The Art of Finding Needles in Haystacks
Event Selection in ATLAS

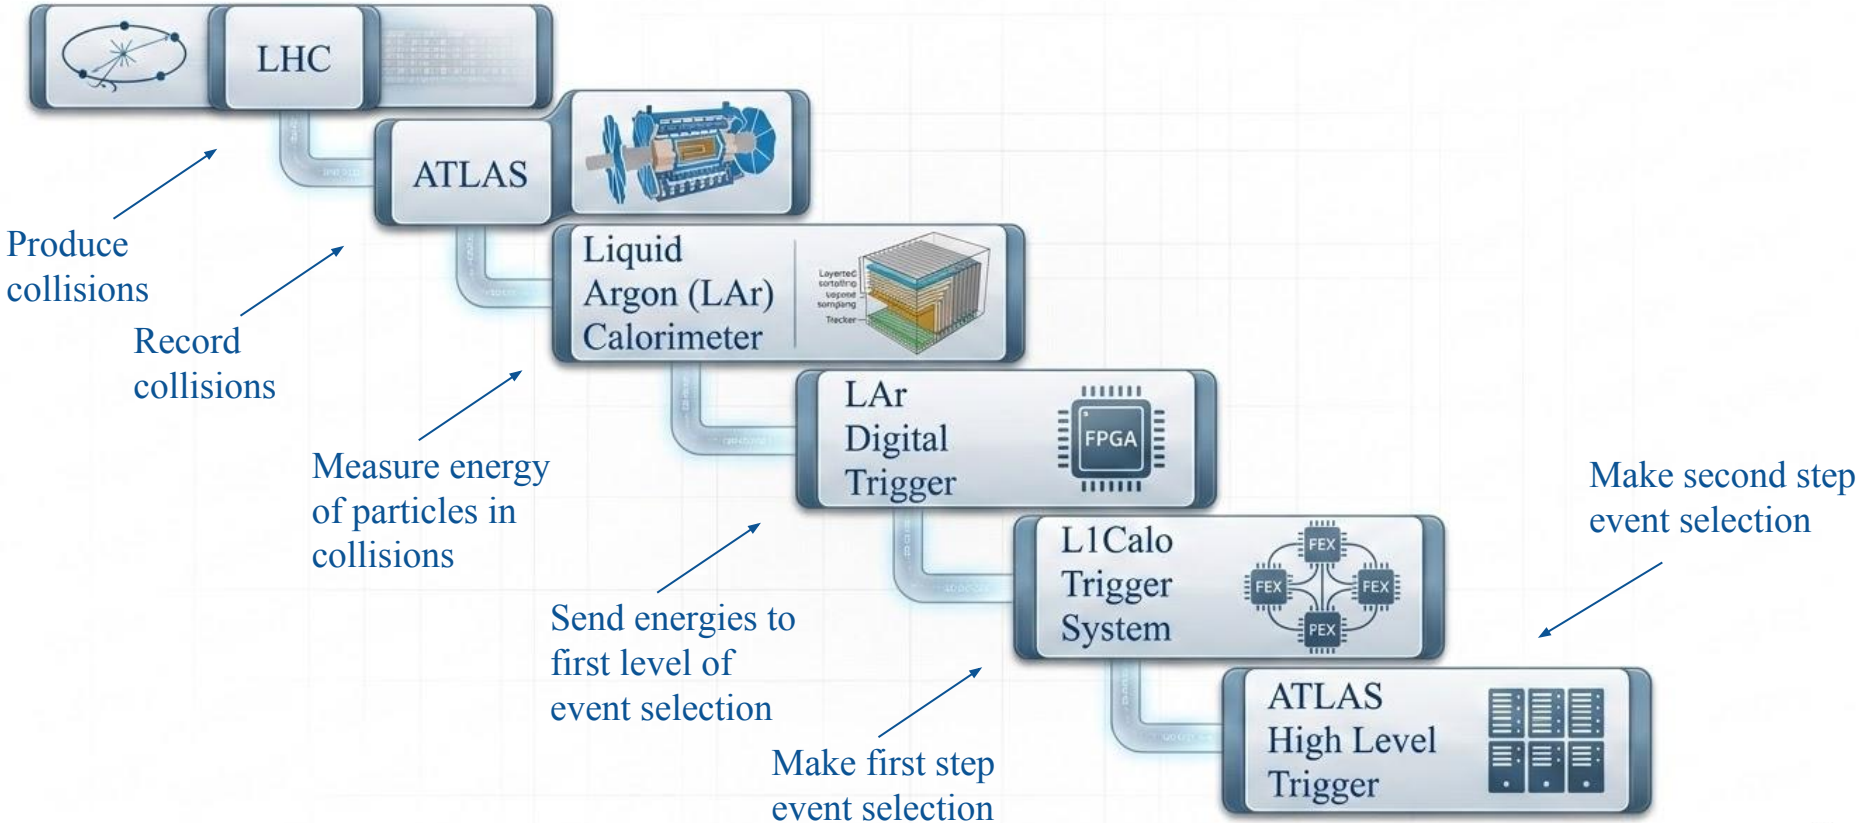
October 23, 2026

Adriana Milić





Overview


LHC
ATLAS
Liquid Argon (LAr) Calorimeter
LAr Digital Trigger
FPGA
L1Calo Trigger System
FEX
FEX
FEX
FEX
ATLAS High Level Trigger
Produce collisions
Record collisions
Measure energy of particles in collisions
Send energies to first level of event selection
Make first step event selection
Make second step event selection

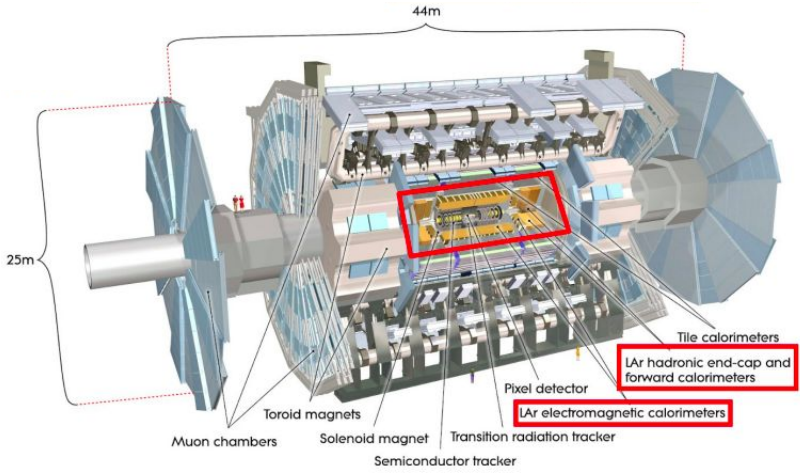
LHC & ATLAS

ATLAS and LHC at CERN

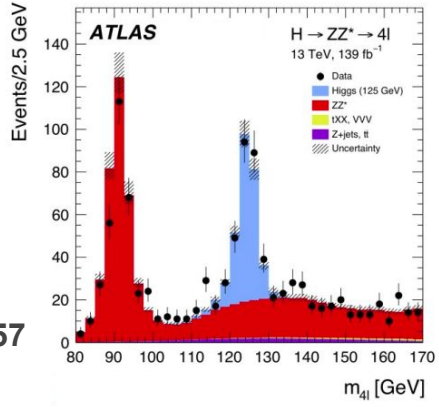
- **Large Hadron Collider (LHC)**
  - Largest and highest energy particle accelerator
  - Since 2010 in operation

- **ATLAS** (~5000 members) is one of the two **large, general-purpose experiments at the LHC.**

- **Main physics goals:**
  - Understand origin of mass: search for a **Higgs boson** ✓
  - **Precision measurements** of the Standard Model (SM) of particle physics
  - Search for **physics beyond the SM**
    - Often very rare processes
  - Sometimes **branching ratios of** $\mathcal{B} < O(10^{-7})$!



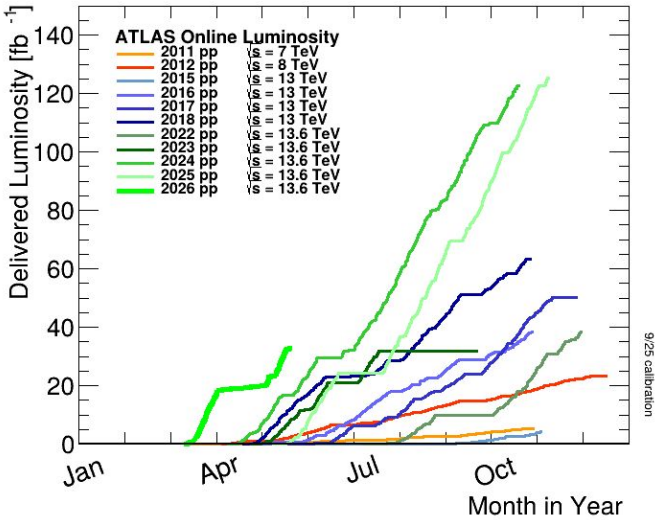
JINST 3 (2008) S08003

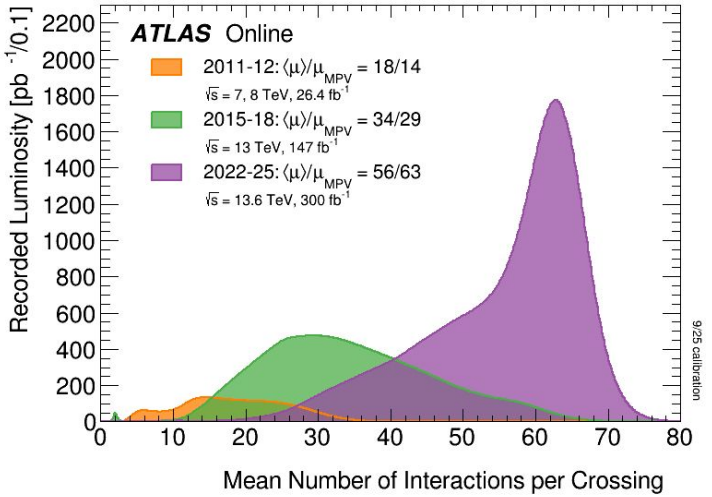


Eur. Phys. J. C80 (2020) 957

ATLAS and LHC at CERN

- **Proton-proton collisions** at **13.6 TeV or ion collisions at 5.46 TeV**
- Nominal LHC **bunch crossing frequency: 40 MHz** (25 ns spacing)
  - 40 Million events per second!
  - Need to pick interesting events and reduce data by factor ~5000 within microseconds.
- **Pileup:** multiple interactions per bunch crossing
- **Extremely challenging environment** for the experiments



