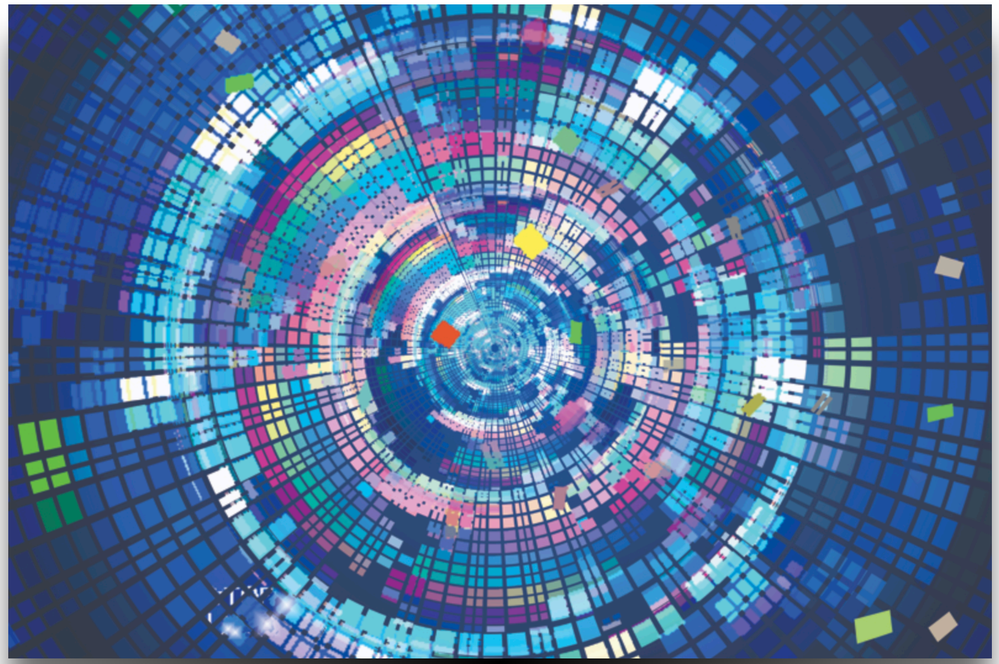

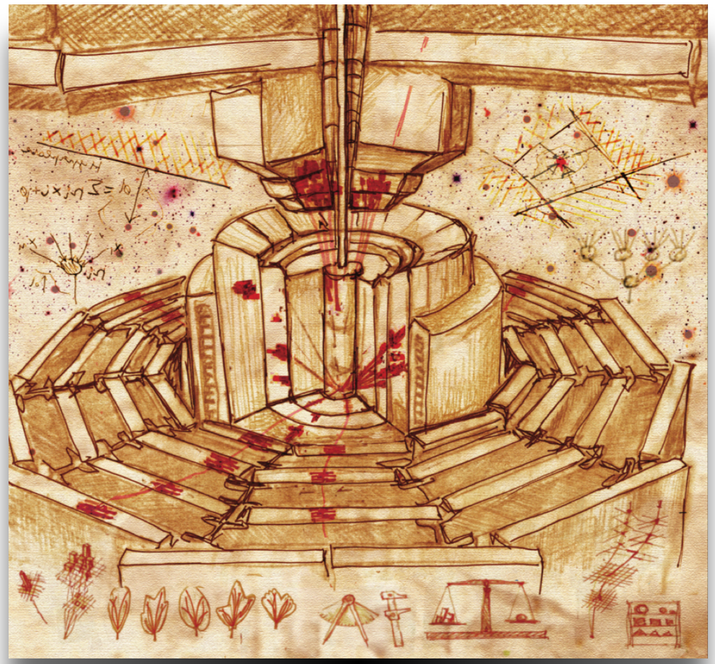

The Upgrade of the CMS Level-1 Muon Trigger for the HL-LHC

J. Konigsberg / TIPP 2023

Outline

- Landscape
  - The High Luminosity LHC (HL-LHC)
  - The CMS HL-LHC upgrade
    - Tracking at Level-1
    - The Muon Detector System upgrade
  - The CMS Level-1 Trigger Upgrade
    - Motivation / Concept / Key features
- The L1 Muon Trigger upgrade
  - Concept and implementation
    - Architecture
    - Electronics
    - Algorithms and firmware
    - What comes next
- Caveats
  - Lots of material to cover, here just a flavor of what is
  - Work in progress, so all preliminary of course

