Final prototyping of trigger and readout electronics ATLAS+CMS

17th Terascale Detector Workshop

20 March 2025

Viacheslav Filimonov

PRISMA Cluster of Excellence

BMBF-Forschungsschwerpunkt ATLAS-EXPERIMENT FSP 103 ATLAS

Physik bei höchsten Energien mit dem ATLAS-Experiment am LHC

JOHANNES GUTENBERG UNIVERSITÄT MAINZ JG|U

Phase-II Upgrade of ATLAS and CMS

- The Phase II upgrade is planned during the Long Shutdown 3: 2026-2030
- The following Runs 4 and 5 will be in operation: 2030 to 2041
- Peak luminosity of 7.5×10^{34} $cm^{-2}s^{-1}$

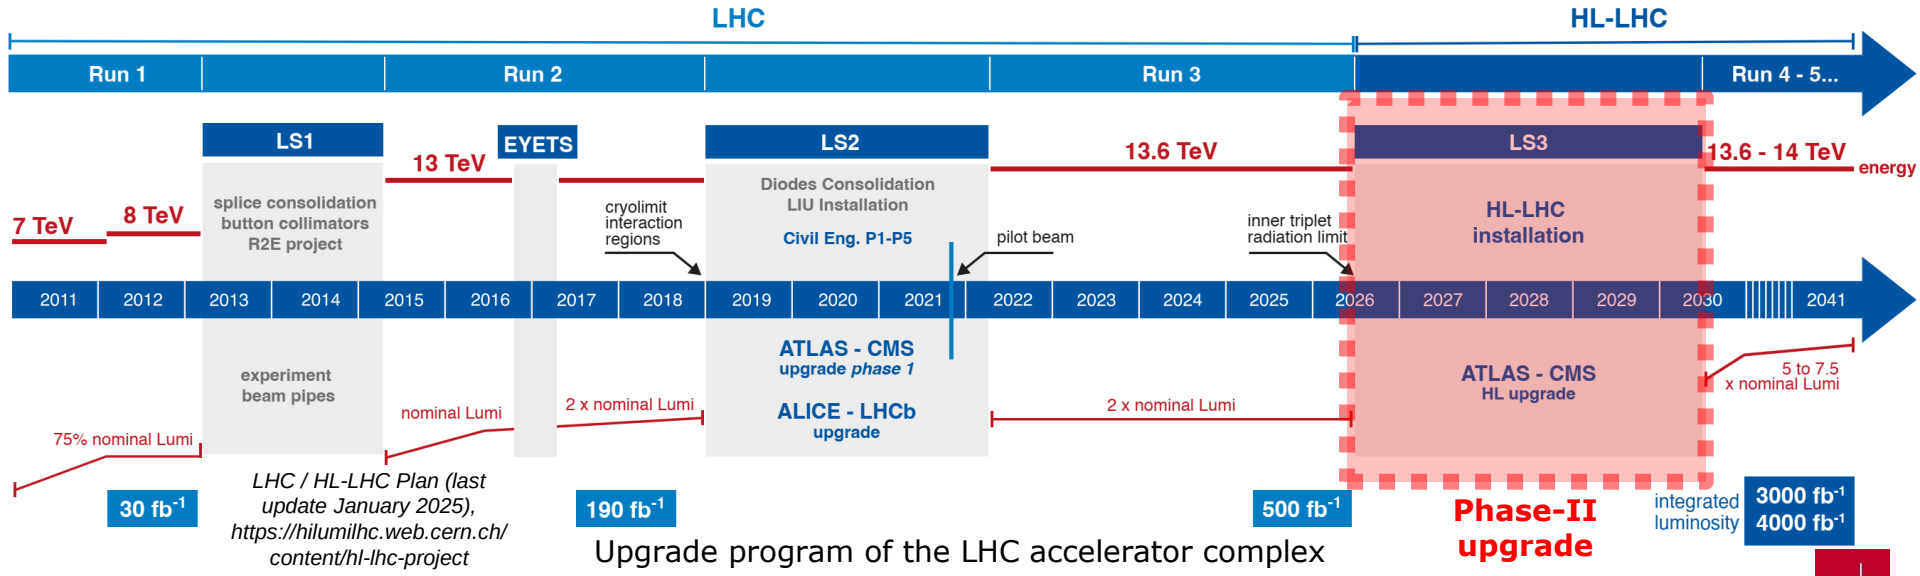

LHC / HL-LHC Plan (last update January 2025), https://hilumilhc.web.cern.ch/content/hl-lhc-project

Upgrade program of the LHC accelerator complex

ATLAS Phase-II Level-0 Trigger system

- The instantaneous luminosity of the LHC in Run 4 is significantly increased → upgrade of the ATLAS trigger system is necessary
- Phase-II Level-0 Trigger System
  - Performs real time event selection
  - Reduces the event rate: 40 MHz → 1 MHz (100 kHz in Run 3)
  - Staying below the maximum readout rate of the ATLAS detector
  - Overall system latency budget 10 μs (2.5 μs in Run 3)

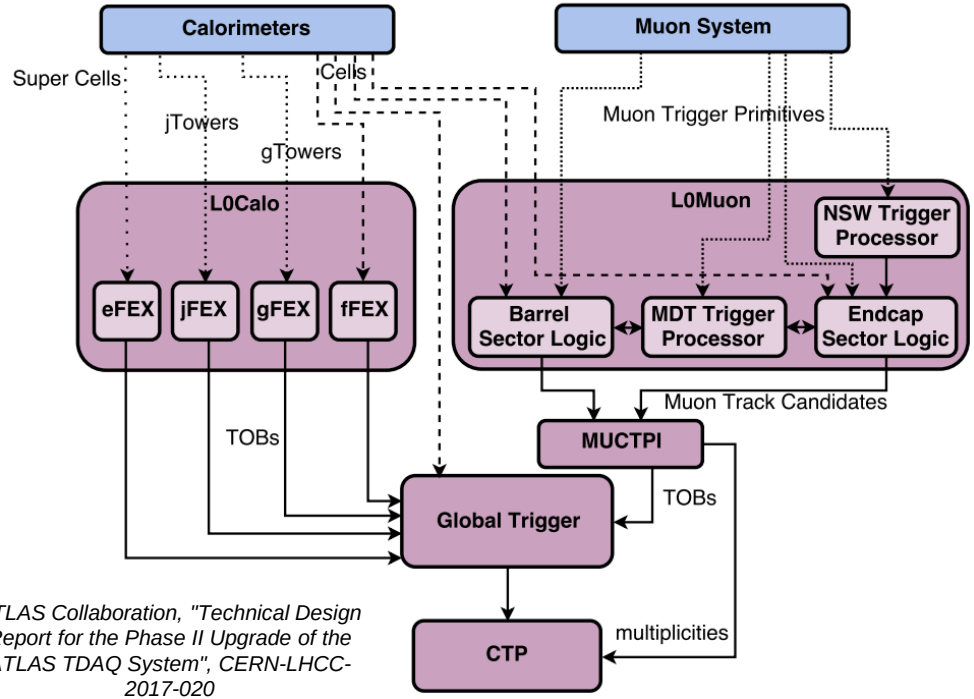

ATLAS Collaboration, "Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System", CERN-LHCC-2017-020

A block diagram of the ATLAS Level-0 trigger system after the Phase-II upgrade

L0Calo

- Processes LAr and Tile calorimeter data to build Trigger Objects (TOBs)
- Includes Phase-I legacy hardware: e/j/g FEXes
- Adds new fFEX for forward EM, jets and taus

- All systems in ATCA form factor
- High end FPGAs and optical modules
- Occupies 5 ATCA shelves

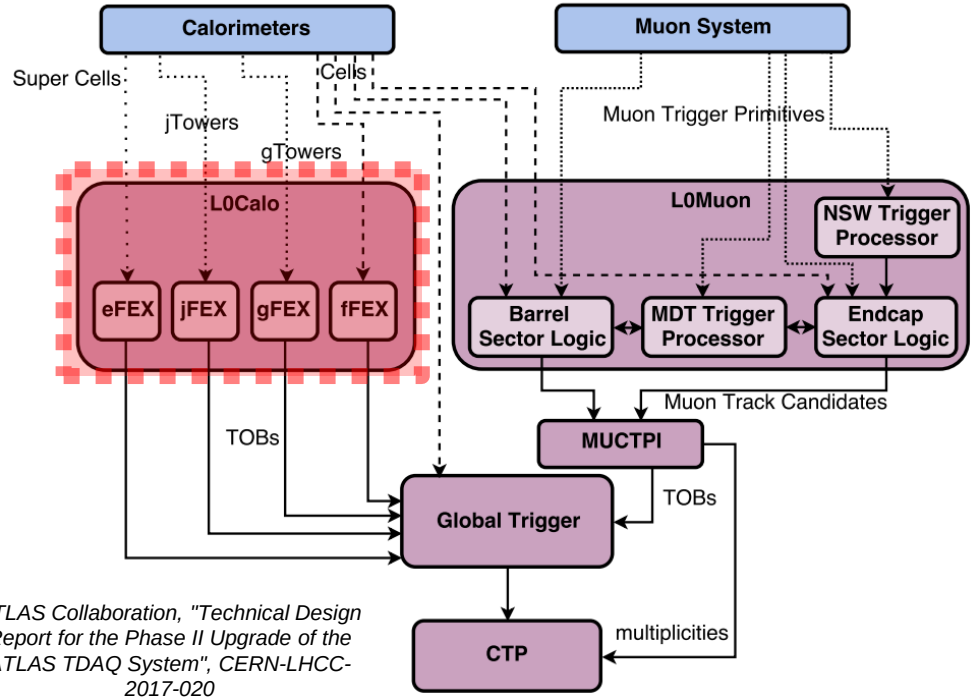

ATLAS Collaboration, "Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System", CERN-LHCC-2017-020

A block diagram of the ATLAS Level-0 trigger system after the Phase-II upgrade

Typical hardware example: fFEX

- The fFEX system consists of four ATCA modules
  - Each hosting two processor FPGAs: AMD Ultrascale+ 13P
- High-speed optical transceivers: Samtec FireFly
  - Support data transmission at speeds up to 25.8 Gb/s per link
  - Real time data path and readout
- Zynq based control mezzanine
  - Configuration, monitoring, slow control
- FPGA power mezzanine modules support rails of up to 100A

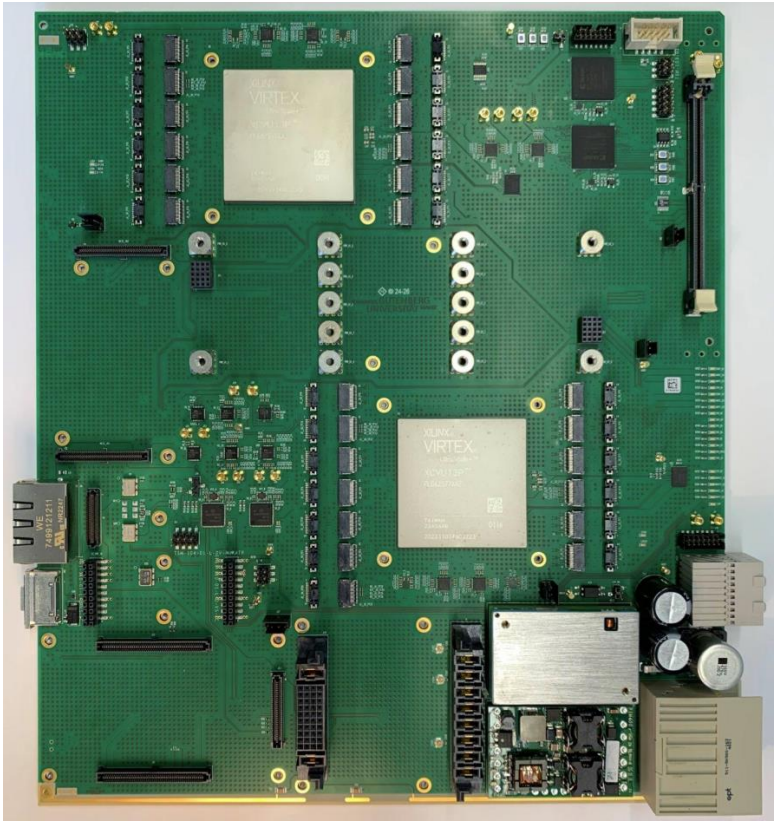

fFEXv1 prototype module hardware overview

fFEX hardware: Signal Integrity

- High-speed PCB design routing techniques
  - High-speed differential pairs adhere to strict physical and spacing constraints
  - Staying within the phase tolerance limit: phase tuning performed
  - Achieving a required differential impedance: in-pair spacing and trace width controlled
  - Minimizing the crosstalk: spacing sufficiently larger than the in-pair spacing used across all pairs

