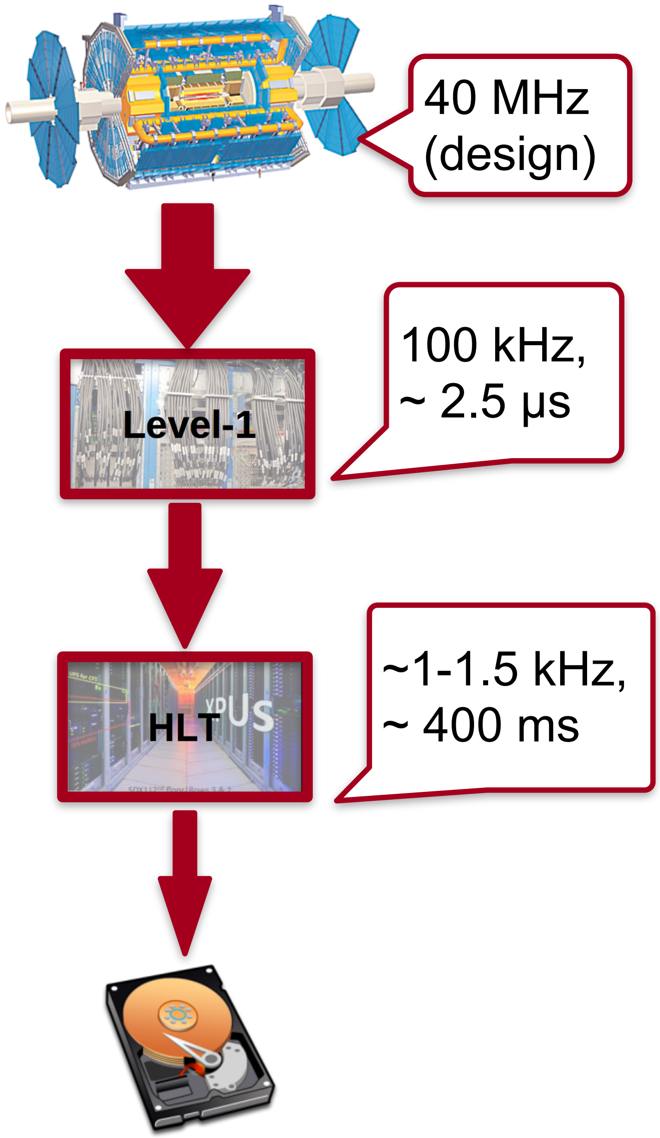

- **Level-1 (L1)**
  - Hardware-based
  - Coarse selection based on limited input from calorimeter & muon systems
  - Rate and latency limit set by detector & trigger hardware
- **High-Level Trigger (HLT):**
  - Software-based
  - Average processing time limited by HLT farm size
    - Commodity hardware; ~40k processing applications
  - Networking based on commercial technologies (Ethernet)

Run 1 → Run 2 → Run 3

- Evolution of the TDAQ system somewhat dictated by evolution of LHC performance

Peak	√s [TeV]	Peak luminosity [$cm^{-2}s^{-1}$]	Peak pile-up
Run 1	8	~ 7 x 10^{33}	~ 35
Run 2	13	~ 2 x 10^{34}	~ 60 (~ 80 in 2017)
Run 3	13 - 14	~ 2 x 10^{34} (luminosity levelling)	~ 60

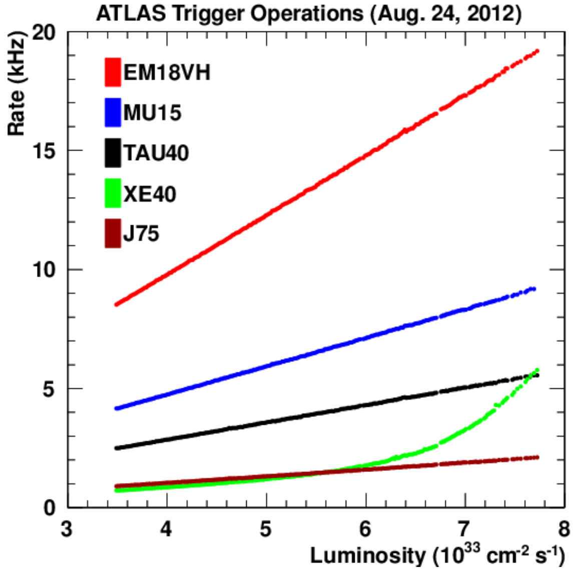

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

- In Run 2, trigger rates increased by a factor of ~ 5
  - A factor of ~2 due to the energy increase
  - A factor of 2-3 due to the luminosity increase

- Options to cope with the increase in trigger rates:
  - Increase output rate → challenge for offline computing
  - Increased trigger thresholds → potential loss of interesting physics
  - Increased trigger rejection / identification power → **improved hardware/software**