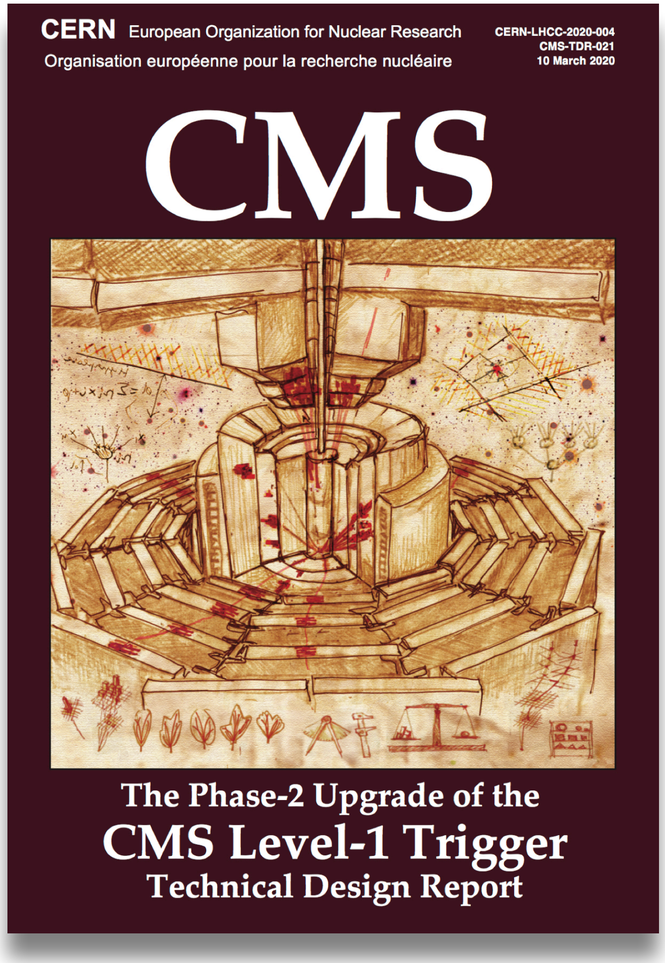

https://cds.cern.ch/record/2714892

~350 pages / ~ 3yrs in the making / ~ 3yrs since

The HL-LHC

- An ambitious project for an ambitious physics program
  - 14 TeV p-p collisions
    - total integrated luminosity of ~3000-4000 fb^{-1}
    - instantaneous ~5-7x10^{34} $cm^{-2}s^{-1}$
    - ~200 simultaneous collisions (pileup)
  - An enormous challenge
    - Accelerator
    - Detectors
    - Trigger and data acquisition
    - Computing
    - Data Analysis

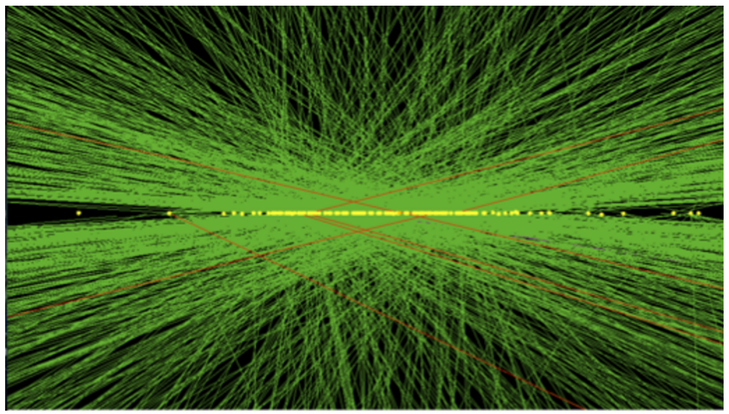

HL-LHC "pileup" challenge

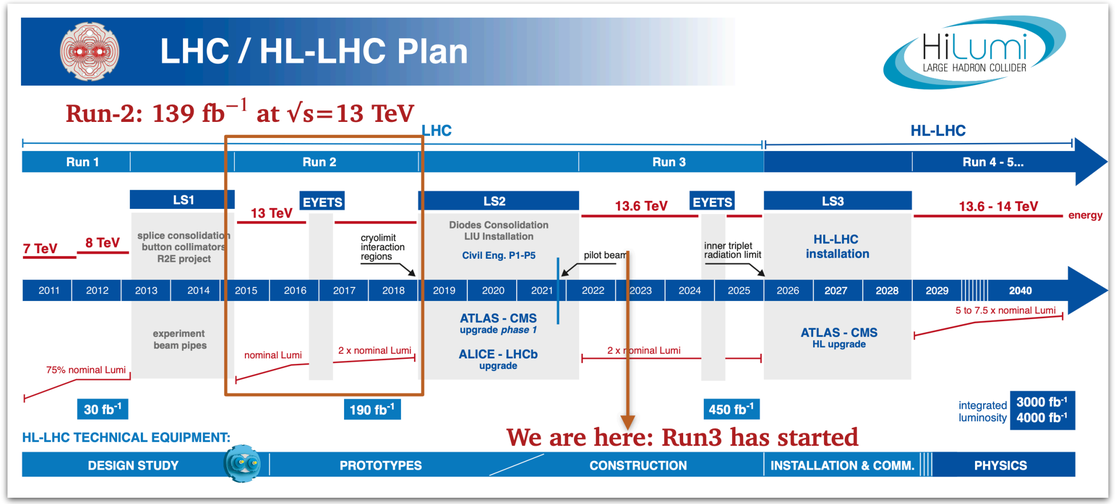

Accelerator Complex Upgrades

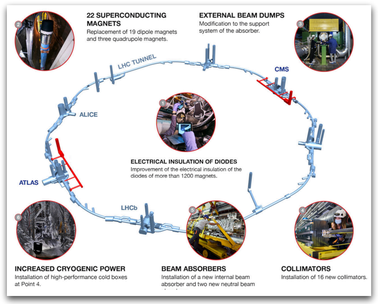

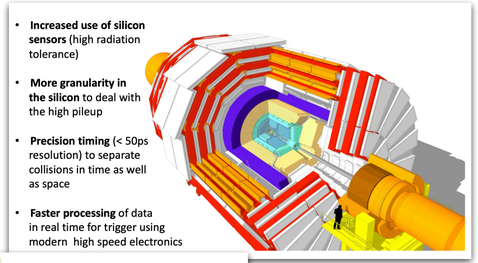

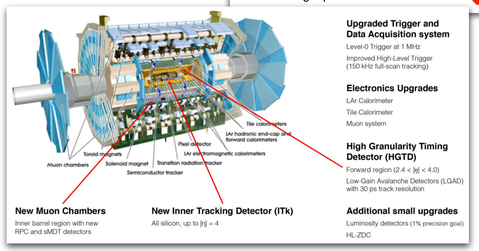

Detector Upgrades

The HL-LHC

- An ambitious project for an ambitious physics program
  - 14 TeV p-p collisions
    - total integrated luminosity of ~3000-4000 fb^{-1}
    - instantaneous ~5-7x10^{34} cm^{-2}s^{-1}
    - ~200 simultaneous collisions (pileup)
  - An enormous challenge
    - Accelerator
    - Detectors
    - Trigger and data acquisition
    - Computing
    - Data Analysis