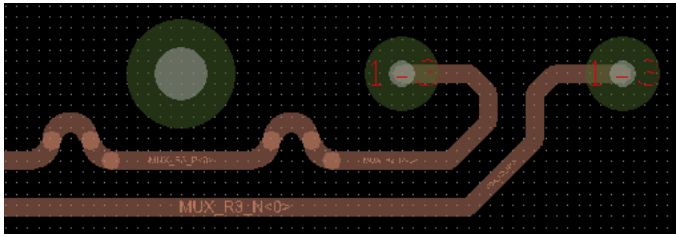

Bump style arc phase tuned high-speed differential pair routed between the optical module and the FPGA

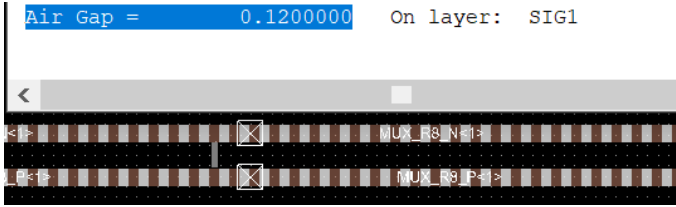

The in-pair spacing (shown as the air gap value) between the traces of a single high-speed differential pair.

Minimum kept spacing (shown as the air gap value) between the neighboring high-speed differential pairs.

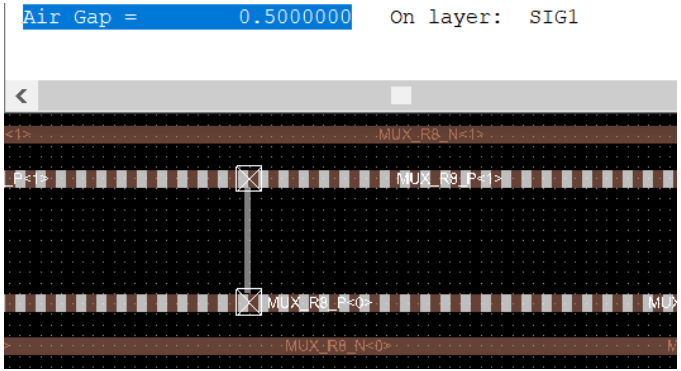

fFEX hardware: Signal Integrity

- High-speed PCB design routing techniques
  - Ground vias in the vicinity of the point where the high-speed track changes its layer from the inner to the outer (or vice versa) in the interface areas

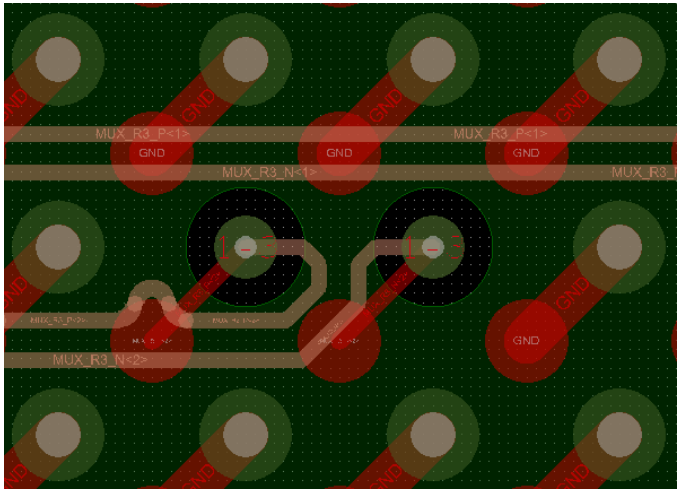

Ground vias surround the point where a high-speed differential pair changes its layer from the inner to the outer

- The high speed stack-up design
  - Minimizing the crosstalk: signal planes shielded by the ground planes
  - Avoiding stubs on the signal lines: high-speed signals occupy the top and bottom inner layers and use microvias
  - Good dielectric constant, ultra-low dielectric loss for high frequencies, highly heat resistant PCB material (EMC EM-890K-89BK) used

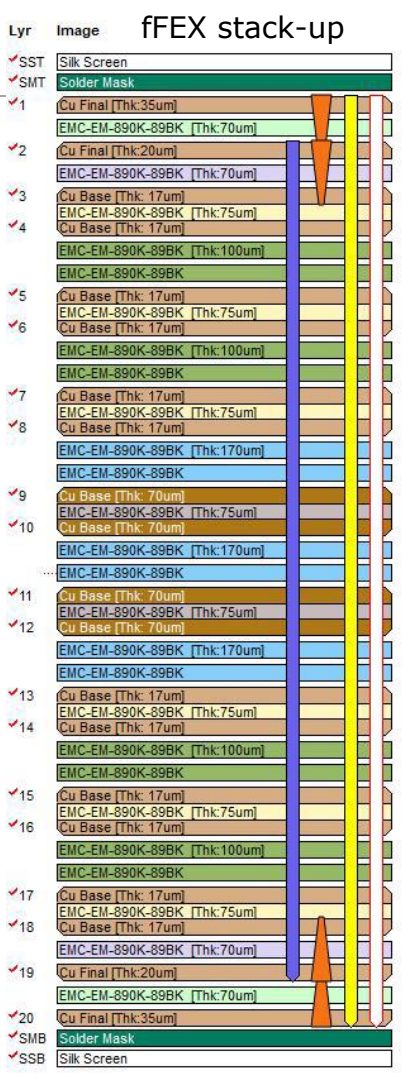

L0Muon and Central Trigger

- L0Muon
  - Forms muon candidates in the barrel and endcap
  - Uses new Processors for the NSW and MDT
  - Occupies 15 ATCA shelves

- Central Trigger
  - Trigger information from barrel and endcap is merged in the MUCTPI before being sent to the Global Trigger and CTP
  - Final L0 decision is made in the CTP based on a trigger menu

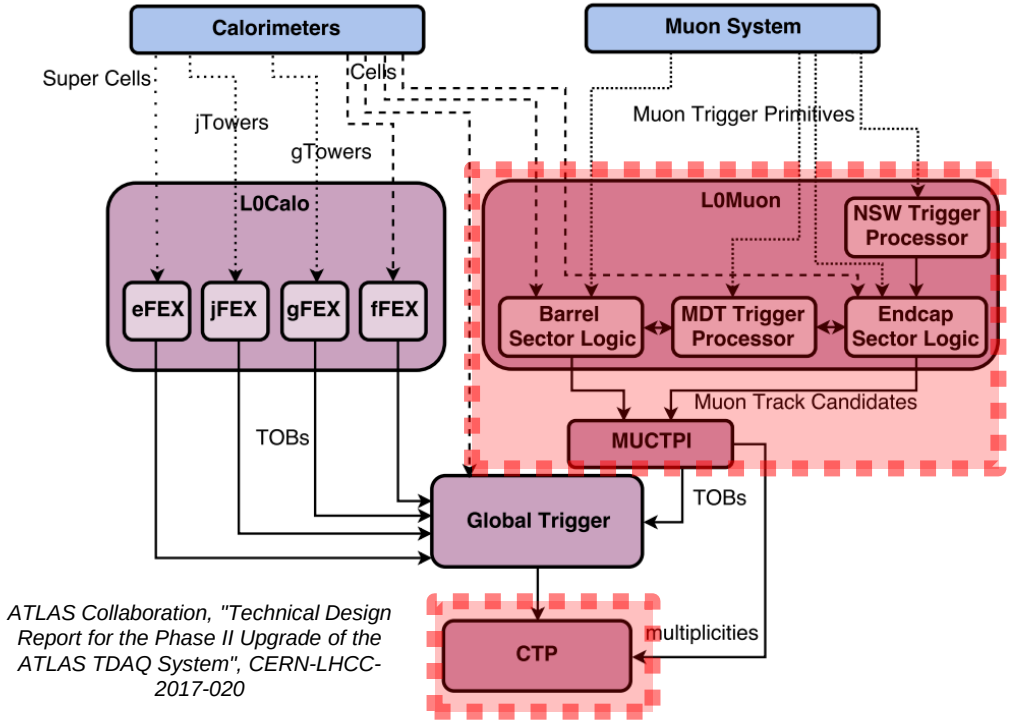

ATLAS Collaboration, "Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System", CERN-LHCC-2017-020

A block diagram of the ATLAS Level-0 trigger system after the Phase-II upgrade

Global Trigger system: Functionality

- As part of the Phase-II Level-0 Trigger System, the Global Trigger **replaces** the Phase-I Topological Processor
- The Global Trigger will absorb the functions of the Phase-I Topological Processor and **significantly extend** them
  - Uses full-granularity calorimeter cells to perform offline-like algorithms
  - Identifies topological signatures
  - Processes the trigger information from the Run 3 hardware systems
  - Transmits the processed trigger information to CTP for final decision

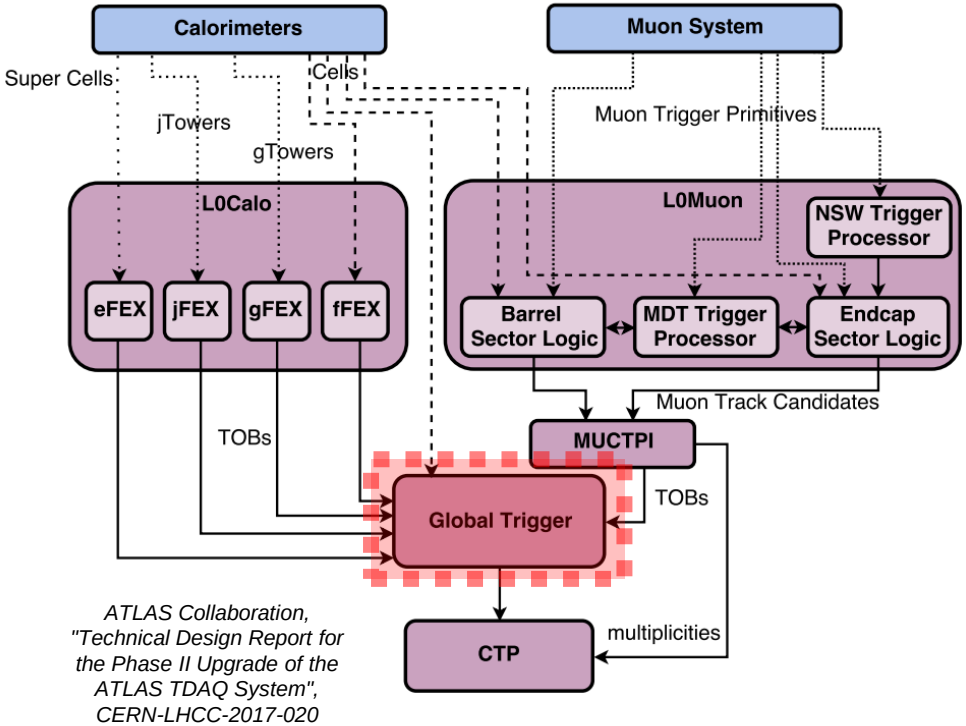

A block diagram of the ATLAS Level-0 trigger system after the Phase-II upgrade

Global Trigger system: Implementation

- Time-multiplexed system concentrates data of full event into a single processor
- Composed of 3 main layers
  - Multiplexing (MUX) layer
  - Global Event Processor (GEP) layer
  - Demultiplexing Global-to-Central Trigger Processor (CTP) Interface
- Will provide a synchronous interface to the rest of ATLAS detector

