Science & Technology Facilities Council

ATLAS Level-1 Trigger for HL-LHC, Requirements, Architecture and Issues

- Requirements
- MDT Readout logic
- Rates with/without Track Trigger
- Architecture
  - L0Calo
  - L0 Muon: RPC, TGC, MDT
  - L0CT
  - L1Calo
  - L1Track
  - L1CT
- Latency
- Issues
- Common Features
- Summary

Much of this material is taken from Norman Gee, Stefan Haas, Yasuyuki Horii, Francesca Pastore, Robert Richter, Osamu Sasaki, Andre Schoening, Mark Sutton, Mark Thomson and Riccardo Vari, with thanks.

Physics Requirements

- Design driven by physics goals
- Strong desire to trigger on leptons at electroweak scale
- Aim:
  - maintain current thresholds for single isolated leptons
  - Maintain trigger efficiency for 20 GeV electrons & muons
  - Manage rates to include jet, missing ET, …
- Need flexibility

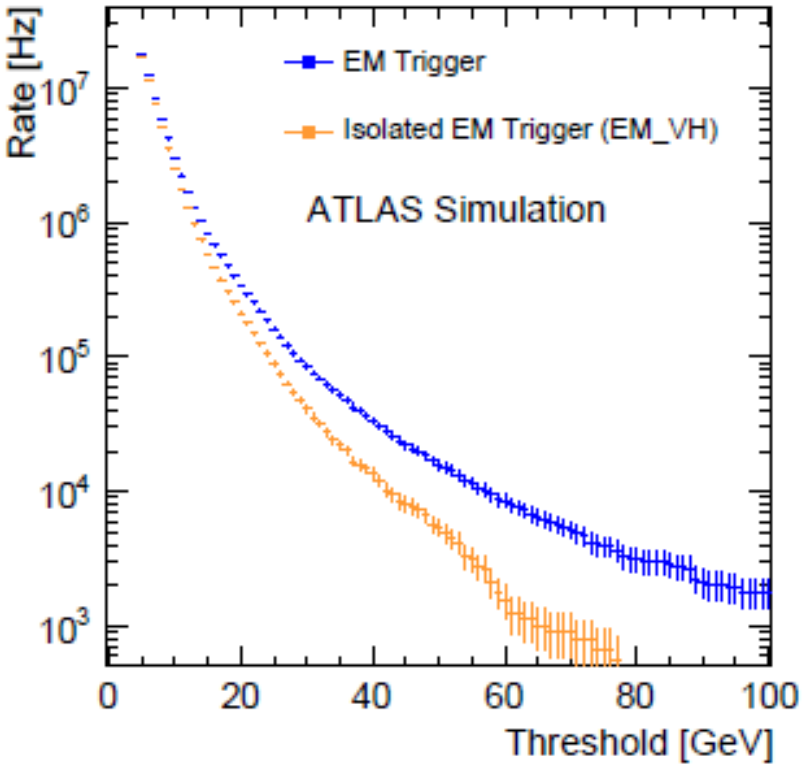

EM20 with Phase-I h/w at Phase II
(L ~ 4 x $10^{34}cm^{-2}s^{-1}$, 95% efficiency)
Rate → ~ 200 kHz

Hardware Requirements (1)

- These are determined by readout electronics at front end
- Constraints published in ATLAS Phase-II Upgrade LoI:

	Max Rate	Max Latency
MDT	~200 kHz	~20 μs
LAr	any	Any
Tile	>300 kHz	Any
ITK	>200 kHz	< 500 μs

- MDT system is limiting case in rate and latency

Hardware Requirements (2)

- Most can be replaced
  - Readout capability determined primarily by cost
- Exception is Muon MDTs
  - 30 % of electronics very hard to change or inaccessible