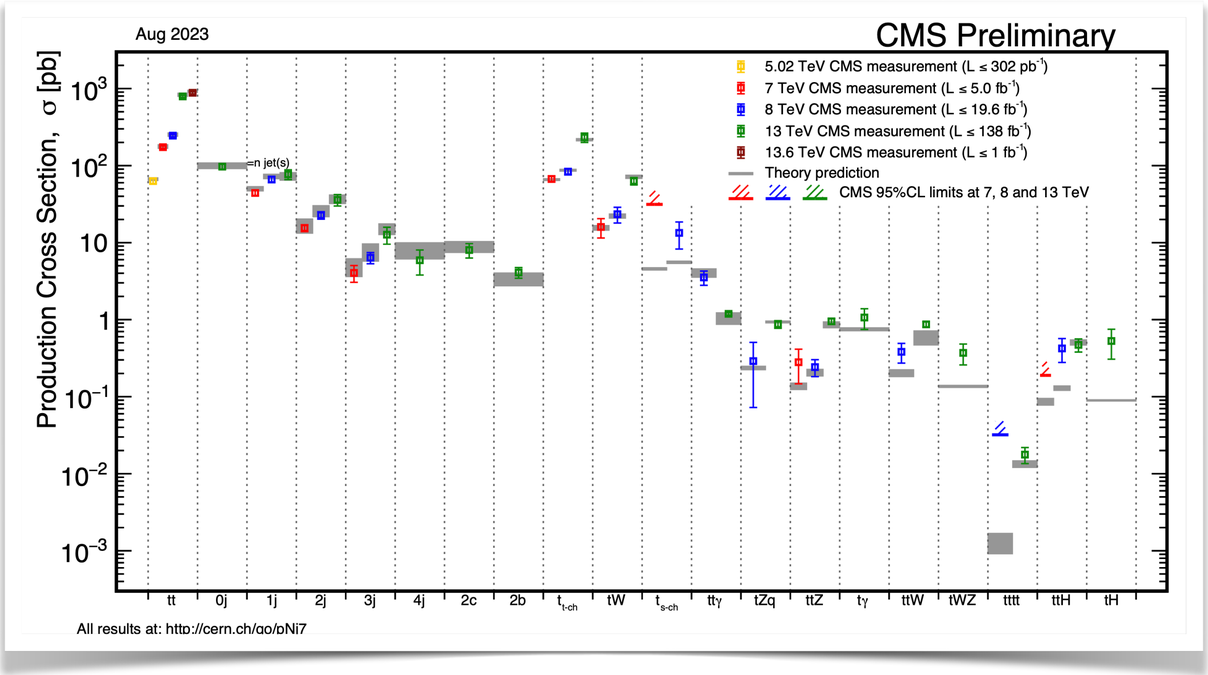

~ x10 more data @ HL-LHC

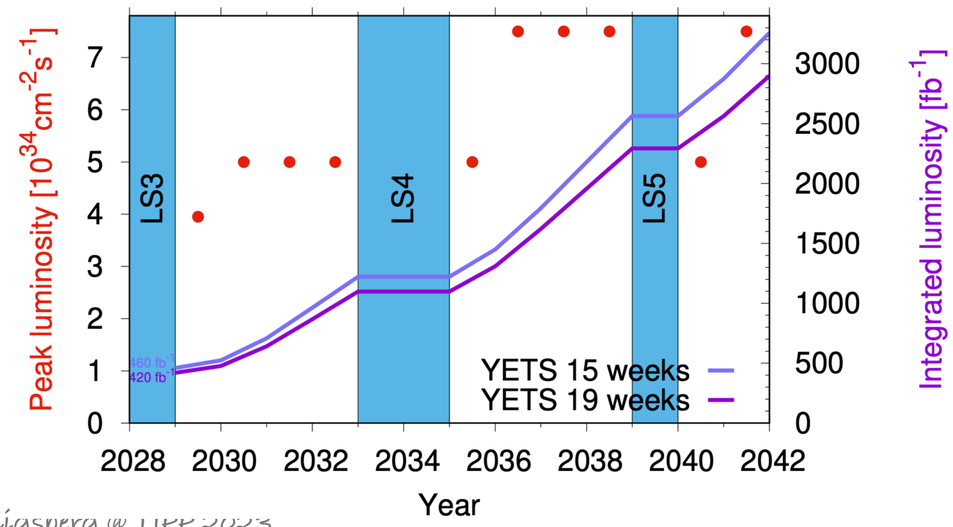

More data:

Rarer processes
More precision
Discovery potential

- Run 1 [2011-12]: 5/20 fb-1 @ 7/8 TeV
- Run 2 [2015-2018]: ~ 140 fb-1 @ 13 TeV
- Run 3 [2022-2025]: started, ~ 300 fb-1 @ 13.6 TeV

