ATLAS TRIGGER UPGRADES

Stephen Hillier (University of Birmingham)
on behalf of the ATLAS Collaboration

ATLAS EXPERIMENT

UNIVERSITY OF BIRMINGHAM

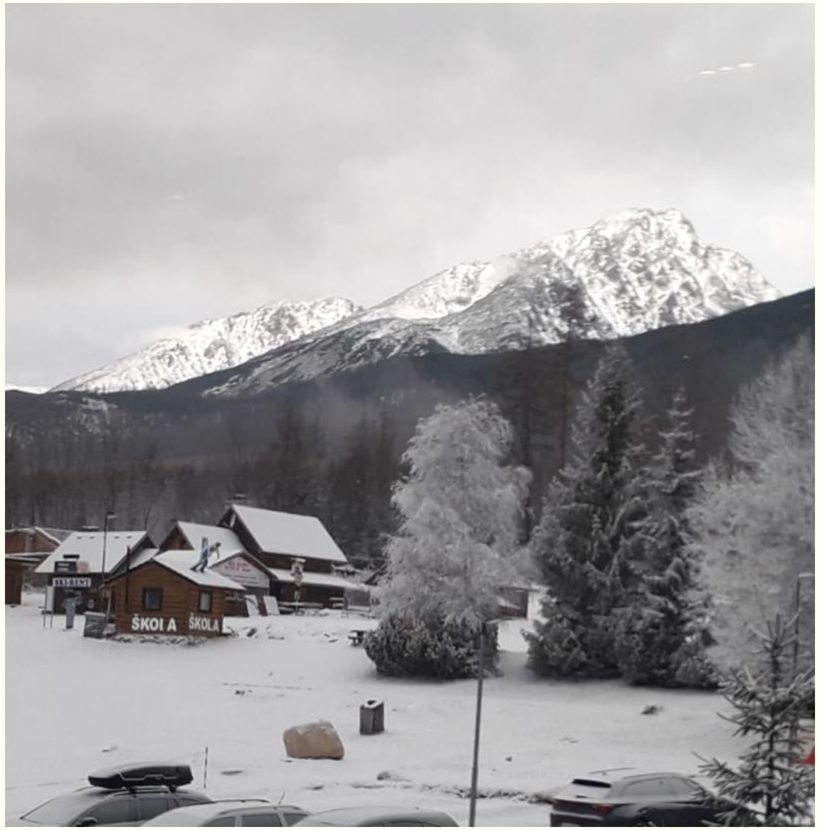

Overview of talk

Challenges of HL-LHC

- Physics Requirements
- Detector Upgrades
- Upgrade Trigger/DAQ Generalities

ATLAS Trigger, DAQ Upgrades

- Level-0 Trigger
- Event Filter
- Readout and DAQ

ATLAS Level-0 Global Trigger

- Architecture
- Hardware

Unless otherwise noted, plots and diagrams are taken from ATLAS TDAQ Phase-2 Upgrade TDR

CHALLENGES OF HL-LHC

And a generic guide to solutions

How did we get here?

- Never forget that triggering at LHC is already a great achievement
  - Average of 65 proton-proton collisions every 25 ns
    - Already at 2+ times nominal LHC instantaneous luminosity
  - Synchronisation of trigger and readout to select correct data associated to interesting event
  - Data spread over multiple time-slices at different points of detector (and buffers)

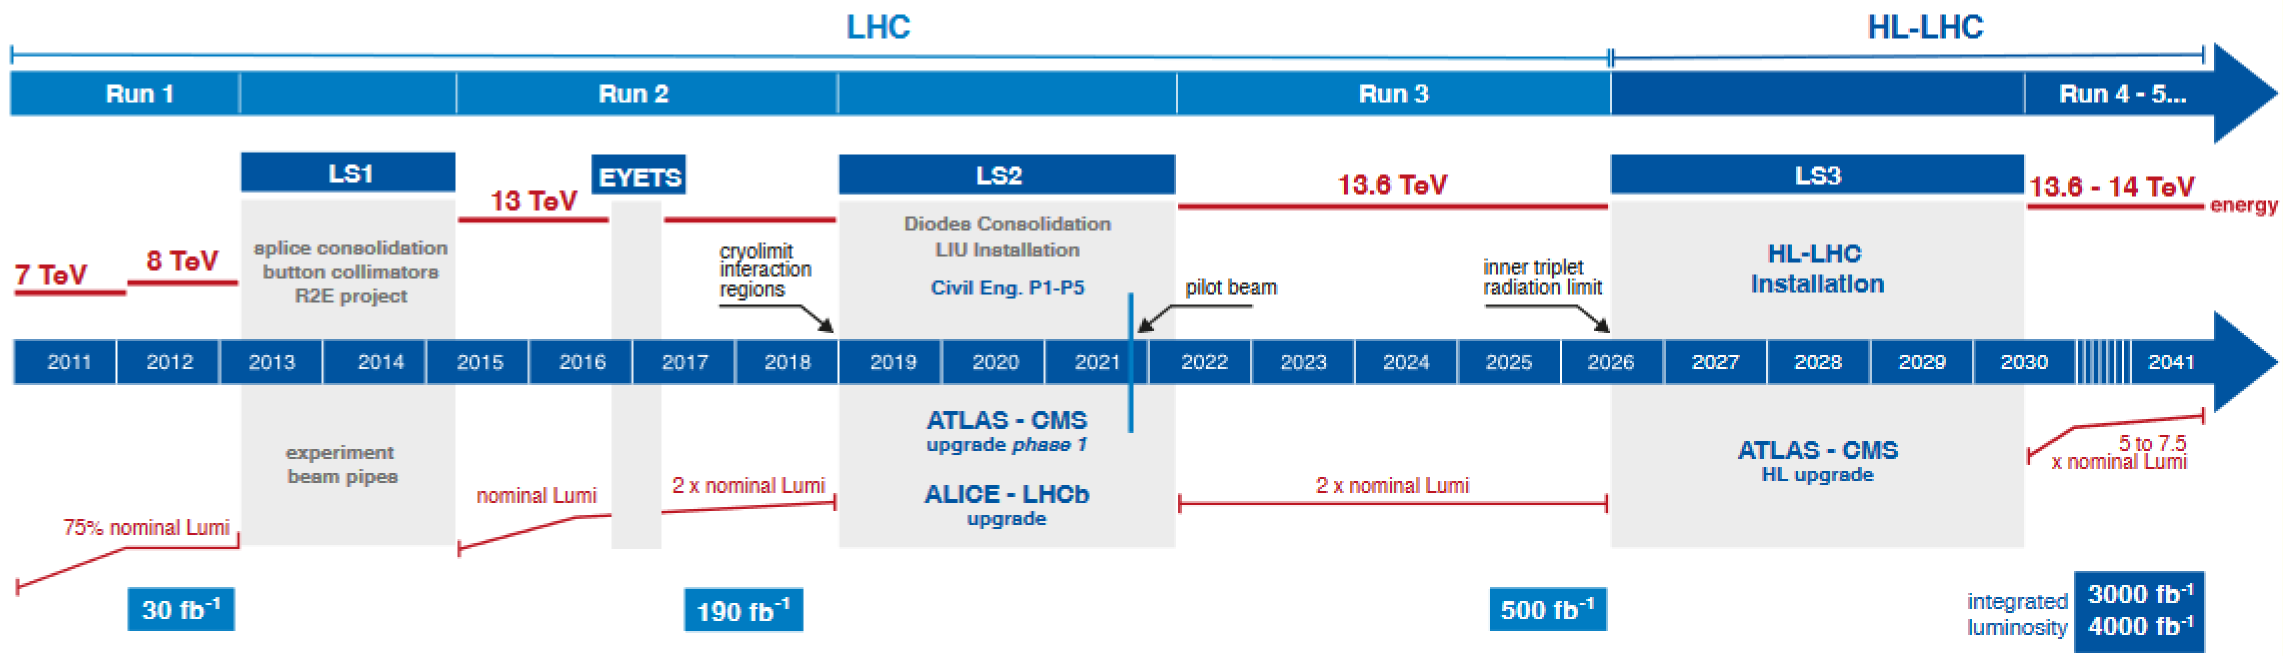

Triggering Challenges of High-Luminosity LHC

- What makes HL-LHC harder?

- Luminosity increases by ~3.5
  - For the same physics (same trigger menu) record at least 3.5 times rate
- Energy potentially increases to 14 TeV
  - Small increase in cross-sections
- Physics signatures become harder to distinguish
  - Denser particle flow environments
  - Greater pile-up effects and ambiguities
- Physics goals more challenging
  - Easier channels already well covered in initial LHC running periods
  - Ambition to push to lower thresholds and more complex physics signatures

Triggering Challenges of High-Luminosity: Mitigations

- How do we solve these challenges?

- Luminosity increases by ~3.5
  - For the same physics (same trigger menu) record at least 3.5 times rate

  Record more events, x 3.5

- Energy potentially increases to 14 TeV
  - Small increase in cross-sections

  Record more events, say x 4

- Physics signatures become harder to distinguish
  - Denser particle flow environments
  - Greater pile-up effects and ambiguities

  Increase detector (and trigger) granularity

- Physics goals more challenging
  - Easier channels already well covered in initial LHC running periods
  - Ambition to push to lower thresholds and more complex physics signatures

  Record more events, x 5
  Be more selective

High Pile-up and Tracking

- Pile-up (μ) increases from current ~65 to ~200 interactions per bunch crossing
- Requires higher granularity tracking
  - Old tracker reaching end of lifetime
- New Inner Tracker (ITK) entirely Silicon detector
  - Shorter strips, more channels
  - 5.1G channels pixels, 60M channels strips
    - c.f. current 92M and 6M
  - Also covers larger angular range
    - η up to 4.0 rather than 2.5
  - Higher radiation tolerance
- Consequence for Trigger/DAQ
  - Increased data size (about 2.5 x current)
  - Large combinatorial problem to extract tracks - High Level Trigger (EF) only