Run-3 ATLAS TDAQ System

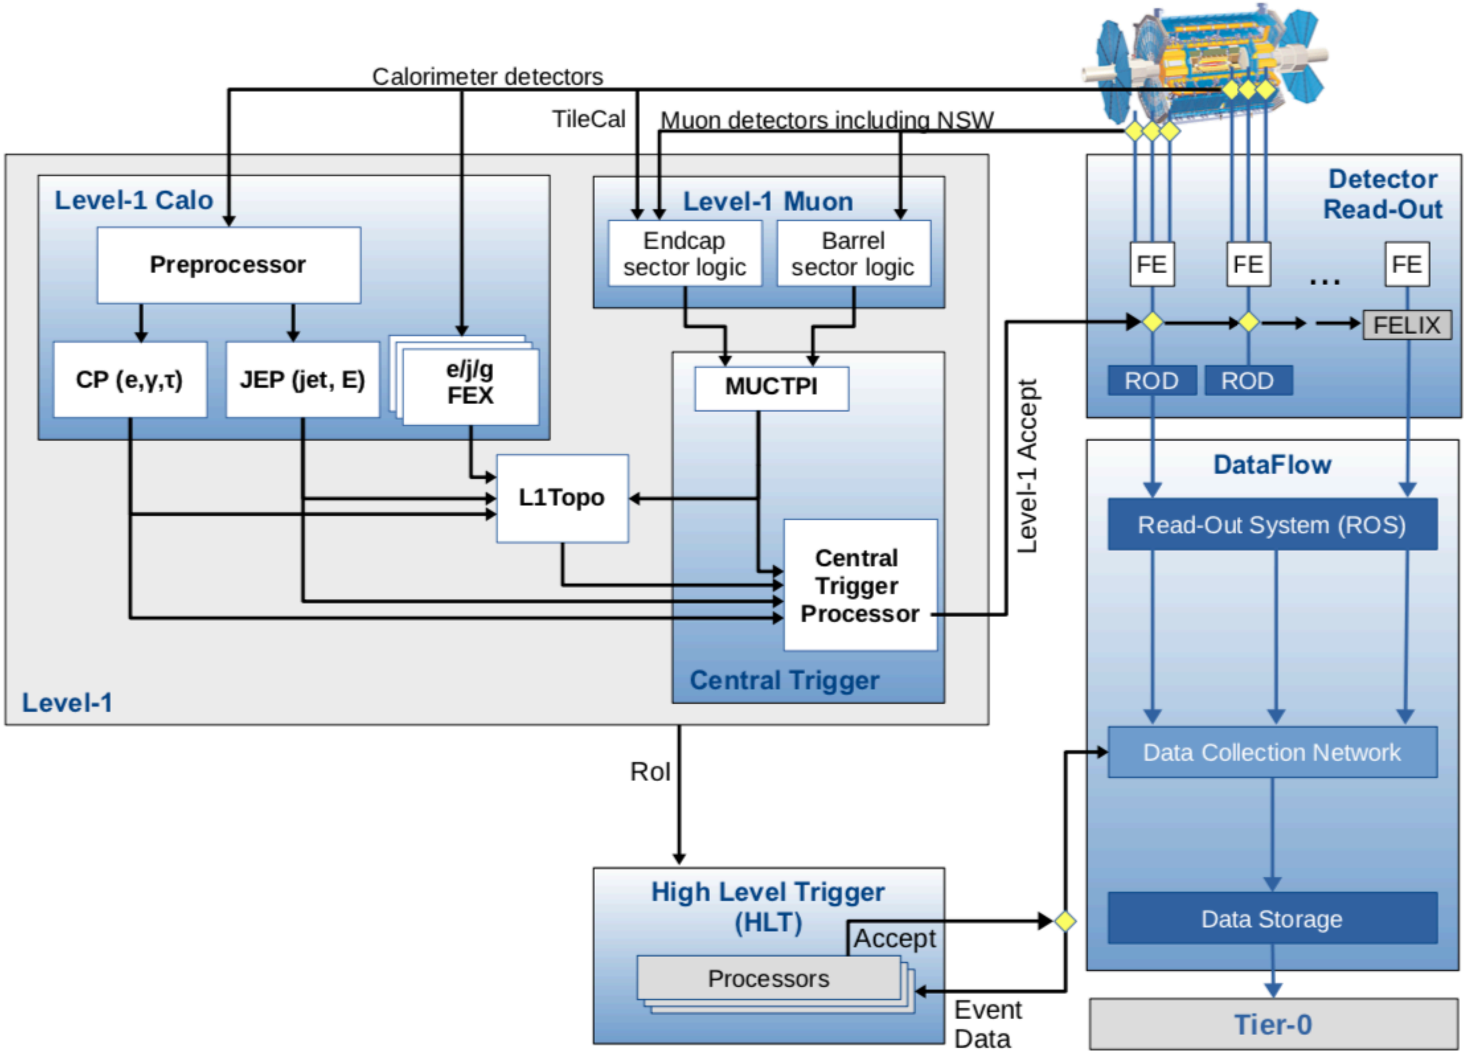

ATLAS TDAQ System - with FTK

SLAC

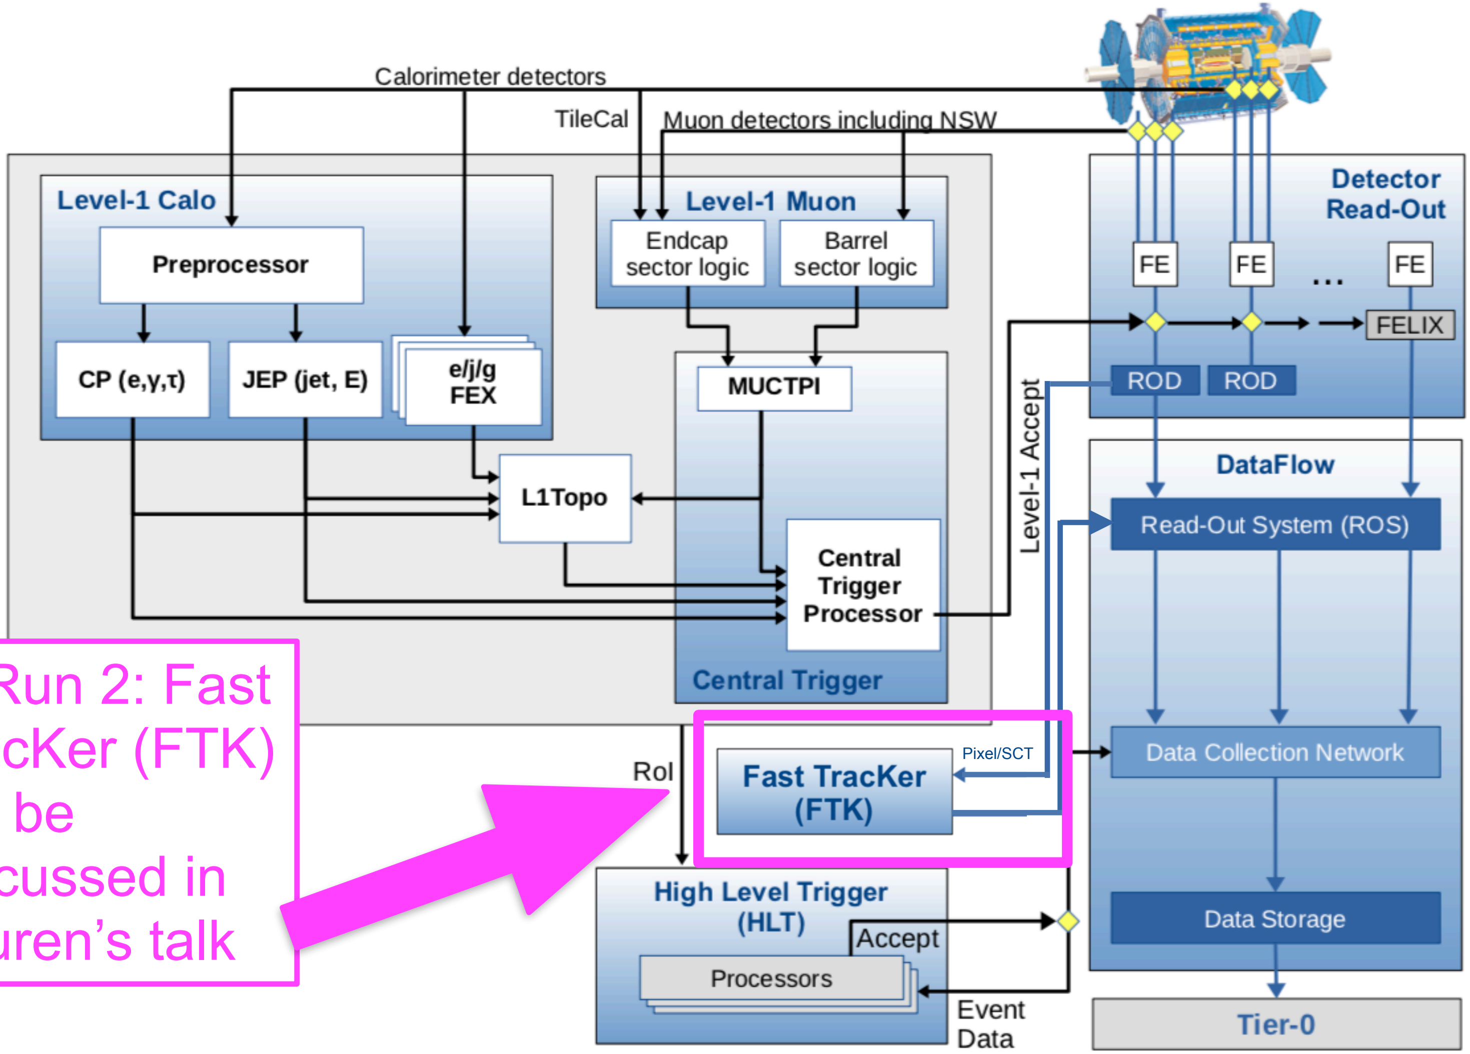

In Run 2: Fast TracKer (FTK) will be discussed in Lauren's talk

Run-3 ATLAS TDAQ System - Level-1 Calo

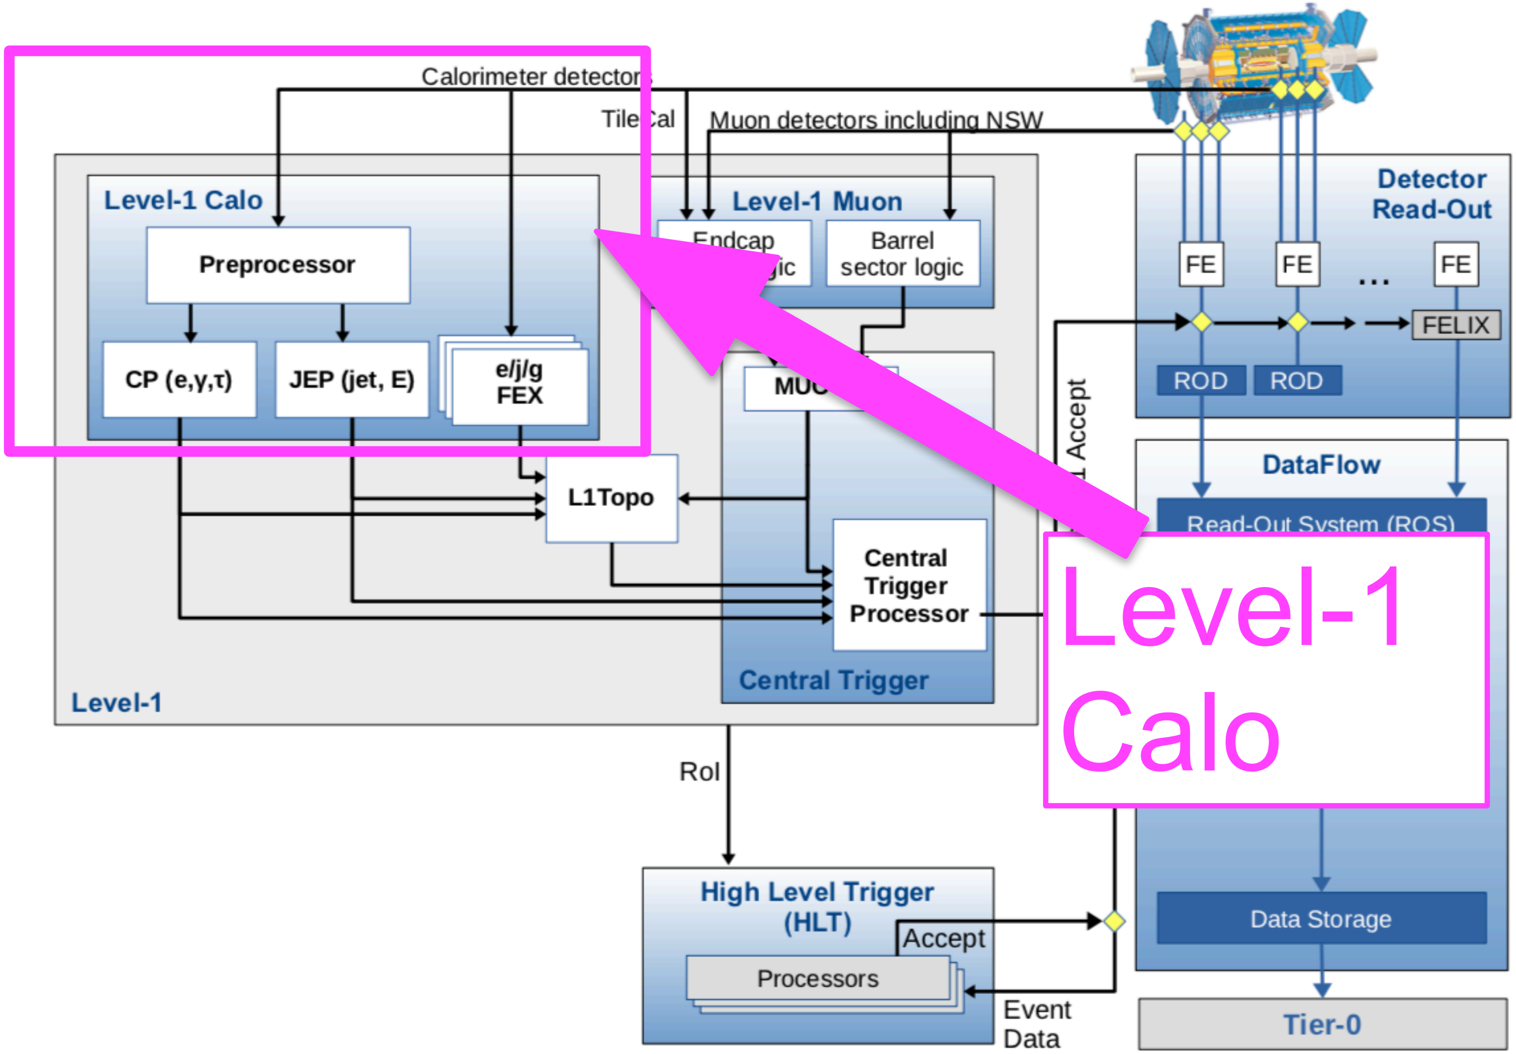

Level-1 Calorimeter Trigger

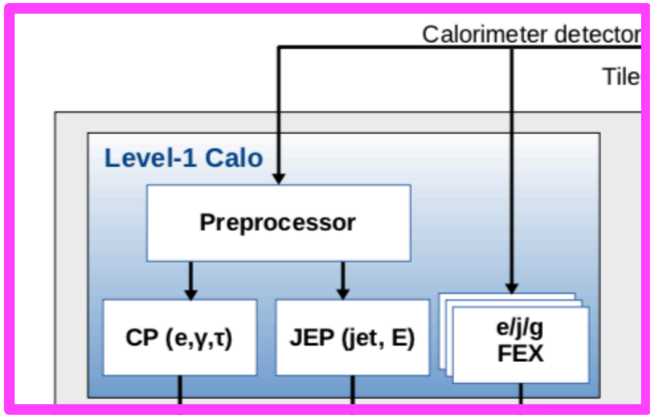

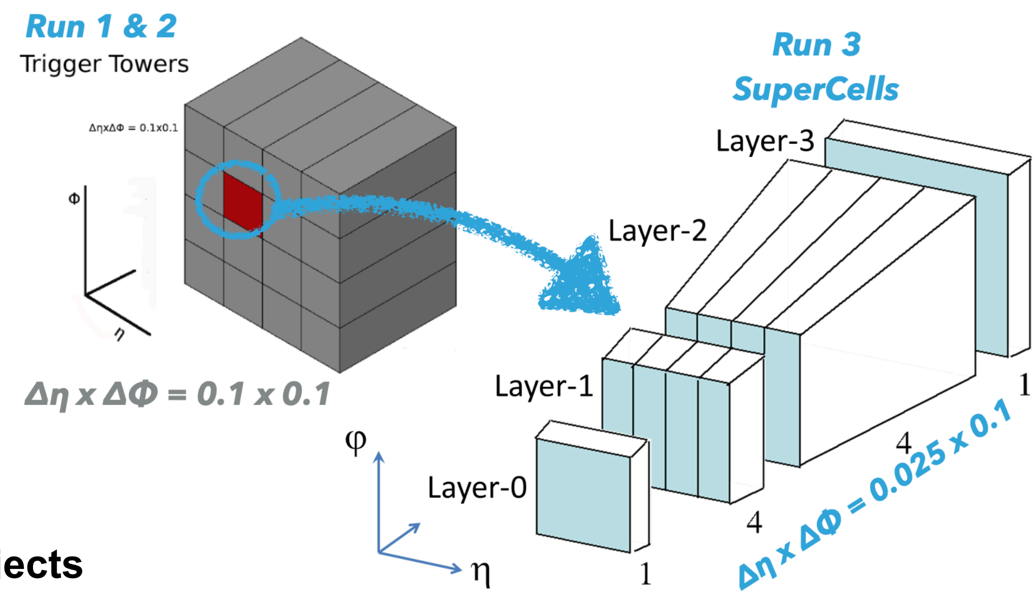

ATL-DAQ-SLIDE-2020-310

- Identifies **calorimeter-based physics objects**
  - Electrons, photons, jets, hadronically decaying taus
  - Global event quantities (missing transverse energy (MET, E_T sums)

- Upgrade for Run 3:
  - Higher granularity LAr calorimeter inputs (**SuperCells**) provide improved resolution compared to TriggerTowers in Run 1/2

Level-1 Calorimeter Trigger

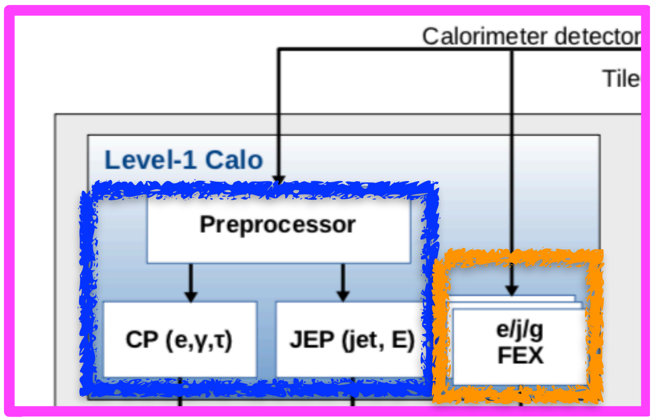

- New **ATCA-based Feature EXtractors (FEX)**
  - More sophisticated algorithms
  - Improved isolation and pileup robustness
  - Reduced rates while keeping thresholds low
- Run-2 and Run-3 systems will run in parallel during commissioning

- **electron FEX (eFEX)** to identify isolated electron/photon and tau objects
  - Significant rate savings at L1 with similar efficiency compared to Run-2 performance
- **jet FEX (jFEX)** to identify jets, MET, E_T sums, hadronically decaying taus
  - Noise suppression and pileup subtraction algorithms to mitigate impact of pileup on ME_T and multi-jet triggers
- **global FEX (gFEX)** to exploit algorithms for computing global event quantities