The CMS HL-LHC upgrade

- Major upgrade of every system !
- + New systems added…
  - New trackers
    - Tracking at Level-1
  - New and upgraded calorimeters
  - New timing layers
  - New and upgraded muon systems
  - New trigger systems
    - Level-1
    - High Level Trigger
  - Luminosity detectors

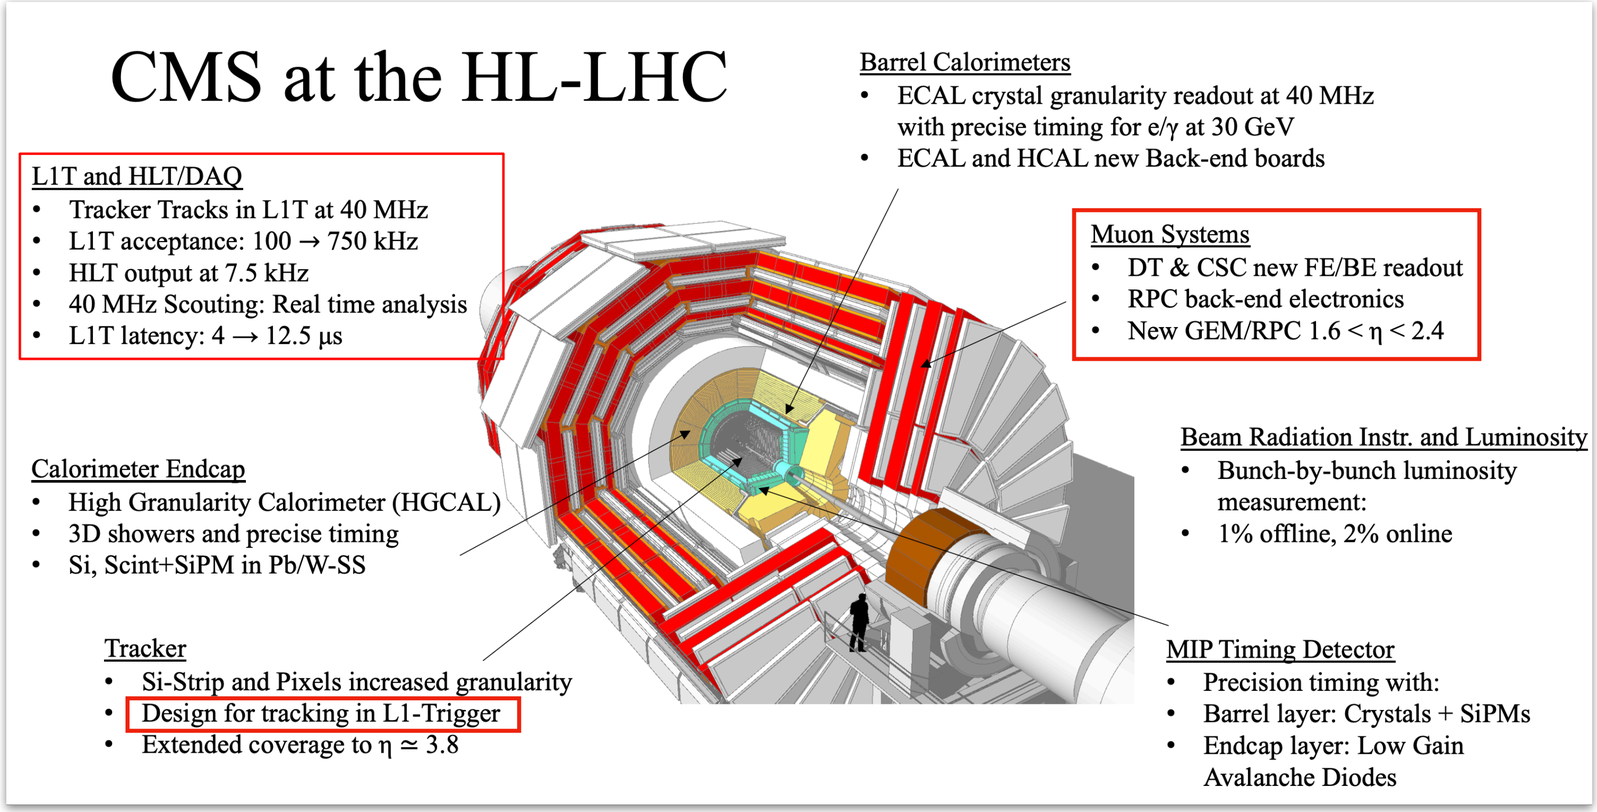

See Anne Dabrowski's presentation on Tuesday

Phase II Outer Tracker

- replacement of entire outer tracker
- 13 200 modules, 190 m^2 silicon, 213 M channel (legacy: 10 M channel)
- 6 barrel layer plus 2×5 endcaps discs

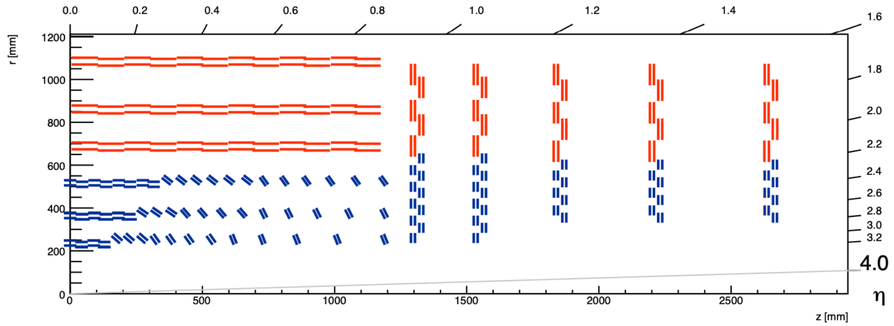

- two kinds of p_T-modules:
  - strip-strip (2S): 5 cm×90 µm
  - pixel-strip (PS): 1.6 mm / 2.5 cm×100 µm