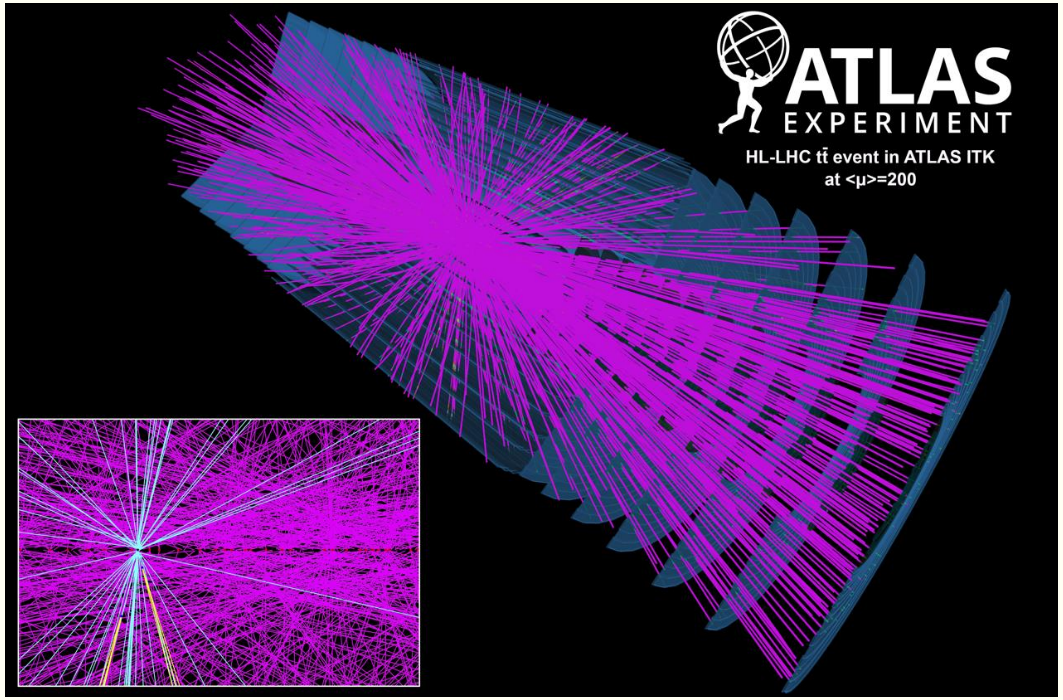

Simulated Top Event: UNSG-2021-58

High Pile-up and Calorimetry

- Consequences of pile-up for calorimetry are not so obvious
  - Occupancy not as high as tracking… mostly (see Forward Calorimeter)
  - Orders of magnitude fewer channels

- However, it's not just about in-time pile-up
  - Calorimeter signals are shaped
  - Stretched over several bunch-crossings
  - Shaped to give best energy resolution *at nominal luminosity*

from ATLAS
Liquid Argon TDR

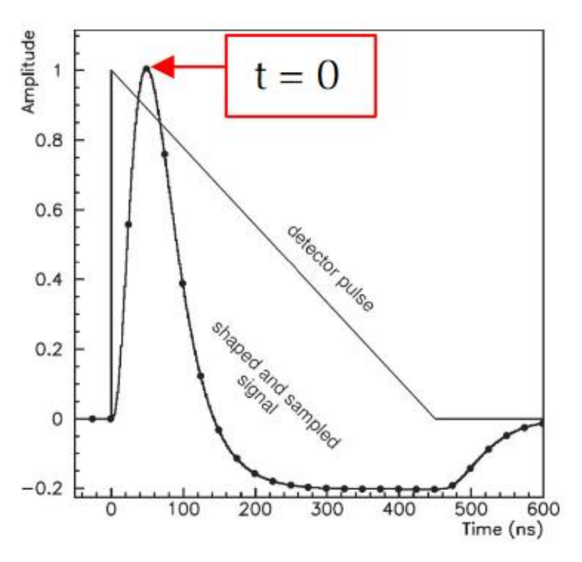

- Out-of-time pileup degrades energy resolution, increases noise
  - Affects all levels, trigger and offline reconstruction
  - Particularly for energy sums used in first level trigger

- Detector won't be upgraded
  - The most that can be done is improve electronics

calorimeter pile-up
'noise' at μ= 200

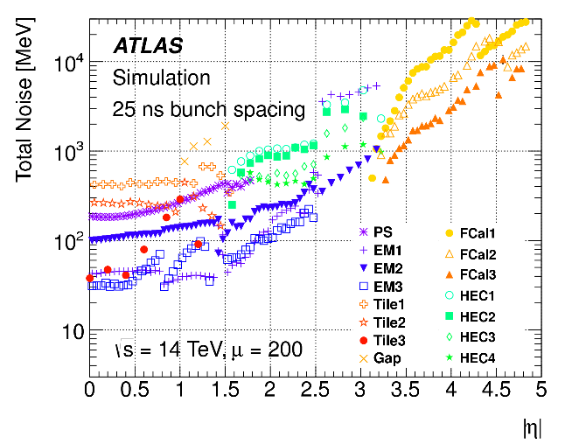

High Pile-up and Muon Detection

- Again, physics occupancy is not high
- Issues limiting trigger performance at HL-LHC
  - Coverage and chamber aging
  - Running with reduced voltage, loosened coincidence
  - Beam background generating fake triggers
- Requires additional chambers and improved logic
  - Fill in existing coverage gaps
  - Expected efficiency x acceptance increase 65% to 95%
  - Better background rejection, particularly in endcap
- Current muon system will mostly be in place
  - Including Phase-1 New Small Wheel upgrade and trigger logic
  - Original MDT (precision) chambers included in Level-0 trigger

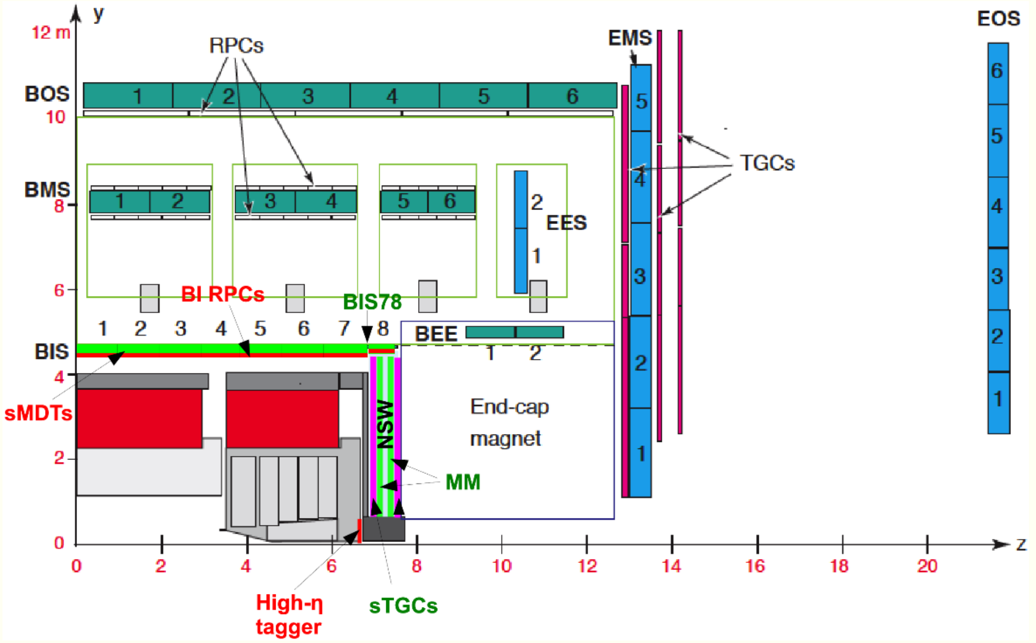

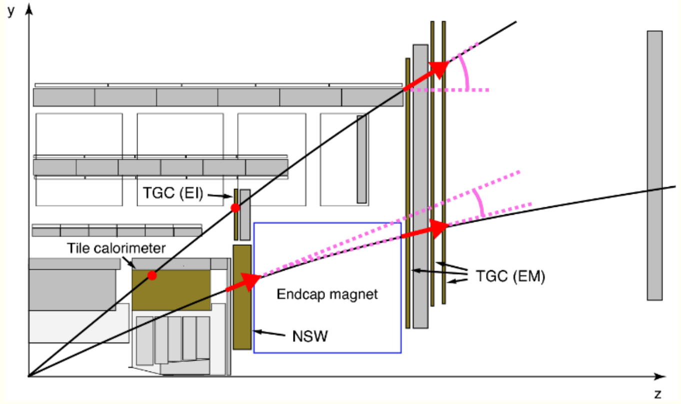

Triggering Challenges of High-Luminosity: Phase-2 Upgrade

- How do these solutions affect Trigger/DAQ?

- Luminosity increases by ~3.5
  - For the same physics (same trigger menu) record at least 3.5 times rate

Record more events, x 3.5

- Energy potentially increases to 14 TeV
  - Small increase in cross-sections

Record more events, say x 4

- Physics signatures become harder to distinguish
  - Denser particle flow environments
  - Greater pile-up effects and ambiguities

Increase detector (and trigger) granularity

- Physics goals more challenging
  - Easier channels already well covered in initial LHC running periods
  - Ambition to push to lower thresholds and more complex physics signatures

Record more events, x 5
Be more selective

Triggering Challenges of High-Luminosity LHC: DAQ

- How do these solutions affect DAQ?

Increase DAQ data rates (to record, approximate)

- Luminosity increases by ~3.5
  - For the same physics (same trigger menu) record at least 3.5 times rate

current x 3.5

- Energy potentially increases to 14 TeV
  - Small increase in cross-sections

current x 4.0

- Physics signatures become harder to distinguish
  - Denser particle flow environments
  - Greater pile-up effects and ambiguities

current x 4.0 x 2.5 = 10

- Physics goals more challenging
  - Easier channels already well covered in initial LHC running periods
  - Ambition to push to lower thresholds and more complex physics signatures

current x 5.0 x 2.5 = 12.5

4 GB/s increase to 50 GB/s

Triggering Challenges of High-Luminosity LHC: L0 to EF

- How do these solutions affect Level-1 Trigger?