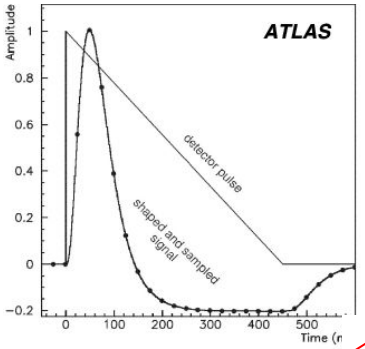


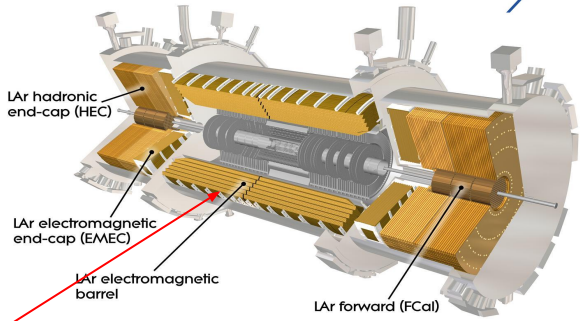
[LuminosityPublicResultsRun3]

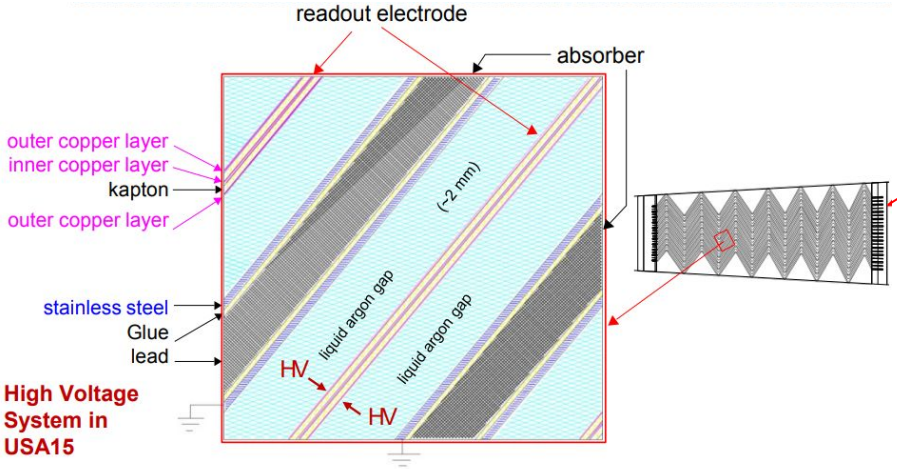
Liquid Argon calorimeter

LAr calorimeter design

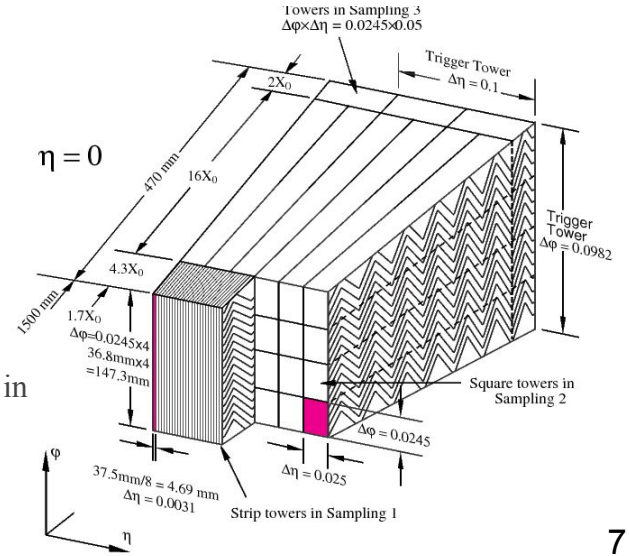
- LAr is a **sampling calorimeter** → layers of absorber plates initializing particle showers are alternated with layers of active material (LAr) which perform energy measurements.
- Particle showers produce a triangular ionization signal.
- Signal is amplified, shaped, and digitized.



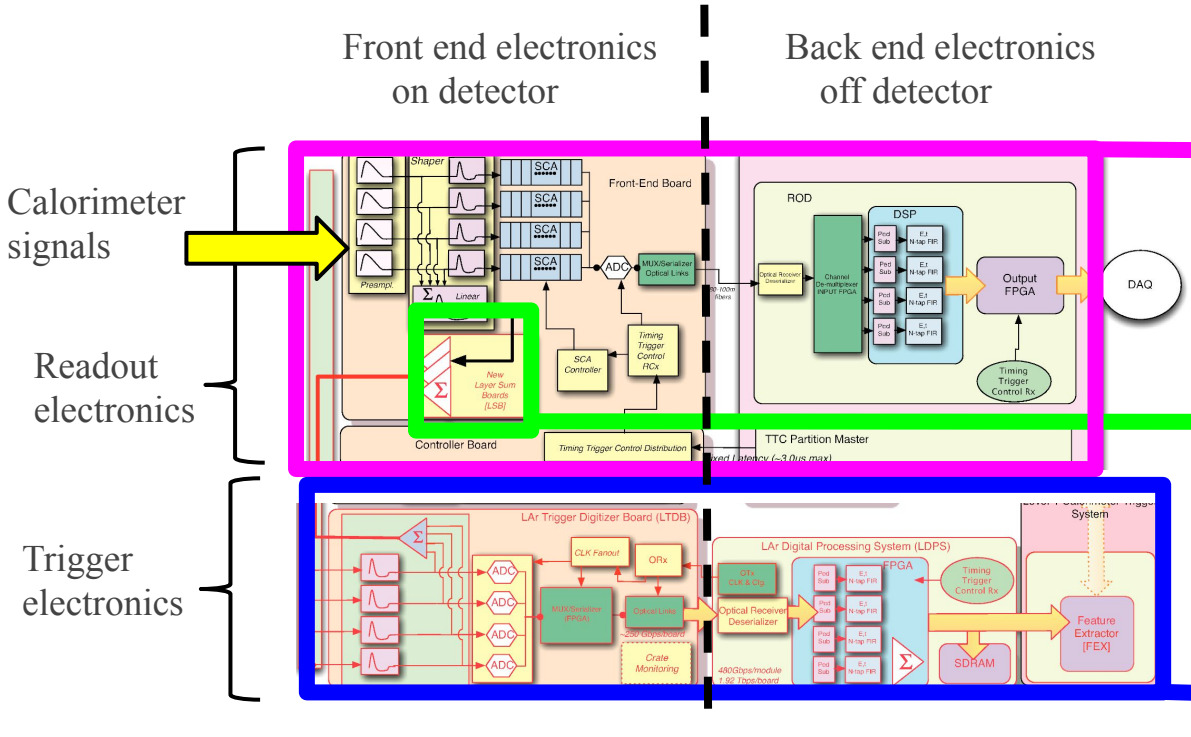




Calorimeter is longitudinally separated in layers with different granularity.



Readout and trigger electronics



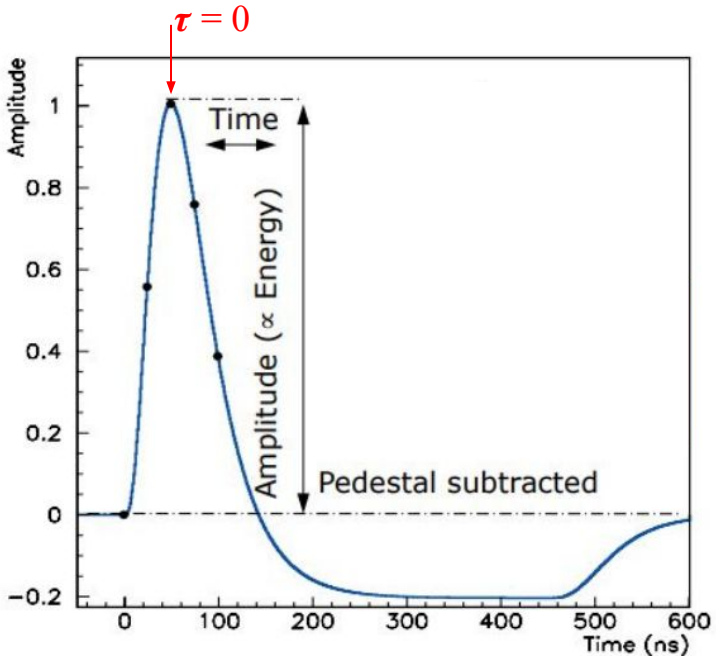
Operating at <2µs latency!
Links ~11 Gbps

Calorimeter cell signals (~200k) are shaped, digitized in the front end. Energy and timing computed in the backend, FPGA based.

Cell signals are summed to trigger signals (~34k) and sent to the trigger path.

Trigger signals are digitized in front end. Energy and timing computed in back end, and sent to the first level trigger (all FPGA based).