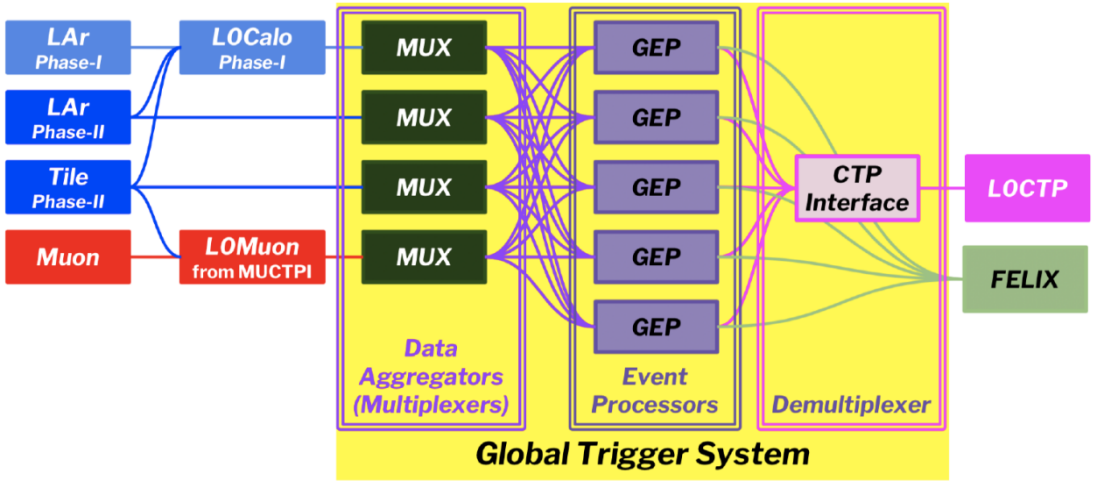

ATLAS Collaboration, "Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System", CERN-LHCC-2017-020

Global Trigger system: Time multiplexing

- 48 MUX nodes receive **real time** data from L0Calo, Calorimeter and MuCTPi every BC and transmit a full event to a single GEP node every BC
- As a result, the latency budget for each event processor is 1.2 µs

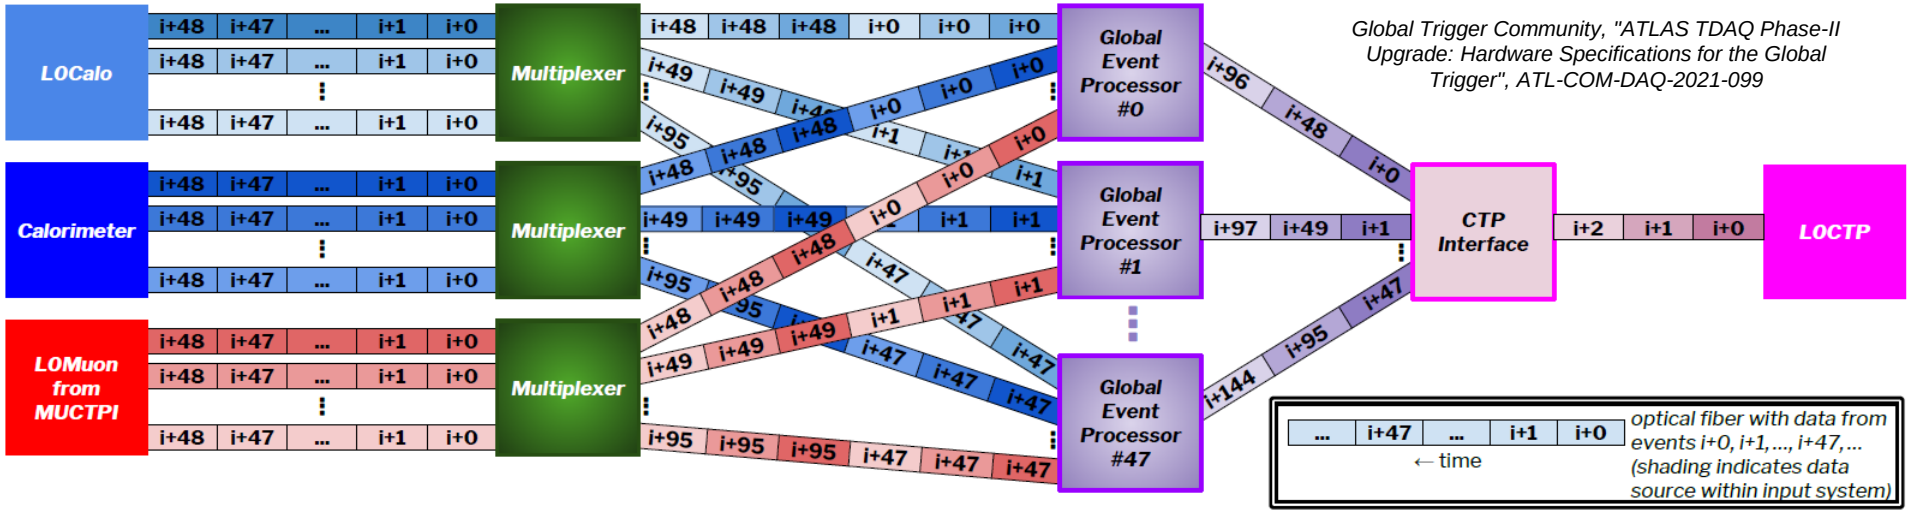

Firmware evolution example

- Abundance of algorithms within the GEP node in order to process the full event data
- Even though each algorithm block in Phase-II has significantly higher latency budget, it also has significant resource limitations

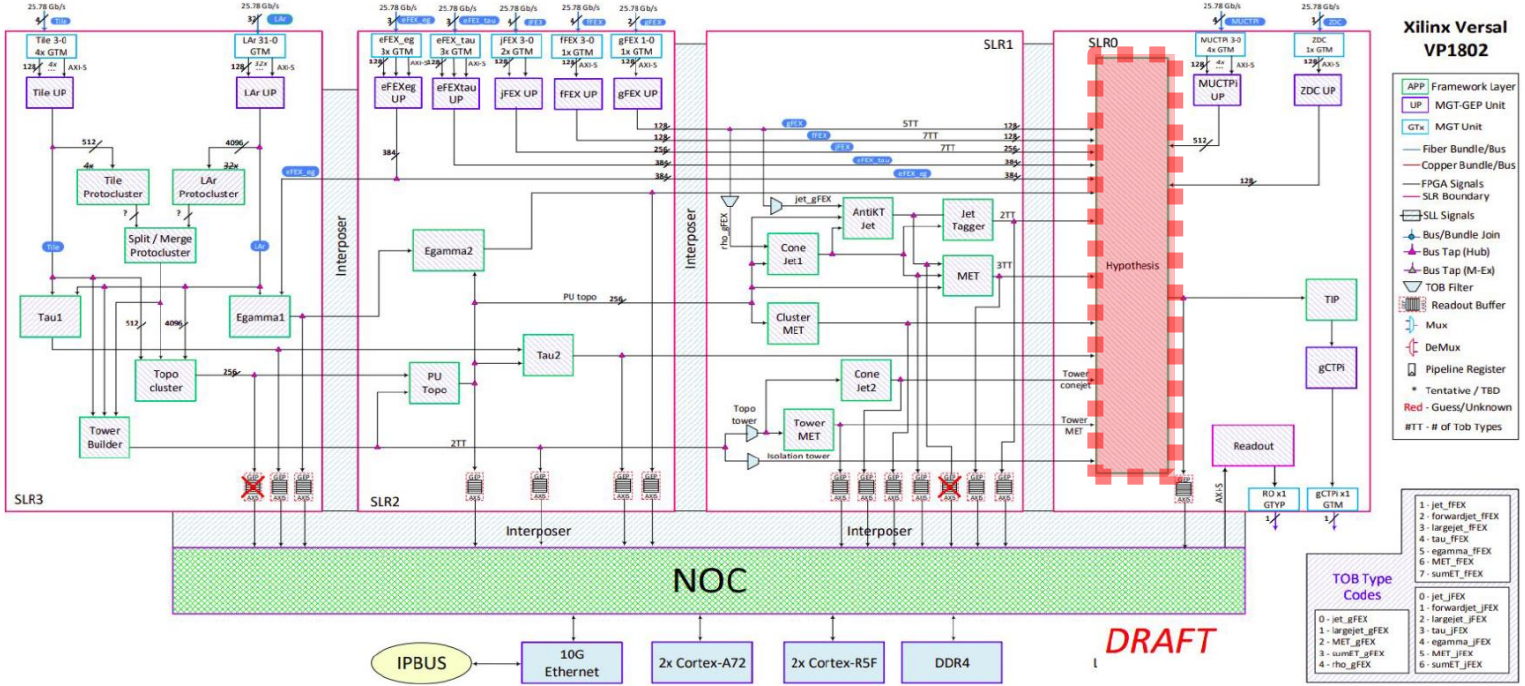

Jeff Eastlack

Recap: Hypothesis (Topo) firmware in Phase-I

- LHC Bunch Crossing synchronous firmware – new event data every 25 ns
- “Select” algorithms select all Trigger Objects (TOBs) passing configurable parameter-based threshold
- “Sort” algorithms output a list of the leading TOBs with the highest Transverse Energy (ET), that pass the thresholds, and sort them by ET
- “Decision” algorithms perform calculations for one or more lists of TOBs, including angular differences, invariant masses, large jet reclustering, and missing transverse energy

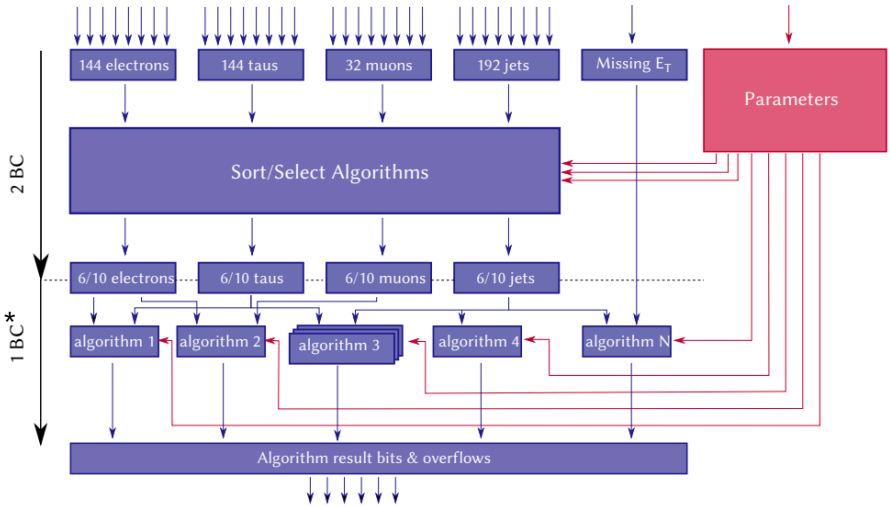

L1Topo “Sort/Select” and “Decision” algorithms structure

* 1 BC ~ 25 ns
BC (Bunch Crossing) – collision between the particle beams

J. Damp, “Search for Dijet Resonances with the Level-1 Topological Processor at ATLAS”, PhD thesis: Johannes Gutenberg-Universität Mainz, 2020