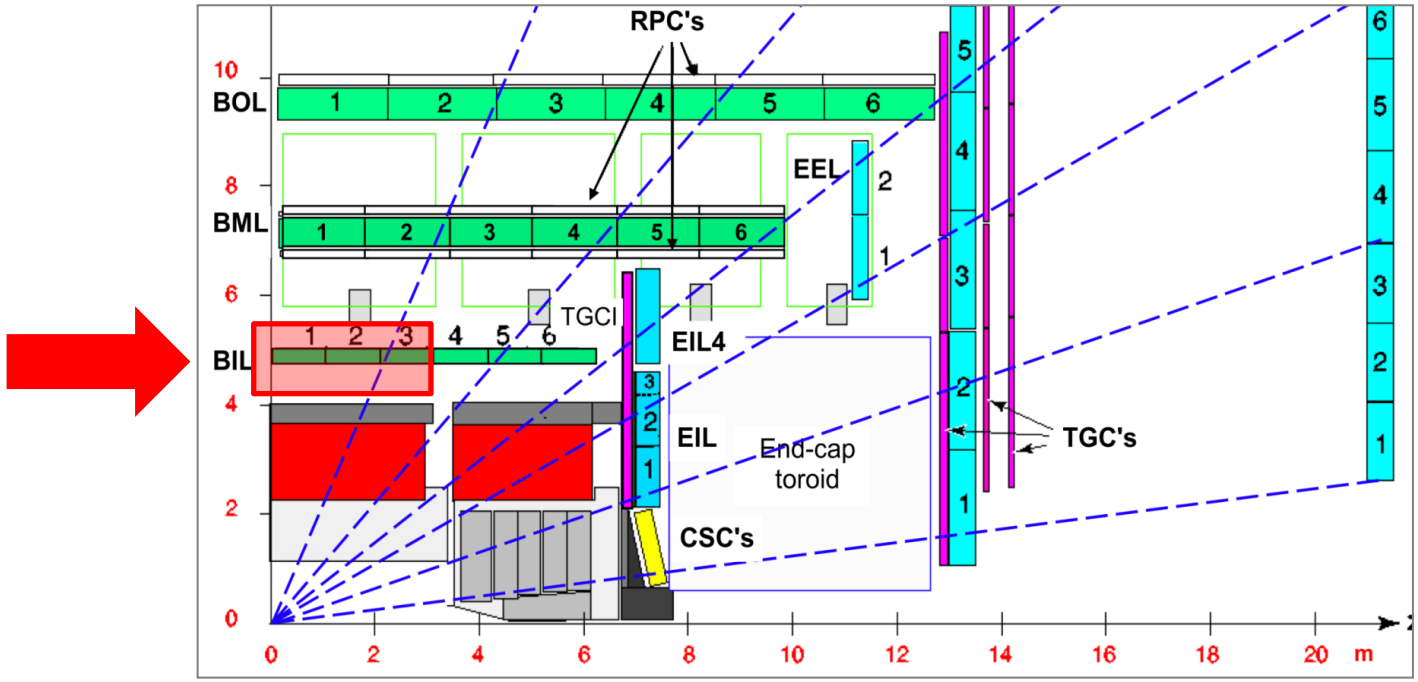

MDT Readout

- Front end of present MDT ("mezzanine card"):

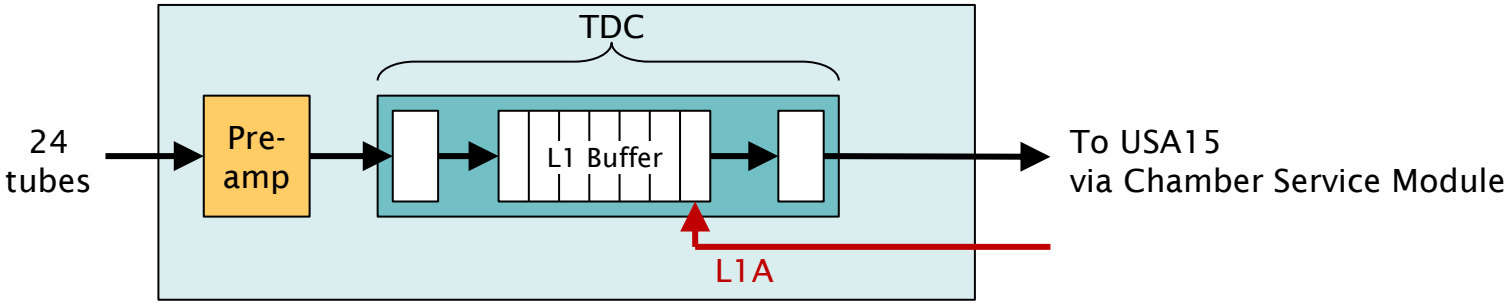

- Hit data comprise leading & trailing edges of pulse
  - Digitised, time stamped…
  - Data-driven pipeline, held for latency of L1A
- Length of pipeline →latency constraint
- Max. output rate across link →trigger rate constraint
  - Exceed these, data lost

MDT Data Output

- 2 modes of reading out MDT:
- Edge Mode
  - Data

Status	Tube ID	Not Used	Absolute Time, Leading Edge
Status	Tube ID	Not Used	Absolute Time, Trailing Edge

  - Used in LoI calculations
- Pair Mode
  - Data

Status	Tube ID	Pulse Width	Absolute Time, Leading Edge

  - Not currently used because TDC fails to recognise trailing edge of short (noise) pulses – locks channel
    - ~ once in tens of minutes
  - Proposal to overcome by issuing periodic reset every few seconds
- With no loss of information, pair mode halves occupancy of buffer + readout bandwidth

MDT Constraints

Readout latency

- Average occupancy of L1 buffer

Latency / μs	3 μs	20 μs	30 μs
Edge mode	6%	38%	56%
Pair mode	3%	19%	28%

- Average Occupancy of readout link
  - (80 Mb/s)

L1A Rate	100 kHz	200 kHz
Edge mode	35%	70%
Pair mode	24%	49%

- 60–70% is occupancy limit in buffer and readout link
  - Beyond this data get lost at instantaneous peaks

Trigger latency

- Trigger requirements summary: **20 GeV isolated leptons, 200 KHz, ~20–25 μs**
  - (cf 20 us LoI)

Rate Estimates

- Estimate of HL-LHC rates based on phase 1 system
  - Trigger rate at least 500 kHz

Object(s)	Trigger	Estimated Rate No L1Track
e	EM20	200 kHz
γ	EM40	20 kHz
μ	MU20	> 40 kHz
τ	TAU50	50 kHz
ee	2EM10	40 kHz
$\gamma\gamma$	2EM10	as above
$e\mu$	EM10_MU6	30 kHz
$\mu\mu$	2MU10	4 kHz
$\tau\tau$	2TAU15I	40 kHz
Other	JET + MET	~100 kHz
Total		**~500 kHz**

Rate Estimates

- Estimate of HL-LHC rates based on phase 1 system
  - Trigger rate at least 500 kHz without track trigger

Object(s)	Trigger	Estimated Rate No L1Track	With L1Track
e	EM20	200 kHz	40 kHz
γ	EM40	20 kHz	10 kHz
μ	MU20	> 40 kHz	10 kHz
τ	TAU50	50 kHz	20 kHz
ee	2EM10	40 kHz	< 1 kHz
$\gamma\gamma$	2EM10	as above	~5 kHz
$e\mu$	EM10_MU6	30 kHz	< 1 kHz
$\mu\mu$	2MU10	4 kHz	< 1 kHz
$\tau\tau$	2TAU15I	40 kHz	2 kHz
Other	JET + MET	~100 kHz	~100 kHz
Total		**~500 kHz**	**~200 kHz**

Meets MDT requirement

Track Trigger Options

- Self-seeded L1Track
  - Front end initiates own readout
  - Pros:
    - Reduces impact on Level 1 architecture
  - Cons:
    - Challenging:
      - Data volume
      - connection between stave layers