Energy computation



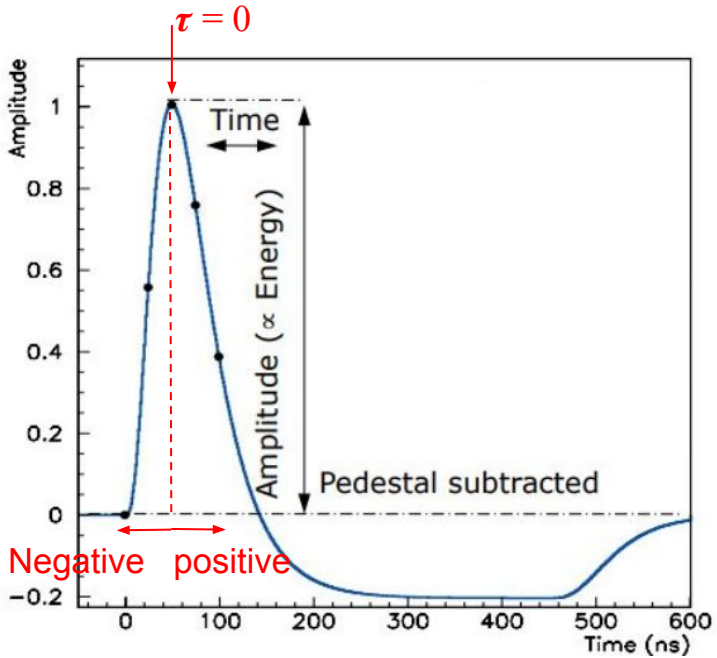
- Optimal Filtering Coefficients (OFCs, A_i, B_i) are derived from the ideal pulse shape & noise autocorrelation matrix
- Applying OFCs gives
  - Amplitude ($\propto$ energy) of the pulse
  - Time-shift between pulse & ideal pulse shape

$$E_T = \sum_{i=0}^{3} A_i (S_i - p - b)$$

$$E_T \tau = \sum_{i=0}^{3} B_i (S_i - p - b)$$

A_i Coefficients for energy reconstruction
S_i Sampled signal
B_i Coefficients for timing
p Pedestal level
b Baseline correction

Timing computation



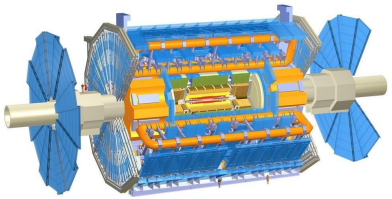
- Electronic boards send samples at 40 MHz → Not trivial to know at what time energy deposit occurred.
- E_T and timing τ computed for 34k trigger units every bunch crossing (@ 40MHz) with a fixed latency
- τ criteria applied to assign correct bunch crossing

$$\begin{cases} -8\ \text{ns} < \tau < 14\ \text{ns}, & \text{if } E_T \geq 10\ \text{GeV} \\ |\tau| < 8\ \text{ns}, & \text{if } E_T < 10\ \text{GeV} \end{cases}$$

- Send E_T for each trigger unit upstream to first level trigger only if τ criteria are satisfied, otherwise $E_T = 0$

ATLAS trigger system

ATLAS trigger system



Raw data from
detector at 40 MHz

Level 1 trigger (hardware):
Muon detectors and calorimeter cells with reduced granularity used,
Regions of Interest (RoI) defined

100 kHz

Level 2 trigger:
Seeded by Level 1 RoI
Full granularity of the detector used

2-3 kHz

Event Filter:
full event building

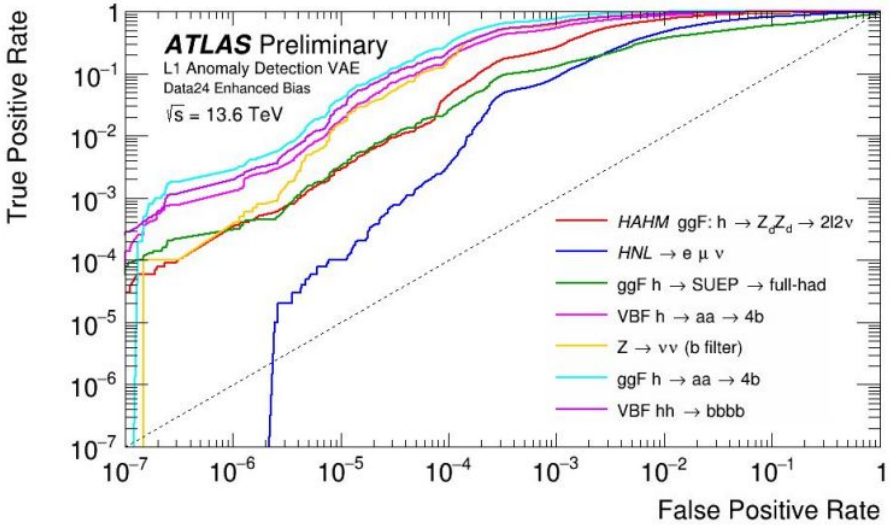
2.5 µs ~10 ms ~1 s

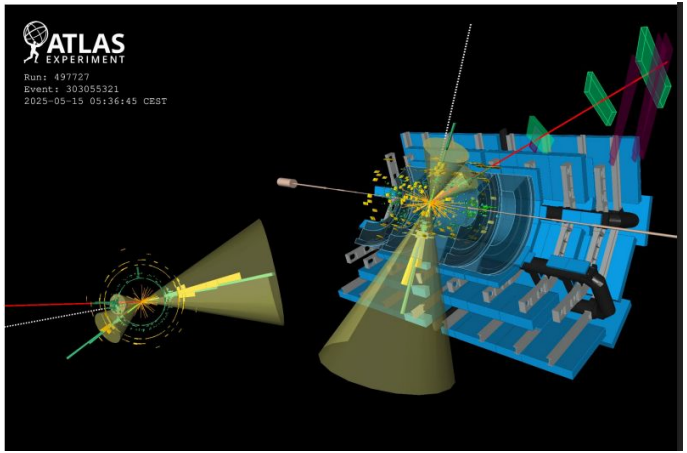
Execution time

Machine learning and neural networks in first level trigger

- **Already deployed in the past data taking period (“Run 3”, 2022-2026)**
- **First hardware based ML algorithm in ATLAS: BDT in eFEX**
  - Algorithm implemented in FPGA for triggering on hadronic τ leptons using BDT

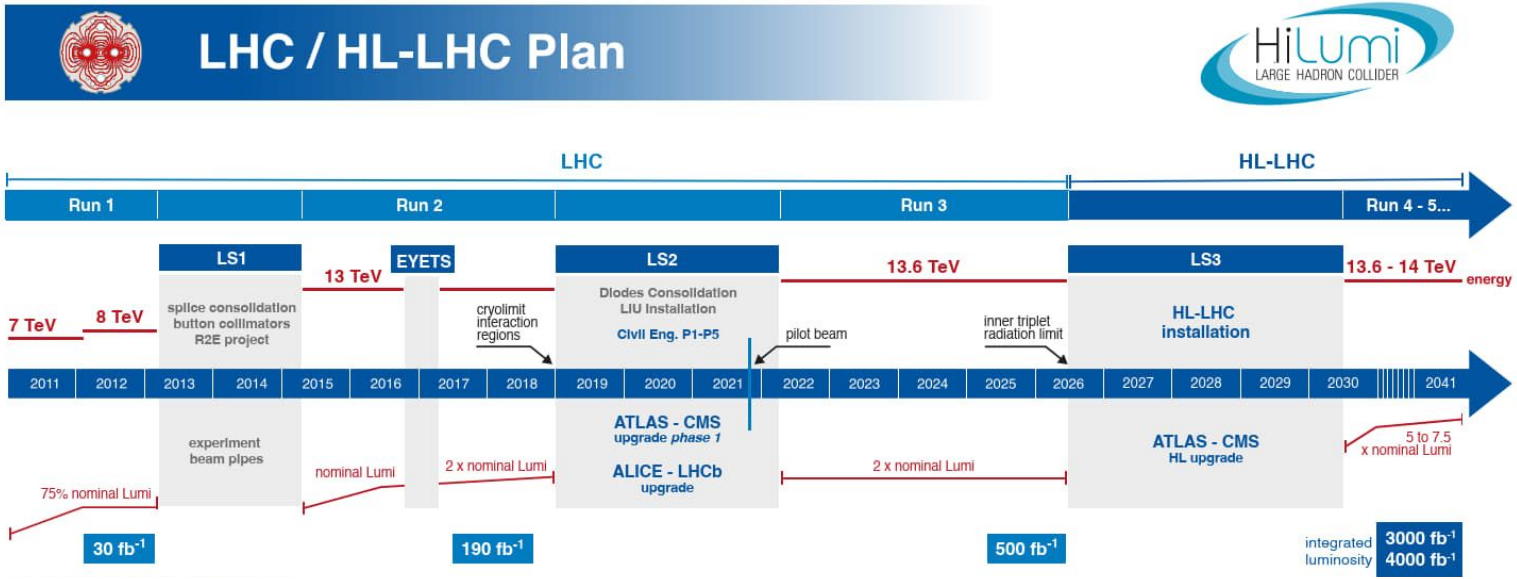
- **Anomaly trigger**
  - General idea: let data guide you!
  - Regular triggers: based on anticipated signatures. But are we creative enough to guess the next discovery?





Future upgrades

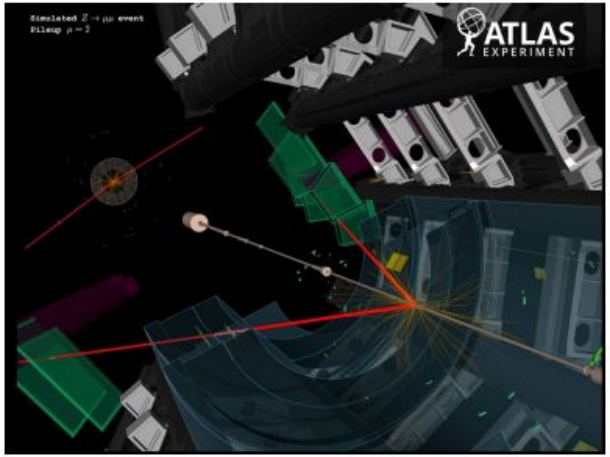
LHC Upgrade



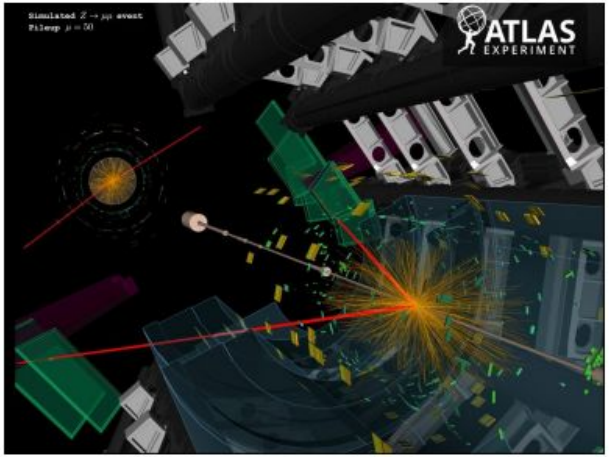
- The LHC accelerator is going through an upgrade that will last three years.
- The HL-LHC period will start in ~2030
  - More than 5 times the initial nominal instantaneous luminosity (5-7.5 10^{34} $cm^{-2}s^{-1}$)
  - This will increase the average pile up from μ~70 to μ~200