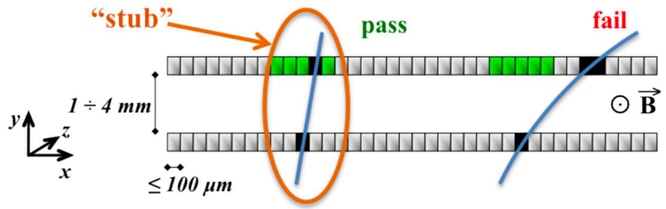

p_T>n GeV tunable window

- new modules allowing rough p_T estimate
- cut at 2 GeV reduces data rate by ~ one order of magnitude
- p_T-information is useful during track reconstruction

Tracking at Level-1

Tilted TBPS: precision assembly of Layer 3 ring

O(x10) reduction in hit multiplicity for p_T>2 GeV hits

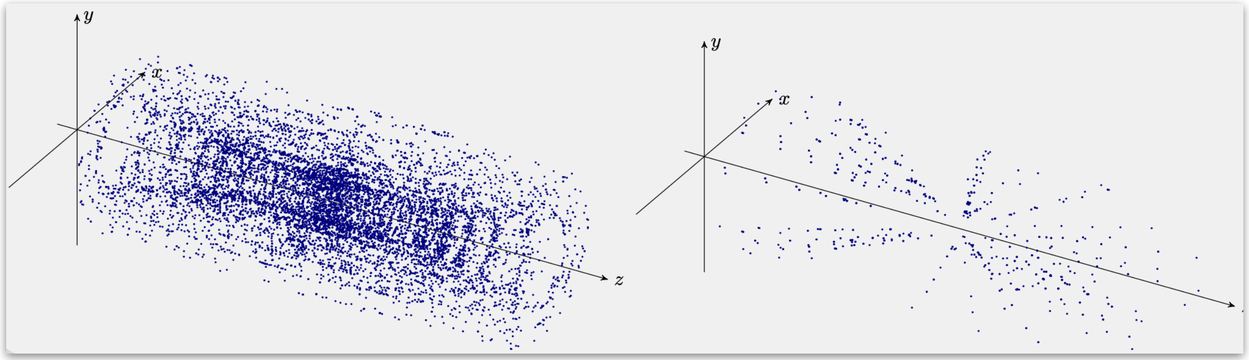

CMS L1 track trigger algorithm

Use p_T>2 GeV hits for track finding

Kalman filter for track fitting

- track finding through road search based on tracklet seeds
- multiple seeding layer combinations used for redundancy and high efficiency
- found tracks sharing stubs get merged (duplicate removal)
- track fit uses kalman filter
- BDT provides qualifiers for improved fake reduction

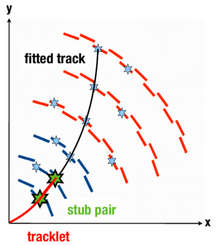

Tracking at Level-1

- Two-stage processing architecture for Track Finding : Layer 1 (DTC), Layer 2 (TF)
- Outer Tracker divided into 9 ϕ sectors
  - Track Finding also segmented into 9 ϕ sectors, with time-multiplexing period of 18

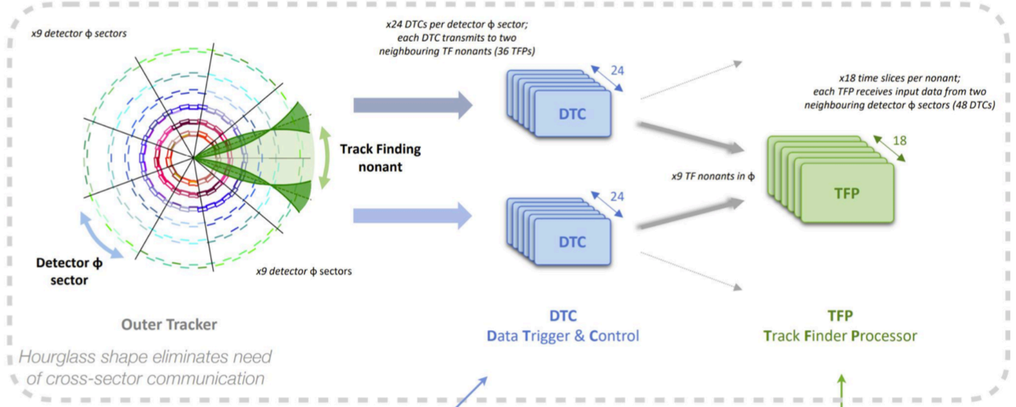

Stub data pre-processing stage

Track finding stage

- Tracks are described by a *96-bit track word*
  - Up to two tracks every 3 link frames (64b)
  - Up to 104 tracks per sector/half-η/event

link to TF-L1 specification

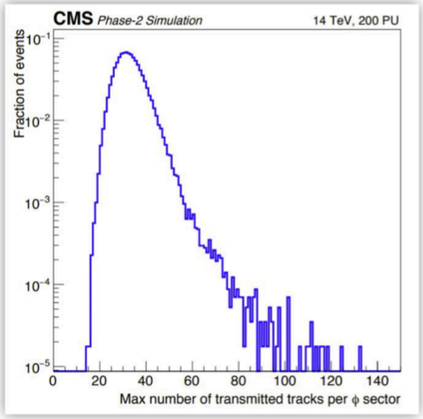

Track word

Bits	Field
0-5	Other MVAs
6-8	MVA quality
9-15	hit pattern
16-18	bend χ^2/d.o.f.
19-31	d_0
32-35	$\chi^2(r\text{-}z)$/d.o.f.
36-47	z_0
48-63	$\tan(\lambda)$
64-67	$\chi^2(r\text{-}\phi)$/d.o.f.
68-79	ϕ_0
80-94	$1/R$
95	isValid