- RoI-based
  - Seeded by RoIs from Muon + Calo triggers
  - Pros:
    - Reduces impact on tracker
  - Cons:
    - More complicated trigger: L0/L1 split
    - Latency
  - ATLAS Baseline

Phase-II Trigger Architecture

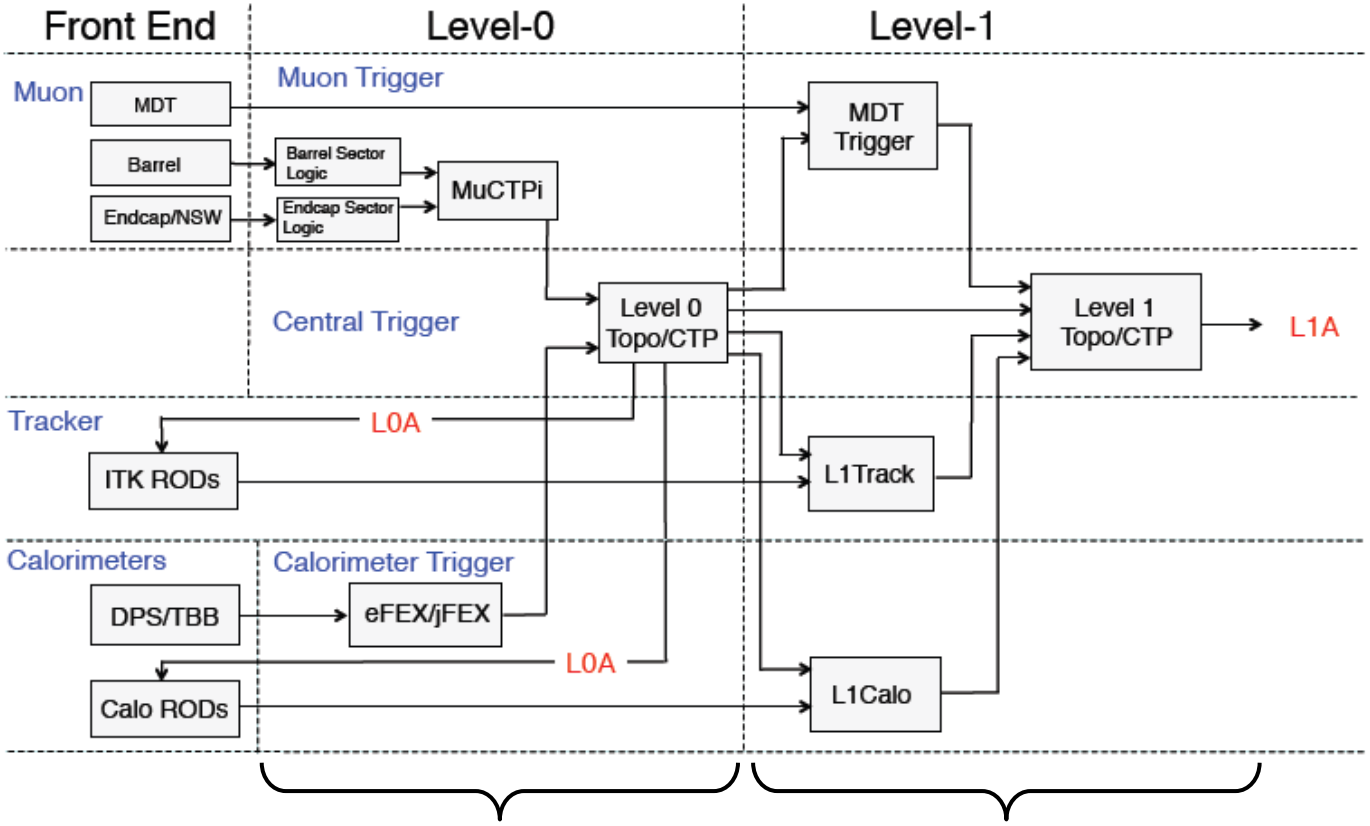

~6μs latency (← MDT trigger) ~20–25 μs (← MDT readout)

500 KHz rate (← Phase-I L1) 200 kHz rate (← MDT readout)

Buffering

- 2 Options for Detector electronics:
  - Single L1A:

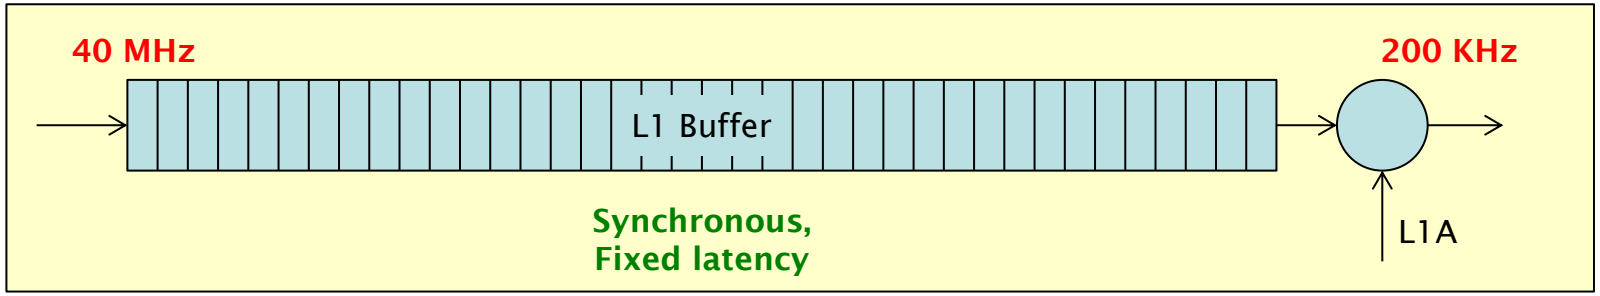

  - L0A + L1A:

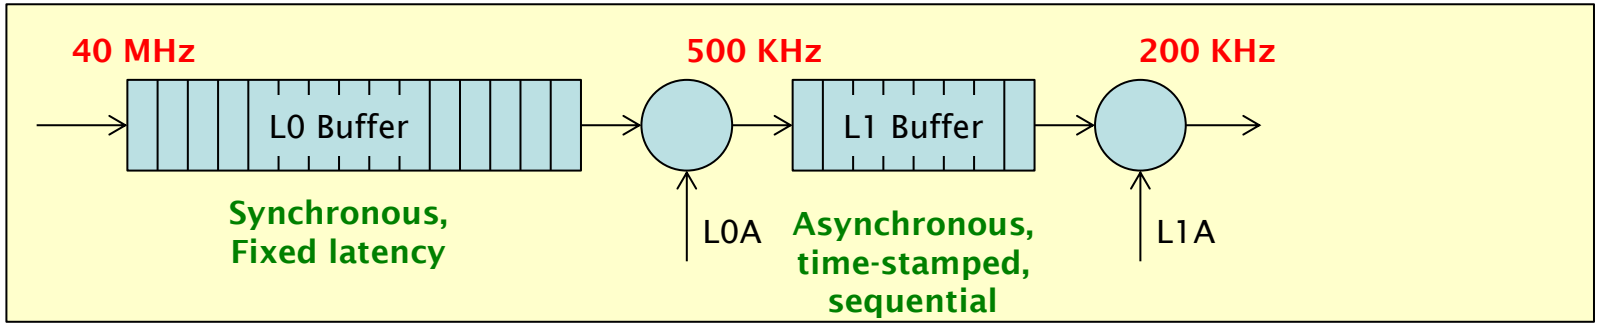

    - More complex
    - Data sparsified: total buffer ~1/4 single L1 scheme

L0Calo

- Uses hardware installed in Phase-1 (eFEXs, jFEXs, gFEX(?), RODs)
  - Digital input from all calorimeters
  - Current h/w retired
- Minor changes
  - Timing & Control Interface:
    - L0A, L1A
    - Daughter board
  - Firmware
    - FEX algorithms?
    - ROD
      - RoI data to L1Calo (+ L1Muon?)
      - Regional Read-out Request (R3) to L1Track
- R3 path
  - Critical latency path
  - Near real-time
    - asynchronous, low latency within defined envelope
  - Fast, dedicated point—point links

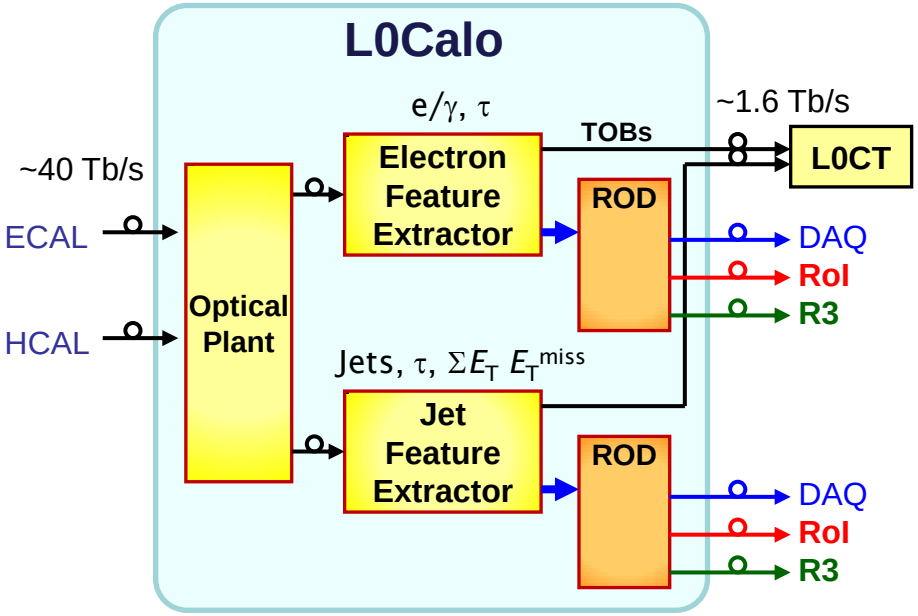

Muon Triggers

- Currently, many fake triggers in muon system
  - Muon tracks not originating at point of interaction
  - Muons of P_T < threshold
    - ← Slope of turn-on curve
    - ← P_T resolution
    - ← Spatial resolution
- Muon systems
  - TGC (Endcap)
  - RPC (Barrel)
  - MDT (Endcap + Barrel)
    - Not used in L1 before Phase-II Upgrade

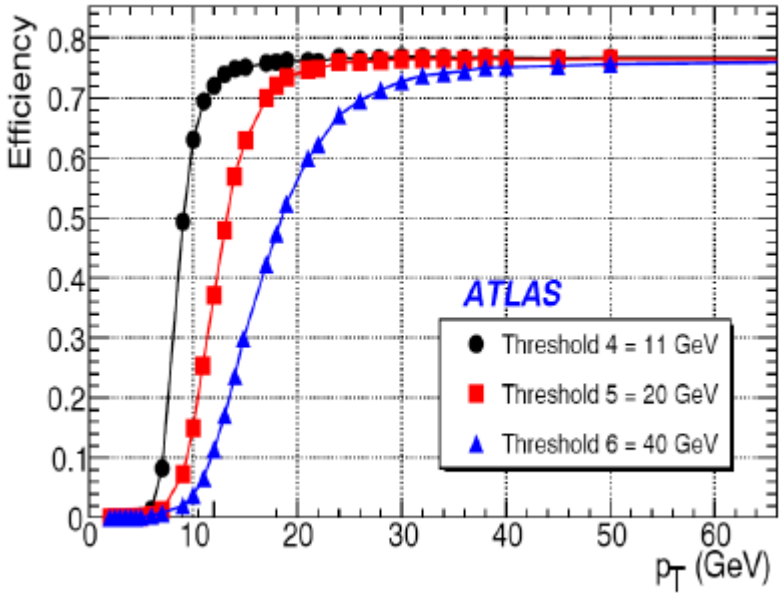

Muon acceptance at L1

L0Muon Barrel (RPC)

- New inner layer of RPCs?
  - Improves redundancy
  - coverage 73 → 86%
- Readout electronics upgraded
- Trigger logic
  - On detector:
    - FE board: TOT:
      - Pulse width
        → charge distribution
        → position
    - DCT (Data Collector & Transmitter)
      - Digitise, sync to LHC clock, zero suppress
    - 832 GBT links @ 10 Gb/s to USA15
      - (+ 416 for new inner layer?)
      - Trigger + ROD share link?
- Majority of trigger logic moved off detector – FPGAs
  - Track finding, P_T calculation...