Recap: Hypothesis (Topo) firmware in Phase-I

- **1 BC (25 ns) – very tight latency budget** for Decision algorithms → full parallelization required
- High resource usage as a consequence: **2.5M LUTs** across 6 FPGAs

Algorithm Example

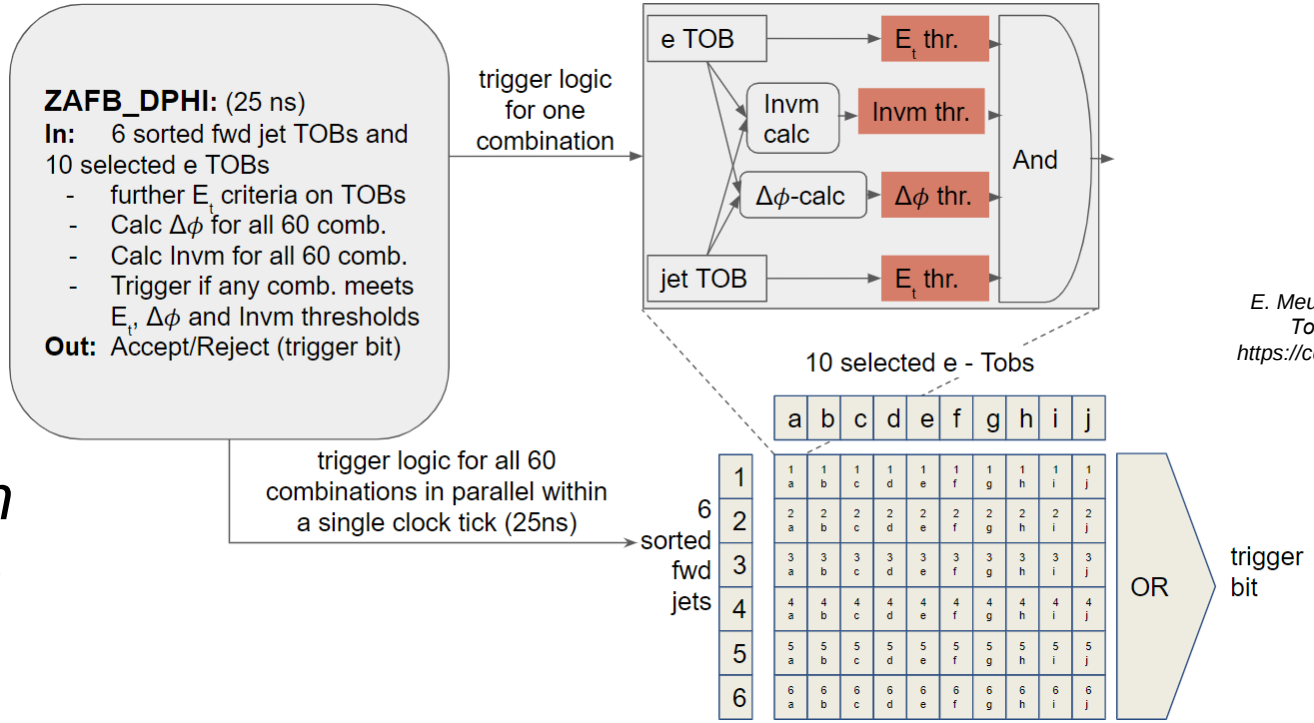

E. Meuser, "The ATLAS Level-1 Topological Processor", https://cds.cern.ch/record/2869237

Hypothesis firmware: Phase-II

- Main strategy – fitting within the tight resource budget at a cost of higher latency
- Instead of processing all combinations in parallel in a single clock tick, processing them sequentially
- Significant resource reduction

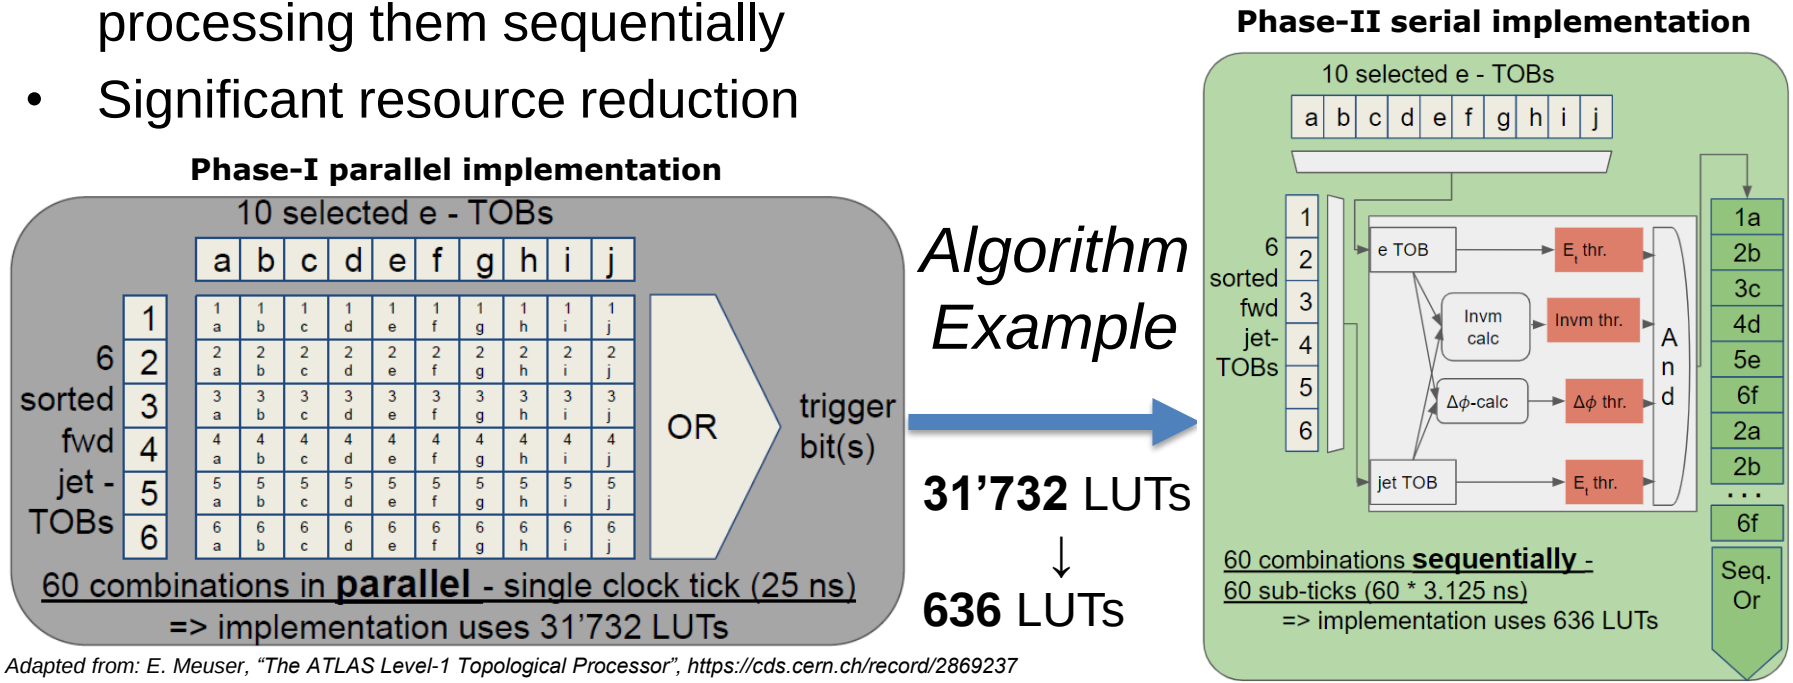

Adapted from: E. Meuser, "The ATLAS Level-1 Topological Processor", https://cds.cern.ch/record/2869237

Global Trigger system: GCM Hardware

- The Global Common Module (GCM) - the main hardware element of the global trigger
- Compose every layer of the Global Trigger

- Standard ATCA form factor
- Two AMD Versal Premium XCVP1802 devices (MUX + GEP / gCTPi)
- 20 Samtec FireFly optical modules
  - Real time data path
  - Readout

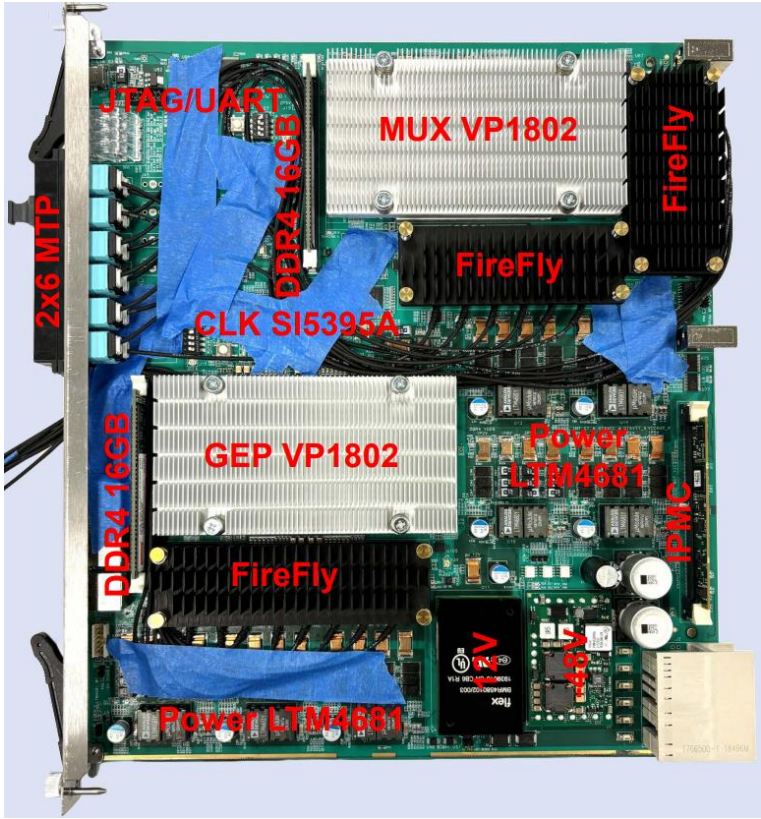

GCMv3 prototype module hardware overview

CMS Phase-II Level-1 Trigger system

- Upgrade of the CMS trigger system is necessary as well due to the significant increase of the instantaneous luminosity of the LHC in Run 4

- Phase-II Level-1 Trigger System
  - Performs real time event selection
  - Reduces the event rate: 40 MHz → 750 kHz (1 MHz in ATLAS)
  - Overall system latency budget 12.5 μs (10 μs in ATLAS)

The Phase-2 Upgrade of the CMS Level-1 Trigger. (2020). [Techreport].