Level-0 rate increase, currently ~100 kHz (Level-1)

- Luminosity increases by ~3.5
  - For the same physics (same trigger menu) record at least 3.5 times rate

350 kHz

- Energy potentially increases to 14 TeV
  - Small increase in cross-sections

400 kHz

- Physics signatures become harder to distinguish
  - Denser particle flow environments
  - Greater pile-up effects and ambiguities

500+ kHz due to non-linearities

- Physics goals more challenging
  - Easier channels already well covered in initial LHC running periods
  - Ambition to push to lower thresholds and more complex physics signatures

1MHz = 10 x current

High Level Trigger input 0.2 TB/s increase to 5 TB/s

Higher trigger rates, less rate reduction, more dataflow

- The fraction of events passing each trigger level will be higher than current system
  - Necessarily, since interesting physics is more common
  - But not necessarily easier to distinguish
    - EF has extra help from improved tracking
    - L0 does not use tracking data, relies on improvements elsewhere
- Dataflow at all levels needs to be increased – enhanced readout/dataflow paradigm

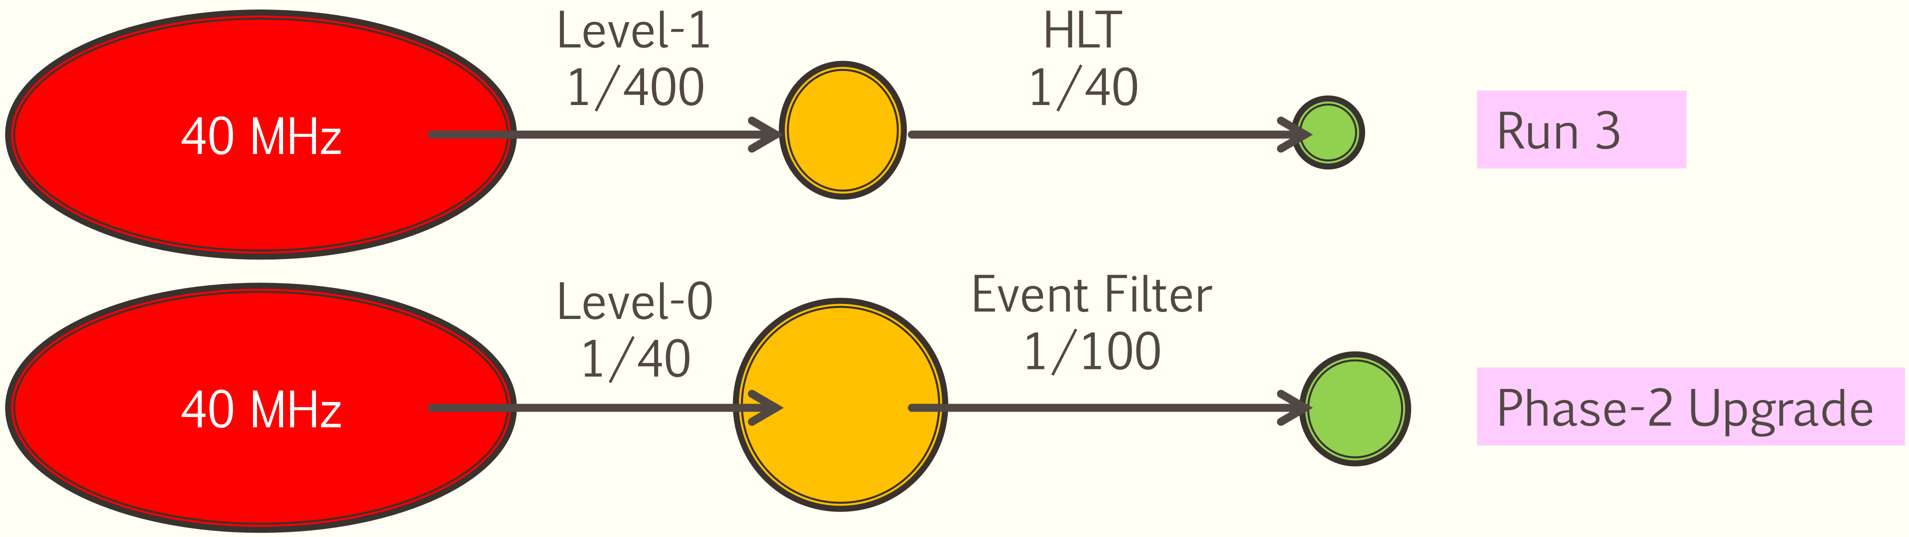

Time to think: more latency, more buffering

- To improve trigger algorithms, need new hardware and more time
- Solution, extend first level latency
  - Current (Level-1) latency 2.5 μs (100 LHC clock ticks)
  - Phase-2 (Level-0) latency 10 μs
  - Increase is even more significant since signal delays eat up at least 1 μs
- Event Filter also requires more powerful algorithms and CPUs
  - May be able to cope with traditional CPU farm
  - Explore other architectures, GPU, FPGAs, AI engines
- All this requires more buffering
  - Detector front-end buffers larger quantities of data for longer periods
  - Readout system requires larger event data buffering

ATLAS TRIGGER/DAQ UPGRADES

Quick overview of all system plans

ATLAS TDAQ Architecture for Phase-2

- Overall picture quite similar to the past
  - But all components updated or reprogrammed
  - Handles higher data and trigger rates
  - Some new components to improve trigger decisions
- Hardware Level-0 Trigger handles full granularity detector data
  - Latency 10 μs, input rate 40 MHz, output rate 1 MHz
- Software Event Filter
  - Input rate 1 MHz, Output rate 10 kHz
- DAQ handled by upgraded FELIX
  - All detectors move to low deadtime readout at 1 MHz
  - Event building, data distribution via Dataflow

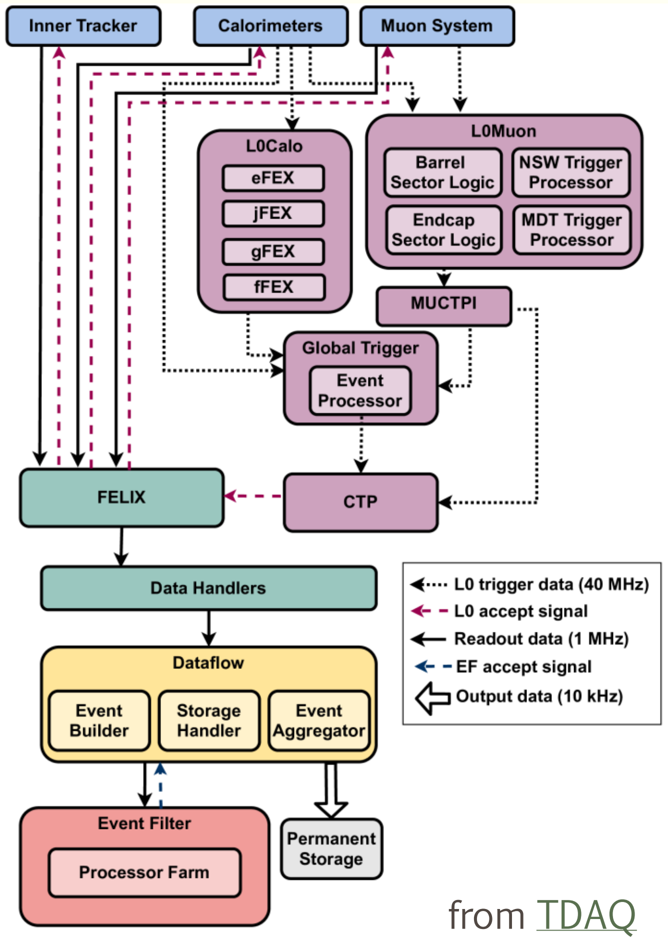

from TDAQ TDR amendment

Level-0 Trigger

- Custom built hardware trigger
  - ATCA based modules
    - c.f. legacy VME based
  - Largely optical signals
    - c.f. legacy largely electrical
  - Optical data speeds up to 25 Gb/s per link
    - c.f. legacy typically up to 10 Gb/s
- ATCA modules with multiple FPGAs
  - Typically FPGAs with System on Chip
  - Potential for including AI engines
- Essentially four separate parts
  - L0Calo: calorimeter signal processing, Phase-1 system plus new module and firmware
  - L0Muon: muon detector processing, all new logic plus additional MDT information
  - L0Global: whole event processing, including full granularity calorimeter and muon data
  - L0CTP: final combinatorial decision, increased number of input items, new timing distribution

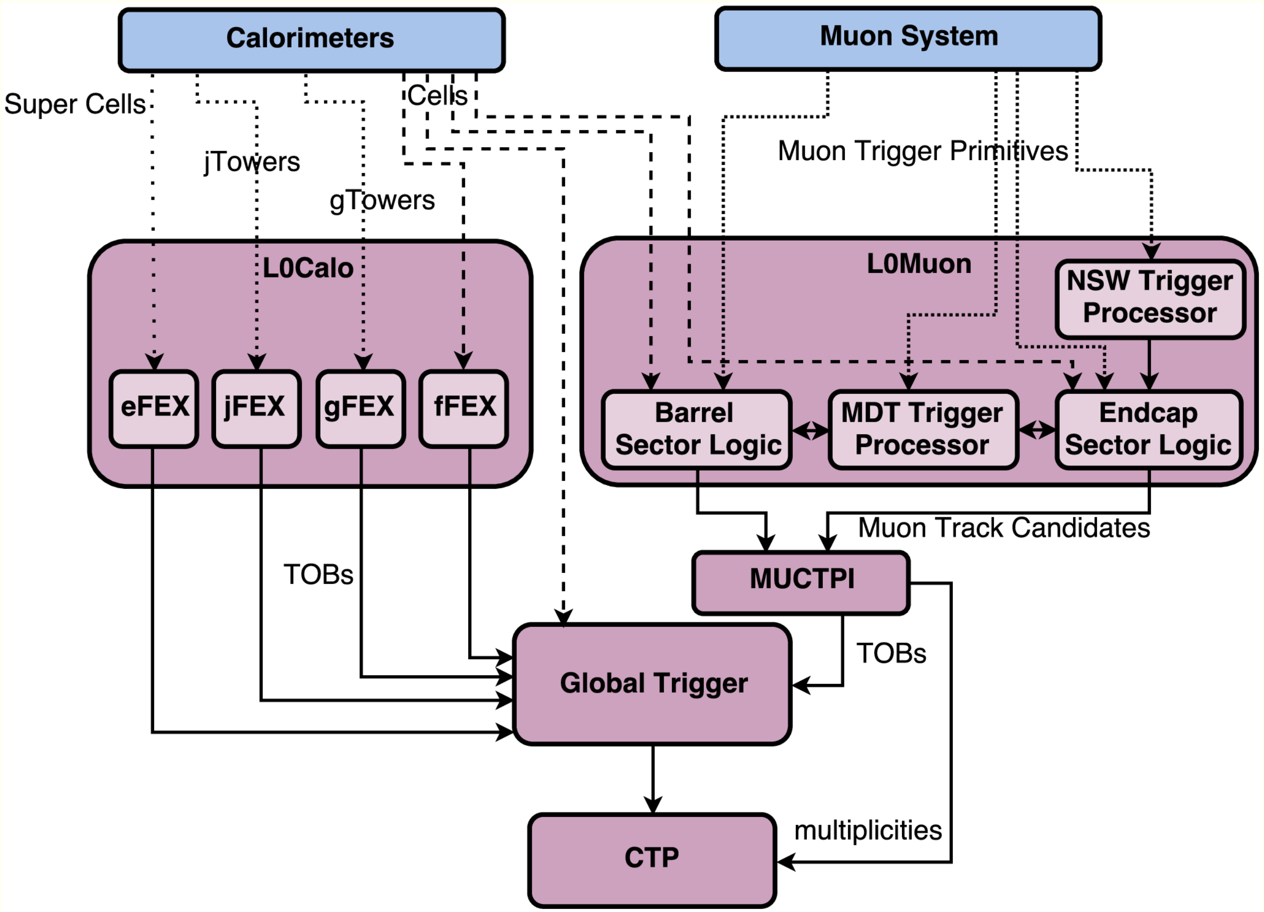

Level-0 Calorimeter Trigger

- Consist of four Feature Extractors (FEX)
  - Input consists of reduced granularity digital calorimeter data
  - Three FEX systems already in place in Run-3
- Forms Trigger Objects (TOBs) for candidate physics signals
  - Electrons, Photons, Taus in eFEX
  - Jets, Taus, Forward Electrons, Missing Energy in jFEX
  - Large-R Jets, Missing Energy in gFEX
  - Improved Forward objects in fFEX (Phase-2)
- Inputs from Liquid Argon and Tile Calorimeters
  - Complex Optical plant required for signal routing
- Outputs sent as list of TOBs to L0-Global
  - TOB = Trigger OBject