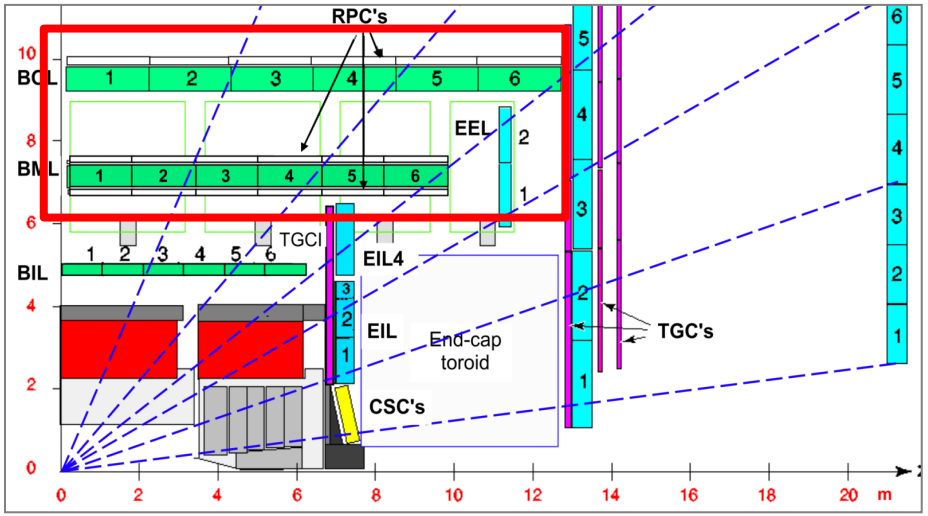

L0Muon Endcap (TGC)

- Phase I:
  - nSW TGC chambers:
    - vectors eliminate non-PI tracks
- Phase II:
  - Replace inner-layer Big Wheel TGCs with high-resolution chambers developed for nSW?
    - Not feasible for Outer Big Wheel
      - Too many chambers
      - Rely on MDT trigger to increase resolution
- Readout electronics upgraded
- Trigger logic moved off detector
  - On detector: digitise, zero suppress
  - 5000 GBT links @ 6.4 Gb/s to Sector Logic
    - Carry Trigger + RO data

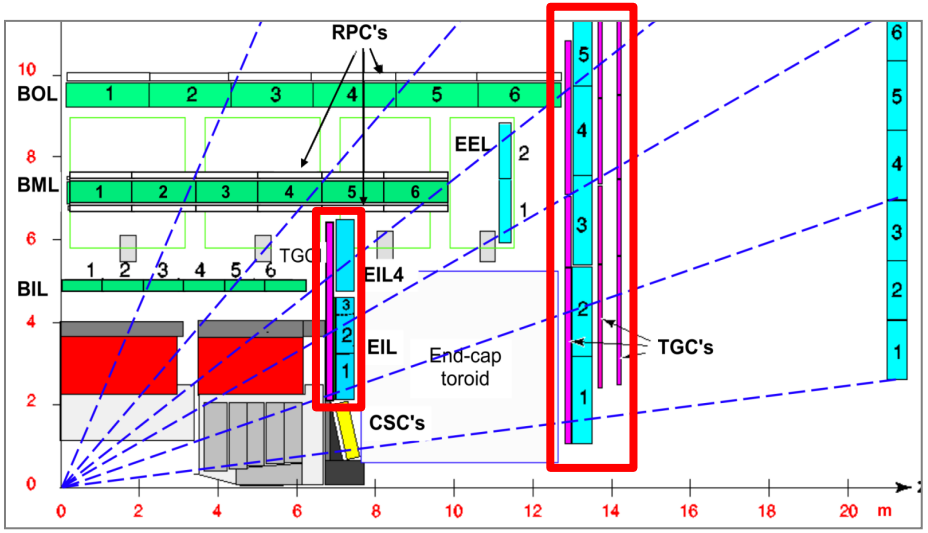

Muon MDT Overview

- Upgrade of readout electronics: add parallel fast readout path

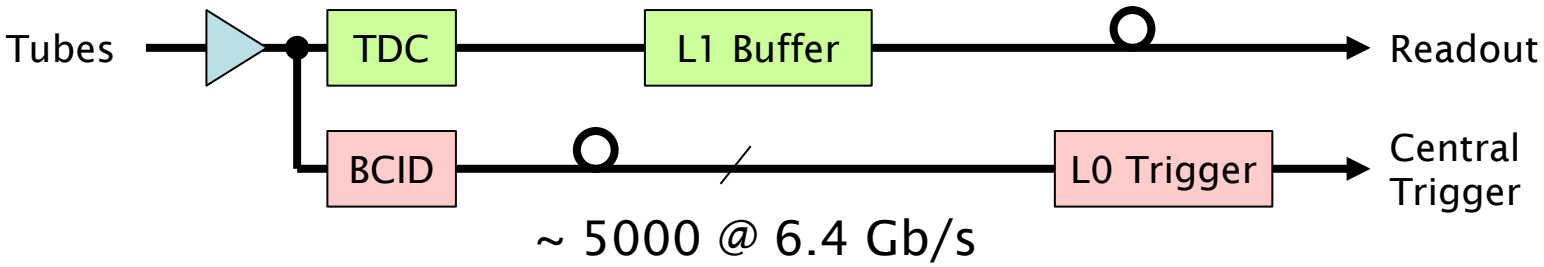

- Front End
  - Synchronous, fixed, low latency
  - TDC clock 40 (80) MHz
    - cf 1.28 GHz in slower path
    - 0.5 (0.25) mm precision
    - × 10 better than TGC & RPC
  - Digitise, BCID
- USA15
  - Decode, fit tracks, calc. P_T…
  - FPGAs
  - Quality ~ current L2 MuFast algorithm