- Output to L1 trigger
  - Each track finder board outputs tracks over 6 pairs of links (η≥0 vs η<0)
    - *3 pairs to Correlator Trigger (layer-1), 2 pairs to GTT, 1 pair to GMT, operating at 25 Gb/s*
    - *Identical data content on each link pair*

Latency of about 5μsec, O(1k) tracks @ 25 Gbps

Muon Detector System upgrade

- Motivation
  - More modern, faster, electronics
  - Added redundancy & robustness, for higher collision rates at larger rapidities
- Barrel region
  - Drift Tubes (DT) and Resistive Plate Chambers (RPC)
    - No new detectors
    - Upgrade of Front End/Back End electronics and readout
- Overlap and Endcap Region
  - Electronics upgrade of several existing systems
    - Cathode Strip Chambers (CSC) & RPC
  - Build additional detectors @ the larger rapidities
    - Improved RPCs (iRPC)
    - New Gas Electron Multiplier detectors
      - GEM
      - ME0

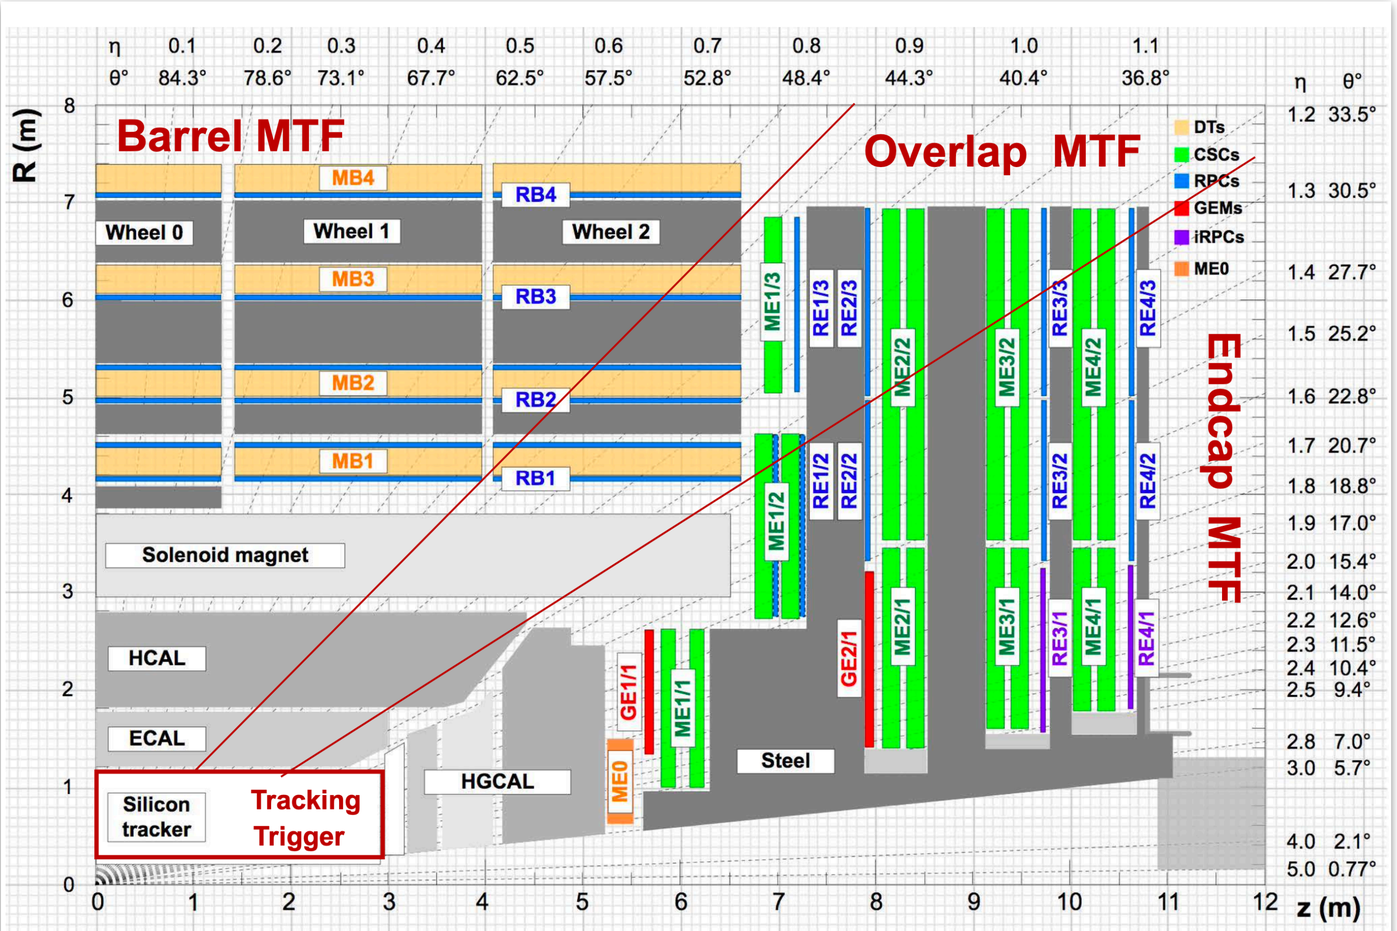

Longitudinal view of a quarter of CMS with highlighted muon detectors and tracker . The approximate boundaries between three muon sub-triggers are indicated.