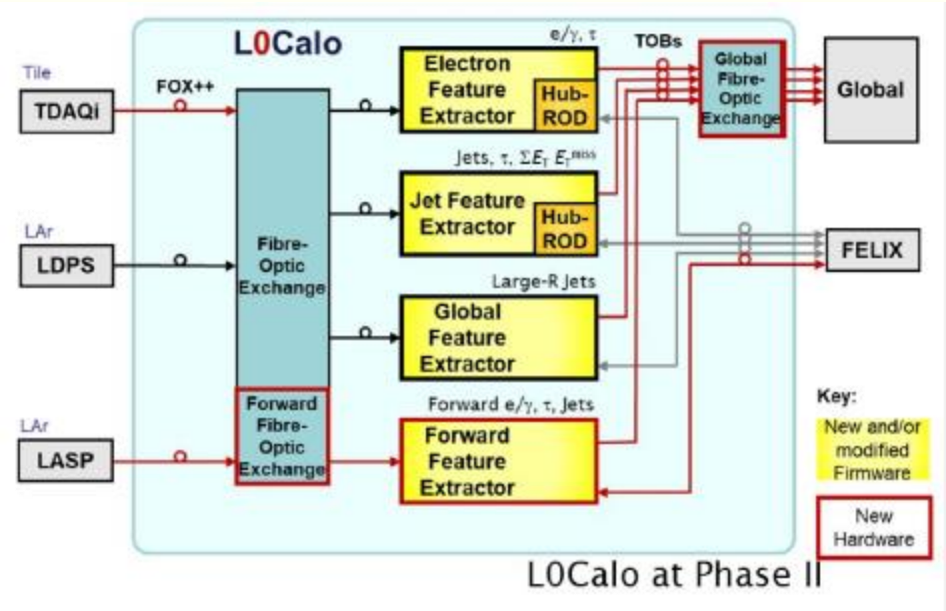

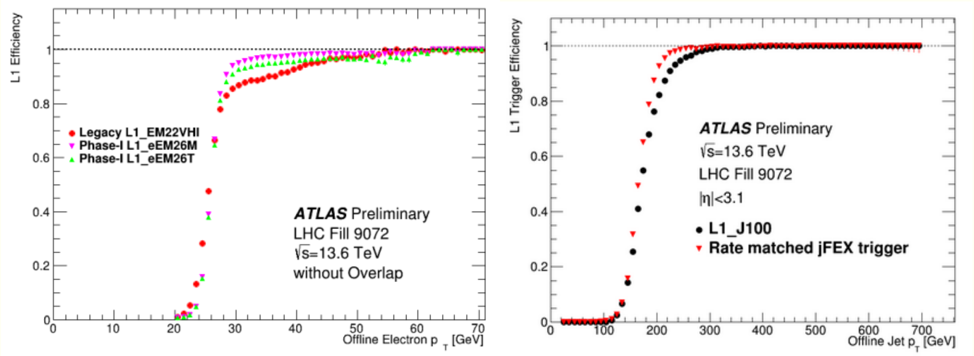

Electron and jet efficiencies using Phase-1 L1Calo

Level-0 Muon Trigger

- Several sequential and parallel trigger processors for each part of muon detector
  - New in Phase-2: MDT trigger
  - MDT processing seeded by other detectors
  - New detectors plugging gaps in coverage
    - sMDT, RPC regions
- Improved trigger processing in all regions
  - High granularity data streamed out optically
  - Trigger logic moved off detector
  - Allows for more precise and programmable decisions
  - Reduction of fakes, better efficiency
- Common processing sector logic platform
  - Programmability via FPGAs
- Prototyping and testing well underway

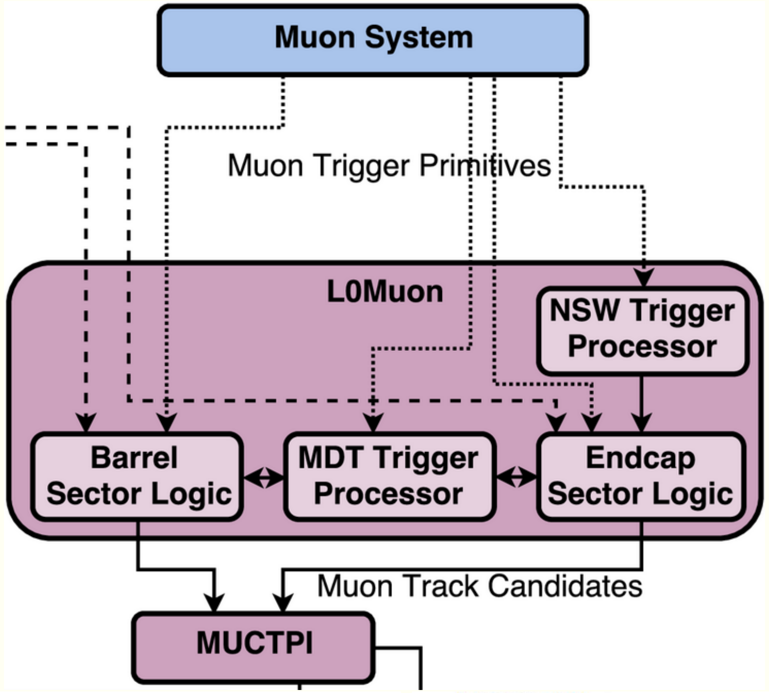

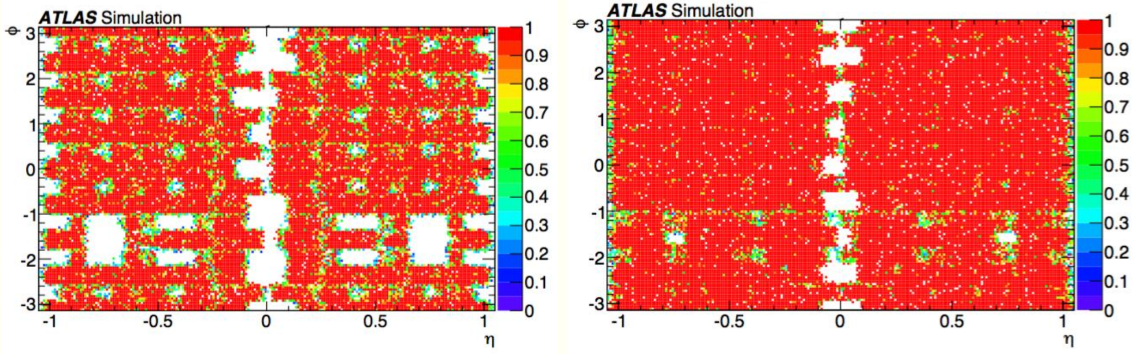

Barrel muon acceptance, Run-3 vs Run-4

Level-0 Global Trigger

- Time-multiplexed full event processor
  - Input data essentially full calorimeter granularity
- Multiple Global Common Module (GCM) boards behave as Level-0 farm
  - Need to maintain fixed latency
  - Each GCM processes 1/N events
    - N ~ 50
- Replaces and enhances current Topological Processor functionality
  - Refines input TOBs with more detailed data and complex algorithms
- Described in detail later

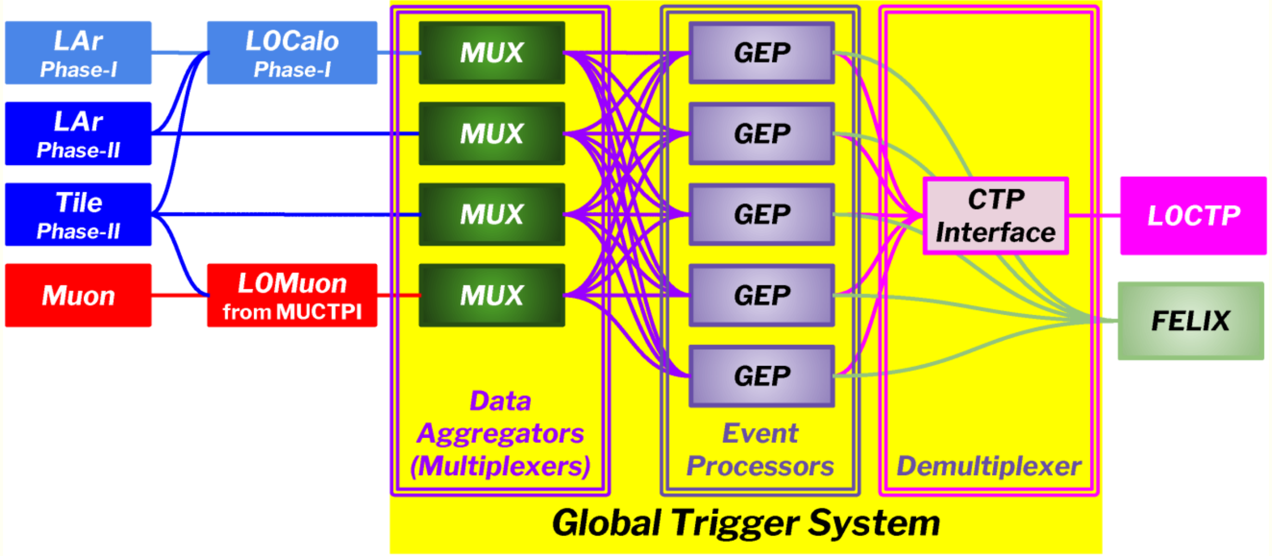

L0 Global v3 prototype

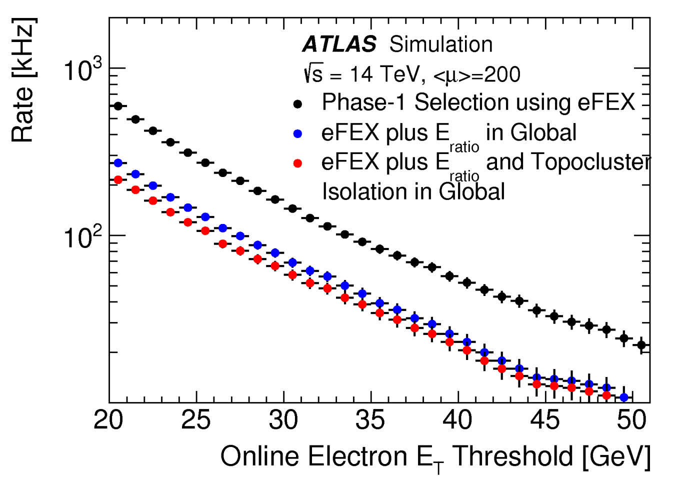

Level-0 Central Trigger

- Responsible for several over-arching functions
  - Final Muon processor combination in MUCTPI
  - Formation of L0Accept based on inputs and Trigger Menu, Central Trigger Processor (CTP)
  - Distribution of clock, synchronisation and decisions Local Trigger Interface (LTI)
- MUCTPI and CTP share common ATCA platform
  - Functionality refined by FPGA load
  - Design based on current MUCTPI ATCA module
- CTP governs overall L0A rate and deadtime
  - More individual trigger inputs, more logic
  - Deadtime requirements looser than current system
    - Bigger demand on detector readout
    - But lower overall deadtime, even at 1 MHz L0A rate