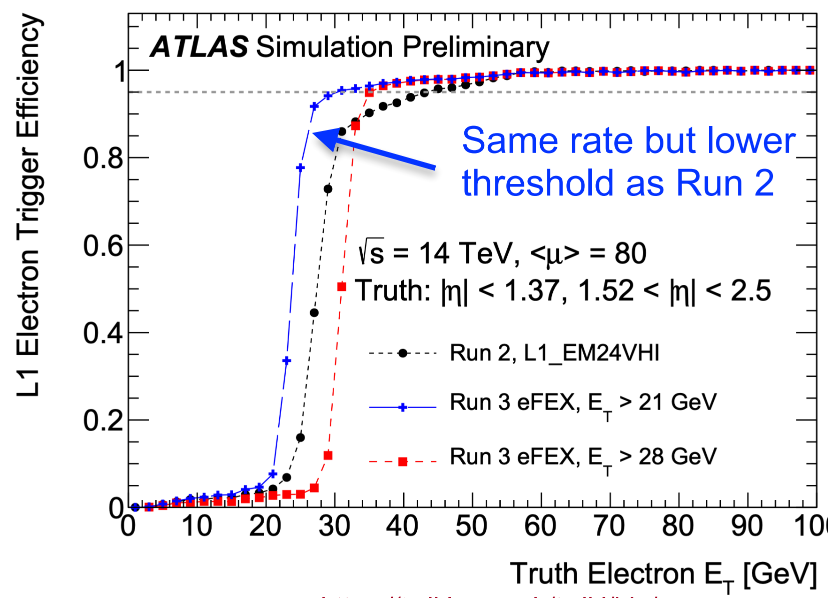

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults

Run-3 ATLAS TDAQ System - Level-1 Muon

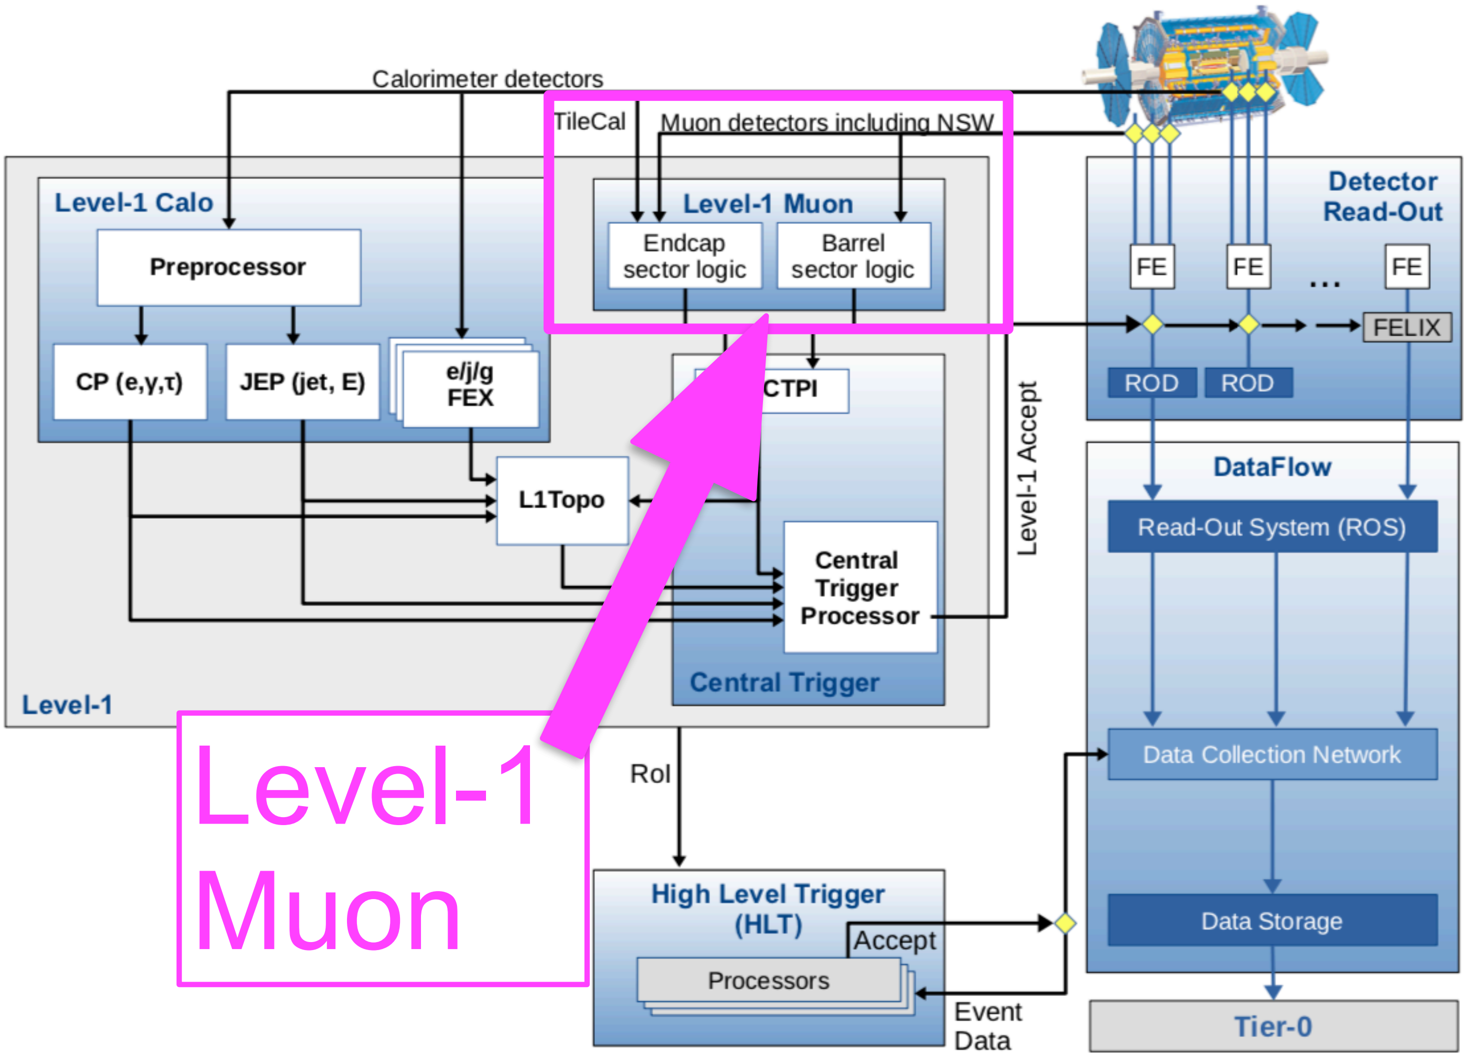

Level-1 Muon Trigger

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

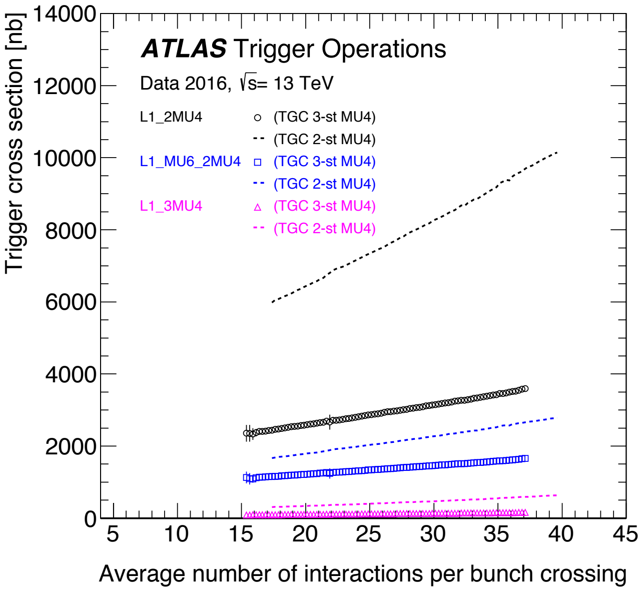

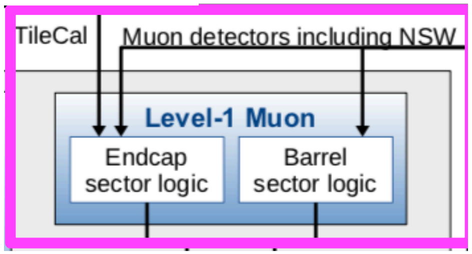

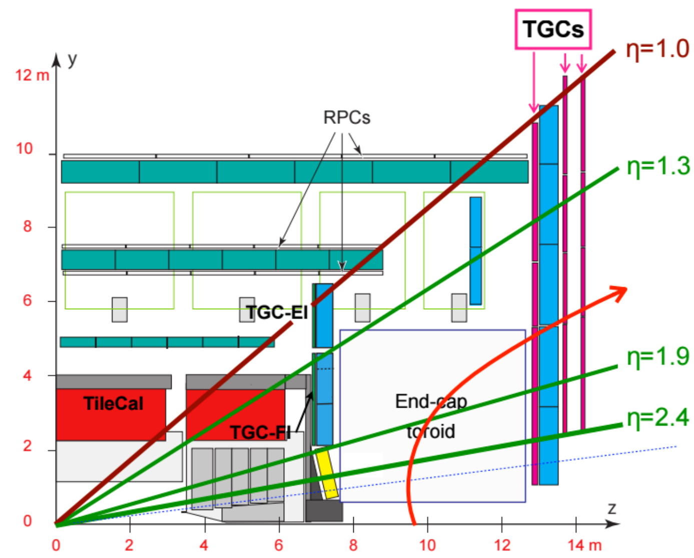

- **Barrel (**Resistive Plate Chambers**, |η| < 1.05)** and **Endcaps (**Thin Gap Chambers**, 1.05 < |η| < 2.4)**
- Coincidences to reduce fakes and therefore rate
  - Barrel: Two- (low-p_T) and three-station (hight-p_T) coincidence triggers
  - **Coincidence between TGC chambers** to reduce fakes (up to 60%) from proton background
  - Additional coincidence with tile calorimeter to reject slow charged particles (typically protons)
- No p_T measurement; Only threshold passed (and multiplicities)
- **New in Run 3: New Small Wheel (NSW) (endcap)**
  - L1 rate reduction expected already with only one working NSW while maintaining same efficiency

Level-1 Muon Trigger

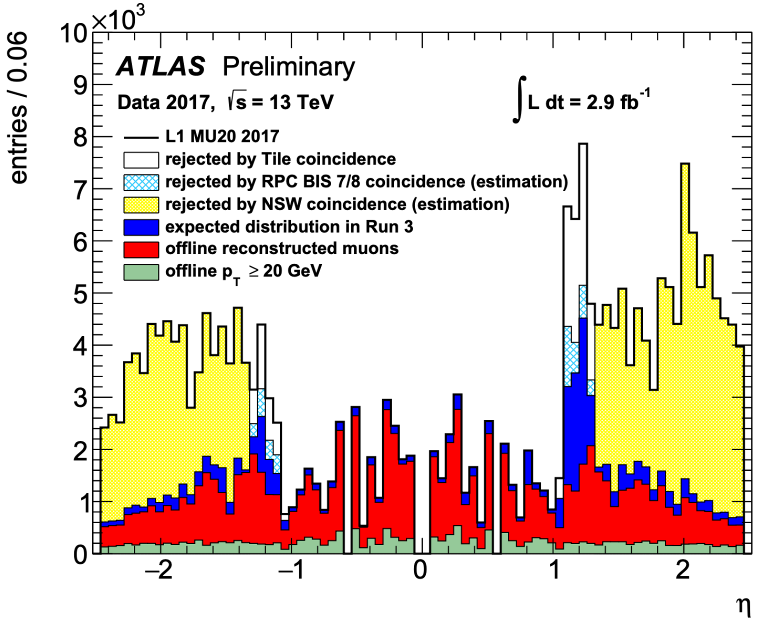

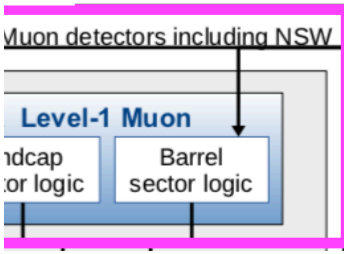

https://twiki.cern.ch/twiki/pub/AtlasPublic/L1MuonTriggerPublicResults

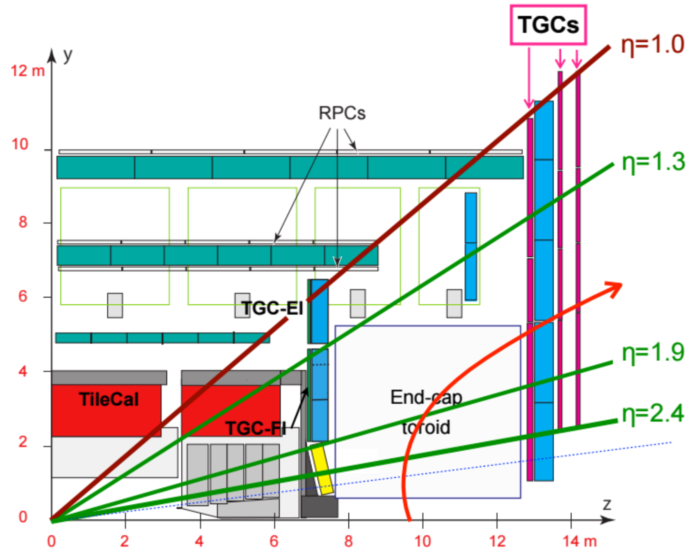

- **Barrel (**Resistive Plate Chambers**, |η| < 1.05)** and **Endcaps (**Thin Gap Chambers**, 1.05 < |η| < 2.4)**
- Coincidences to reduce fakes and therefore rate
  - Barrel: Two- (low-p_T) and three-station (hight-p_T) coincidence triggers → rate reduction with only minimal efficiency loss
  - Coincidence between TGC chambers to reduce fakes (up to 60%) from proton background
  - Additional coincidence with tile calorimeter to reject slow charged particles (typically protons)
- No p_T measurement; Only threshold passed (and multiplicities)
- **New in Run 3: New Small Wheel (NSW) (endcap)**
  - L1 rate reduction expected already with only one working NSW while maintaining same efficiency