Muon MDT Options

- RoI based
  - RoI from RPC + TGC
  - Minimises bandwidth off detector & data volume to process
  - Latency ~ 4 μs
- Not RoI based
  - Latency ~ 2.4 μs
  - Increased volume of data to transport & process
  - Viable in endcap
    - Toroid field
      - Little deflection for high-P_T muons
        → Narrow search cones
- L0 or L1?
  - LoI (& hence most diagrams) shows as part of L1, but
  - L0 latency of 6 μs allows MDT to be part of L0 trigger

L0 Central Trigger

- Requirements not yet well defined
- Functionality will include
  - Concentrating muon data
  - Topological processing
    - Bottleneck?
  - Flexible, fast combination of objects → triggers
  - Pre-scaling
  - Deadtime
    - Simple – 1 BC? (~1%, Peak rate 20 MHz)
    - Complex (leaky bucket)
      - Detector requirements?
  - Generation of L0A
- FPGA-based implementation

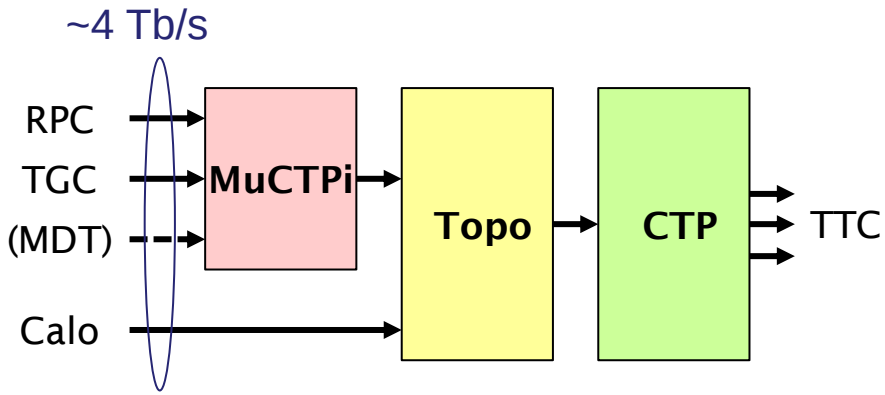

L1Track Implementation (1)

- Matching tracks with Calo + Muon TOBs
  - 3–10 reduction in rate of TOBs
- Can't read all data – no. links, detector mass, power
- Self Seeded
  - Filter data on cluster size (momentum)
    - Algorithm implemented on ABCn130 prototype
  - Identify coincident hit pairs (inner/outer stave layers)
  - Output ~ 4 Gb/s stave → USA 15
  - Reconstruct all high-P_T tracks using subset of strip layers
- RoI Driven
  - Baseline
  - RoI from L0Muon + L0Calo
    → Regional Readout Request [R3]
  - Buffer all data on detector (ABCn130)
  - R3 data ~10% data in L1 pipeline
    - 1 GBT per stave/petal

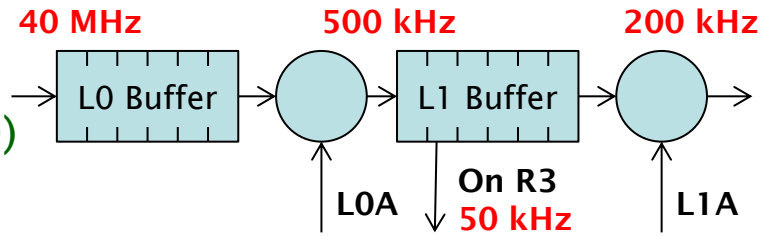

L1Track Implementation (2), RoI-Driven, Cont.

- ID Readout (to trigger)
  - Limited by strip readout
    - Daisy chain
    - Queuing problem
  - Barrel: 95% data within 6 μs
  - Endcap: some areas problematic
    - Spare GBT bandwidth & redundant FE links
      ⇒ 95% data within 5 μs
  - Links shared with L1A data – must prioritise
- Track Finder (USA15)
  - Latency constraint ~ 6 μs
    → highly parallel pattern matching
  - AM ASICs as in FTK?
    - Many more patterns required
      → development needed
  - Other technologies to be explored