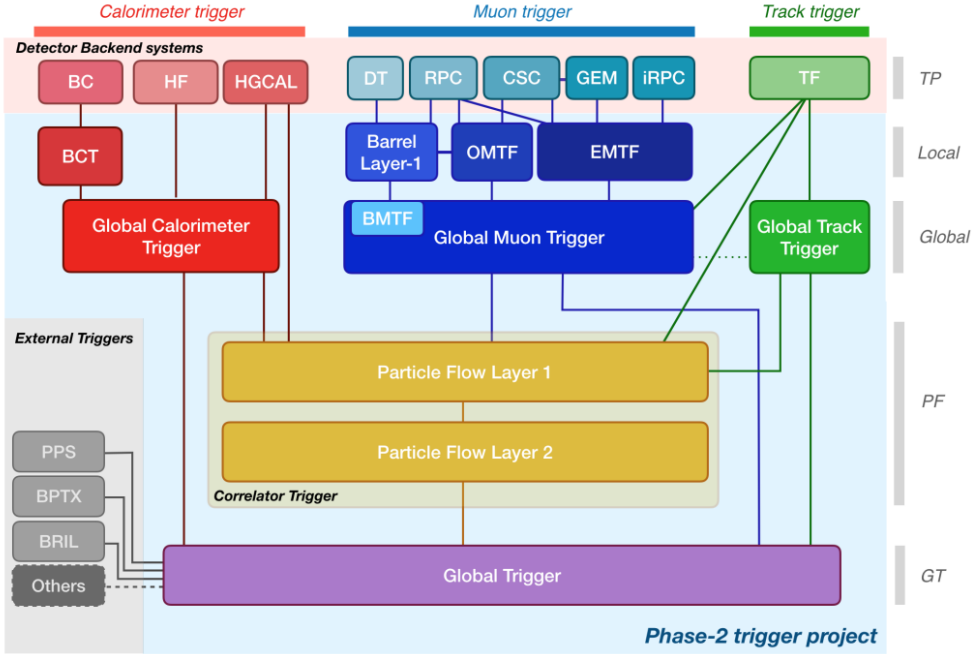

A block diagram of the CMS Level-1 trigger system after the Phase-II upgrade

Calorimeter and Muon triggers

- The muon trigger processes data from both the muon system and tracking system to create the muon trigger objects

- The calorimeter trigger processes information from the electromagnetic and hadronic calorimeters and builds calorimeter trigger objects
  - Includes Regional and Global Calorimeter Triggers

The Phase-2 Upgrade of the CMS Level-1 Trigger. (2020). [Techreport].

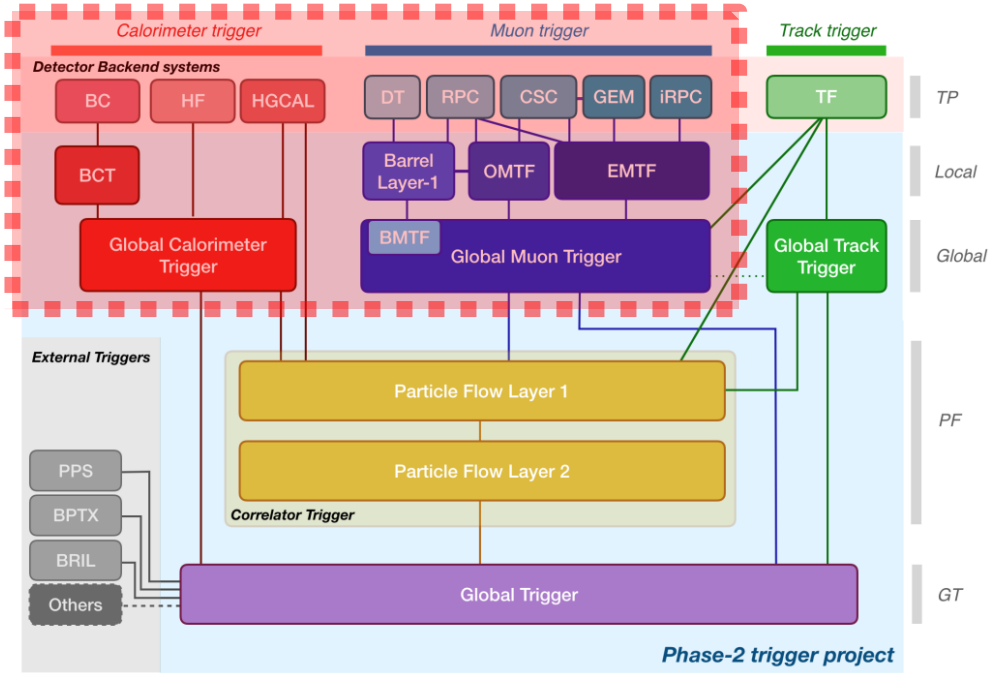

A block diagram of the CMS Level-1 trigger system after the Phase-II upgrade

Track Trigger

- Hardware-based system
  - Uses dedicated pT modules in the outer tracker to form correlated hit pairs (stubs) that are rapidly processed in the track finder modules
  - The GTT processes data from the track finder modules to build high-level track objects
- ATLAS: Track reconstruction is integrated into the Event Filter
  - Uses software-based and heterogeneous computing rather than a dedicated hardware system

The Phase-2 Upgrade of the CMS Level-1 Trigger. (2020). [Techreport].

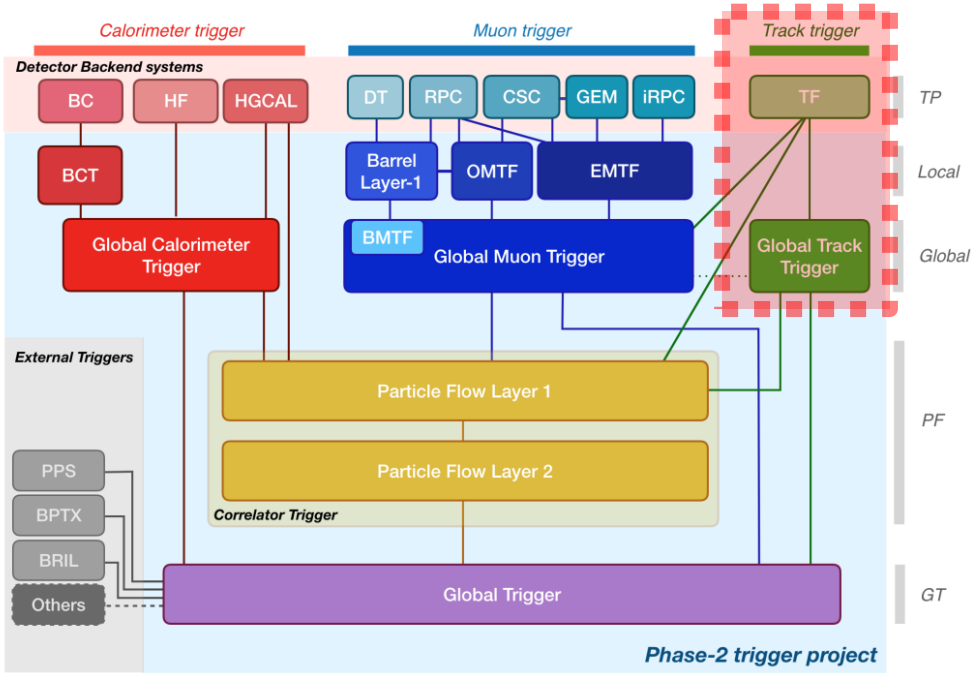

A block diagram of the CMS Level-1 trigger system after the Phase-II upgrade

Correlator and Global Triggers

- The correlator trigger aggregates all the data processed in the calorimeter, muon, and global track triggers to build more accurate trigger objects

- The global trigger receives inputs from all upstream trigger systems and implements the trigger menu
  - Parallel evaluation of O(1000) trigger algorithms that each select a specific event signature

The Phase-2 Upgrade of the CMS Level-1 Trigger. (2020). [Techreport].

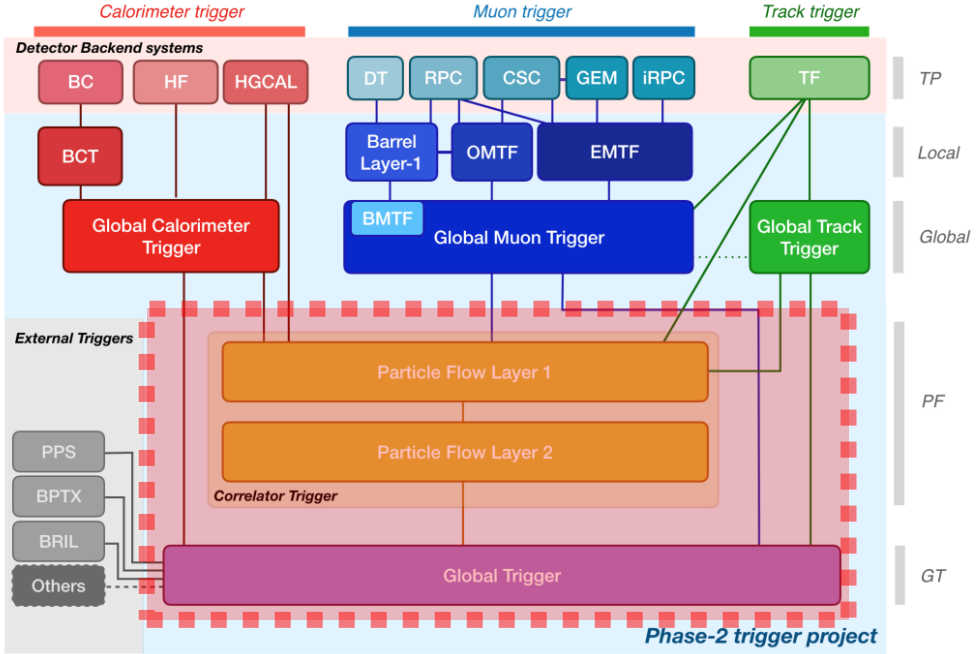

A block diagram of the CMS Level-1 trigger system after the Phase-II upgrade

Scouting System

- The global trigger is reprogrammed several times per year → necessary to detect any data inconsistencies as fast as possible
- The scouting system runs in parallel to the trigger system and takes in a subset of trigger primitives and objects at 40 MHz to verify correct system operation
- New algorithm developments can be prototyped and tested rapidly in the scouting system using the global trigger inputs

- Dedicated hardware platform
  - 2 AMD VU35P FPGAs
  - FireFly optical modules: 24x 25 Gb/s input bandwidth
  - QSFP: 5x 100 Gb/s output bandwidth

APx: Generic Trigger Processing Board

- All subsystems are composed of several types of generic trigger processing boards
- Standard ATCA form factor
- Single Xilinx Virtex UltraScale+ VU13P device
  - Speed grade -2
- Samtec FireFly optical modules
  - 120 TX/RX lanes at 25 Gb/s – real time data path
  - 4 TX/RX lanes at 28 Gb/s – readout
- ZYNQ-7000 based IPMC
- Control mezzanine
  - ZYNQ MPSoC-based

Isobel Ojalvo, "The APx Board for the CMS Phase 2 L1 Calorimeter trigger: Testing and Performance", TWEPP 2024

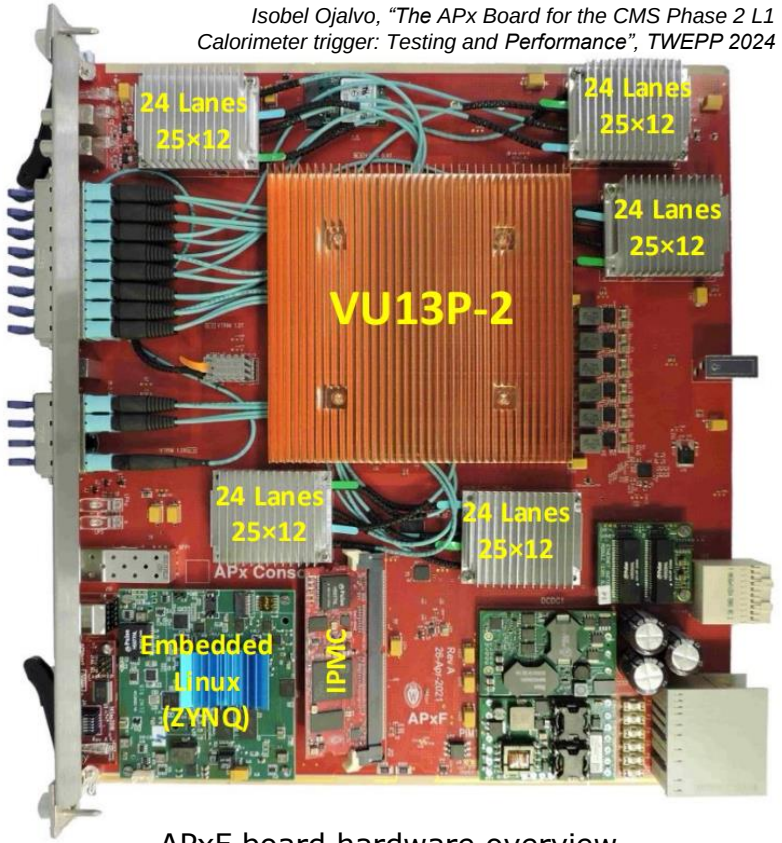

APxF board hardware overview

ATLAS Readout: Architecture

- Front-End LInk eXchange (FELIX) - a data router that serves as an interface between the detector electronics and commodity computing
  - Readout, configuration, trigger, clock distribution, monitoring

- Serves all sub-detectors
- Consists of 350 servers with new custom PCIe FELIX cards
- Handling data at 1 MHz readout rate with a total throughput of 4.6 TB/s.