DT RPC CSC
GEM ME0 iRPC

Muon Detector System upgrade

GE1/1 (Installed in LS2)

- ✓ 36 superchambers (SC) per endcap
- ✓ 2 triple GEM detectors per SC (Gas gap: 3/1/2/1 mm)
- ✓ 10° in φ
- ✓ $1.5 < \eta < 2.0$

GE2/1 (installation before LS3)

- ✓ 18 superchambers (SC) per endcap
- ✓ 2 triple GEM detectors per (SC)
- ✓ 20° in φ
- ✓ $1.6 < \eta < 2.4$

ME0 (installation during LS3)

- ✓ 18 stacks per endcap
- ✓ 6 layers of triple GEM detectors per stack
- ✓ 20° in φ
- ✓ $2.0 < \eta < 2.8$

RE3/1, RE4/1 (installation before LS3)

- ✓ 72 new iRPC chambers
- ✓ 20° in φ
- ✓ $1.8 < \eta < 2.5$

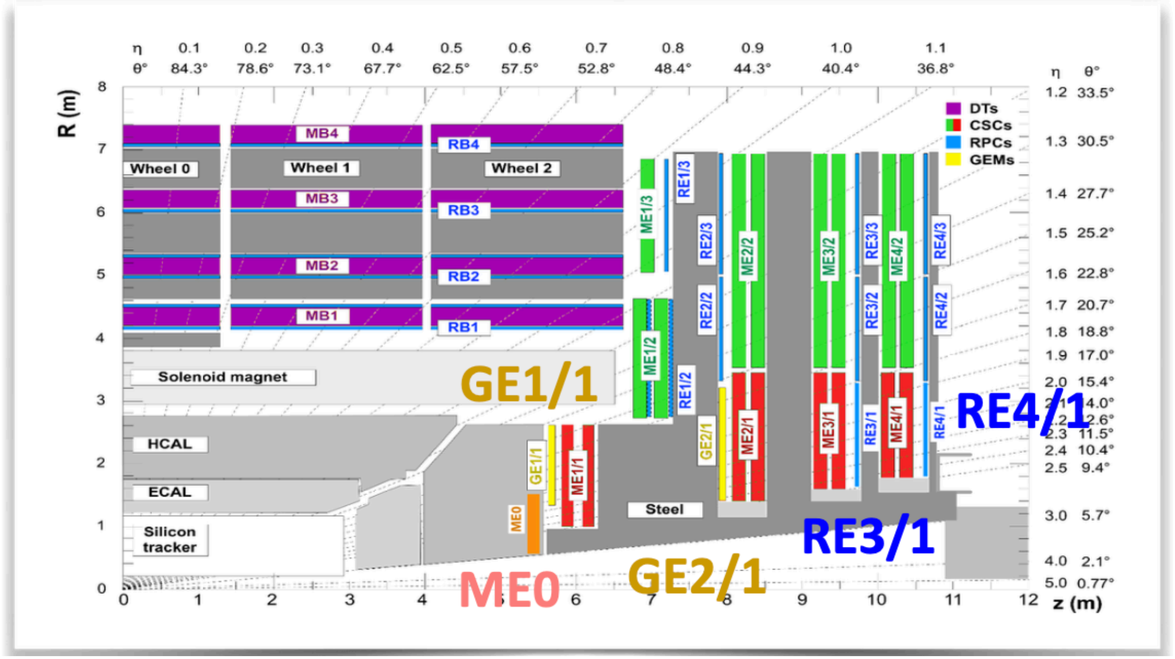

Installation of GE1/1 muon station in 2020

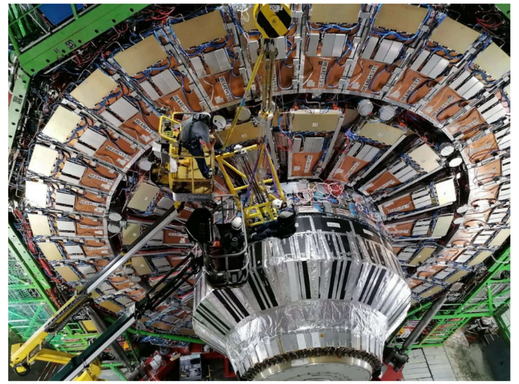

DT electronics slice test

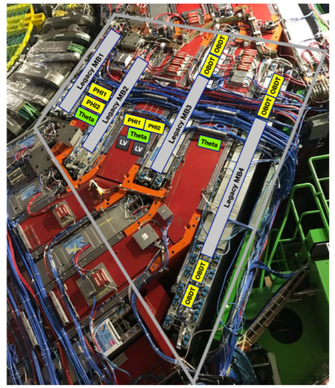

GE1/1 taking data in Run 3

RE4/1 demonstrator installed

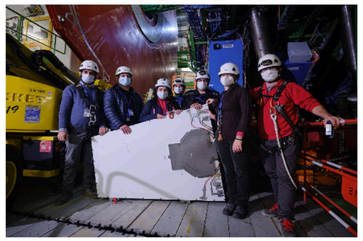

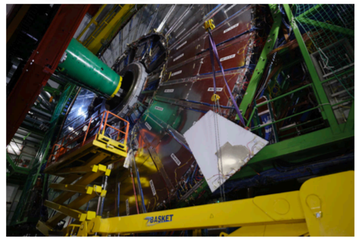

GE2/1 demonstrator

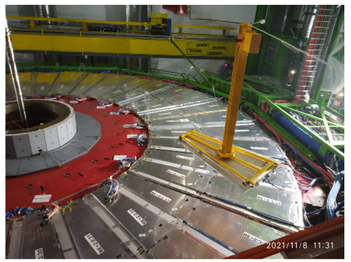

ME0 chambers

The CMS Trigger upgrade

- For HL-LHC CMS will keep its two-tier trigger concept

➢ *Level-1 (hardware based) Trigger*