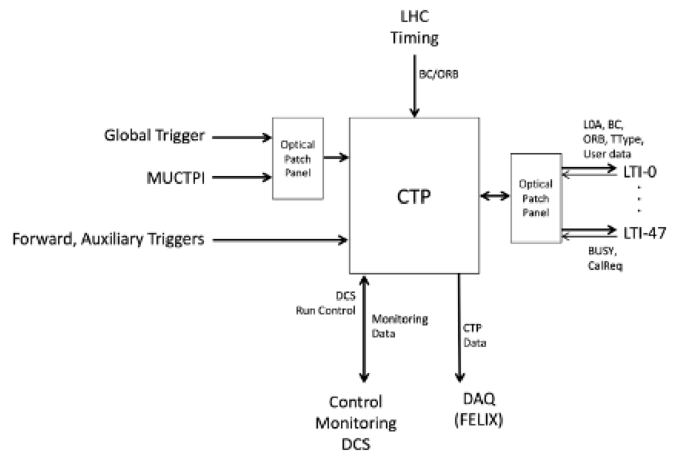

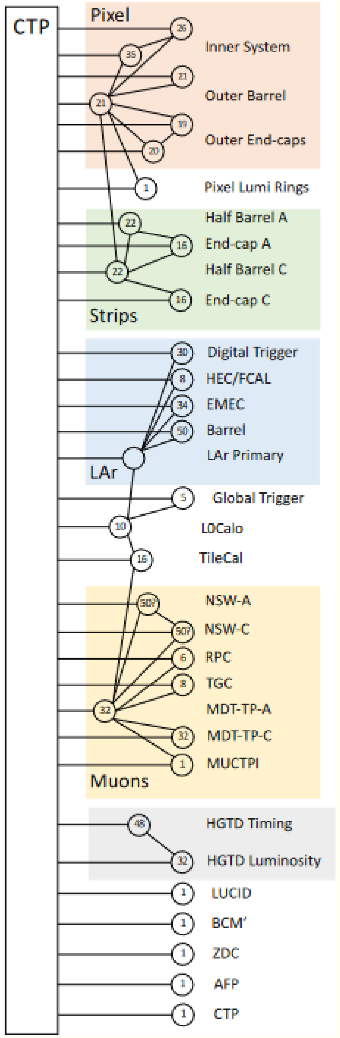

Event Filter

- Event Filter forms the software level of trigger
  - Reduces 1 MHz L0A rate to 10 kHz recorded data
  - Possibility for higher rate streams with reduced data size
  - Full event data used at 1 MHz
- Typically performed by custom, fast algorithms in CPU farms
  - Exploring possibility of other (hybrid) architectures
    - GPU enhancement for some algorithms
    - FPGA accelerators being studied
  - Increased applicability of Machine Learning techniques
  - Larger data size provides opportunities (and challenges)
    - Particularly in Inner Tracker
    - Processing time very sensitive to pile-up levels

CPU optimisation from TDAQ TDR amendment

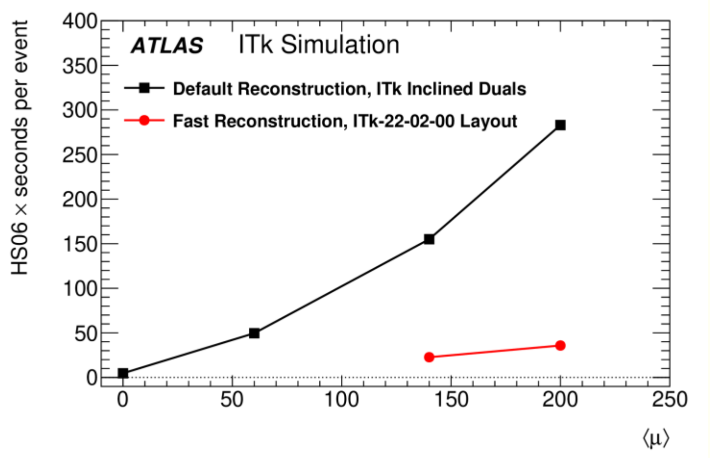

GPU speed-up from ATL-DAQ-PROC-2022-002

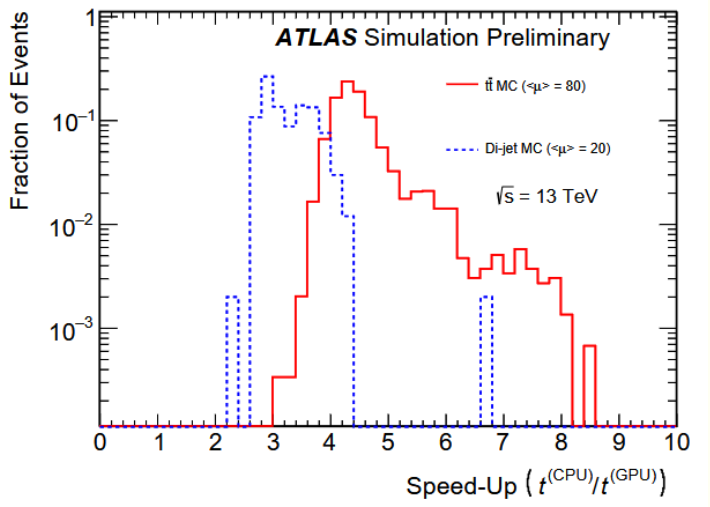

Event Filter Tracking Evaluation

- Inner Tracker has vastly more data
  - And also far more tracks
- Fast tracking has dedicated R&D task-force to evaluate architecture
  - CPU vs GPU vs FPGA, accelerator technologies
    - Focussed on AMD Xilinx FPGAs
  - Design of common language/interfaces for each possibility
- Multiple demonstrators benchmarking tracking algorithms
  - Seeding, track finding/fitting, pattern recognition, ambiguity removal
  - Neural network options considered
  - Exploring usage of High Level Synthesis (HLS)
- Common experiment independent tracking project, ACTS
  - A Common Tracking Software

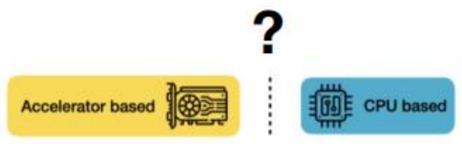

Readout System: FELIX

- Original Readout completely replaced
  - Required for higher bandwidth
  - Complete system runs at 4.6 TB/s throughput
- Common FELIX hardware solution used for all detectors
  - First version of FELIX already used for new detectors in Run-3
  - Used in combination with commodity servers and network
- Data handlers used to collect data fragments and process for Event Building
- Prototyping well underway

FLX-182

FLX-155

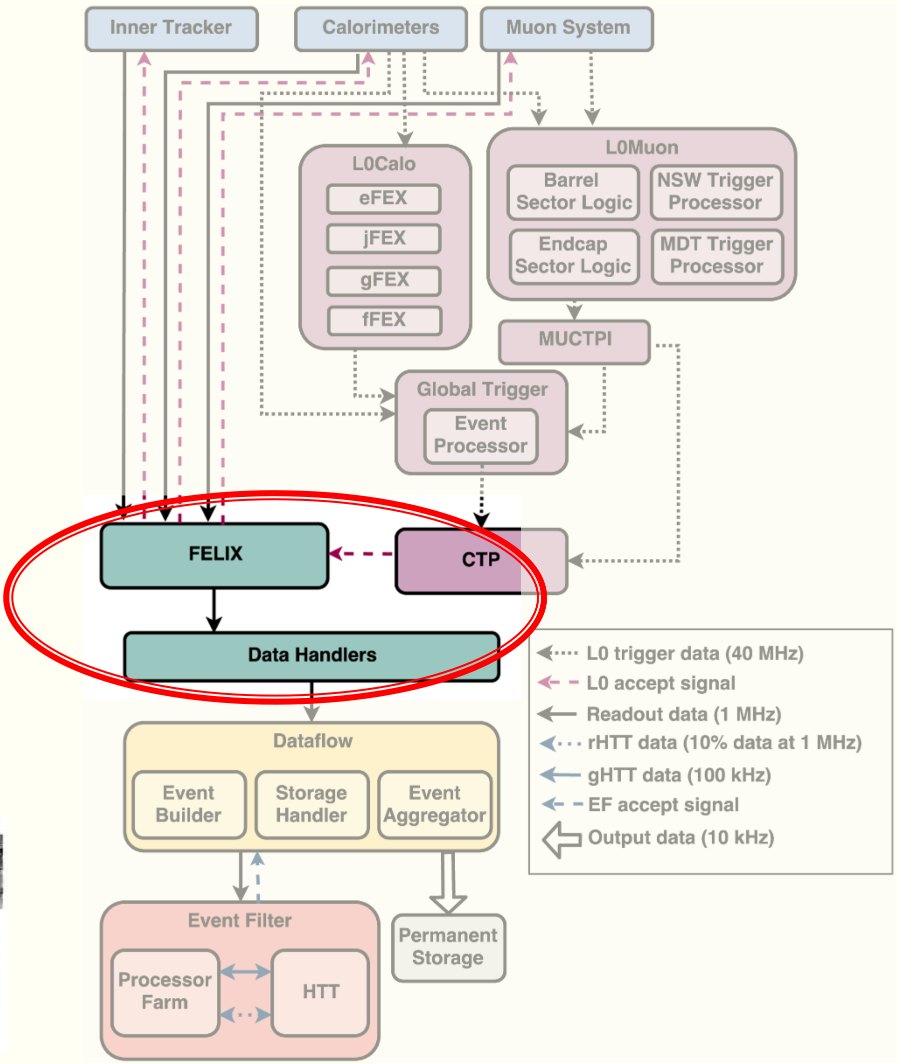

Dataflow, Network and Online Software

- These provide the necessary glue to connect and organise data taking
  - Dataflow: managing buffering and provision the right data promptly wherever needed
    - Based on L0 and Event Filter decisions
  - Network: High speed communication between FELIX, Event Filter, event building, recording etc
  - Online Software: configuring TDAQ and detectors to coordinate activities
- Detailed simulation of dataflow/network to identify and mitigate bottlenecks
  - Buffer sizes
  - Network limitations
  - Needs model of detector data with respect to likely physics events and rates