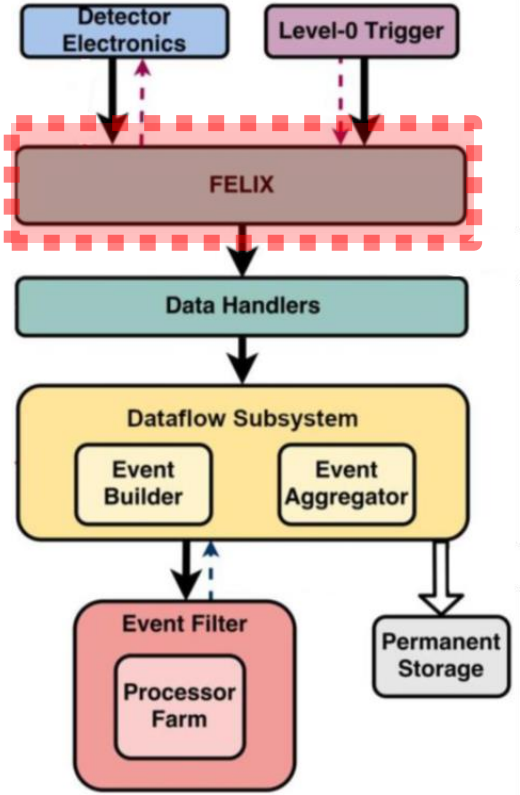

A block diagram of the ATLAS TDAQ system after the Phase-II upgrade

ATLAS Readout: FELIX Hardware

- FLX-182
  - AMD Versal Prime VM1802
  - PCIe Gen 4 x 16 (240Gb/s)
  - 24 bidirectional optical links (25Gb/s)
- FLX-155
  - AMD Versal Premium VP1552
  - PCIe Gen 5 x 16 (480Gb/s)
  - 48 Bidirectional optical links (25Gb/s)

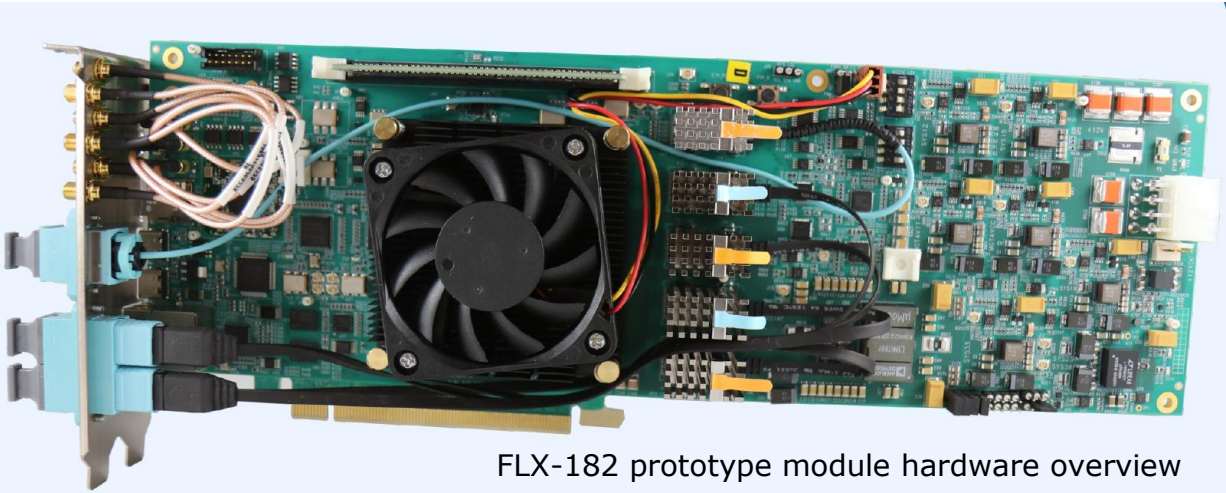

FLX-182 prototype module hardware overview

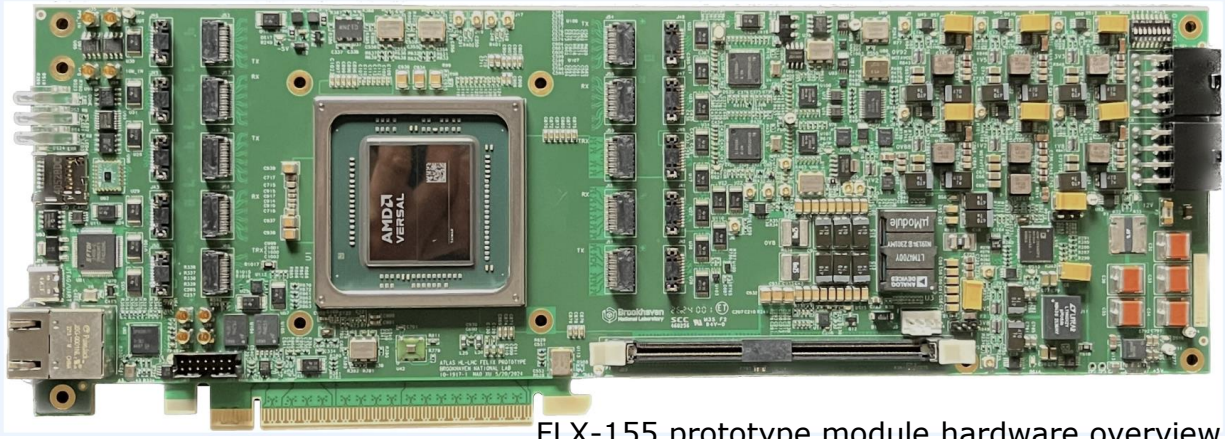

FLX-155 prototype module hardware overview

CMS Readout: Architecture

- Main differences with ATLAS
  - 0.75 first level trigger rate
  - x3 links (lp)GBT: 50000 at up to 10 Gb/s
  - x1.4 data throughput
- The on-detector front-end electronics sends the physics data to off-detector back-ends

The Phase-2 upgrade of the CMS data acquisition and High Level Trigger. (2021). [Techreport].

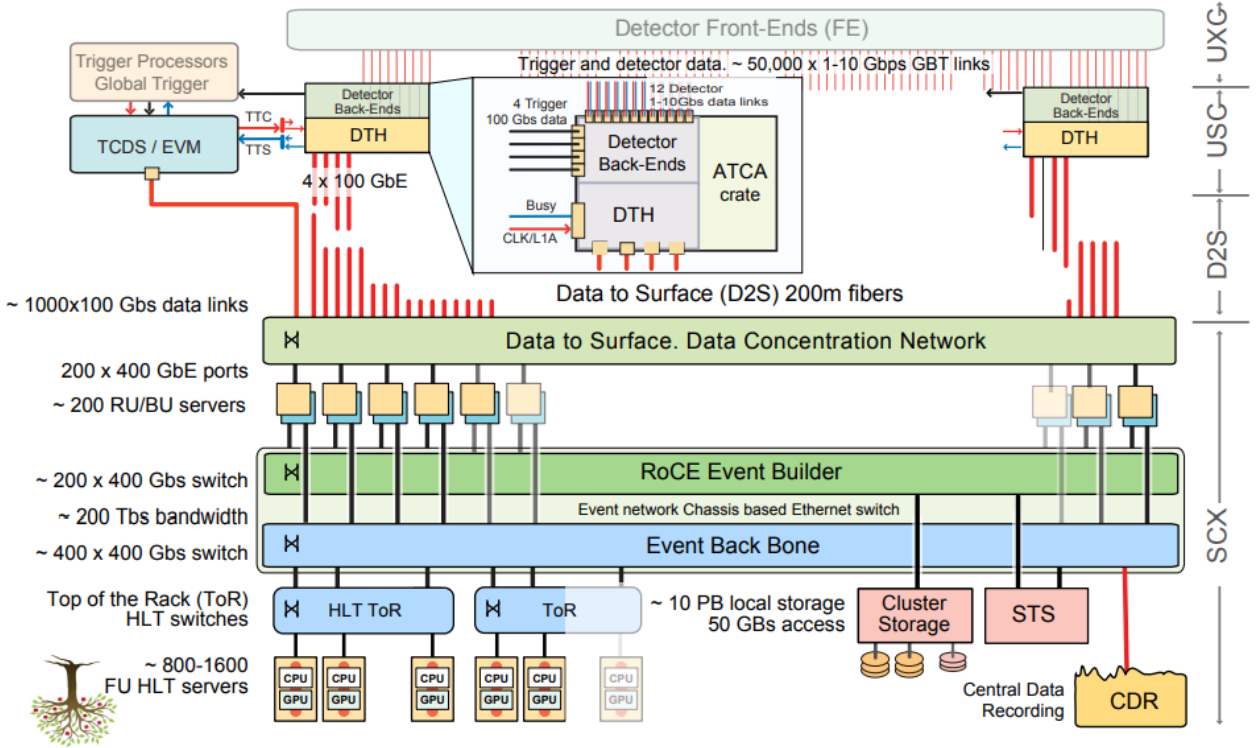

Block diagram of the CMS Plase-2 DAQ

CMS Readout: Hardware

- The back-end boards
  - Standard ATCA form factor
  - Specific to sub-detector
  - Implement communication with the front-ends and forward the received physics data via point-to-point optical links to a common DAQ board: the DAQ and Timing Hub (DTH400).
- DTH400
  - Two AMD VU35P FPGAs
  - Aggregates the data over a full orbit of the LHC
  - Transmits them to the computing centre

The Phase-2 upgrade of the CMS data acquisition and High Level Trigger. (2021). [Techreport].