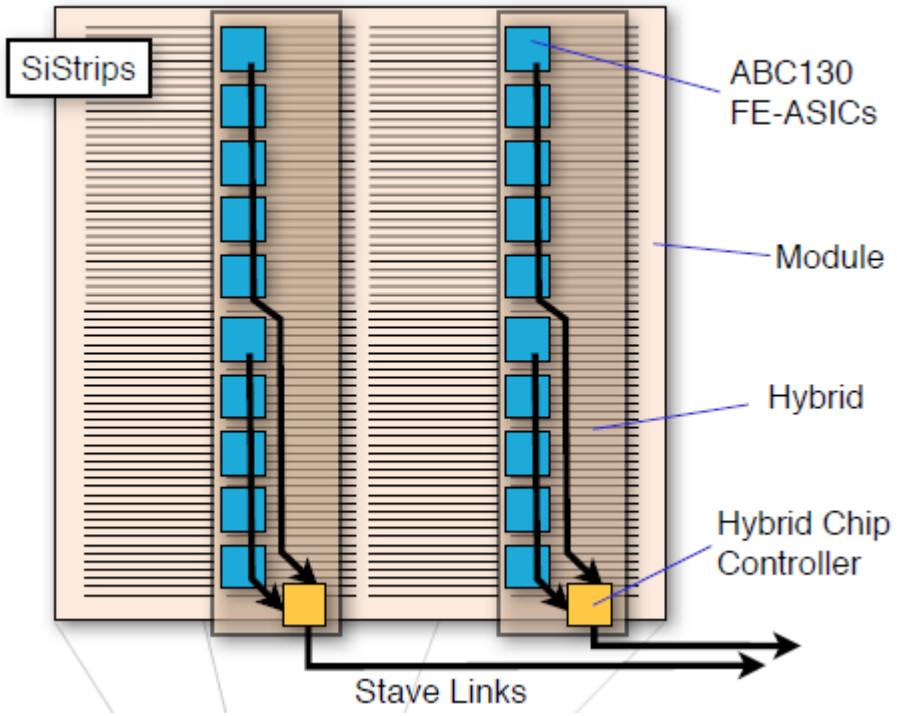

L1Calo

- On L0A, Calo RODs send data for RoIs to L1Calo
  - ~ 2Tb/s (4Mbit @ 500 kHz) on ~ 1000 x 10 Gb/s optical links
- Full granularity cal. Data (e.g., x4 L0Calo in layer 2)

- New cluster algorithms reduce rates
  - Sharper turn-on curves
  - More precise spatial location
- π^0 (calorimeter Layer-1 strips)
- Improved jet algorithms? (e.g., iterative within max. latency)
- Output: TOBs ~2 Gb/s (40 kb @ 200 kHz)

- USA 15
- FPGA based
  - Latency potentially ~ 3 μs

- Send list of rejected L0 RoIs (R3s) to L1track?
- Send RoIs (R3s) to L1Track only after L1Calo?
  - Reduces tracking load

L1 Central Trigger

- Requirements not yet well defined
- Functionality:
  - Match Calo & Muon objects to tracks
  - Topological processing
  - Pre-scaling
  - Deadtime…
  - Generation of L1A
- Share crate with L0CT?
  - Both drive Accepts → distribution
- Use generic hardware with dedicated firmware?
  - High-bandwidth optical links, large FPGAs…

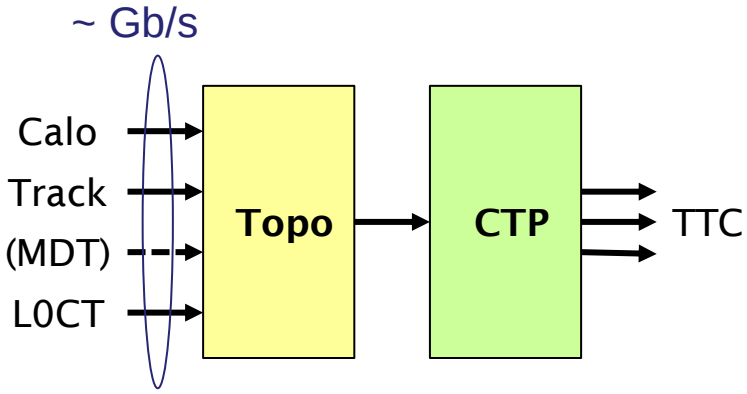

Latency

- Shown is L0/L1 model
- Sparsify L1Track data further to reduce latency?
  - Filter on cluster size as per self-seeded option?
  - L1Calo data?