Run-3 ATLAS TDAQ System - L1Topo and CTP

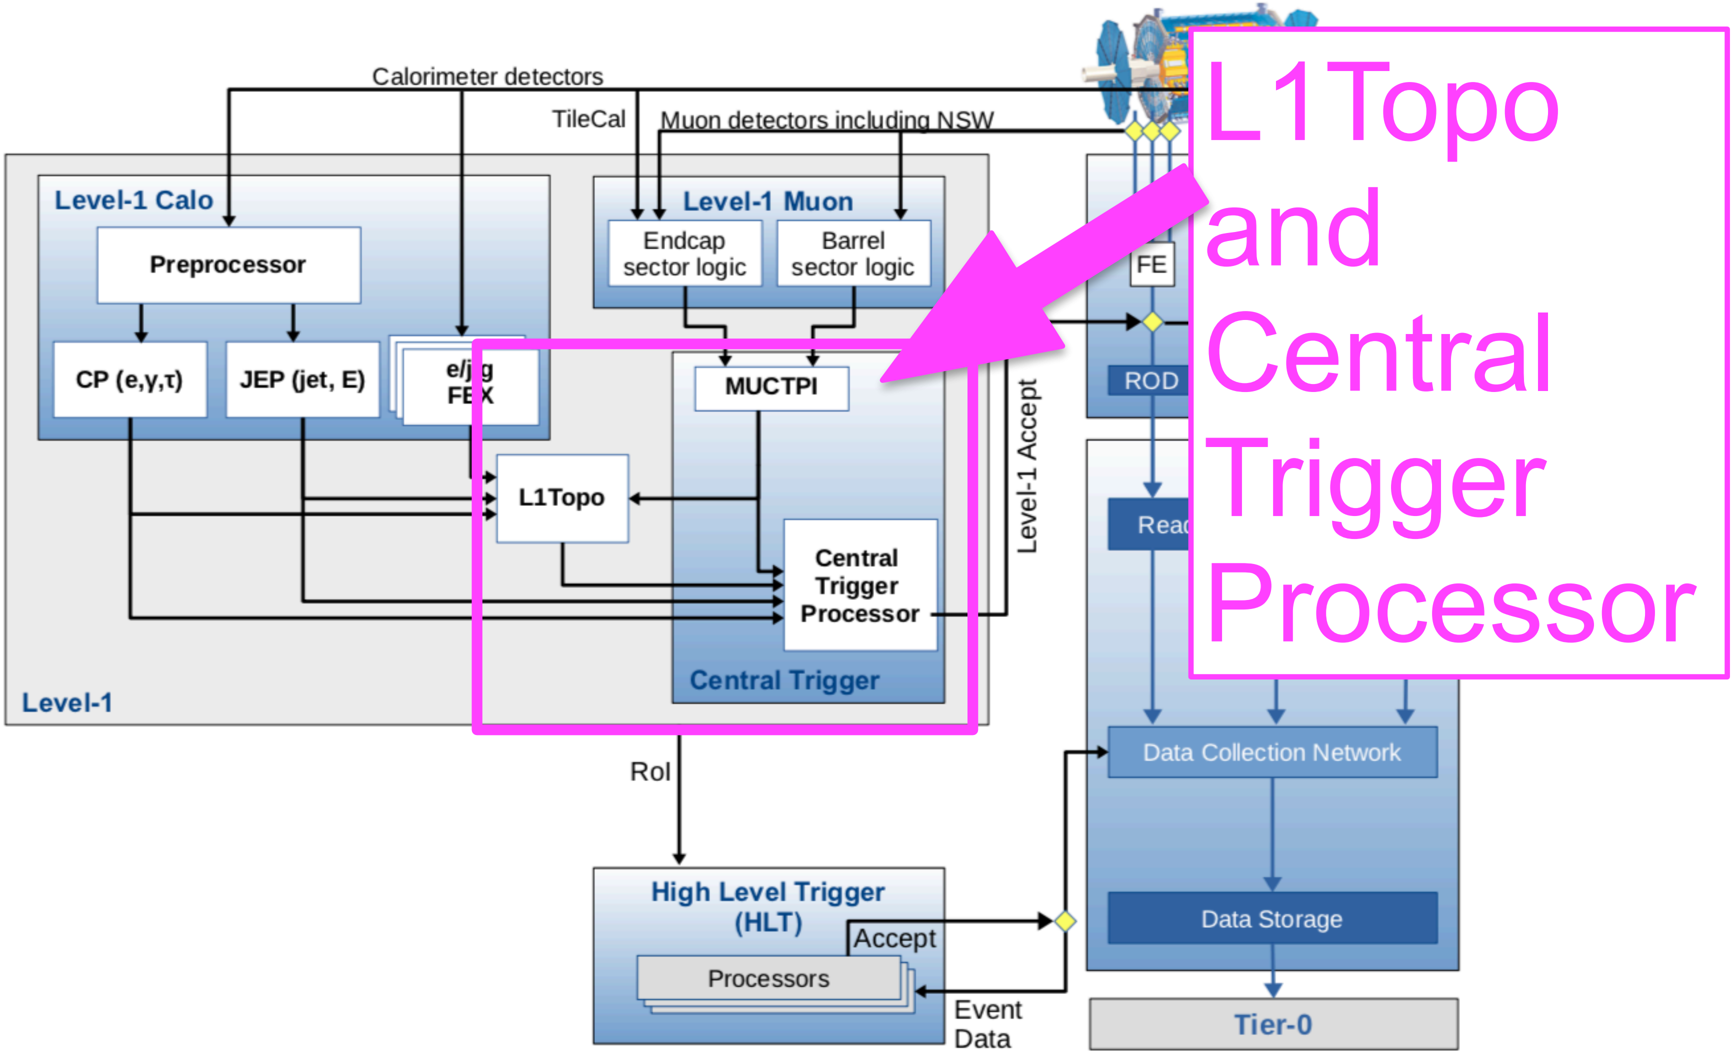

L1Topo and Central Trigger Processor

- **L1Topo** combines information from L1Calo and **MUCTPI** (Muon Central Trigger Processor Interface) into variables that are used for L1 selections (topological, angular, kinematic selections, sums of quantities)
  - Significant rate reduction, increased signal purity without impact in physics acceptance

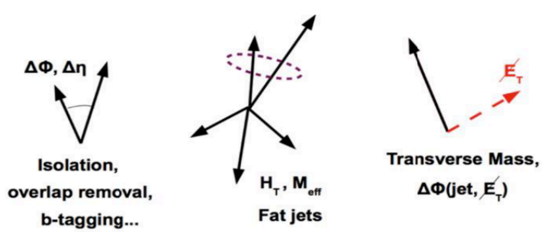

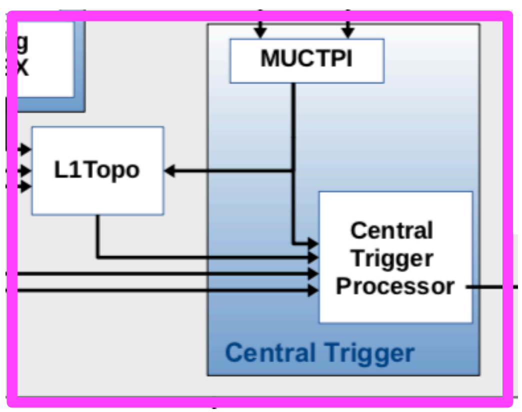

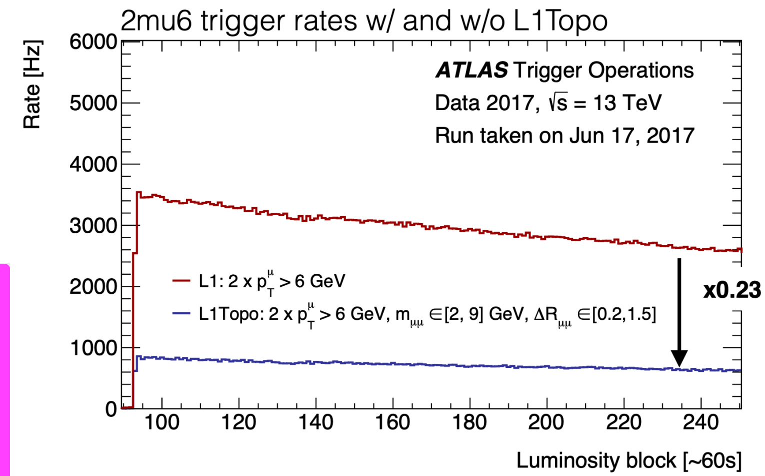

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

- **Central Trigger Processor (CTP)**
  - Time-alignment of trigger signals taking the inputs from the various L1 sub-systems
  - Apply multiplicities, logical selections and topological selections
  - **Final trigger decision**
  - Applies prescales and bunch groups

Run-3 ATLAS TDAQ System - High Level Trigger

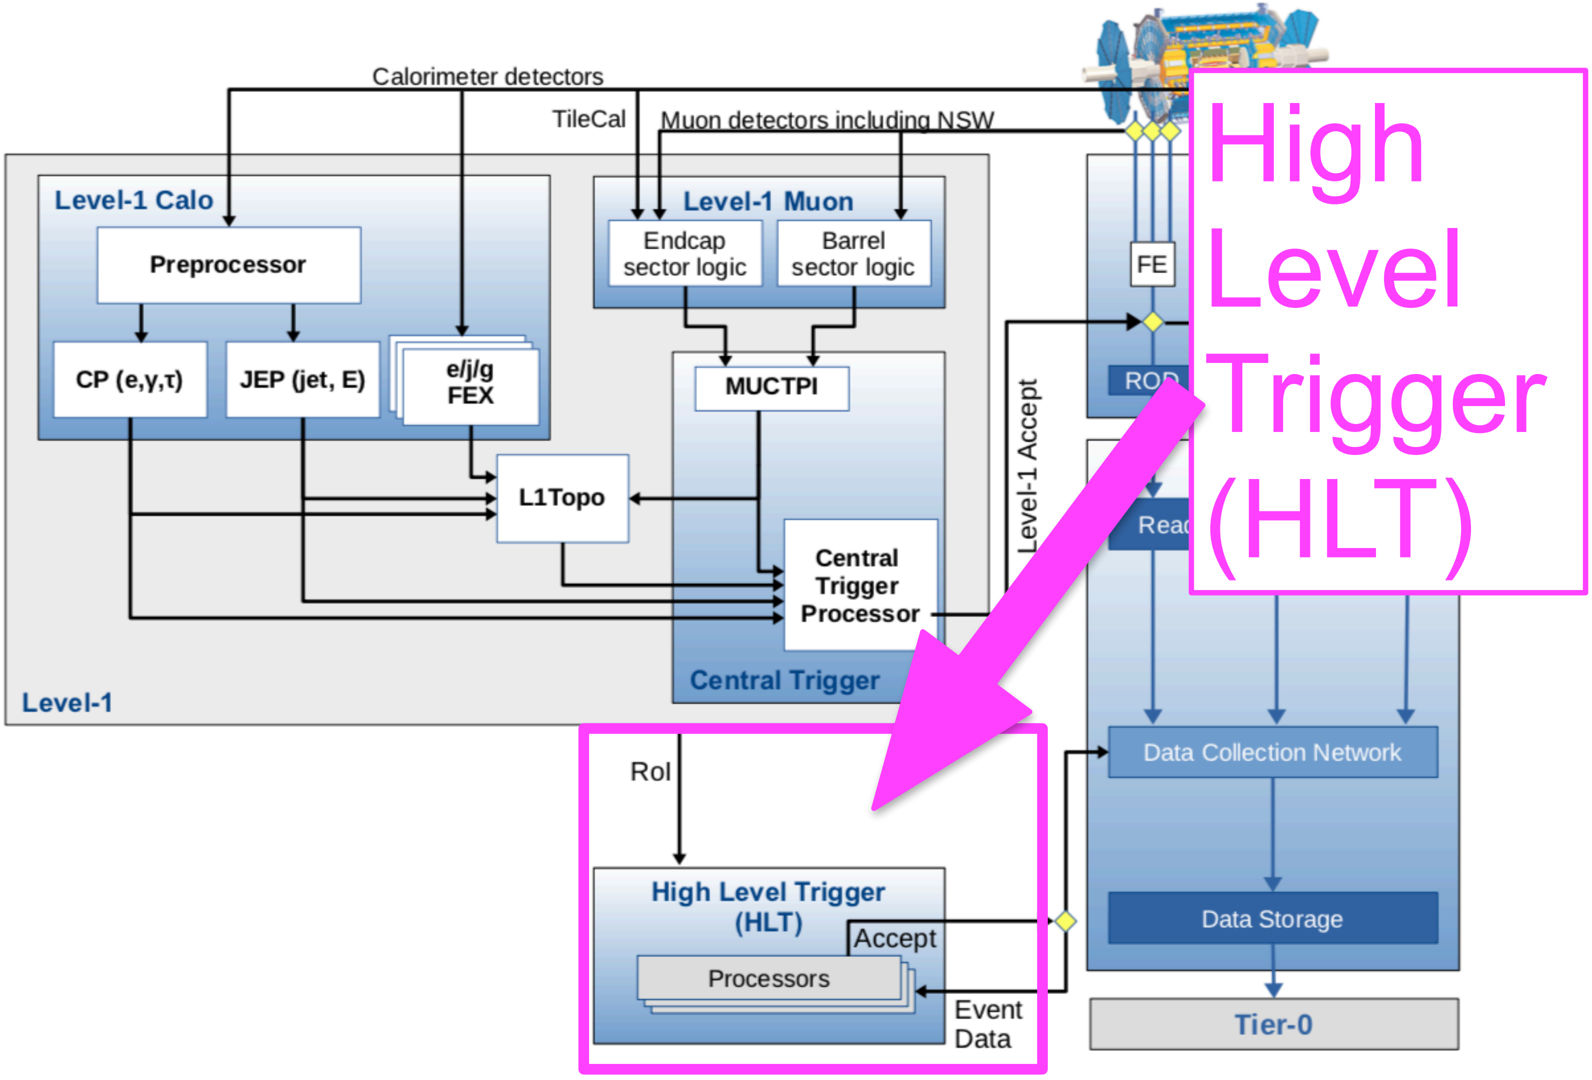

High-level Trigger

- **High-level Trigger (HLT)**
  - Input from Level-1 is regional information, the Region-of-Interest (RoI)
    - Geometrical region in η x Φ with information about type of object (EM, MU, TAU,…) and thresholds passed (p_T, E_T)
    - RoI is the place where more computationally/time expensive reconstruction algorithms can be run, e.g. tracking (more on this slide)
  - Decision about the **final event accept** is made