ATLAS Upgrade

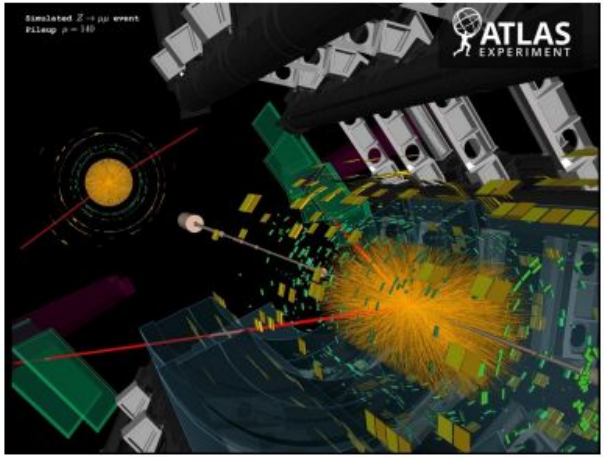
Pileup 2



Pileup 50



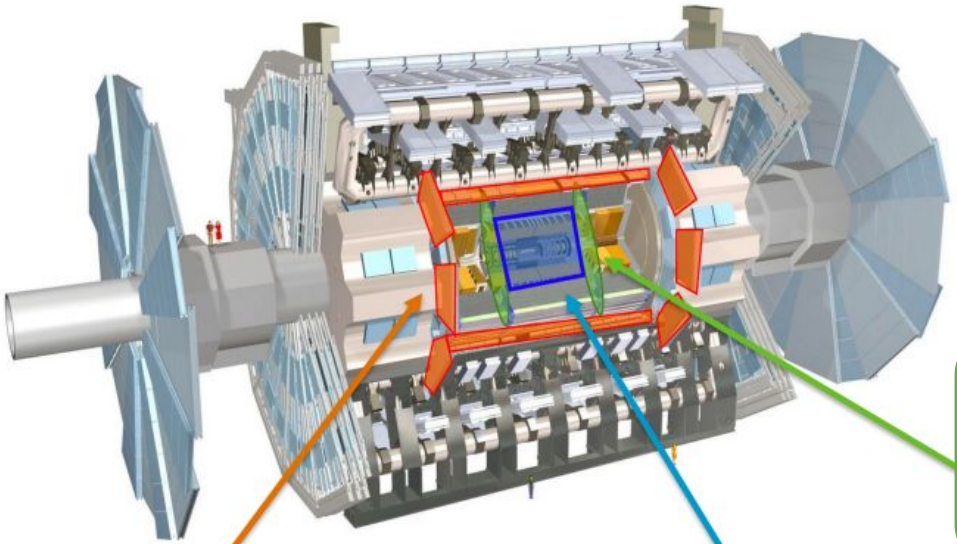
Pileup 140



- Need to upgrade ATLAS to deal with more busy events and higher radiation.
  - Upgrade is underway.

- Trigger latency will increase, but more info to process.
  - From 2.5 μs to 10 μs
  - Allows to run **more complex trigger level algorithms**.

ATLAS Phase-II Upgrade overview



Upgrade of trigger and DAQ system
- L0 trigger at 1 MHz
- Improved High Level Trigger

Electronics upgrades
- Replacement of on/off detector electronics for LAr, Tile, and Muon detectors
- 40 MHz continuous readout with finer segmentation

High Granularity Timing Detector (HGTD)
- Forward region ($2.4 < |\eta| < 4.0$)
- Precision time reconstruction (30 ps) with Low Gain Avalanche Diodes (LGADs)
- Improved pileup separation and bunch-by-bunch luminosity

New Muon Detectors
- Improved trigger efficiency/resolution and reduced fake rate

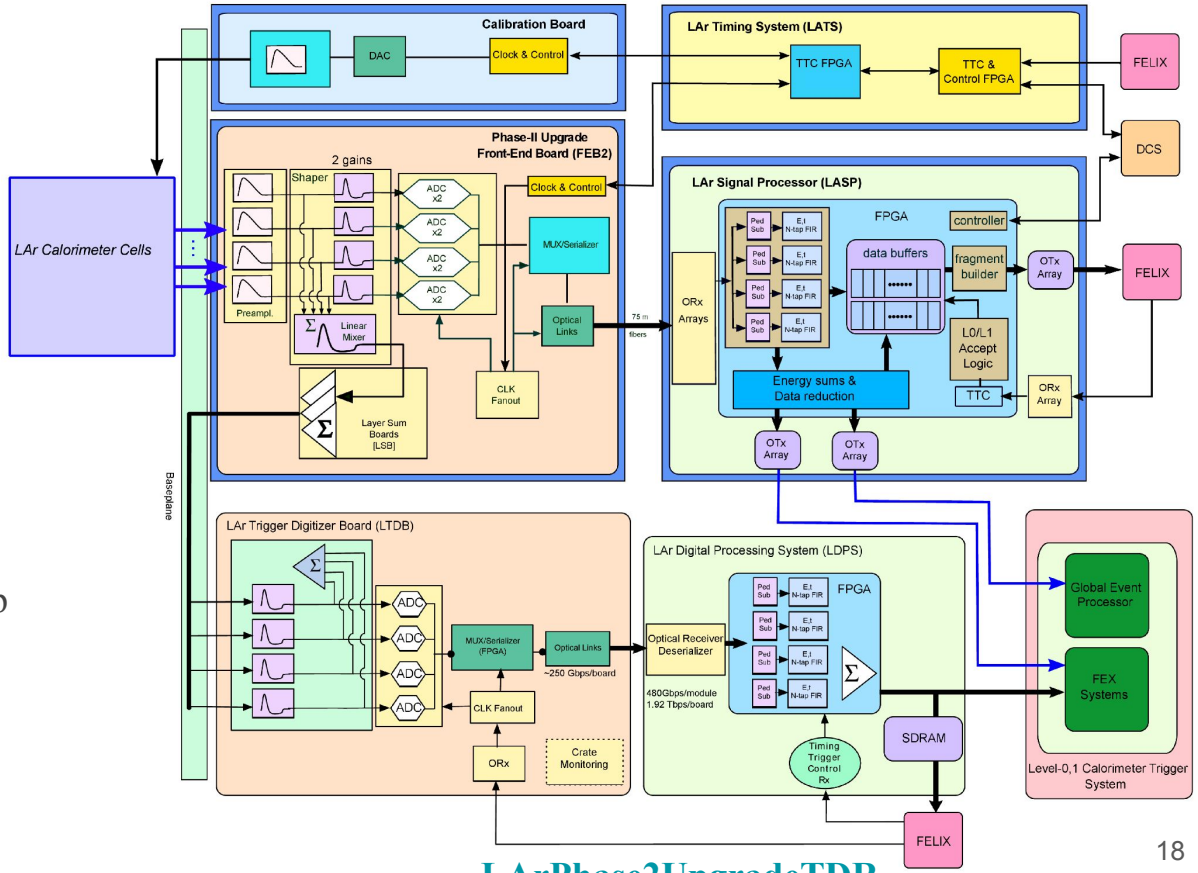
New Inner Tracking Detector (ITk)
- All silicon, coverage up to $|\eta| = 4$
- Less material and finer segmentation

LAr Phase-II Upgrade

- **Readout requires complete replacement** to cope with increased rates and radiation exposure.
  - Thousands of boards with custom ASICs and very powerful FPGAs to be replaced!

- Provide **full calorimeter information** (250 Tbps!) for improved trigger decisions

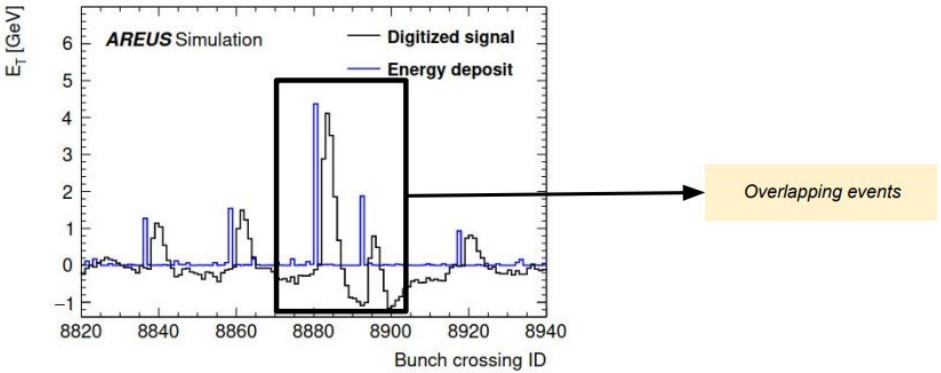
- Many places where ML could help improve physics performance.
  - **How to make them perform within tight latency budget and hardware constraints?**

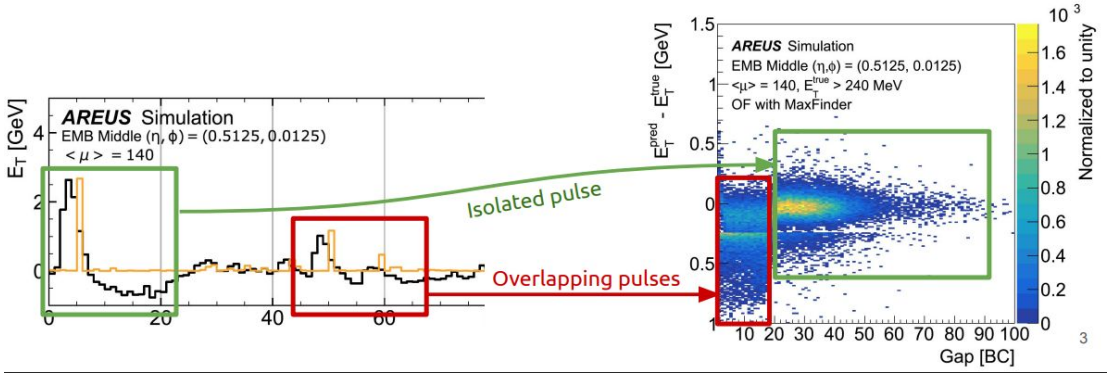


[LArPhase2UpgradeTDR]