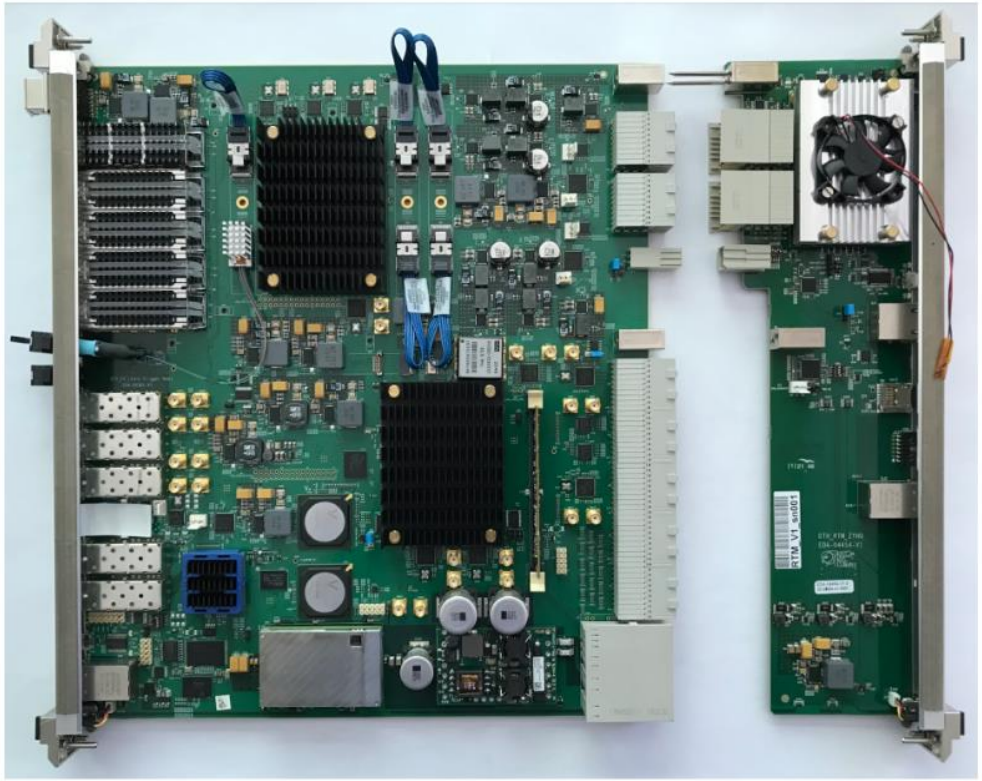

The DTH400 board prototype-2 together with the Rear Transition Module (RTM) housing the embedded controller

Summary

- Trigger and readout electronics for ATLAS and CMS make use of high end FPGAs and optical modules
- A robust, scalable and high-performance platform provided by the ATCA standard supports demanding requirements in terms of power, PCB real estate, modularity and reliability
- PCIe standard choice for FELIX allows moving much of the data processing from custom hardware into flexible software running on standard server platforms, reducing cost and improving scalability and maintainability
- Increasing usage of the common hardware concept shifts the development efforts towards the firmware which often implements various functionalities of the identical hardware modules

Backup

Phase-I Upgrade of the ATLAS detector

- The Phase I upgrade took place during the Long Shutdown 2: 2019-2022
- The following Run 3 is in operation: 2022 to 2025
- Peak luminosity of $2x10^{34}$ $cm^{-2}s^{-1}$

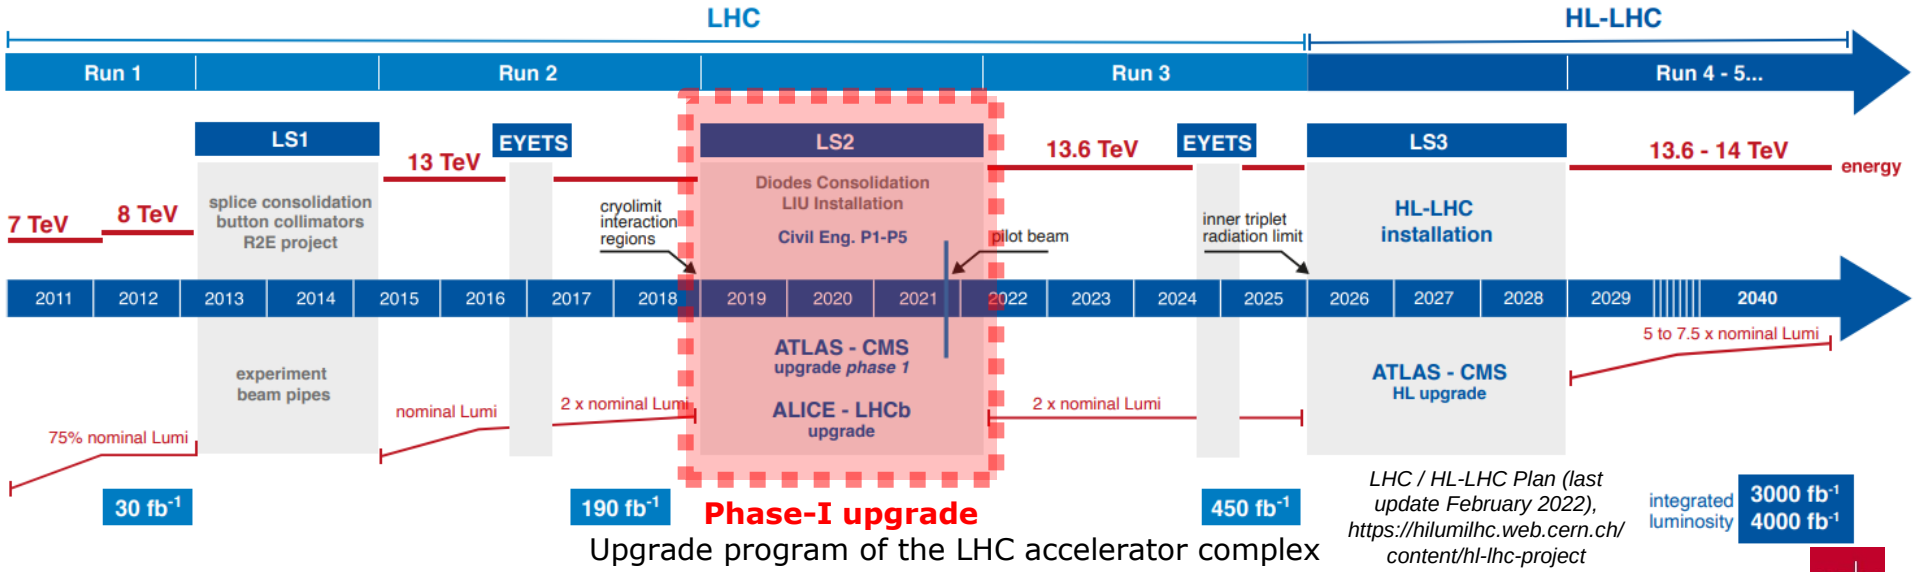

Upgrade program of the LHC accelerator complex

LHC / HL-LHC Plan (last update February 2022), https://hilumilhc.web.cern.ch/content/hl-lhc-project

Phase-I Level-1 Trigger System

- The increased instantaneous luminosity of the LHC in Run 3 → upgrade of the ATLAS trigger system is necessary
- Phase-I Level-1 Trigger System
  - Performs real time event selection
  - Reduces the event rate: 40 MHz → 100kHz
  - Staying below the maximum readout rate of the ATLAS detector
  - Overall system latency budget 2.5 μs

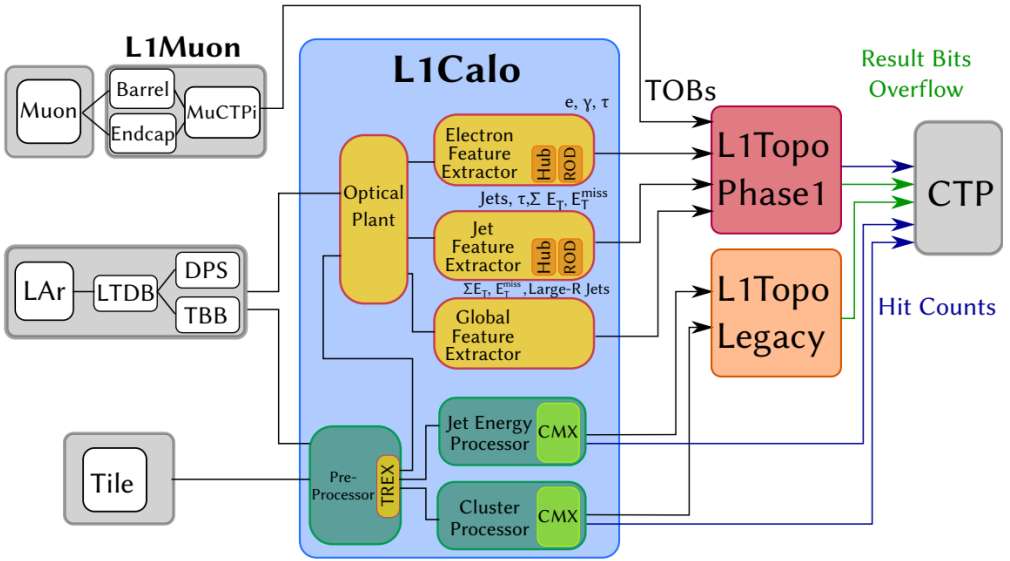

A block diagram of the Level-1 trigger system after the Phase-I upgrade

ATLAS Collaboration, "Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System", CERN-LHCC-2013-018

Phase-I L1Topo system: Functionality

- As part of the Level-1 trigger system, the Level-1 topological trigger (L1Topo) processes data on the **real-time** data path from the individual Feature Extractors (FEXes) and the upgraded Muon to Central Trigger Interface (MUCTPI) to perform topological triggers as well as triggers, counting number of objects
- Provides high processing capabilities in order to make use of the input objects with increased granularity from the new FEXes and the MUCTPI

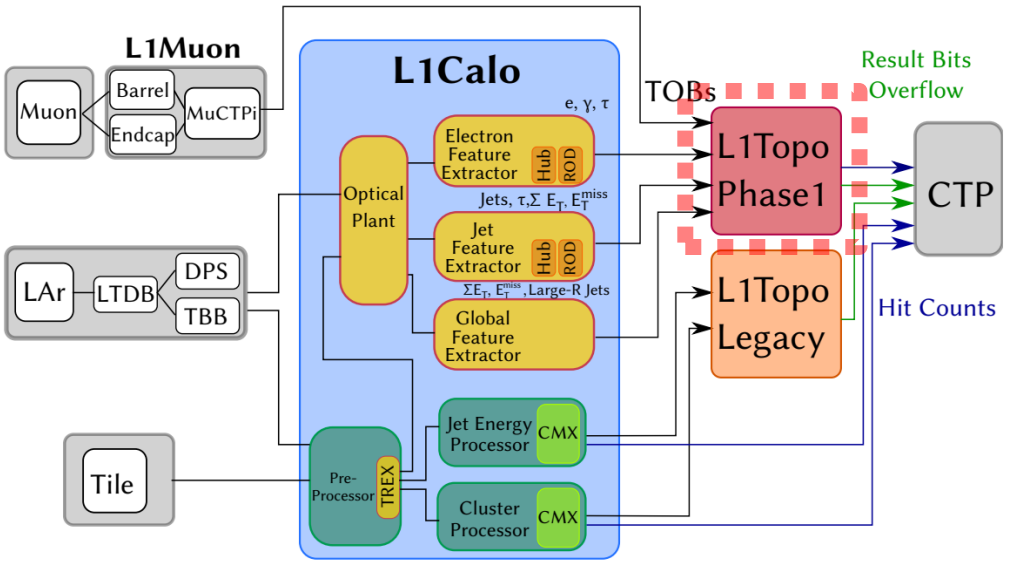

A block diagram of the Level-1 trigger system after the Phase-I upgrade

ATLAS Collaboration, "Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System", CERN-LHCC-2013-018

Phase-I L1Topo system: Hardware

- The L1Topo system consists of three ATCA modules
  - Each hosting two processor FPGAs: Xilinx Ultrascale+ 9P
- High-speed optical transceivers: Avago MiniPODs
  - Support data transmission at speeds up to 11.2 Gb/s per link
- Zynq based control mezzanine
  - Configuration, monitoring, slow control
- Similar hardware building blocks as on the jFEX module