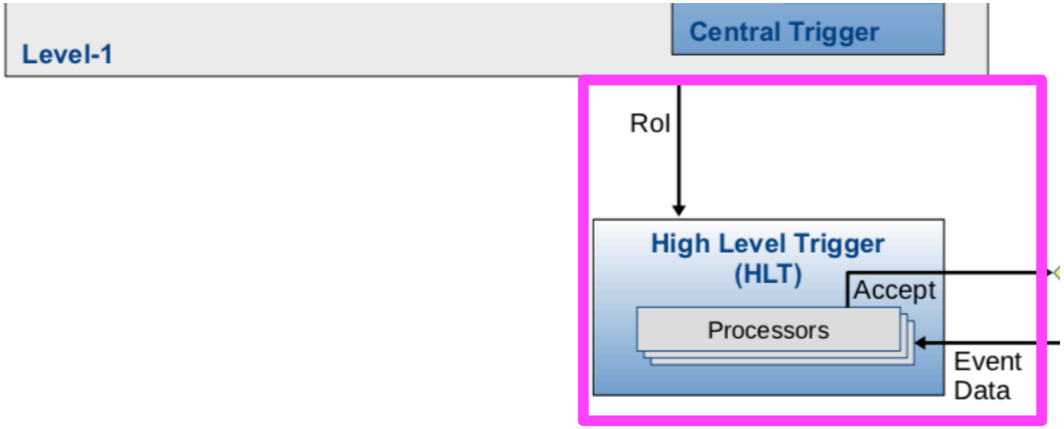

Reconstruction Algorithms

- To stay within HLT rate & processing time limits, the reconstruction of trigger signatures is made up of **several steps** to achieve **early rejection**
  - **Early rejection**: stop processing of algorithms as soon as a step fails
  - Fast reconstruction
    - Often trigger-specific algorithms
    - Either guided by Region of Interests (RoIs) or full detector detector (e.g. jets, MET)
    - e.g. tracking: computationally/readout/network traffic intense to run tracking on complete event, tracking is only run on data from RoIs
  - Precision reconstruction
    - (Very similar) algorithms used also in offline reconstruction
    - Full detector information available

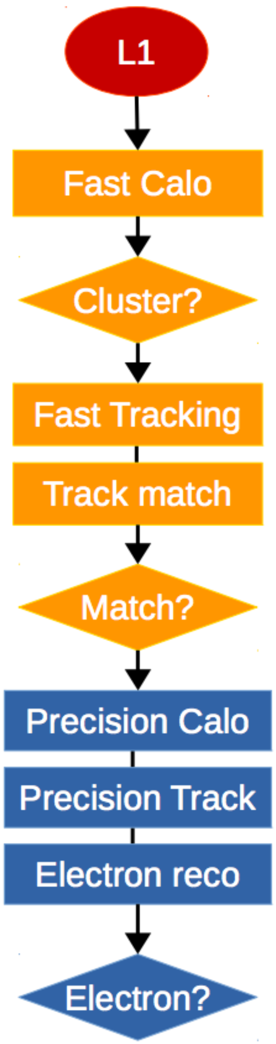

- Types of algorithms:
  - *Feature Extraction*: builds objects (tracks, clusters, …)
  - *Hypothesis*: apply selection cuts (track pT, invariant mass, …)

More info can be found in these talks:
ATL-DAQ-SLIDE-2019-738
ATL-DAQ-SLIDE-2020-336

Run-3 HLT Framework

SLAC

- ATLAS is redesigning its core framework for native, efficient and user-friendly multi-threading support → **AthenaMT**
  - Three kinds of parallelism
    - Inter-event: multiple events are processed in parallel
    - Intra-event: multiple algorithms can run in parallel for an event
    - In-algorithm: algorithms can utilize multi-threading and vectorization
- **Trigger framework is also being reimplemented for Run 3**
  - Will profit from redesign to make more use of offline reconstruction software
  - HLT requirements (partial event reconstruction in RoIs and early rejection) have been considered during initial design-phase of the AthenaMT framework
    - Replacing scheduling and change by native Gaudi Scheduler which is also used offline

Read-Out and Data Flow

- **L1 accept** causes **front-end (FE)** detector electronics to transfer relevant data to the **Read-Out Drivers (RODs)**
  - Detector-specific custom hardwares (mainly VMEbus)
  - Perform initial data processing and formatting
- After ROD stage, data is sent via optical link to the **Read-Out System**
  - First common stage of the DAQ system
  - Data is buffered in custom PCIe I/O card (RobinNP)
- The **Data Collection Network** handles all I/O on the HLT nodes, including RoI requests from the HLT and full event building over 10 Gb Ethernet network
- Events accepted by the HLT are sent to **Data Storage** for packaging and transfer to permanent storage offline

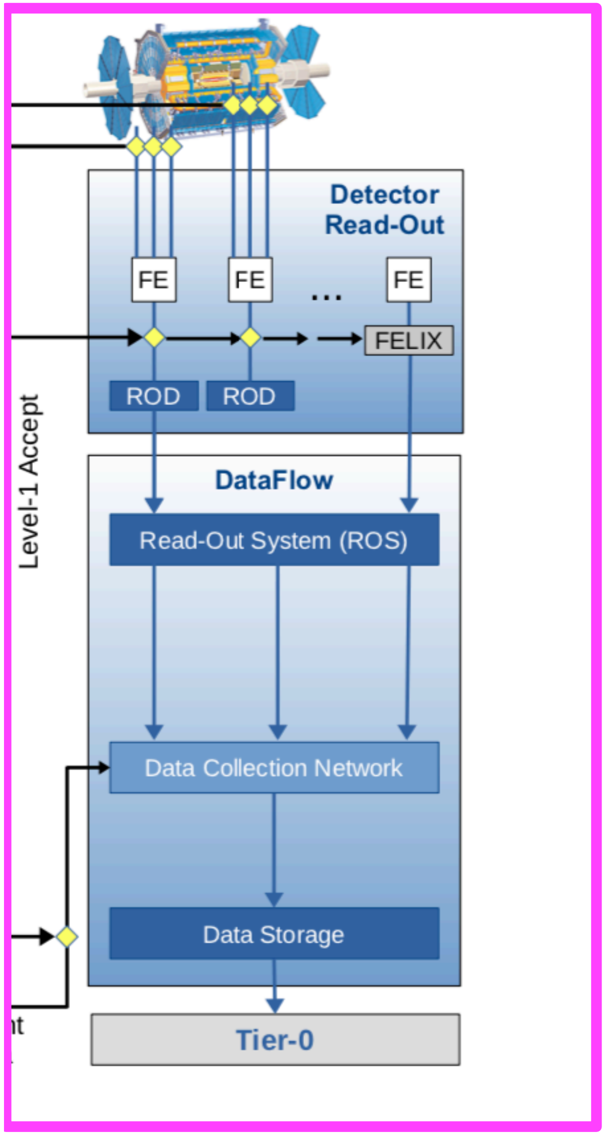

What Objects to Trigger on and How

Trigger Menu

- What to record is defined by the **physics program** of ATLAS
  - Look at what *signatures* the final state (high-p_T) objects of the desired physics leave in the detector
  - Can also look at global event quantities (total energy, missing transverse momentum/energy)
  - Each physics signature can be reconstructed in several steps in form of a *trigger chain / line*
  - *Trigger Menu* is the collection of these trigger chains
  - A typical ATLAS trigger menu is rather complex, contains several hundreds of chains
    - ~ 500 in Run 1,
    - ~ 1800 - 2000 in Run 2
  - Trigger Menu varies with **luminosity and time** (fine-tuning according to running conditions), and is dependent on **rate limitations** at trigger levels and **online resources** (computational needs, bandwidth)

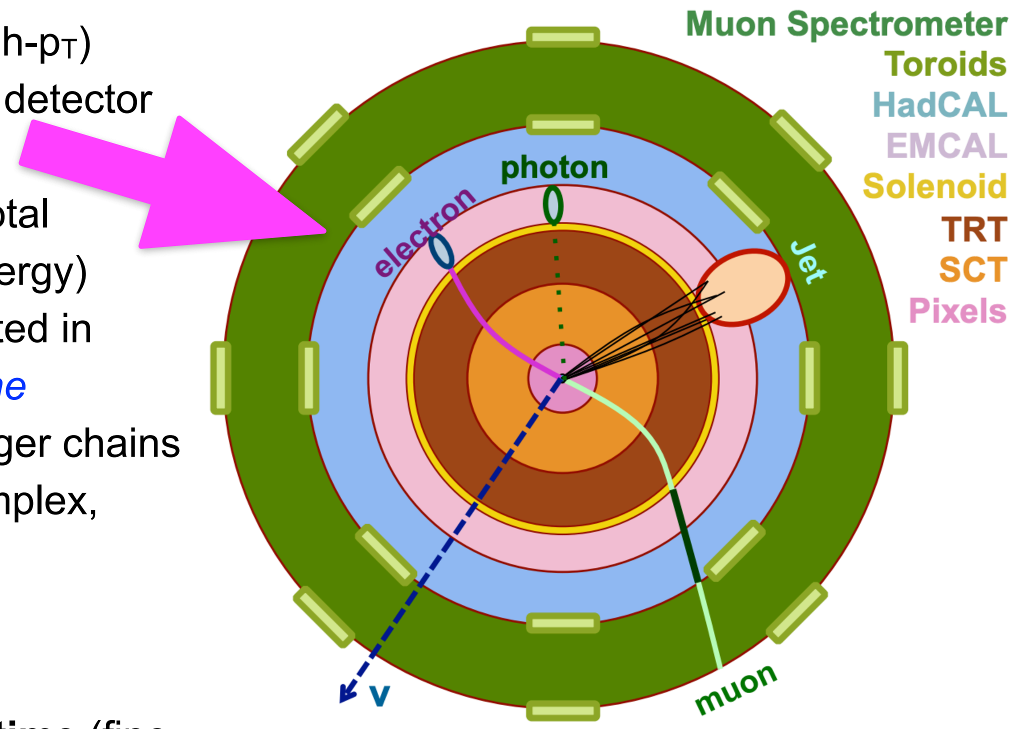

A. Sfyrla, Summer Student Lectures 2018

Trigger Menu at 2e34

- Trigger menu is documented in ATL-DAQ-PUB-2019-001
  - Highest rates at HLT for single muon/electron, di-tau, MET
  - Same p_T thresholds for electron/muon at HLT
- Also documented dedicated menus for heavy-ion, low μ datasets (μ~ 2) and other special runs

Trigger	Typical offline selection	Trigger Selection		L1 Peak Rate [kHz]	HLT Peak Rate [Hz]
		L1 [GeV]	HLT [GeV]	L=2.0×10^{34} cm^{-2}s^{-1}	
Single leptons	Single isolated μ, $p_T > 27$ GeV	20	26 (i)	16	218
	Single isolated tight e, $p_T > 27$ GeV	22 (i)	26 (i)	31	195
	Single μ, $p_T > 52$ GeV	20	50	16	70
	Single e, $p_T > 61$ GeV	22 (i)	60	28	20
	Single τ, $p_T > 170$ GeV	100	160	1.4	42
	Two μ, each $p_T > 15$ GeV	2×10	2×14	2.2	30
	Two μ, $p_T > 23, 9$ GeV	20	22, 8	16	47

ATL-DAQ-SLIDE-2020-320

Trigger Chains

- Not all trigger chains need to run at full rate
  - Rate might be just too high
  - Often just a subsample is enough
  - Can add triggers when luminosity drops to make optimal use of resources
- *Prescales* are used to reduce rate
  - Prescale of N (e.g. N=10): Only accept 1 out of N events

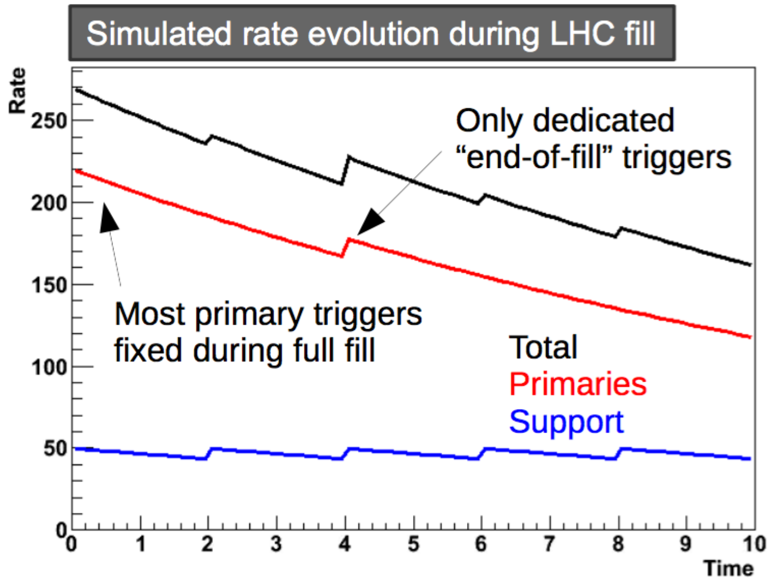

- Trigger chains are classified into
  - Primary/physics chains: chains for physics signals in general (unprescaled)
  - Backup chains: higher thresholds and/or tighter selections (e.g. in case of unexpected luminosity increase)
  - Supporting & background chains: to collect data for auxiliary measurements in physics analyses (e.g. data-driven background extraction, measuring trigger efficiencies), usually prescaled
  - Alternative triggers: using different selection algorithms
  - Monitoring and calibration chains: to monitor the data quality (e.g. to check the performance of tracking by the inner detectors) & calibration purposes for detectors