Machine learning for LAr energy reconstruction

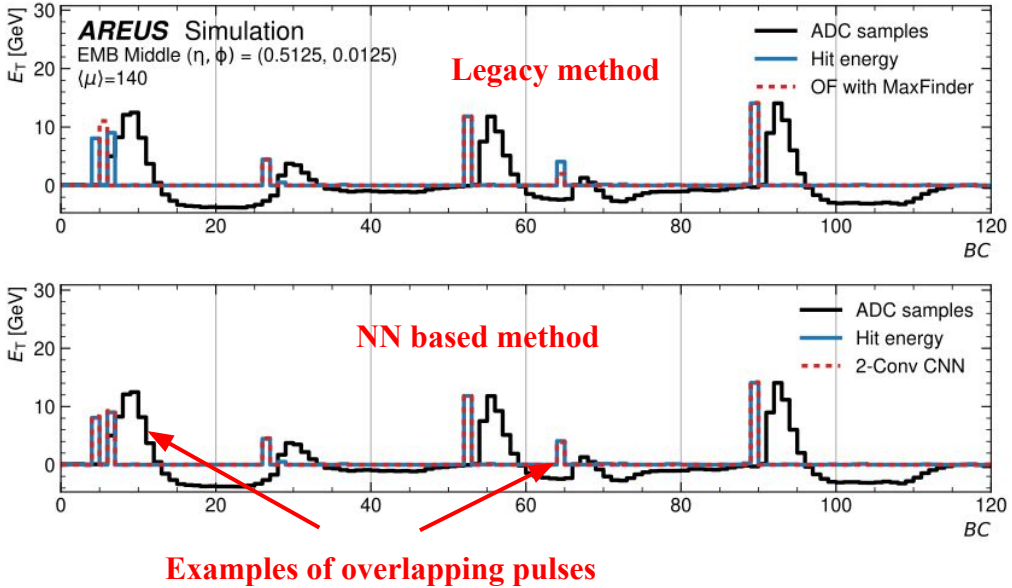
- More simultaneous proton-proton interactions make **pulses overlap.**
- The **traditional Optimal Filtering method can then misreconstruct energy**, especially when deposits occur close together in time.
- Can a neural network recover energy more reliably without exceeding hardware limits?

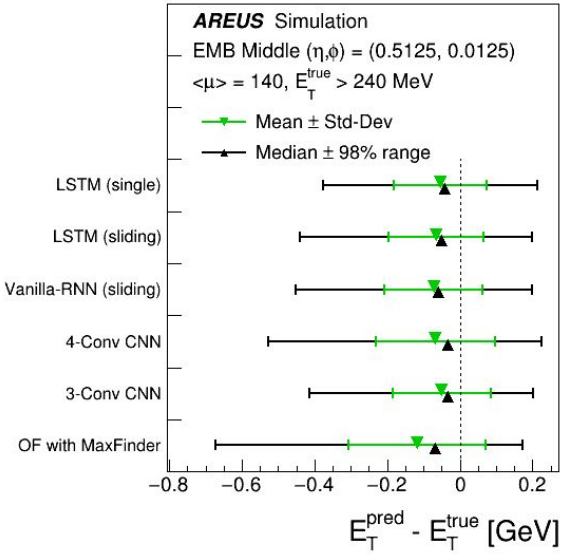




Machine learning for LAr energy reconstruction

- Overall **better energy scale and resolution for all neural networks** with respect to Optimal Filtering.





[LArCaloPublicResultsUpgrade]

Possible candidate: Vanilla RNN

- The produced Vanilla RNN firmware fits the resource limitations
- **HLS4ML** = open source software designed to facilitate the implementation of AI algorithms on FPGAs
- Useful for prototyping, but hand optimised VHDL is needed to meet the constraints.

Stratix10 1SG280HU2F50E2VG	Language	Number of channel	*ALM [%]	** DSPs [%]	Latency [ns]	Max. Frequency [MHz]	Multiplexing
RESULTS	HLS optimized	370	23%	100%	302	414	10
	VHDL optimized	392	18%	66%	121	561	14
SPECIFICATION	-	384	max 30%	max 70%	max 125		

* Adaptive Logic Module (ALM)
** Digital Signal Processing (DSP)

- ALM: adaptive logic modules
- DSP: digital signal processors

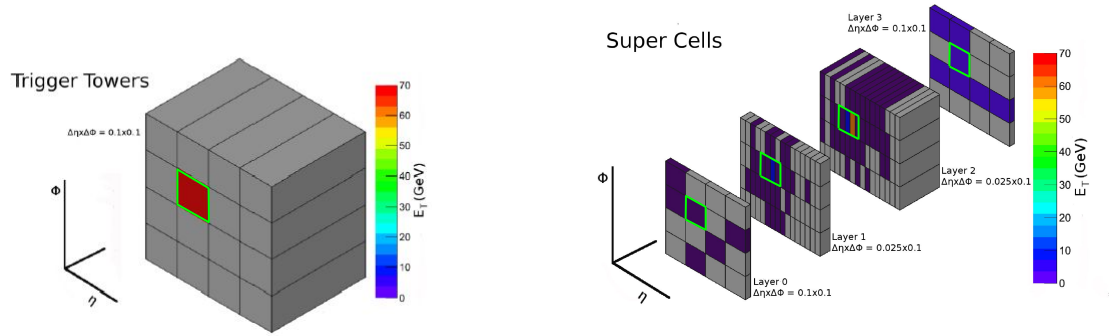
Final thoughts

- **Applications of AI and ML** being used in **several areas of ATLAS and CERN** overall**.**
  - **Part of offline analyses** since decades.
  - Increasingly being used in **real-time system running on FPGAs.**
  - Several task forces being set up to streamline various ongoing efforts
  - Not trivial as ATLAS is a collaboration with **over 5000 members.**

- **AI-driven demand implies sharp cost increase for memory and storage**
  - **Mitigations under investigation.**

Back-up

Trigger Upgrade Overview



Simulation for a 70 GeV electron

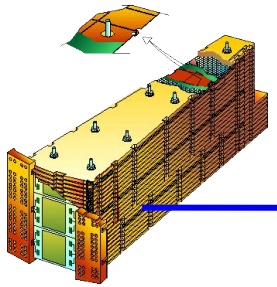
- No longitudinal segmentation
- Fixed size in $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$
- 5.4k trigger towers (TT) from 180k cells

- Longitudinally segmented
- Increased granularity in Front and Middle to $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$
- 34k SCs from 180k cells

Use more information at trigger level to better reject background

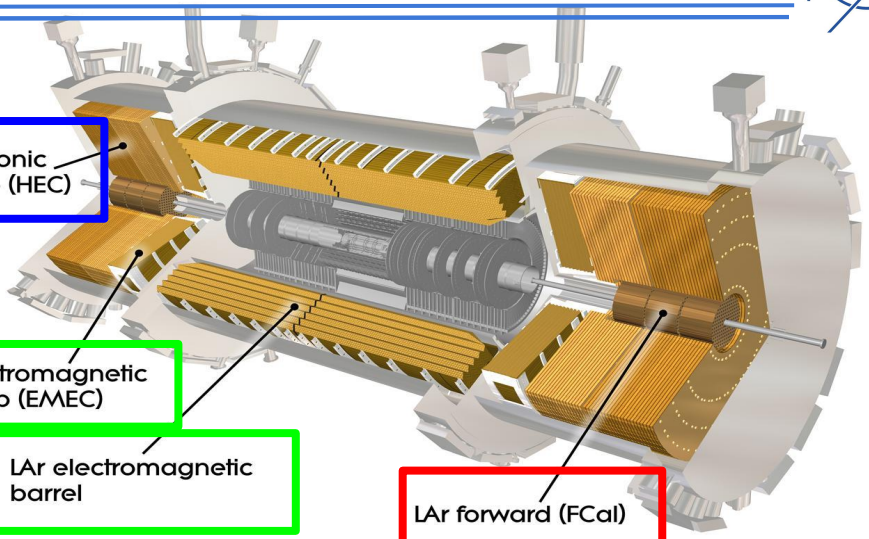
- ~10x more signals transmitted to L1Calo
- Using shower shape information (similar to offline identification) for better pileup rejection