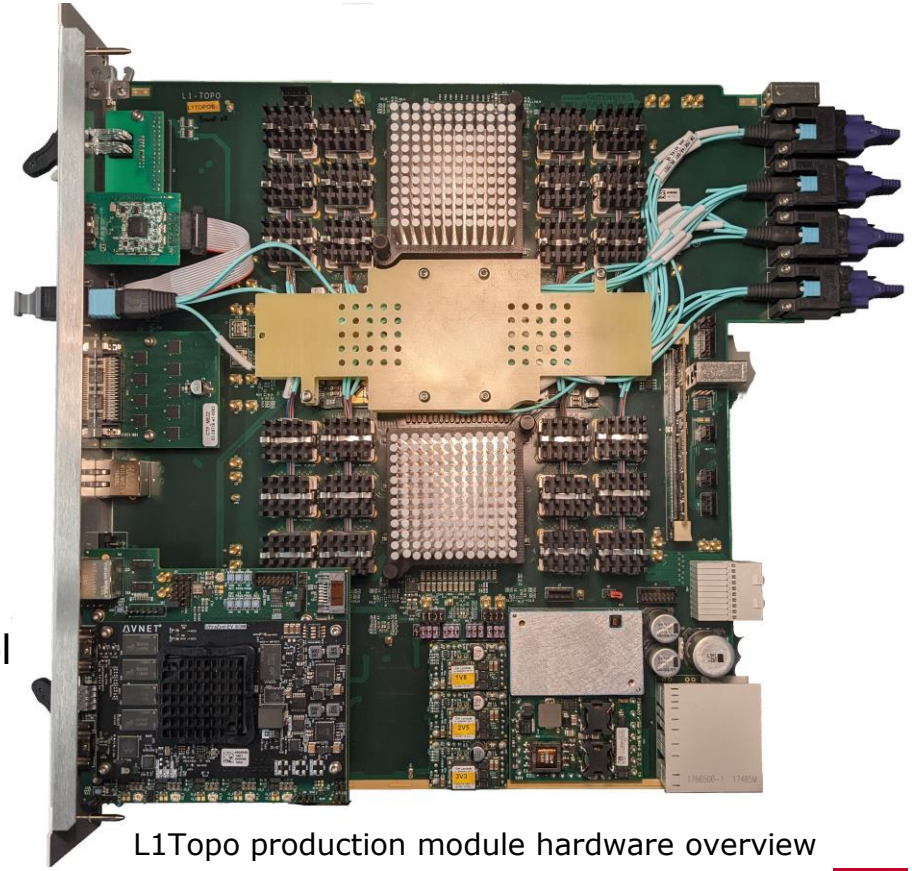

L1Topo production module hardware overview

Phase-I L1Topo system: Firmware

- The algorithms are automatically assembled and configured based on the provided trigger menu
- The algorithm parameters can be set and changed via the IPBus by the Online Software during a Run
- The topological trigger configuration is fully described in

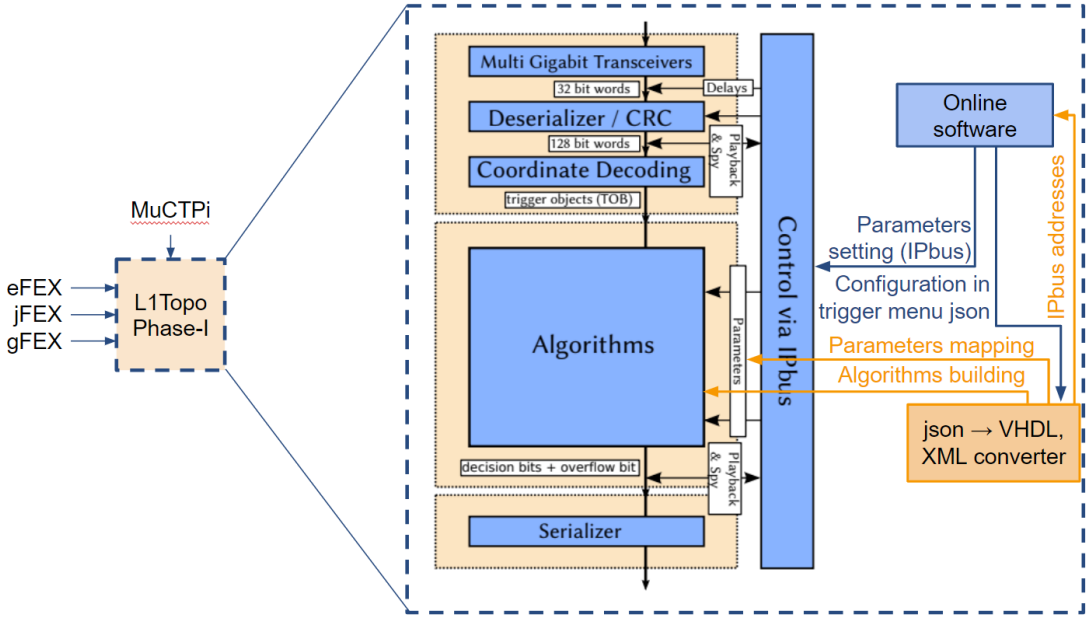

in a single menu-driven json file, from which algorithm VHDL code, as well as IPbus address mapping, are automatically generated

- Consistency between the firmware and the software is ensured
- Menu change may require a firmware rebuild, leaving little time for testing

Phase-I L1Topo system: Commissioning

- The Phase-I L1Topo system has been fully commissioned with the rest of the new L1 trigger systems in ATLAS
- Main commissioning challenges due to 4 different input sources
  - Different input format of TOBs
  - Different granularity of TOB coordinates
  - Complicated detectors' geometry
  - Different time of TOBs' readiness
- Comparison against software implementation
  - Debugging with playback / spy
  - High statistics continuous online monitoring
- The Phase-I L1Topo system has come into routine operation taking data in 2024

Phase-I L1Topo system: First performance results

- L1Topo chains provide about 70 % of unique rate for J/Ψ and ϒ candidates

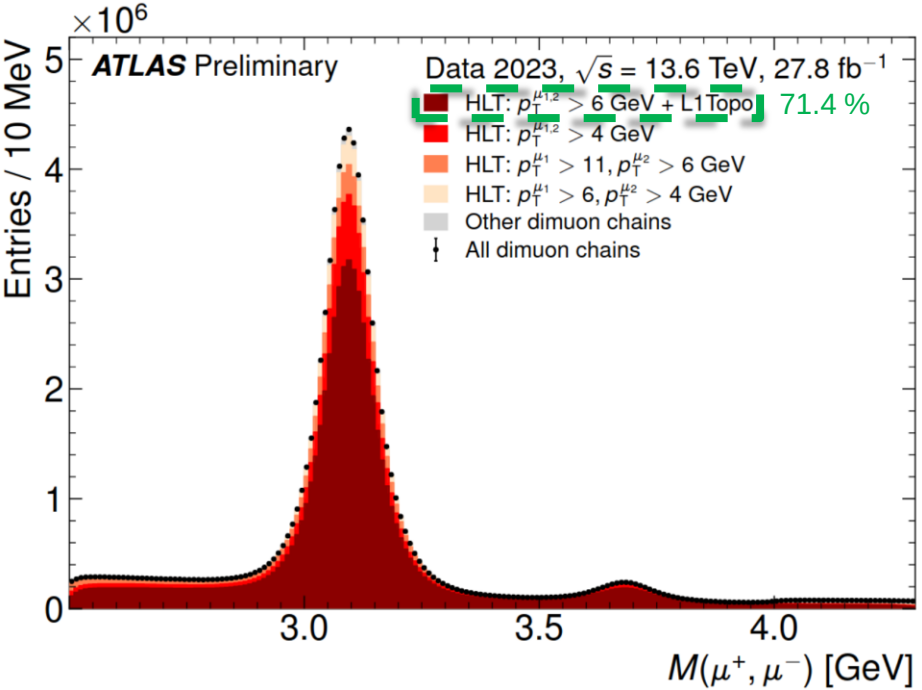

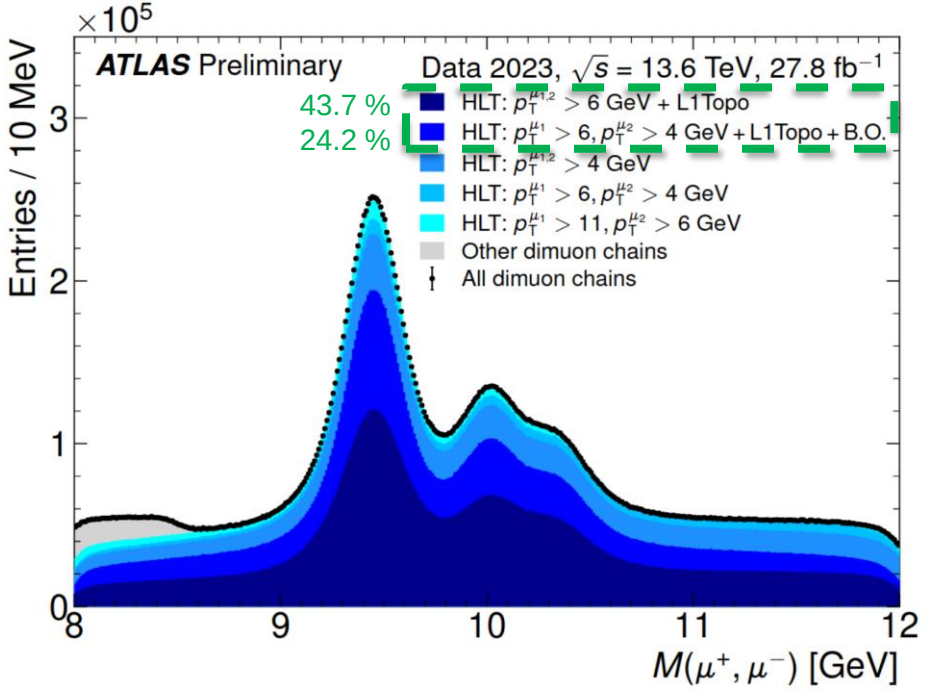

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BPhysicsTriggerPublicResults

ATLAS detector

- Particle beams collide every 25 ns (frequency of 40 MHz)
  - Bunch Crossing – time between the collisions (25 ns)

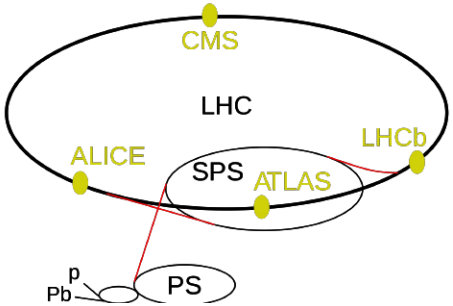

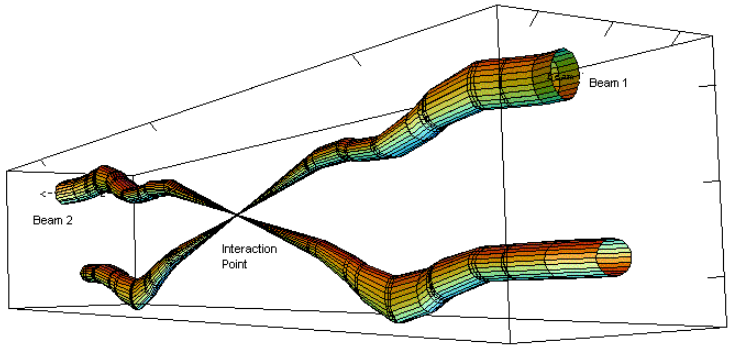

Relative beam sizes around IP1 (Atlas) in collision