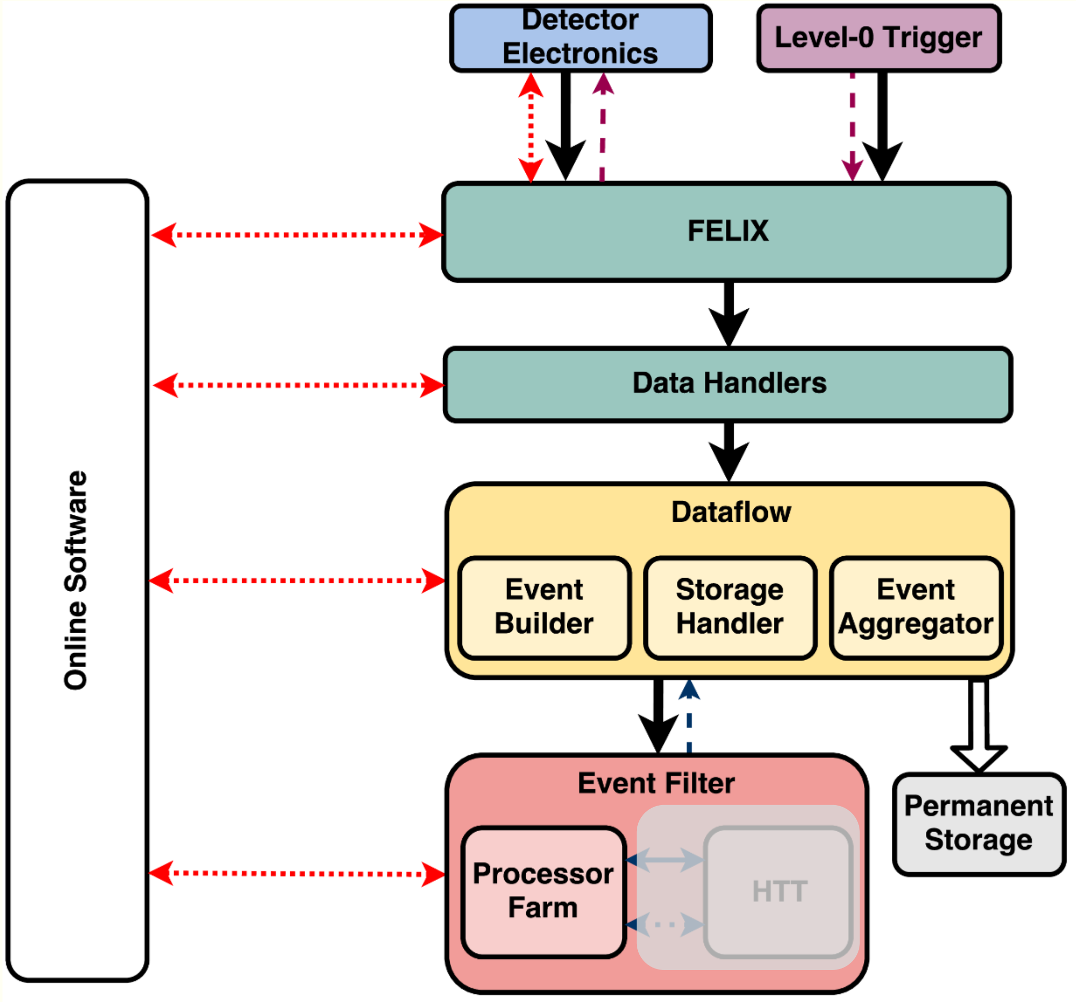

LEVEL-0 GLOBAL TRIGGER

A more detailed case study

Level-0 Global Trigger, concept

- Maximize selectivity by processing full event in a single processor
  - Event Filter-like algorithms at Level-0
  - But limited time (10 μs) and bandwidth
- Direct input from many sources
  - Calorimeters at full granularity
  - L0Calo for calorimeter based objects
  - L0Muon for muon candidates
  - Input 50 Tb/s in total
- Data aggregation to a single node
  - Uses time multiplexing (MUX) to Global Event Processor (GEP)
- Farm of FPGA based processing units
  - Two main FPGAs with multiple functions

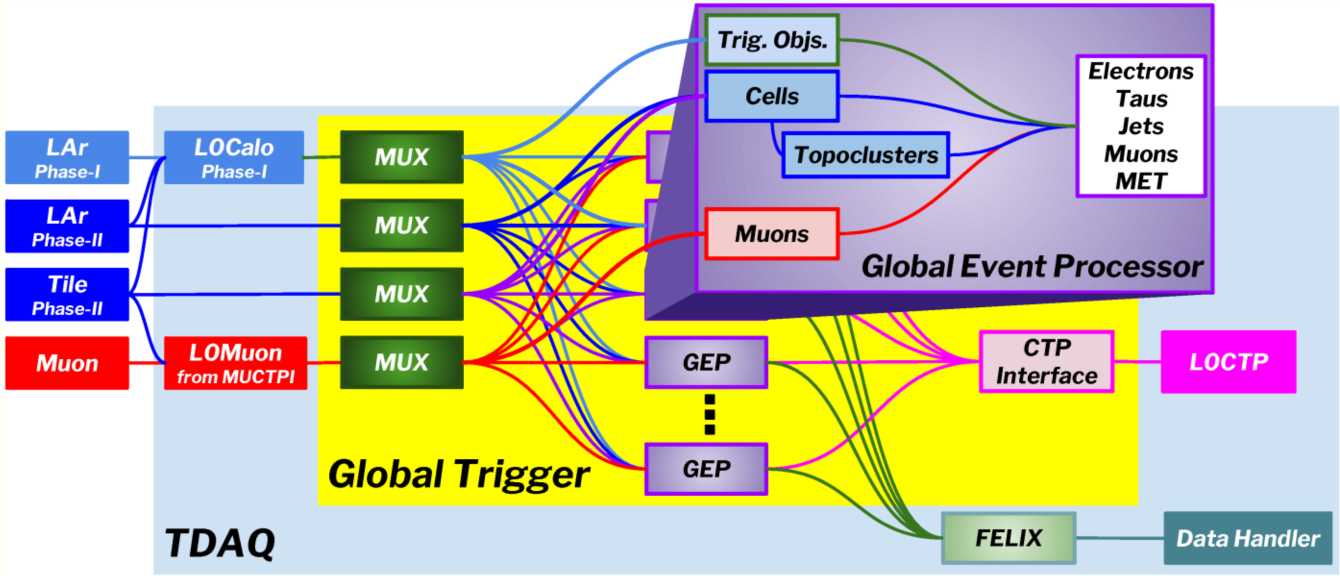

- Processing FPGA refines Level-0 results
  - Localised processing around Trigger Objects (TOBs)
  - Some event level processing algorithms
- Final decisions transmitted to L0CTP
  - On positive CTP result, processing data read out to DAQ

Level-0 Global Trigger, Time Multiplexing

- Nominal 49 node MUX layer
  - Captures incoming data every BC and streams consecutive events to Event Processors in turn (round robin)
- Connected by full-mesh optical 'backplane' to every GEP node
- Nominal 49 node GEP layer
  - Each node receives a new event every 49 BCs, outputs results to single gCTPI node (global/CTP interface)
- gCTPi resynchronises results, and transmits in fixed latency to L0CTP