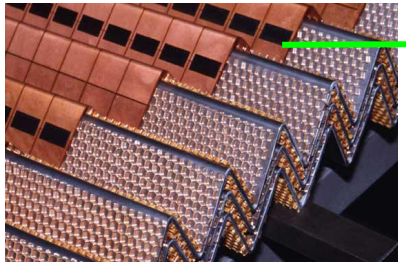
LAr calorimeter design



Cu/Kapton electrode

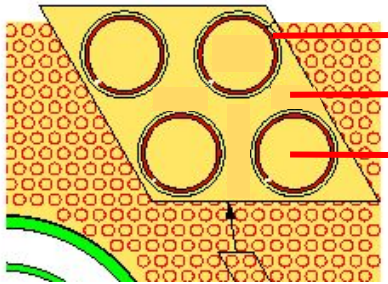
Cu absorber





Cu/Kapton electrode

Pb absorber



LAr gap

Matrix

Rod

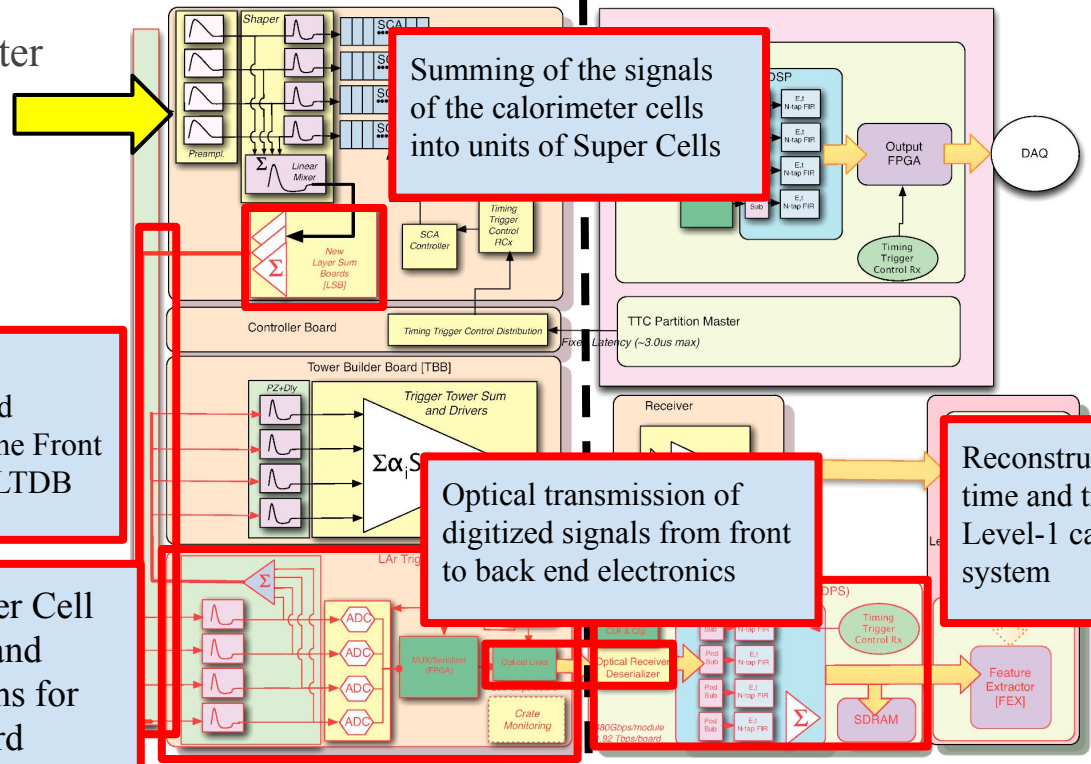
Upgrade trigger and readout electronics

UNIVERSITY OF TORONTO

Front end electronics on detector

Back end electronics off detector

Calorimeter signal



Summing of the signals of the calorimeter cells into units of Super Cells

Baseplane:
Transmission of increased number of signals from the Front End Board (FEB) to the LTDB

Optical transmission of digitized signals from front to back end electronics

Reconstruction of E_T and time and transmission to Level-1 calorimeter trigger system

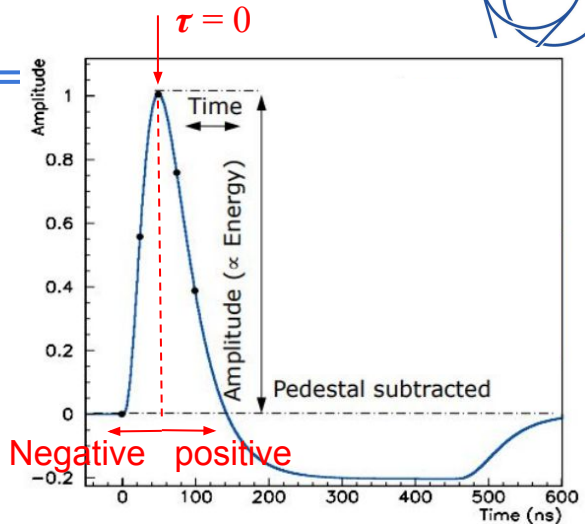
Digitization of Super Cell signals, recreation and feeding back of sums for Tower Builder Board

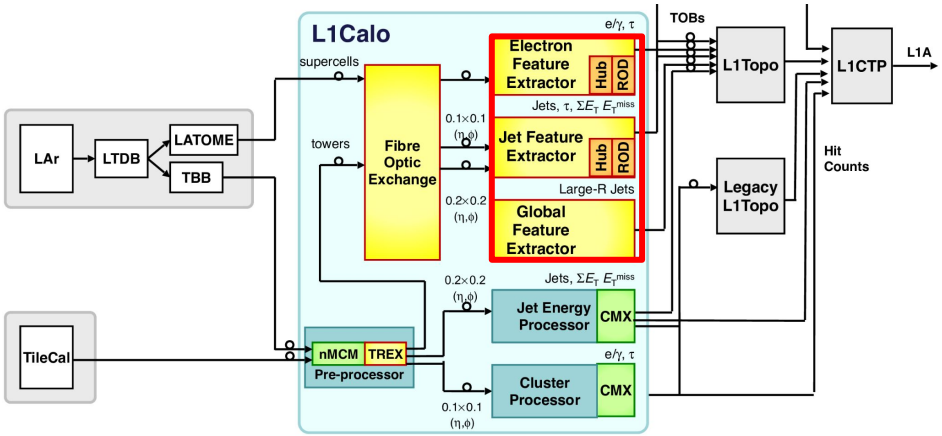
Timing selection

- In LDPB firmware, transverse energy E_T and timing τ calculated for every bunch crossing (@ 40MHz) with a fixed latency
- τ criteria applied to assign correct BCID

$$\begin{cases} -8\ \text{ns} < \tau < 14\ \text{ns}, & \text{if } E_T \geq 10\ \text{GeV} \\ |\tau| < 8\ \text{ns}, & \text{if } E_T < 10\ \text{GeV} \end{cases}$$

- Send E_T for Super Cell to L1Calo only if criteria are satisfied, otherwise $E_T = 0$





LAr calorimeter signals

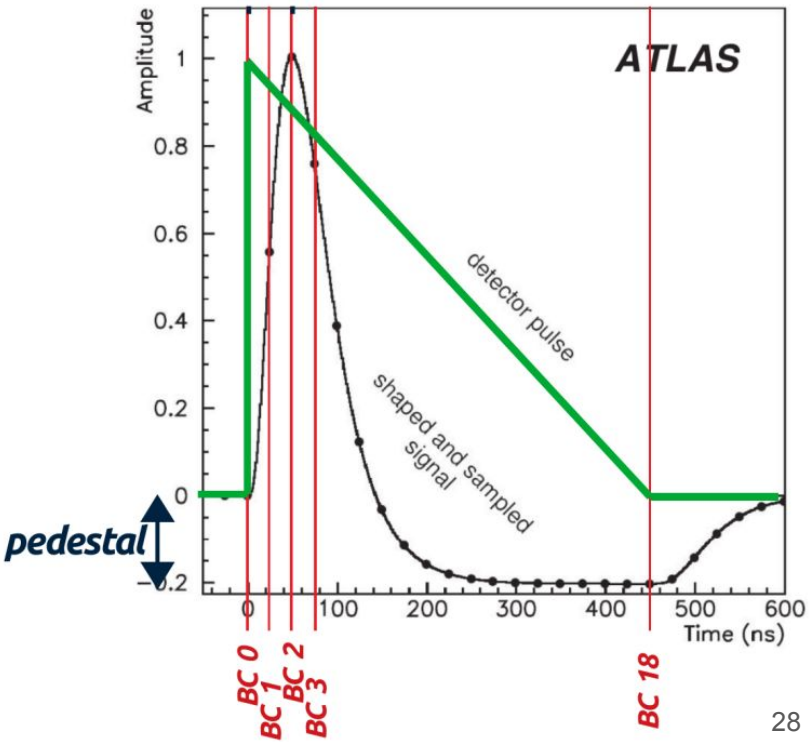
Main readout

- 180k cells are the smallest readout units
  - Cells η-ϕ dimensions and depth depends on layer and position
  - Electric field across cells with the application of High Voltage (from 250 V to 2500 V)
- Particle showers produce a triangular ionization signal
- Signal is amplified, shaped, and buffered waiting for L1 accept in the front-end electronics
- Energy and timing are computed by the back-end electronics with Optimal Filter Coefficients (OFCs)

Trigger

- Signals from close-by cells are summed to form Super Cells (SC)
  - Feed L1 Calorimeter Trigger (L1Calo)

Main readout pulse



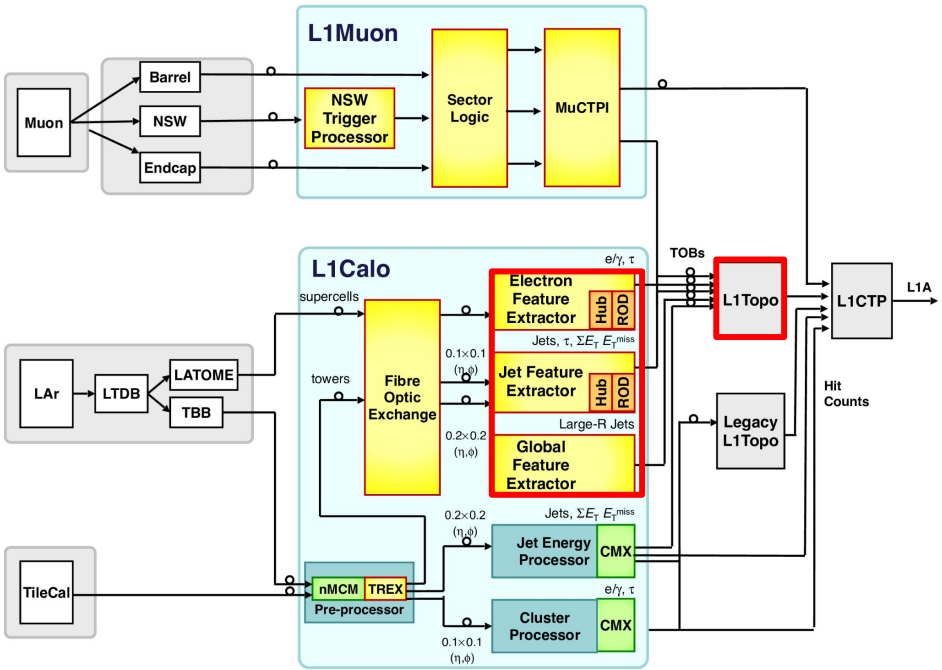
(Super) Cell Energy Computation

$$E_{\text{cell}} = F_{\text{DAC}\to\mu\text{A}} \cdot F_{\mu\text{A}\to\text{MEV}} \cdot \frac{1}{\frac{M_{\text{phys}}}{M_{\text{cali}}}} \cdot G_1 \cdot \sum_{i=1}^{N^{\text{samples}}} a_i \left(s_i - p\right)$$

$F_{\text{DAC}\to\mu\text{A}}$	=	sampling fraction, converts calibration board DAC counts to current
$F_{\mu\text{A}\to\text{MEV}}$	=	factor which converts ionisation current in the calorimeter to total deposited E, from test-beam studies
$\frac{M_{\text{phys}}}{M_{\text{cali}}}$	=	ratio of maxima of physical and calibration pulses with the same input current
G_1	=	cell gain - ADC to DAC from calibration pulse
a_i	=	Optimal Filtering Coefficients (OFCs), derived from predicted pulse shape & noise autocorrelation
s_i	=	samples of the shaped signal digitised in a given electronic gain, measured in ADC counts
p	=	read-out electronic pedestal, measured for each gain