Streaming

- If any trigger chain passes, events are accepted and are written out to **streams**
- Events can be written out to different streams depending on which trigger was passed

Debug Streams
Events for which a trigger decision has not been made, because of failures in parts of the online system

Physics Streams
Data for physics analyses

Express Stream
Full events for fast reconstruction

Trigger-Level Analysis Streams
Contains reconstructed HLT objects

Calibration / Monitoring Streams
Events with minimum of information for detector calibrations at high rate

Delayed Streams
Not promptly reconstructed, only when resources allow it

Streaming - Debug Streams

- If any trigger chain passes, events are accepted and are written out to **streams**
- Events can be written out to different streams depending on which trigger was passed

Debug Streams
Events for which a trigger decision has not been made, because of failures in parts of the online system

- Decrease in events written to Debug Streams throughout Run-2
- Most events recovered and included in physics analyses

arXiv:2007.12539

Stream	Number of Events 2015	2016	2017	2018
Physics	1 694 555 330	5 387 420 813	5 649 311 254	6 400 342 575
Total Debug	411 878	188 860	18 197	1507
Total Debug w.r.t. Physics	2.4×10^{-4}	3.5×10^{-5}	3.2×10^{-6}	2×10^{-7}
Recovered Events	402 671	187 944	18 001	1 455
Recovered Events w.r.t. Total Debug	97.8%	99.5%	98.9%	96.5%

Streaming

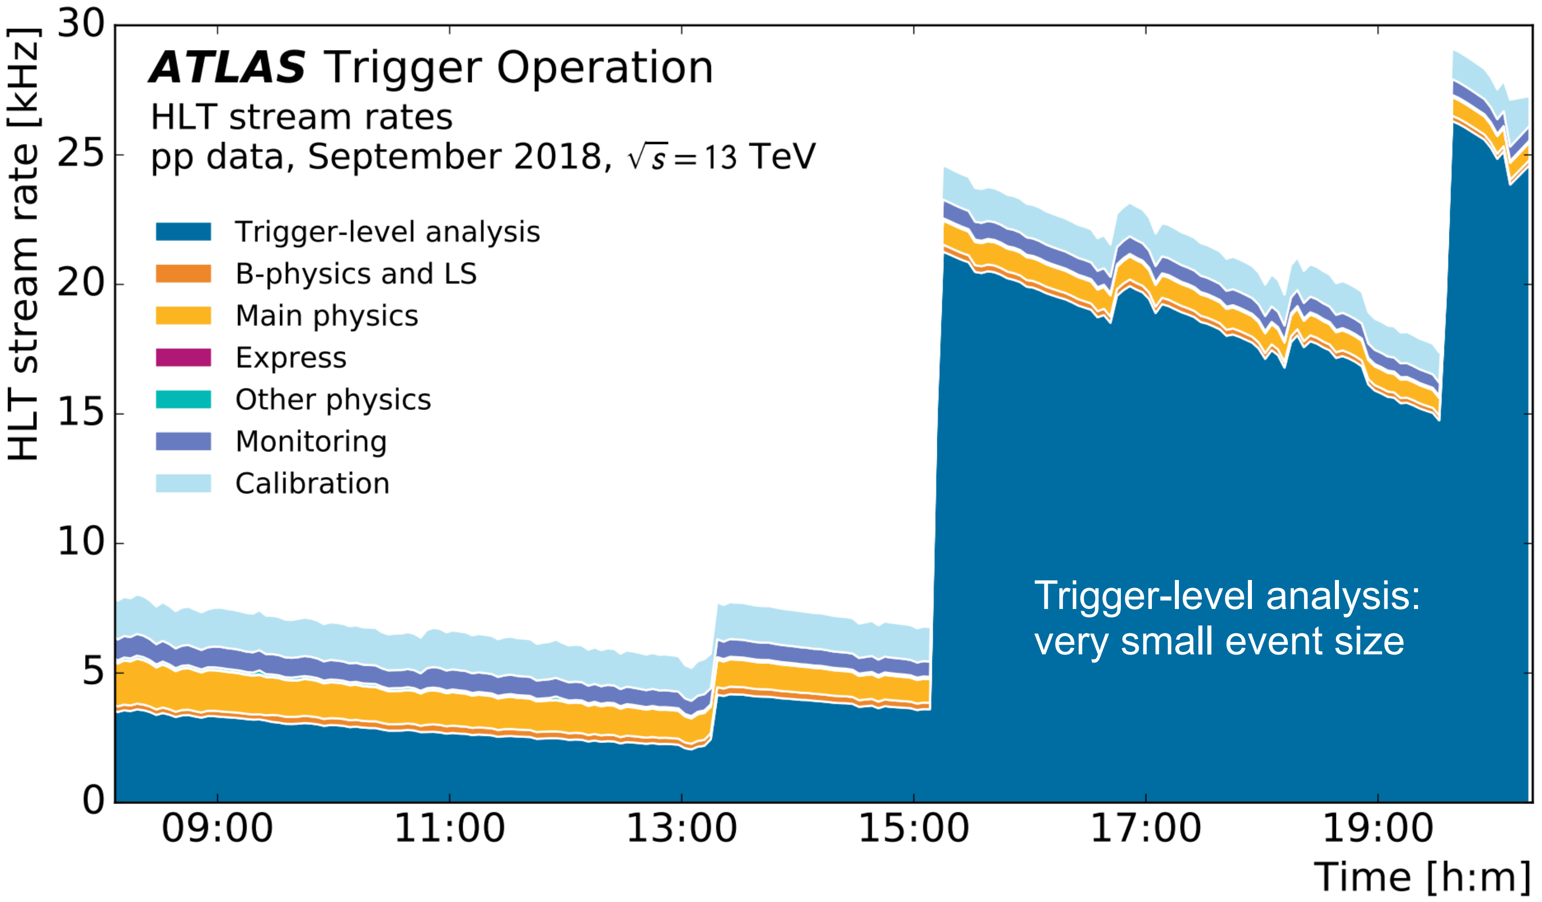

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

Summary

- Experience from Run 1 and Run 2 has shown that the ATLAS TDAQ system is able to efficiently record data while dealing with various constraints and challenging conditions
  - Evolution of the TDAQ system in terms of hardware and software important to maintain physics acceptance and efficiency
  - Improvements for Run 3 are focussing on new L1 hardware and improved HLT algorithms
    - Upgraded L1Calo and L1Muon system
    - Moving closer to offline reconstruction through AthenaMT
  - Versatile trigger menu to record data for a wide range of physics analyses
    - Aim at exploiting the total bandwidth for physics even better and to extend the phase space for physics discovery

- Not shown here: Performance of TDAQ system throughout Run 2, please see https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerPublicResults
- For info on the the HL-LHC ATLAS TDAQ system, see ATLAS-TDR-029

Backup

The ATLAS Detector

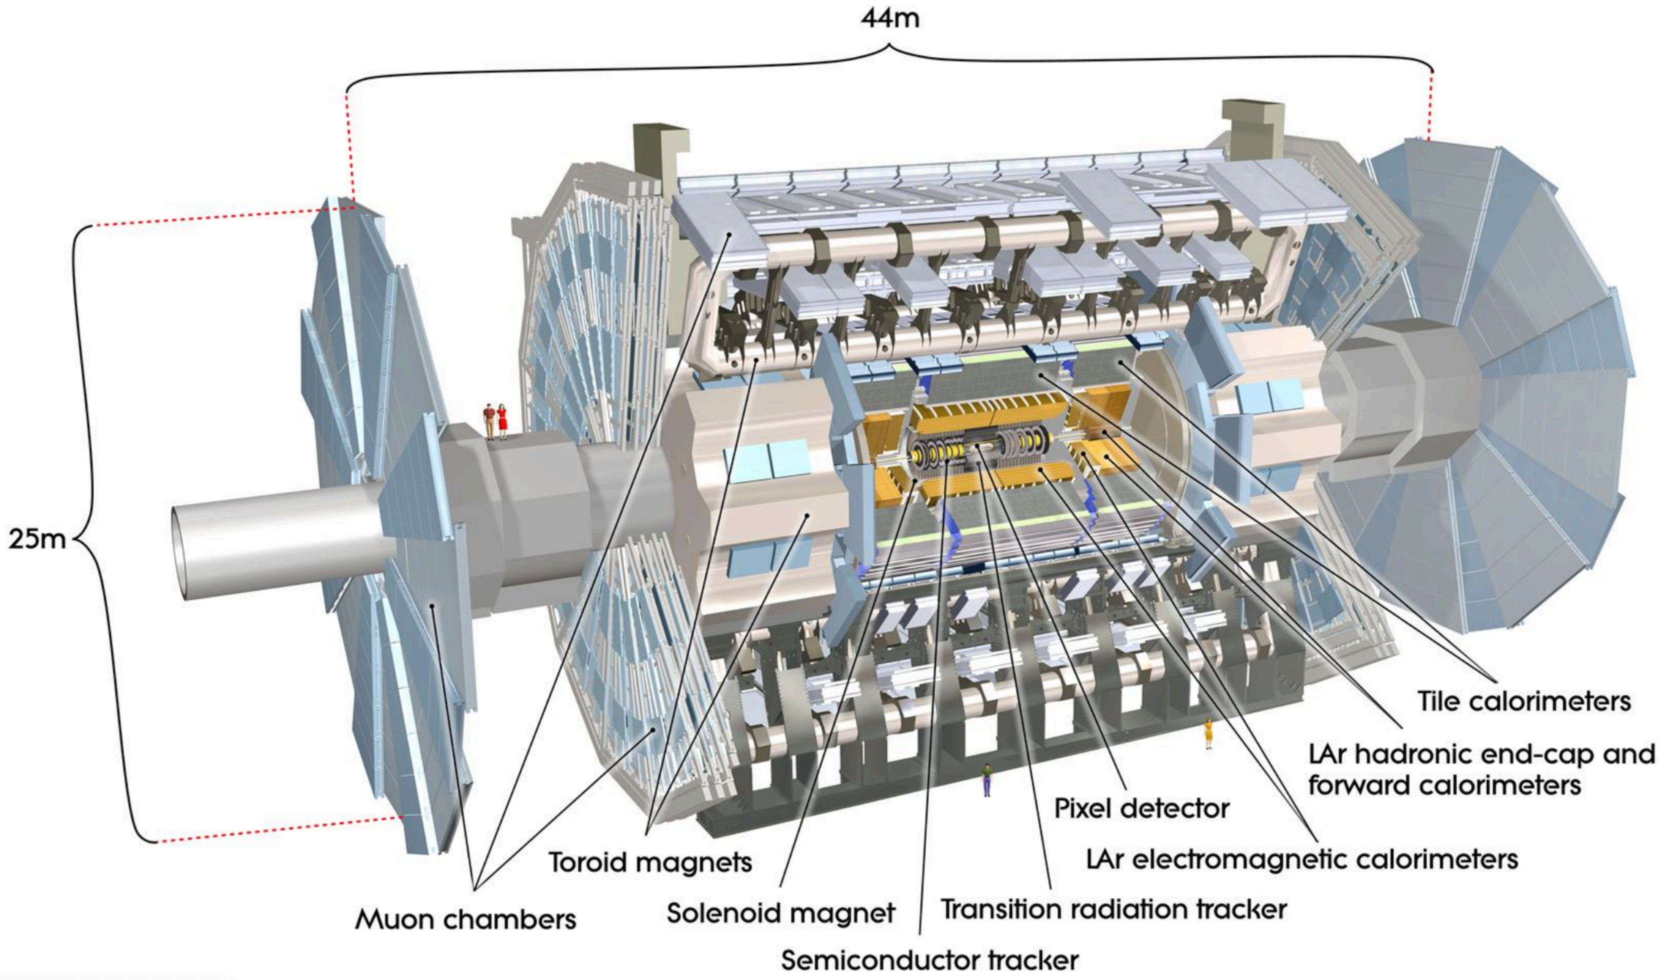

Identifying the "Interesting" Events

- Proton collisions produce mainly hadrons with low transverse momentum
  - Only 2% of all tracks have $p_T > 2$ GeV

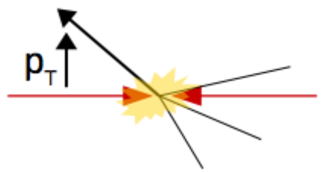

- Interesting physics is usually **high-p_T**
  - $H \rightarrow \gamma\gamma$, $p_T(\boldsymbol{\gamma}) \sim 50\text{-}60$ GeV
  - $W \rightarrow e\nu$, $p_T(\mathbf{e}) \sim 30\text{-}40$ GeV
  - Obvious signatures to use in the trigger
    - Single e/μ triggers used in most analyses

- What if new physics is "soft"?
  - This is where triggering becomes a challenge
  - Upgraded and new features in TDAQ system as well as ideas necessary!