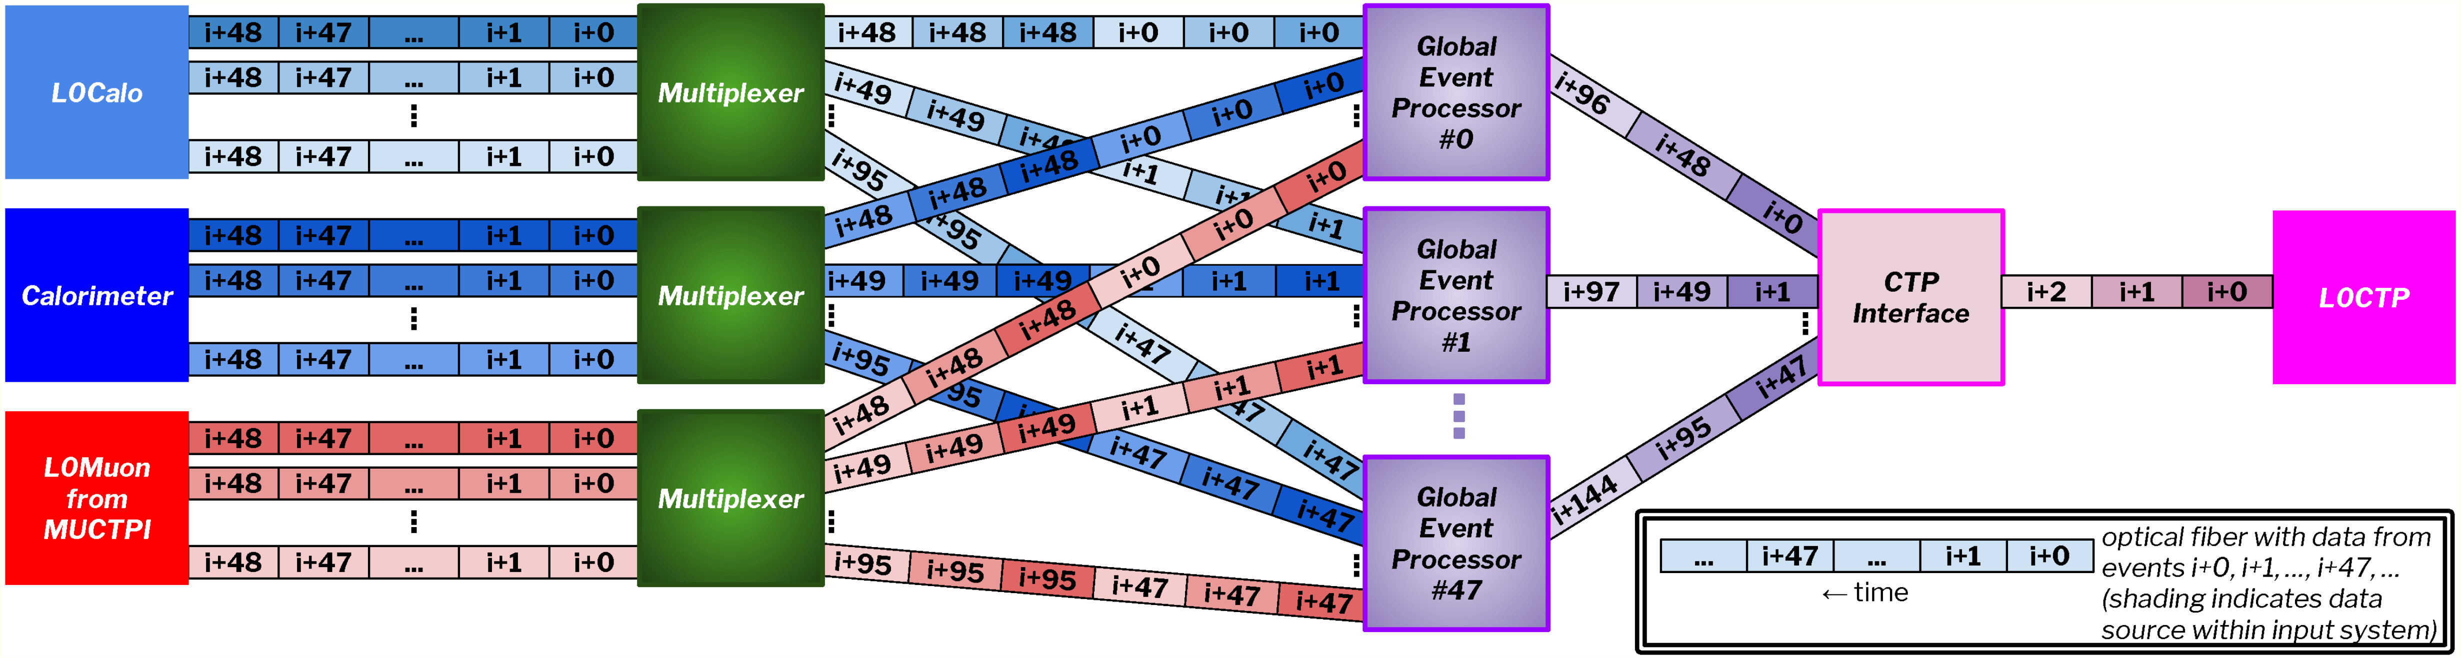

Current Ideas for Global Algorithms

- Algorithms based around current High-Level Trigger algorithms
- Seeded algorithms based on TOBs from FEXes
  - Electron/Tau: improve jet fake rejection with full granularity data
    - Shapes in front EM layer
    - ML algorithms begin investigated
- Topo-clustering for Jets
  - Better energy resolution and pile-up rejection
  - Improved close-by jets and jet substructure measurements
- Global algorithms for full event triggering
  - MET with more sophisticated (ML) pile-up rejection
  - Underlying event subtraction for high-pt objects
  - Isolation energies for muon isolation

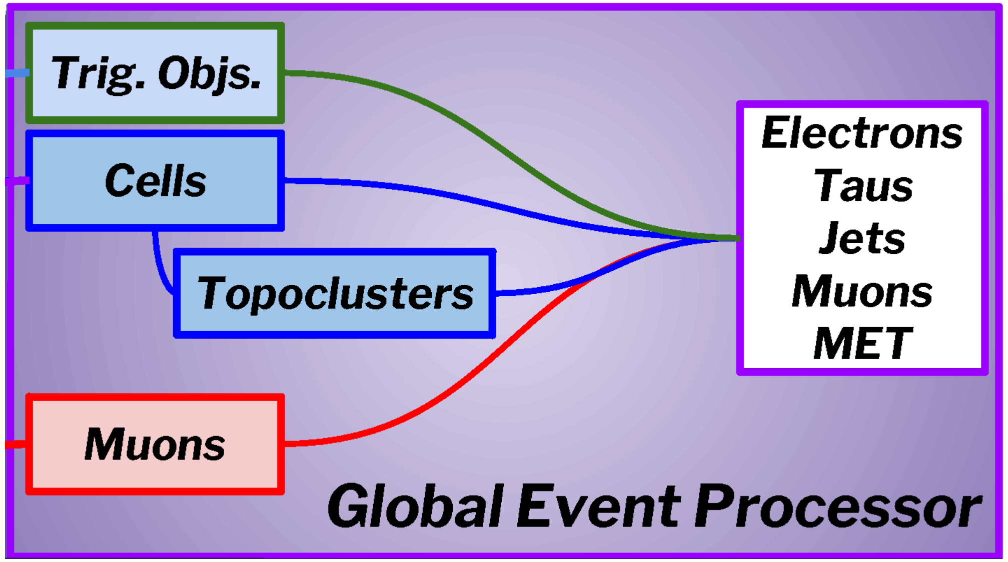

Phase-1 Tau BDT from ATL-COM-DAQ-2024-033

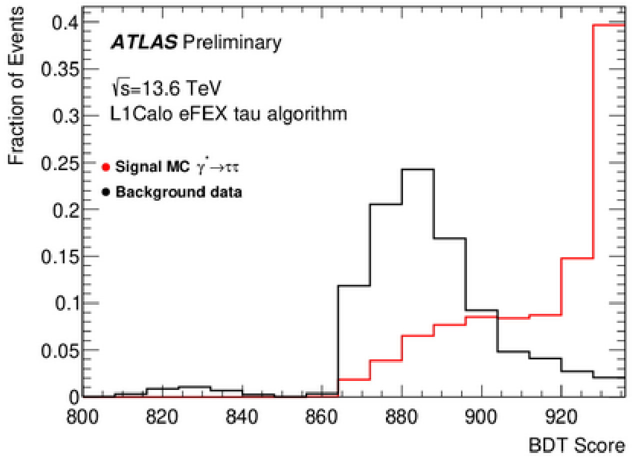

Compromises in Global Architecture

- Is Global the ultimate hardware trigger processor (for non-tracking ATLAS data)?
- Ideal: all inputs into one processor, time to process
  - Essentially a software-like entity
- Reality is not quite so good
  - Still limited latency
    - And some inputs arrive quite late in latency window
  - Bandwidth is not enough for all data
    - 50 Tb/s input bandwidth
    - ATLAS calorimeters have almost 200,000 channels
      - 200000 x 10 bits x 40 MHz = 80 Tb/s
    - Have to restrict input to most significant signals
  - Modern FPGAs are not a single processor
    - Design FPGA contains 4 Super Logic Regions (SLR)
    - Must distribute algorithms between SLRs in a logical fashion
    - Transfer original and intermediate data between SLRs as required

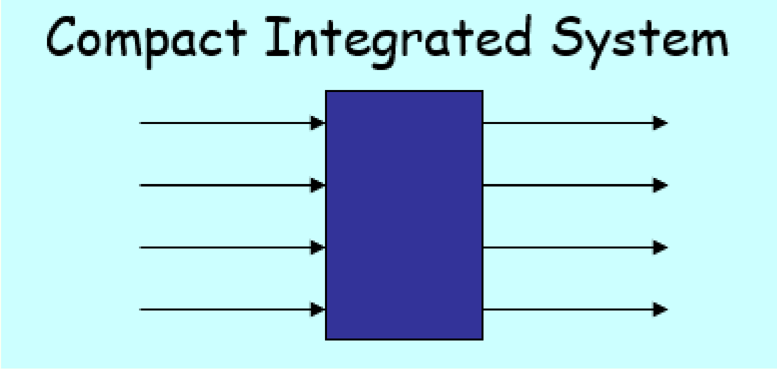

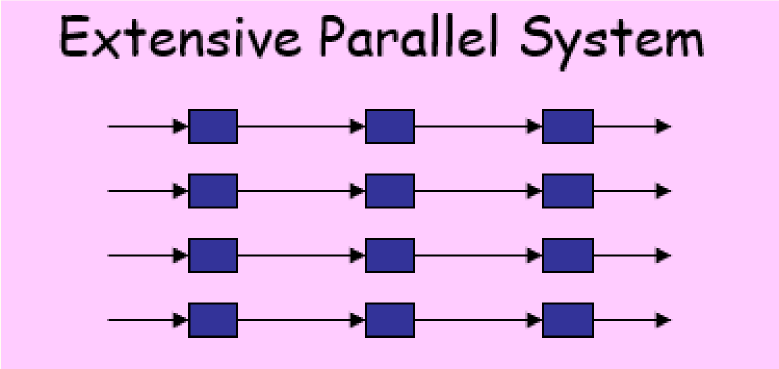

Global Common Module Hardware

- Version 3 of GCM currently in prototyping stage
  - Five modules in full prototype run
  - Already undergoing extensive tests
    - Input and Output interfaces
    - Development of all versions of firmware
- Two Xilinx Versal Premium VP1802 on current version
  - One dedicated as MUX node, balanced input and output
    - Relatively low power
  - One dedicated as GEP node for algorithmic processing
    - Also repurposed as gCTPI in one module
  - Division of labour ensures similar power consumption across all GCM modules
- Final system consists of 50 GCMs (nominal)
  - Spread over 5 ATCA shelves
  - Plus supporting fibre management, PC control and readout