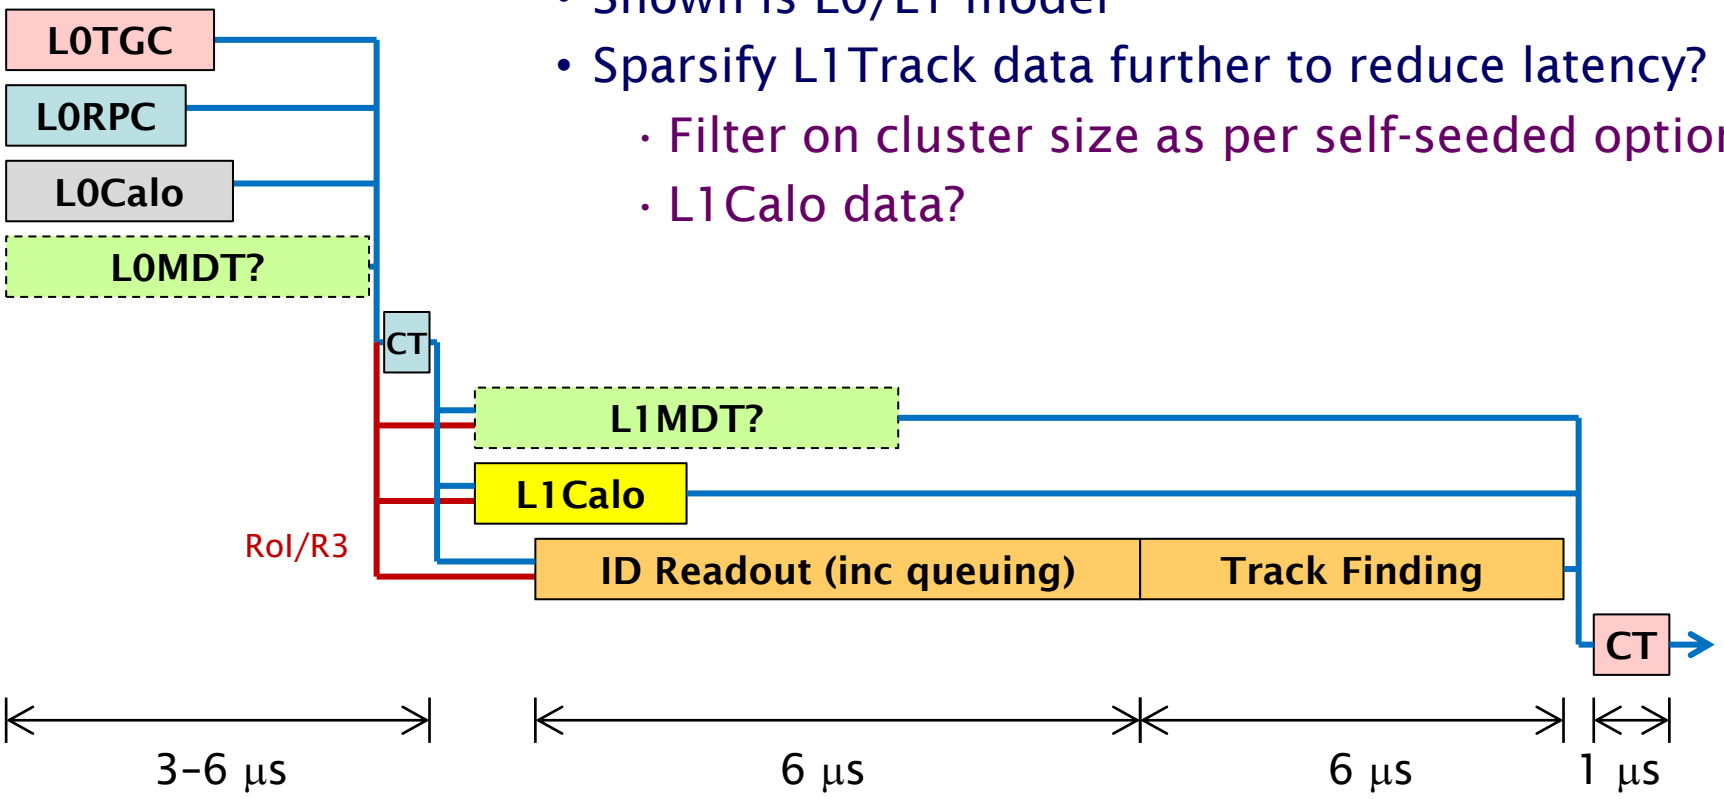

Open Issues

- R&D
  - Accuracy of TOT for RPCs, various Front End ASICs, AM ASICs for L1Track…
- Big questions
  - Architecture
    - Track Trigger RoI-driven ?
    - L0/L1 split?
    - MDT: L0? L1? RoI driven?
  - Latency & rates at L0 and L1
  - Deadtime
  - Algorithms
    - Feature extraction, zero suppression…

} Impact on bandwidth

- Items defining schedule
  - Design of Front End ASICs
  - Trigger hardware installed in Phase I

High Bandwidth Signals

- Motivation: bottlenecks at module & component interfaces
- Optical links ~10 Gb/s (or greater)
  - COTS not a problem
  - Rad hard? Latency?
    - LpGBT required
- High Bandwidth PCB tracks
  - Potentially greater challenge
  - Some systems require more than short, point-point tracks
  - Crosstalk, reflection, attenuation, differential skew…
  - Advanced PCB materials
  - Simulation, manufacturing, testing
    - Issue in Phase I, compounded at Phase II
- Direct optical links to FPGAs?

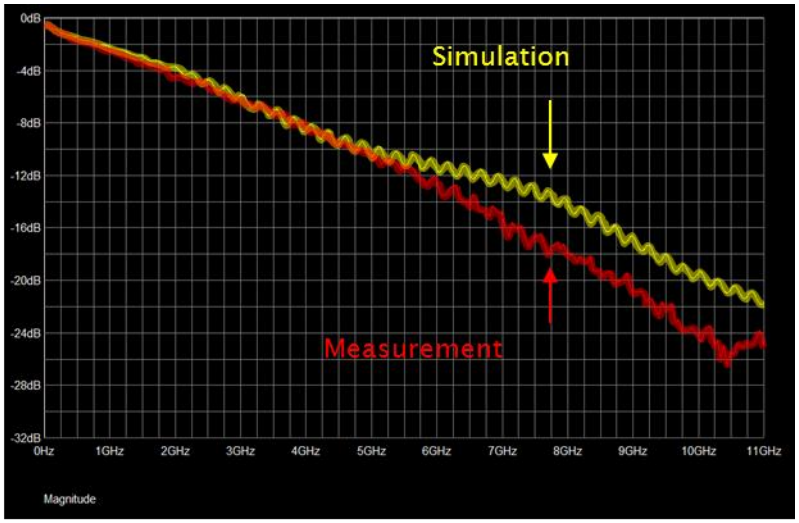

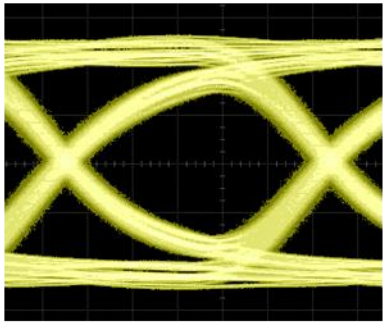

5 Gbps

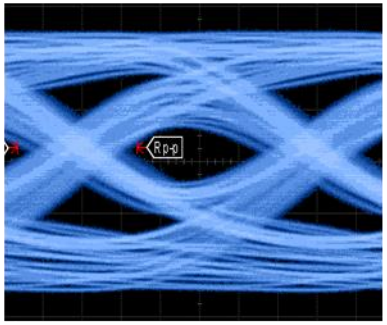

10 Gbps

Common Features

- Movement of trigger logic off detector
  - Relaxes rad-hard requirements
  - Allows FPGA implementation – flexibility
  - Ease of maintenance and operation
- High-bandwidth optical links
- Many of trigger boards will comprise large FPGAs + high-bandwidth optical links
  - Employ common hardware across subsystems?
    - Pro: cost
    - Cons: design compromises, increased complexity, diminished skills base

Summary

- Number of solid proposals of how trigger requirements at HL-LHC can be met
- Track trigger is the key development
- Baseline proposal is L0/L1 architecture
  - L1 RoI driven
- Technology R & D programmes underway
- Key parameters to be fixed