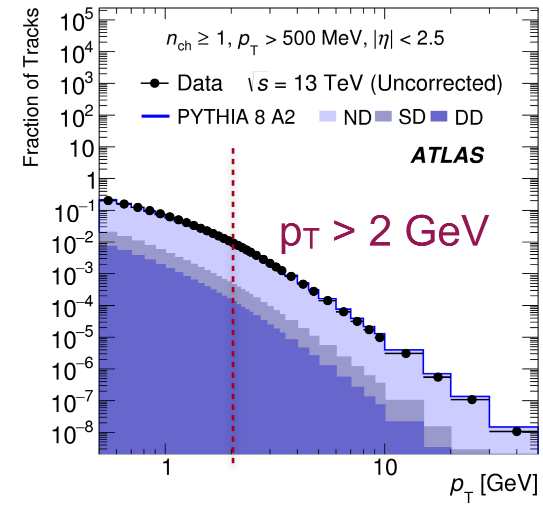

arXiv:1602.01633

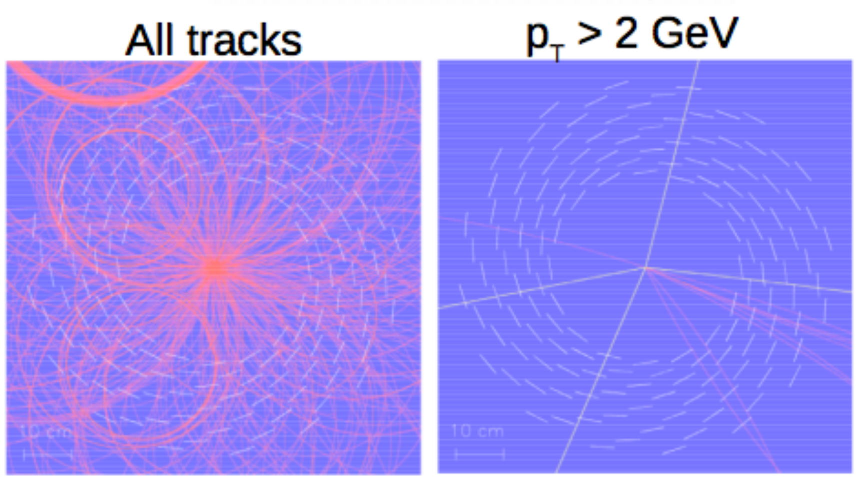

Simulated H→4μ + 17 minbias events

Run 1 → Run 2 → Run 3

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

SLAC

- The ATLAS TDAQ system operated successfully in Run 1 (2009-1012) and 2 (2015-2018)

Peak	√s [TeV]	Peak luminosity [$cm^{-2}s^{-1}$]	Peak pile-up
Run 1	8	~ 7 x 10^{33}	~ 35
Run 2	13	~ 2 x 10^{34}	~ 60
Run 3	13 - 14	~ 2 x 10^{34} (luminosity levelling)	~ 60

- In Run 2, trigger rates increased by a factor of ~ 5
  - A factor of ~2 due to the energy increase
  - A factor of 2-3 due to the luminosity increase
- Options to cope with the increase in trigger rates:
  - Increase output rate → challenge for offline computing
  - Increased trigger thresholds → potential loss of interesting physics
  - Increased trigger rejection / identification power → **improved hardware/software**

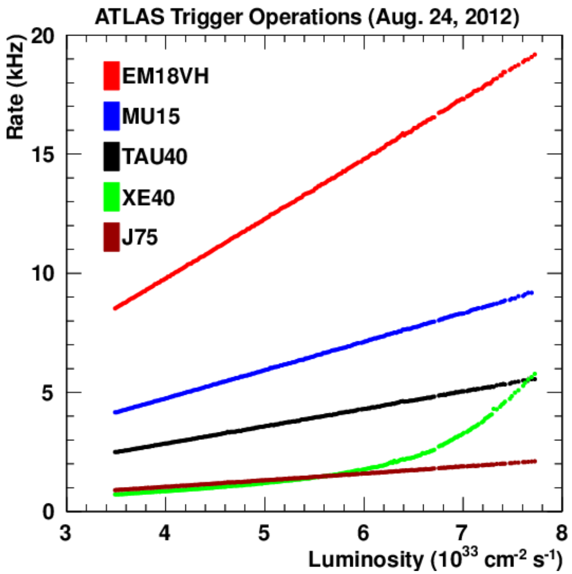

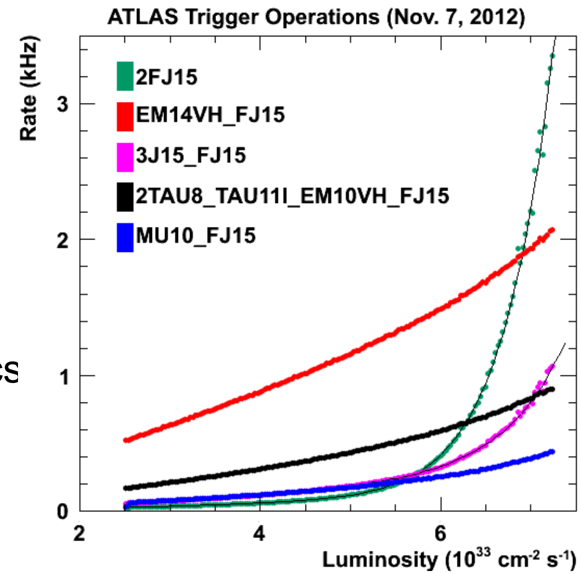

L1Calo - Run 2

- **For triggering on any EM object**
  - **Electrons, Photons, Jets, Taus**
  - **Global event quantities (MET)**

- Run 2 Pre-processor
  - Several **calorimeter cells are summed into trigger towers**
  - Resulting in towers of reduced granularity, e.g. η x Φ = 0.1 x 0.1
    - ~7000 calorimeter trigger towers
- Run 2 Cluster Processor (CP), Jet-Energy Sum Processor (JEP): **Object reconstruction for e, γ, τ, jet, (E)**
  - Find local maximum via sliding window algorithms
    - Apply energy selection based on sum in towers
    - Window size depending on object
    - Electron/Photon 0.2 x 0.2, Jets 0.4 x 0.4
  - Can apply additional selections
    - EM isolation (ring around core)
    - Hadronic isolation (no activity in hadronic layer)

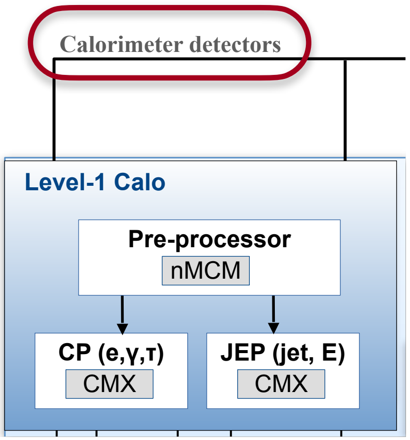

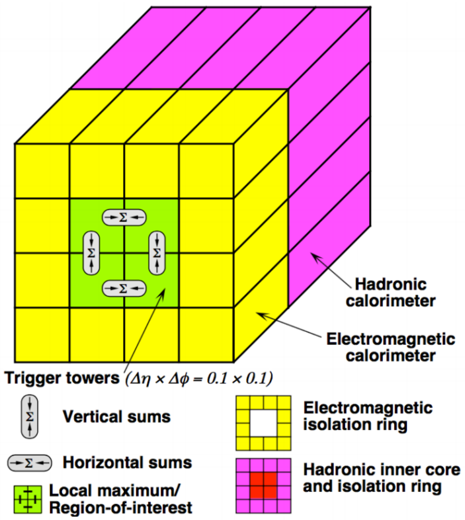

Level-1 Muon

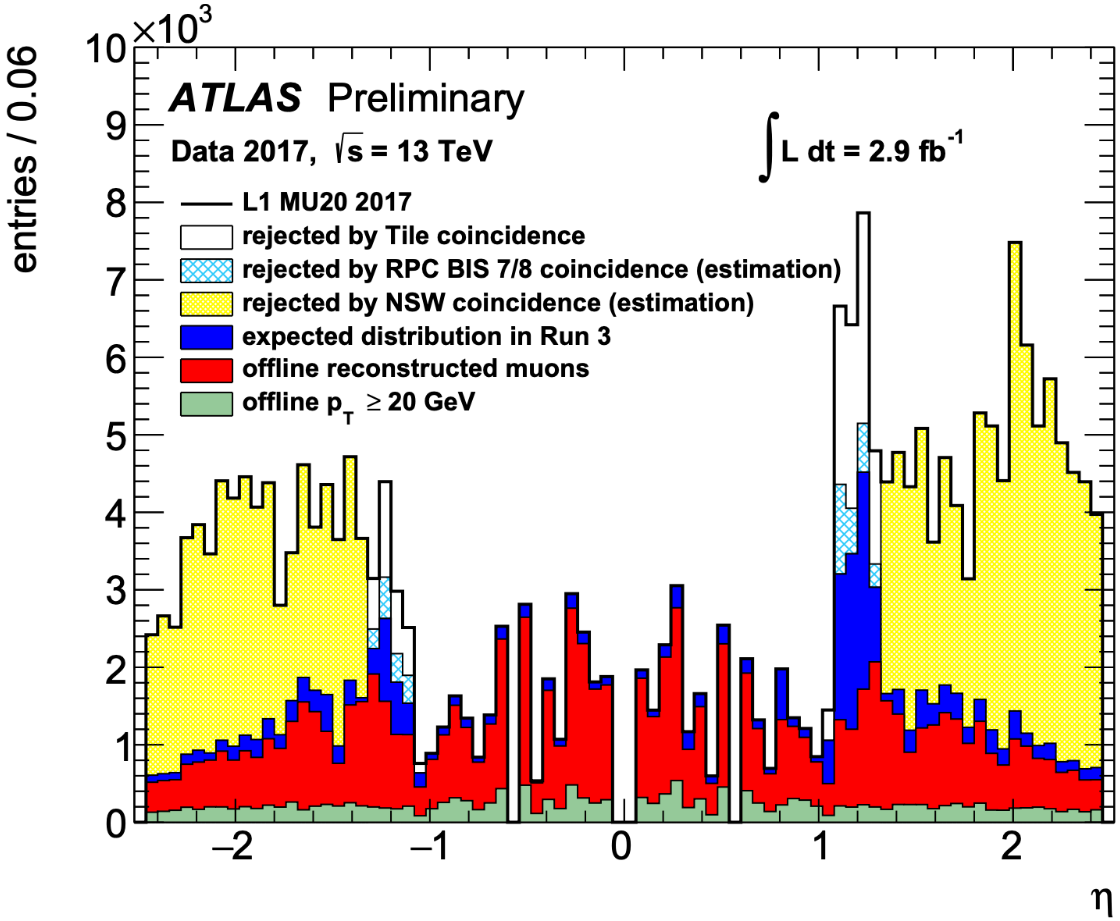

https://twiki.cern.ch/twiki/pub/AtlasPublic/L1MuonTriggerPublicResults

Predicting Trigger Rates and Computing Needs

- Rate and computational needs can be predicted for triggers making use of Enhanced Bias (EB) data
  - EB data: dataset with O(1M) events, enhanced in high-p_T objects, selected using only L1 triggers
  - By knowing the L1 prescales, the selection bias can be removed with event weights
  - This provides an unbiased sample with sufficient statistical precision in the high-p_T/ high-multiplicity regime
- Rate and computational needs predictions are possible through reprocessing a new trigger menu with the enhanced bias dataset

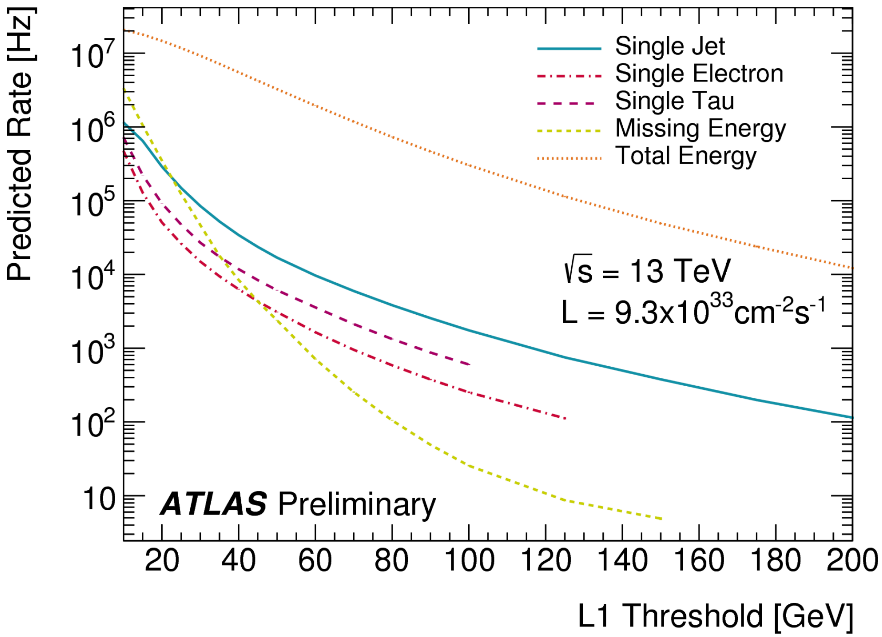

ATL-DAQ-PUB-2016-002

arXiv:2007.12539

Trigger Operation

- Many aspects to consider for and during data-taking:
  - Operating & maintaining the trigger system
    - Shifter and expert on-call organization
  - Planning & getting system ready for intensity ramp ups and special runs
  - Commissioning & integration of trigger sub-systems
  - Dealing with day-by-day operational challenges and limitations
    - Rate limitations due to sub-detectors in certain conditions
    - HLT farm performance
    - Debug stream treatment
    - LHC plan changes
  - Testing, validation & deployment of new menus & software
  - Monitoring during data-taking
  - Data Quality assessment