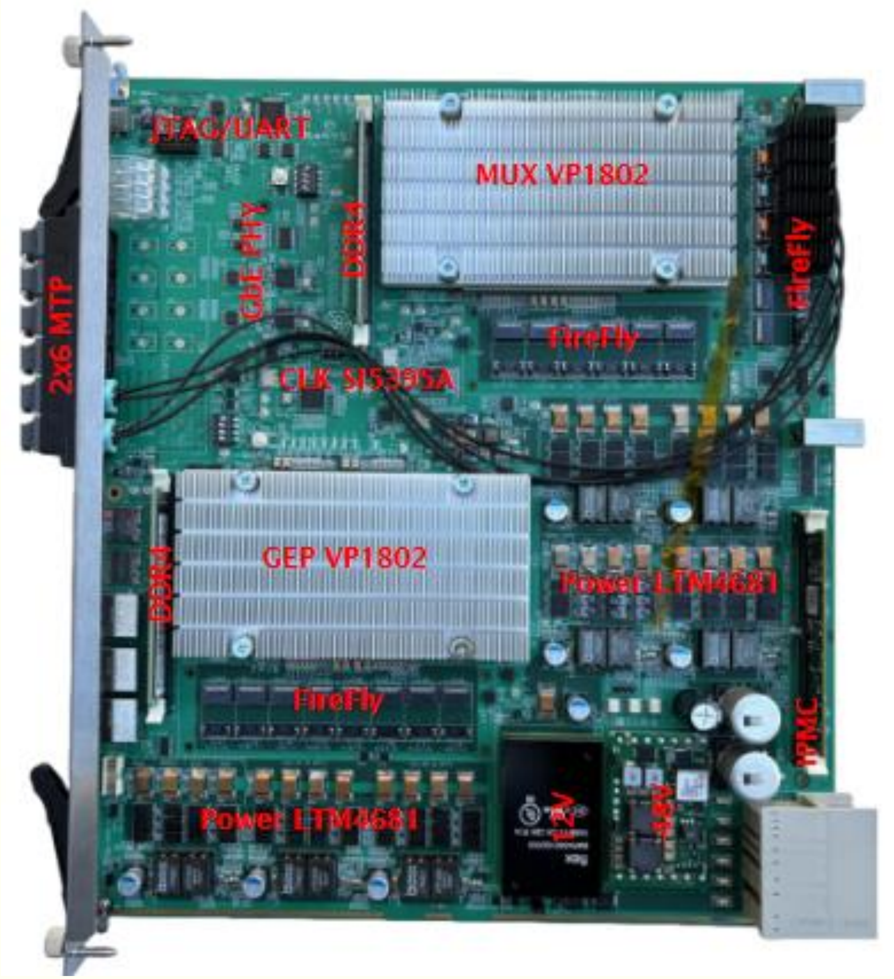

Plans for 10% Global Slice Test

- Next steps
  - Build up full slice through system
  - Test all input/output interfaces in tandem
  - Use FELIX readout to verify functionality
  - Follow testing strategy successfully deployed in previous system commissioning
- Full slice consists of 5 GCM prototypes
  - Need to be fully populated with transceivers
  - Full mesh connectivity as in final system, but only ~10% size
  - One gCTPi to complete path
  - Initially dummy algorithms
    - Once infrastructure works, can start to test with real algorithms and data

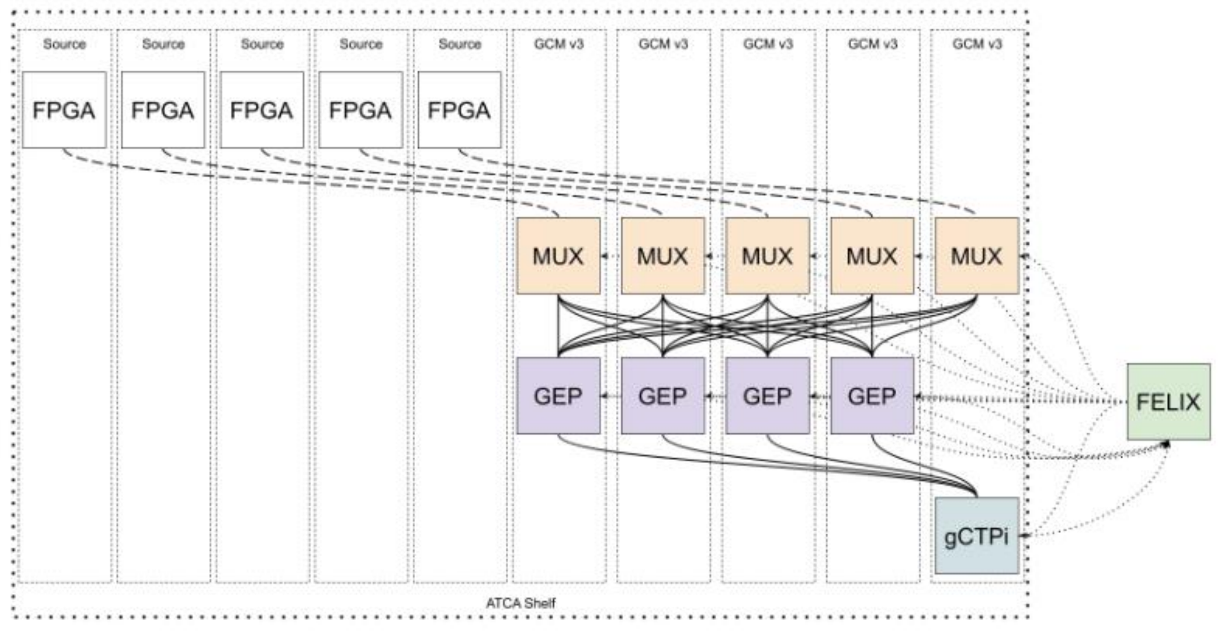

- Tests on individual components/links already well underway
  - Tests performed on FPGA demonstrator boards and prototype GCMs at BNL
- Expect combined slice test in early 2025
  - Initially at BNL, moving to test lab at CERN

Conclusion

- LHC-HL provides a set of challenges at all levels of TDAQ
  - Parameters for ATLAS solutions at LHC-HL have been shown
- ATLAS TDAQ preparations for Run-4 are already well advanced
- Hardware triggering (Level-0) converging on very similar module designs
  - Generic ATCA modules using modern FPGAs typically with SoC control
  - Usage of increasingly fast optical link technologies
  - Module functionality varies by firmware load
  - Investigation of Machine Learning techniques in several areas
- Software triggering (Event Filter) investigating several architectures
  - Traditional CPU based farms
  - GPU processors and accelerators
  - FPGA based solutions for specific problems

Extra slides

- From ATLAS TDAQ Phase-2 Upgrade TDR
  - https://cds.cern.ch/record/2285584

TDR trigger menu

Trigger Selection	Run 1 Offline p_T Threshold [GeV]	Run 2 (2017) Offline p_T Threshold [GeV]	Planned HL-LHC Offline p_T Threshold [GeV]	L0 Rate [kHz]	After regional tracking cuts [kHz]	Event Filter Rate [kHz]
isolated single e	25	27	22	200	40	1.5
isolated single μ	25	27	20	45	45	1.5
single γ	120	145	120	5	5	0.3
forward e			35	40	8	0.2
di-γ	25	25	25,25		20	0.2
di-e	15	18	10,10	60	10	0.2
di-μ	15	15	10,10	10	2	0.2
$e-\mu$	17,6	8,25 / 18,15	10,10	45	10	0.2
single τ	100	170	150	3	3	0.35
di-τ	40,30	40,30	40,30	200	40	$0.5^{\dagger\dagger\dagger}$
single b-jet	200	235	180	25	25	$0.35^{\dagger\dagger\dagger}$
single jet	370	460	400			0.25
large-R jet	470	500	300	40	40	0.5
four-jet (w/ b-tags)		$45^{\dagger}$(1-tag)	65(2-tags)	100	20	0.1
four-jet	85	125	100			0.2
H_T	700	700	375	50	10	$0.2^{\dagger\dagger\dagger}$
E_T^{miss}	150	200	210	60	5	0.4
VBF inclusive			2x75 w/ ($\Delta\eta > 2.5$ & $\Delta\phi < 2.5$)	33	5	$0.5^{\dagger\dagger\dagger}$
B-physics††				50	10	0.5
Supporting Trigs				100	40	2
Total				1066	338	10.4

† In Run 2, the 4-jet b-tag trigger operates below the efficiency plateau of the Level-1 trigger.
†† This is a place-holder for selections to be defined.
††† Assumes additional analysis specific requires at the Event Filter level

Inputs to L0Calo

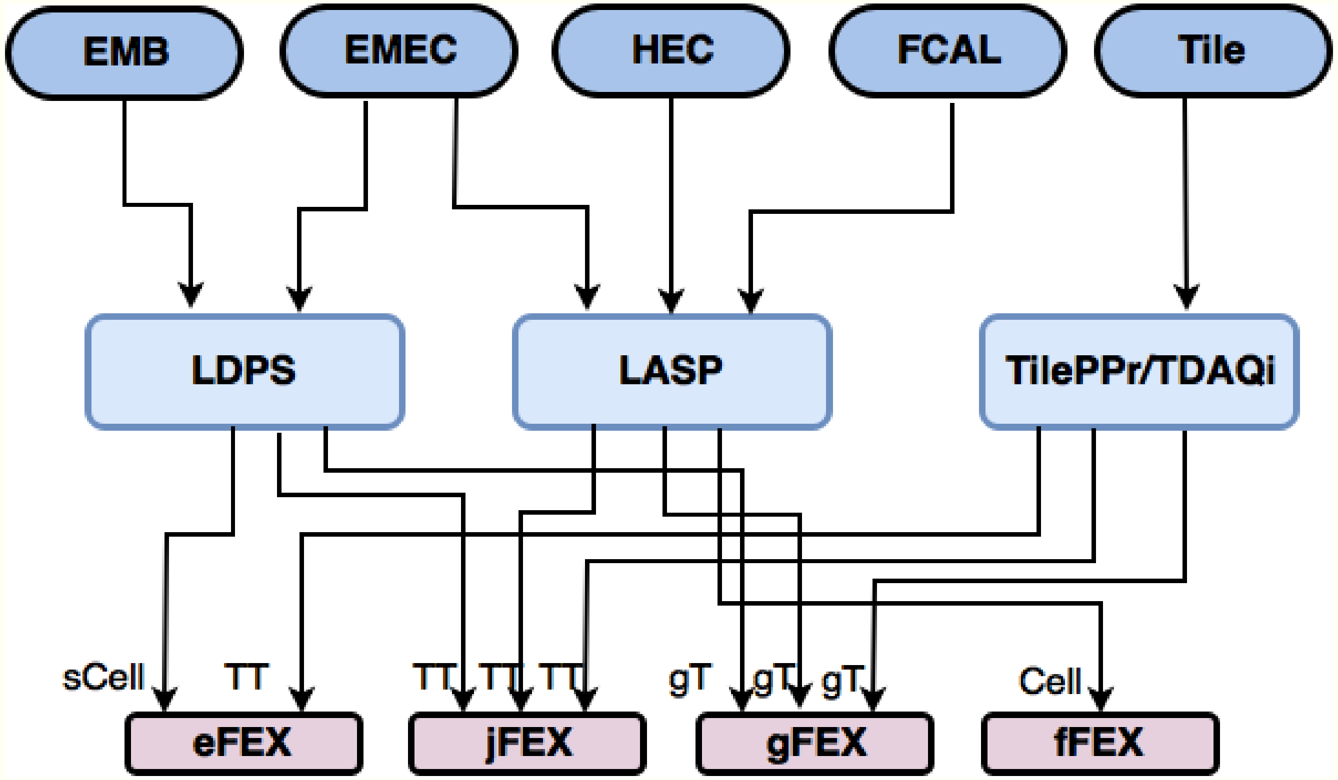

More detail on Global