Daily/weekly calibration runs and quantities measured during test beam needed to obtain input for energy computation

Level 1 trigger system

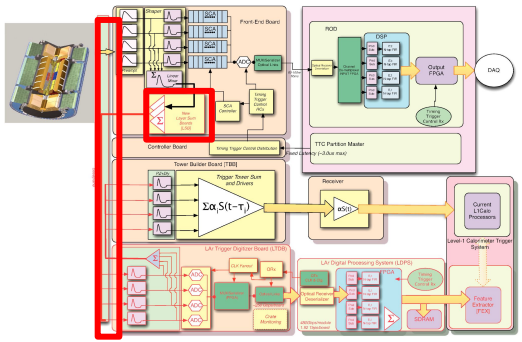


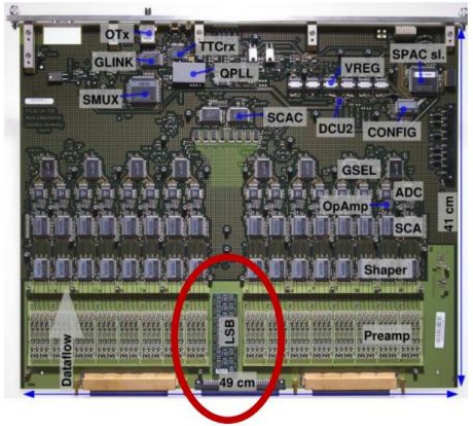
- Muon detectors and calorimeters (LAr and Tile) provide input to first level trigger.
- First level trigger exploits fine granularity information from the LAr trigger signals using Feature EXtractor (FEX) processors.
- Reduced information wrt calorimeter level, possible to implement more sophisticated algorithms.

arXiv:2401.06630

Layer Sum Boards (LSB) and baseplane

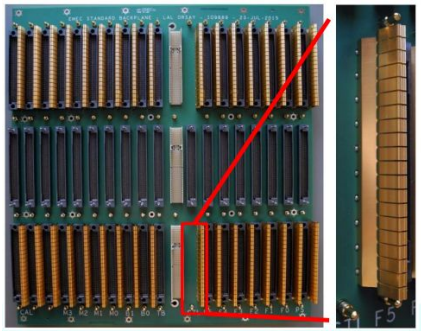
- LSB:
  - 1524 FEBs will be equipped with new LSBs.
  - Production finished, now testing ongoing at Pittsburgh and Arizona.
- Baseplane:
  - Two baseplanes per FEC.
  - Most of baseplanes went into production after being tested in the lab at P1.

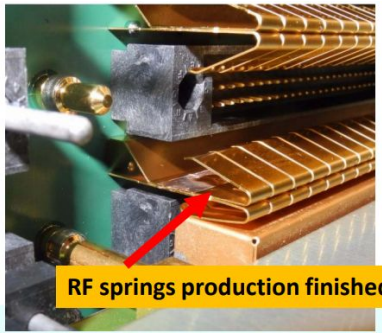




LSB on the currently installed FEBs.

Partition	Pcs needed
EMB	64
EMEC std	32
EMEC sp.	8
HEC	8
FCal	2

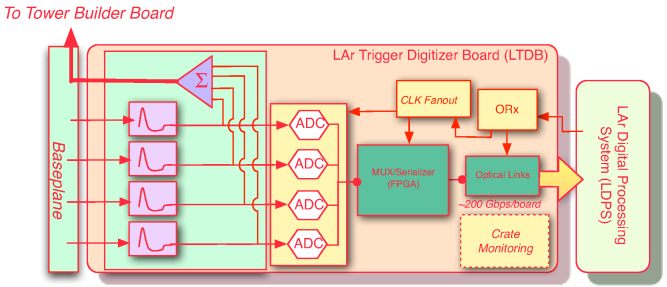


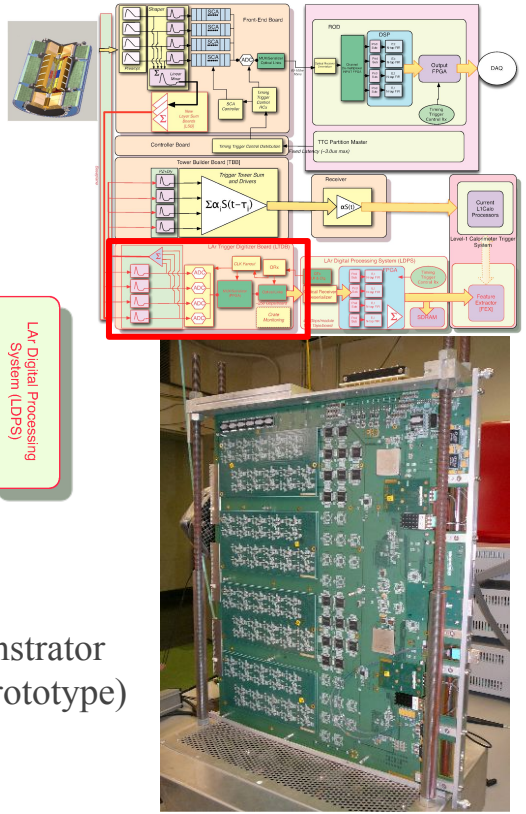


New baseplanes

LAr Trigger Digitizer Board (LTDB)

- Each LTDB will process up to 320 Super Cell signals delivered by the new baseplane from the FEBs.
- The Super Cell signals will be digitized at 40 MHz by high performance ADCs.
- Digitized signals are sent to the back end electronics via optical links (5.12 Gbps).
- 200 Gbps output / board.

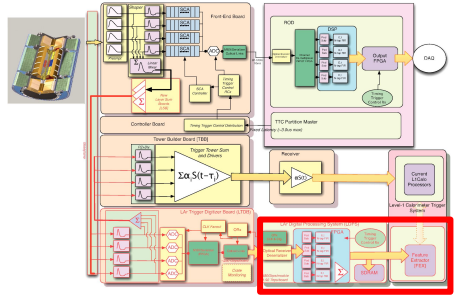


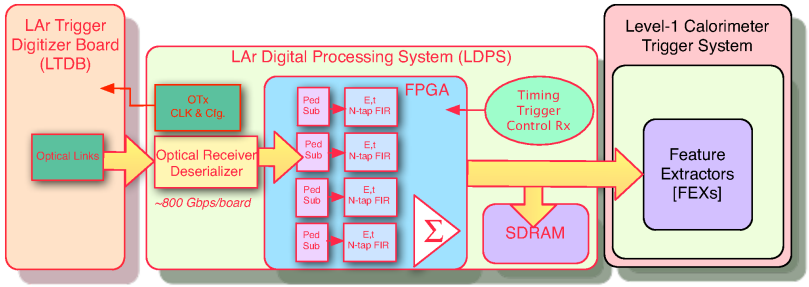


LTDB demonstrator board (pre-prototype)

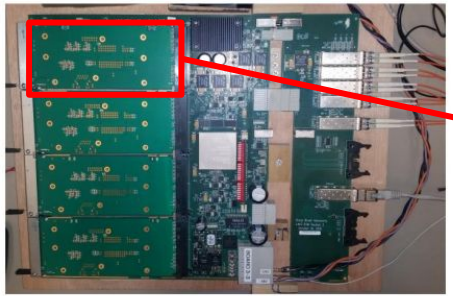
LAr digital processing system (LDPS)

- Each digital processing board (LAr Carrier board) has 4 *LAr Trigger prOcessing MEzzanines* (LATOME).
- Each LATOME receives the data equivalent to one LTDB (320 Super Cells) at 25 Tbps.
- LATOME will reconstruct E_T and time for each Super Cell and transmit the results to the Level-1 calorimeter trigger at 41 Tbps.
- It will additionally calculate global variables (Energy sums, barycenter,..) for the Jet and Global Feature Extractors.

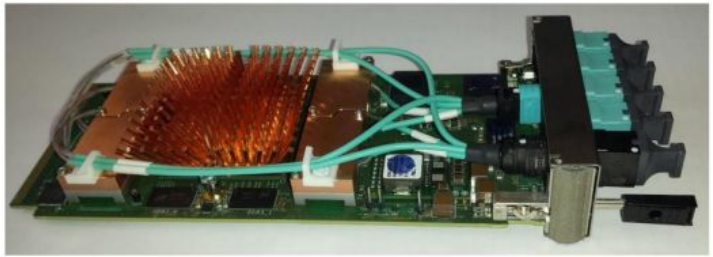




LAr Carrier board



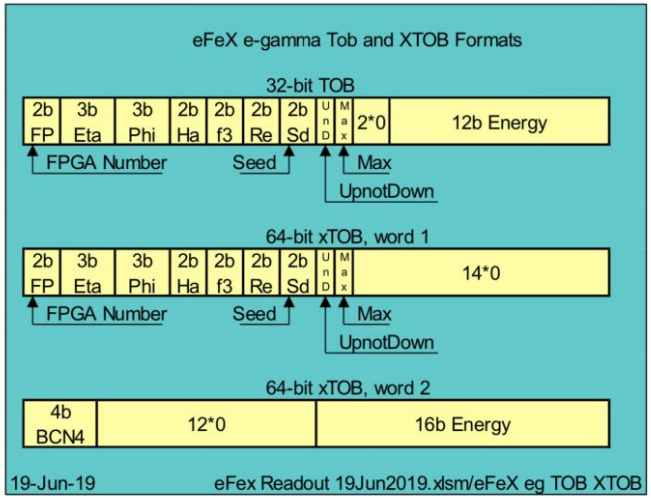
LATOME



L1Calo event format

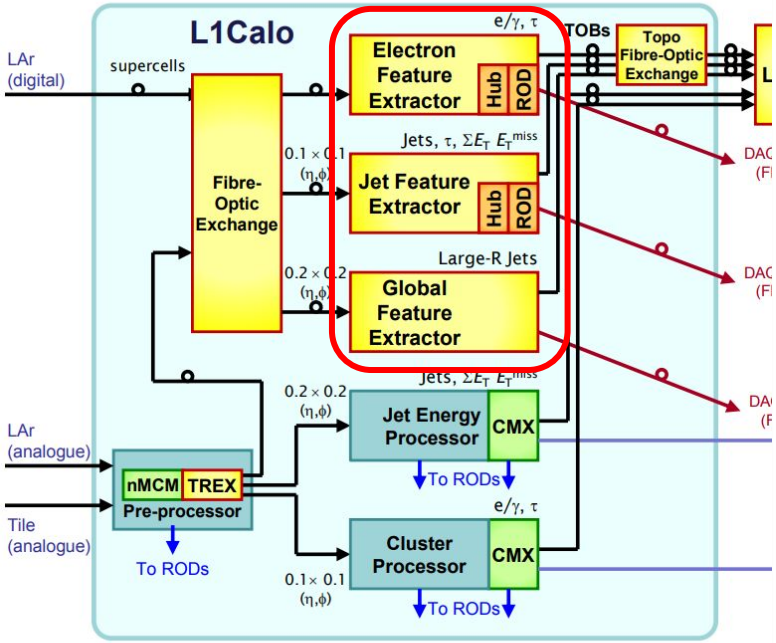
- **L1Calo RAW data**
  - 1 TOB + 1 xTOB = 12 bytes
  - Number of TOBs depends on E_T requirement. Sensible thresholds a few 10s of TOBs
    → < 1kB + overheads /event