Year	Dataset	Trigger DQ Eff.		ATLAS DQ Eff. [%]	Integrated Luminosity
		L1 [%]	HLT [%]		
2015	pp @ 13 TeV (50 ns)	100.00	99.94	88.77	84 pb^{-1}
	pp @ 13 TeV	99.97	99.76	88.79	3.2 fb^{-1}
2016	pp @ 13 TeV	98.33	100.00	93.07	33 fb^{-1}
2017	pp @ 13 TeV	99.95	99.96	95.67	44 fb^{-1}
2018	pp @ 13 TeV	99.99	99.99	97.46	59 fb^{-1}

JINST 15 (2020) P04003

Run 3 and Run 4

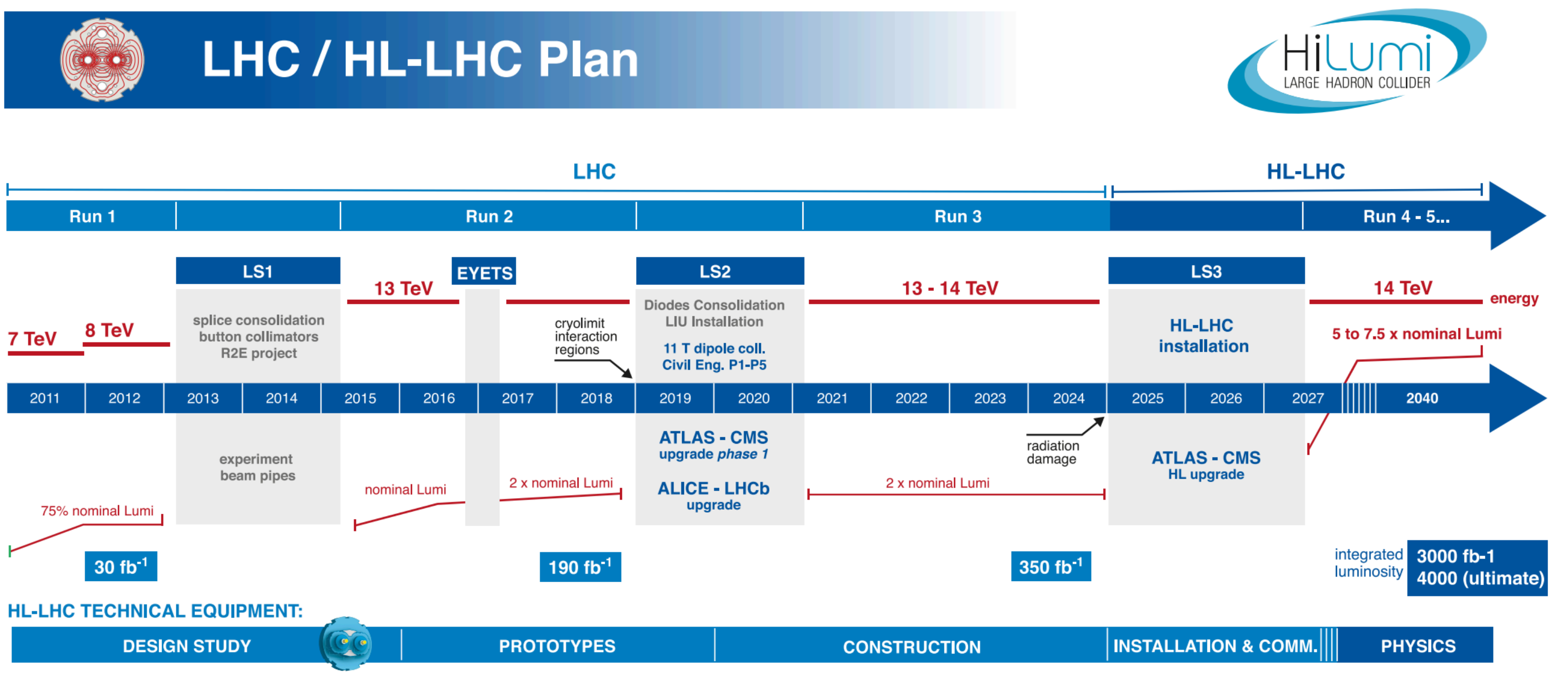

- ATLAS Upgrades for Run 3
  - Partial replacement of muon spectrometer, finer LAr segmentation
- ATLAS Upgrades for Run 4
  - New inner tracker with extended η coverage, trigger and calorimeter upgrades, potentially new calorimeter/timing detector in forward region

Phase-II TDAQ system

SLAC

- Two-Level Trigger and Data Acquisition System
  - hardware-based L0 trigger system
  - software-based Event Filter, aided by dedicated tracking accelerator
- Storage-based data-flow infrastructure
  - decouple real-time domain from software processing
  - enable advanced data processing strategies

Operating points for ATLAS TDAQ:
- L1 latency increase to ~ 10 μs (~2.5 μs today)
- Readout rate increase to 1-4 MHz (100 kHz today)
- Rate to permanent storage ~ 10 kHz (~1 kHz today)

ATL-DAQ-SLIDE-2019-527

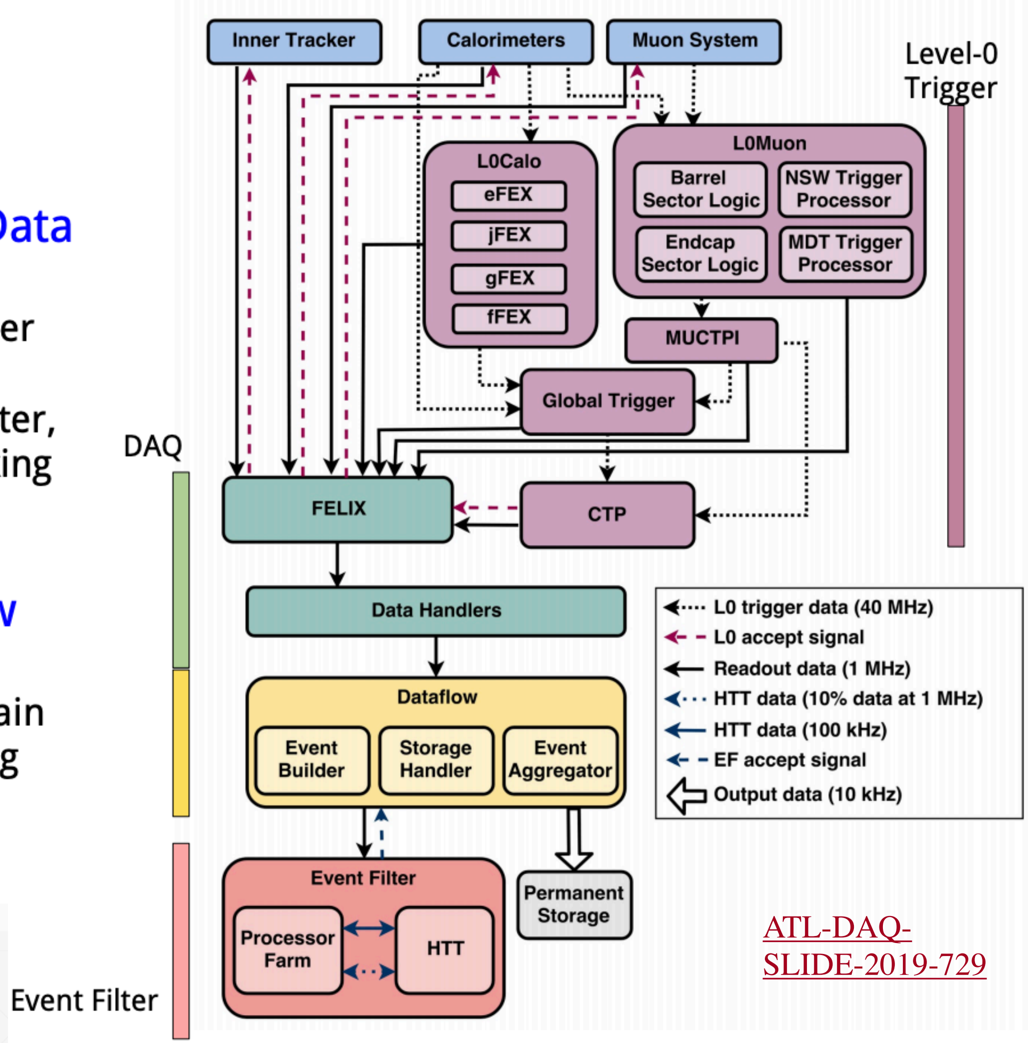

ATL-DAQ-SLIDE-2019-729

Phase-II TDAQ upgrade

Three main systems of the TDAQ Phase-II upgrade architecture:

- Level-0 Trigger
- DAQ (Readout and Dataflow subsystems)
- Event Filter

Single-hardware-level trigger architecture (baseline scenario):

- ➢ capable of evolving into a two-level hardware trigger system (evolved scenario)

The two main **criteria for** an evolution to the **split-level hardware trigger** configuration:

- the hadronic trigger rates
- the inner pixel detector layer occupancies

If either or both are higher than expected, the baseline TDAQ architecture would restrict the trigger menu at the ultimate HL-LHC running conditions.

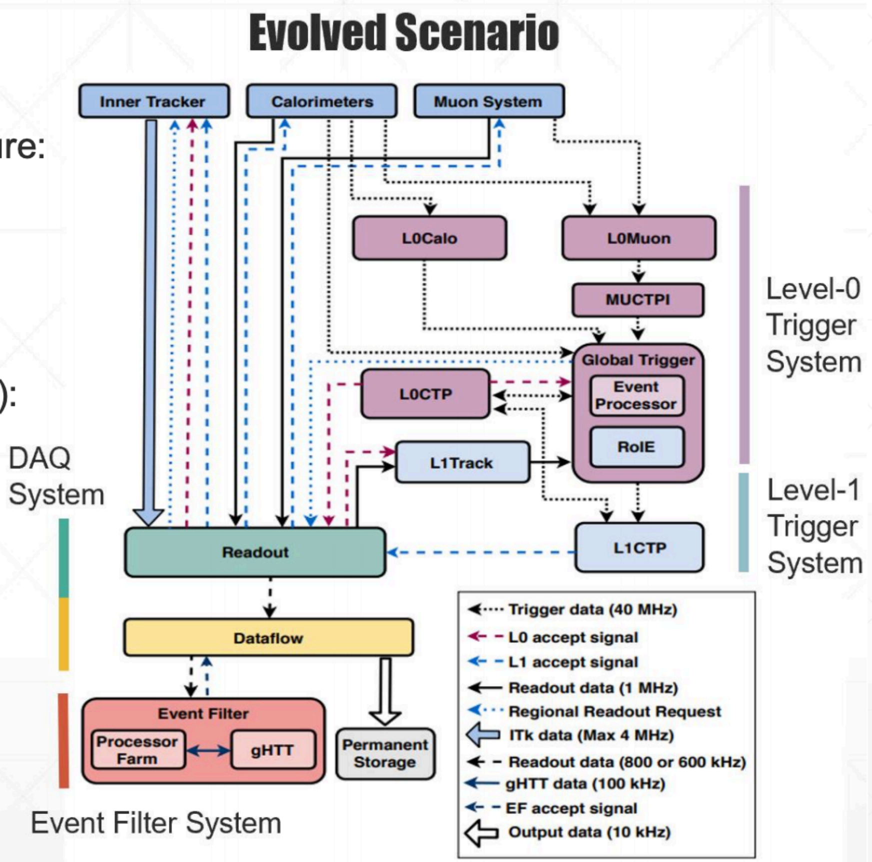

ATL-DAQ-SLIDE-2019-527

Phase-II TDAQ system

- Evolution path to a two-level hardware trigger included in the design
  - L0 – 4 MHz
  - L1 – 1 MHz
  - Event Filter – 10 kHz
- Possible transition from baseline to evolution driven by physics requirements
  - hadronic trigger rates
  - occupancy of inner layers of ITk
- Avoid the baseline TDAQ implementation restricting the trigger menu at the ultimate HL-LHC operating conditions
- Level-1 Trigger combines L0 objects with track information from a dedicated subsystem to discriminate against pileup in the calorimeter

ATL-DAQ-SLIDE-2019-729