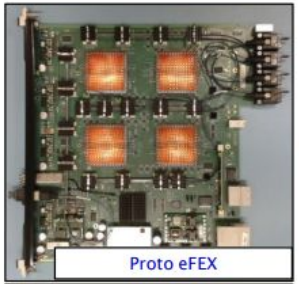
- In comparison: LATOME readout
  ~ 460 kB/event
  - Requested by L1Calo to have E_T and ADC values.
  - Needed for **monitoring, diagnostics, and performance optimization**!

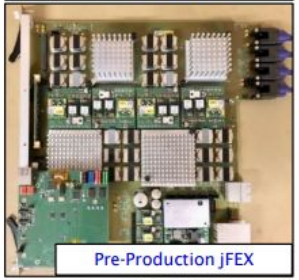


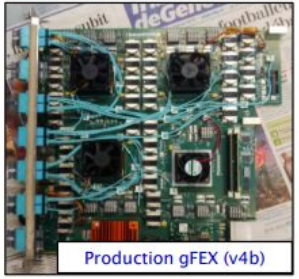
More details here.

The FEXes




Proto eFEX

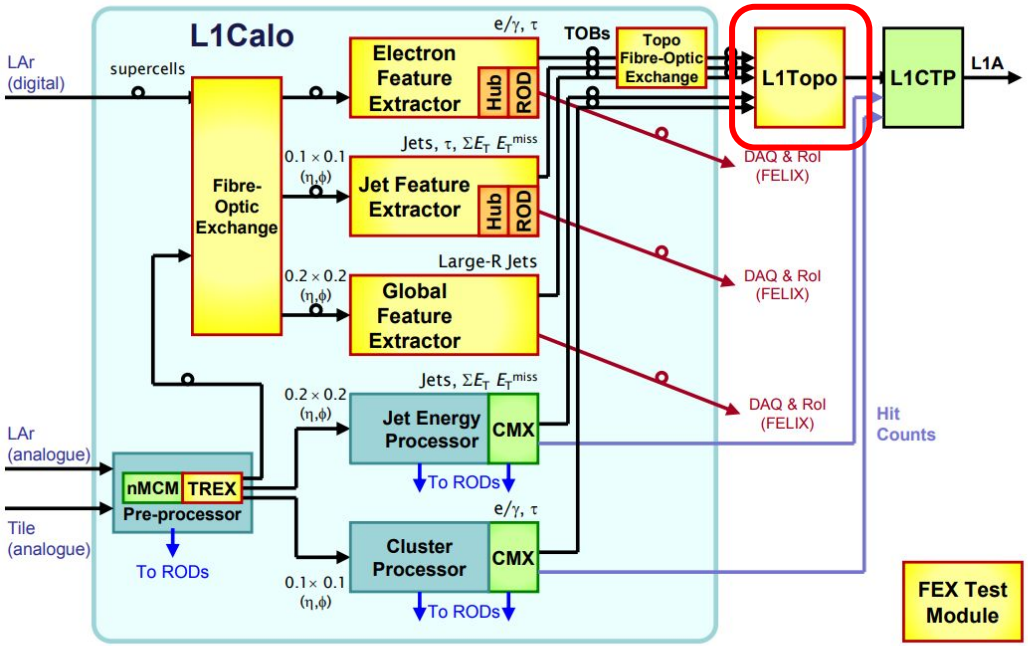

Pre-Production jFEX


Production gFEX (v4b)

	eFEX	jFEX	gFEX
# ATCA modules	24	6	1
# FPGAs per module	4	4	3

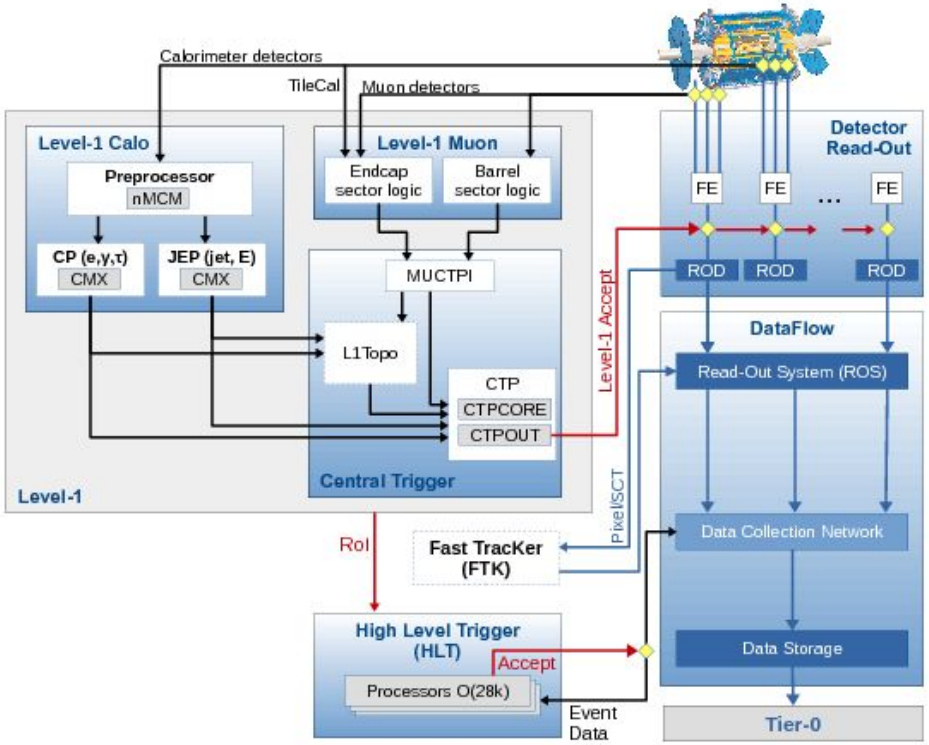
- **eFEX**: Identifies isolated e/γ and τ candidates.
- **jFEX**: Identifies jets, large τs, Σ E_T, E_T^{miss}
- **gFEX**: Identifies large-R jets, computes E_T^{miss}, centrality...

The L1Topo modules



- **Three Topo modules** with 2 FPGAs each
- Different FPGAs with different functionalities
  - **Topo1**:
    Counts multiplicities of objects passing certain thresholds.
  - **Topo2 and Topo3**:
    Trigger decisions based on **event topologies** (angular cuts, mass selections, invariant and transverse masses, combination of objects from different bunch crossings (LLP)).
- **Configuration** of each FPGA specified by **trigger menu**.
  - VHDL config code derived from trigger menu XML files and parameters loaded via IPBus.
  - Few configurations studied to **check timing**.

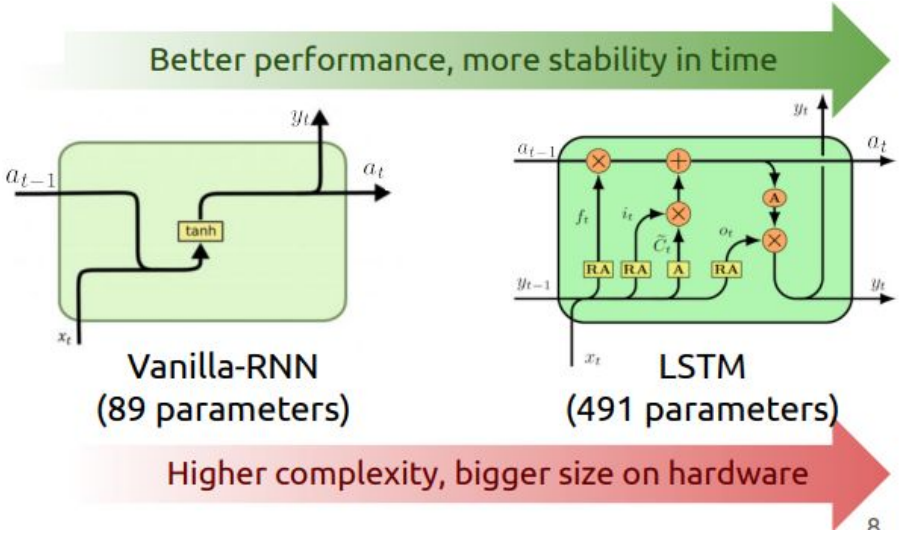
ATLAS trigger principle



- Level-1 Calo and Level-1 Muon get coarse granularity information from detectors.
- By combining that information on CTP level, a L1A is generated and RoI info is sent to HLT.
- FE board readout buffers are read out upon L1A.
- Upon processing info from RoI, HLT issues “accept” on selection of L1Accepted events.

Challenge: Getting model onto an FPGA

- A trained network is only useful if its implementation meets throughput, latency, and resource constraints
- All NNs were optimized to **reduce as much as possible the required computing resources while keeping good performance.**



Algorithm	LSTM (single)	LSTM (sliding)	Vanilla (sliding)	CNN (3-conv)	CNN (4-conv)	Optimal Filtering
Number of parameters	491	491	89	94	88	5
MAC units	480	2360	368	87	78	5