- Extensive use of state-of-the-art FPGAs
- Increase bandwidth: 100 kHz ⇒ 750 kHz
- Increase Latency: 3.8 μs ⇒ 12.5 μs
- Higher granularity (calorimeters & muon systems) and tracking information
- Sophisticated object reconstruction and correlation
- Enhanced physics selection & Scouting system

➢ *High Level (software) Trigger*

- Optimize reconstruction: Improve physics reach, maintain thresholds while increasing efficiency and stay within computing resources.
- Reduction rate (100:1) 1 kHz ⇒ 7.5 kHz
- Data throughput: 2.5 Gb/s ⇒ 61 Gb/s
- Heterogenous architecture – CPU/GPU

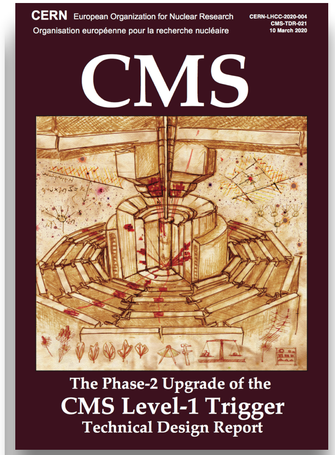

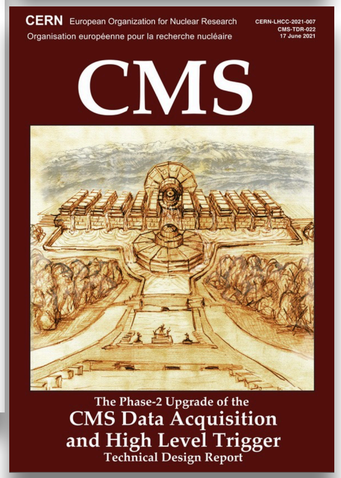

The Level-1 Trigger upgrade

General Features:

- Maintain and expand on current physics capabilities in high-PU environment
- Utilize all information from upgraded detectors
  - Independent **Calorimeter**/**Muon**/**Tracker** paths
  - **Correlator** system for particle flow and mixed-object triggers
  - **Global Trigger** for trigger menu implementation
  - **Data Scouting** at 40 MHz: physics with L1 info only

Output

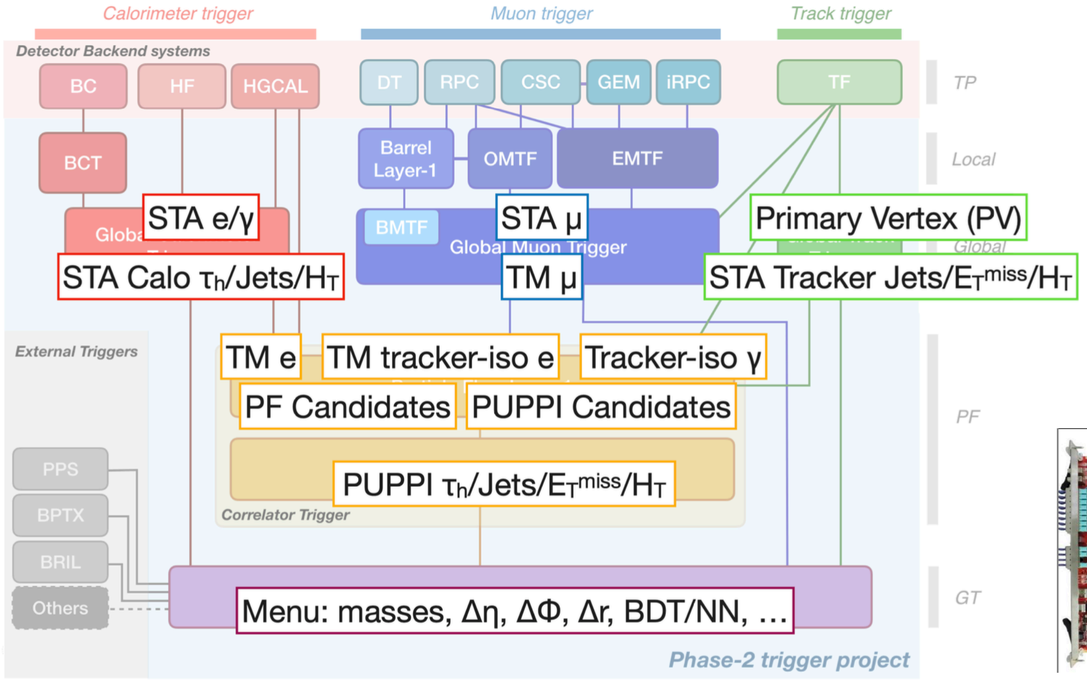

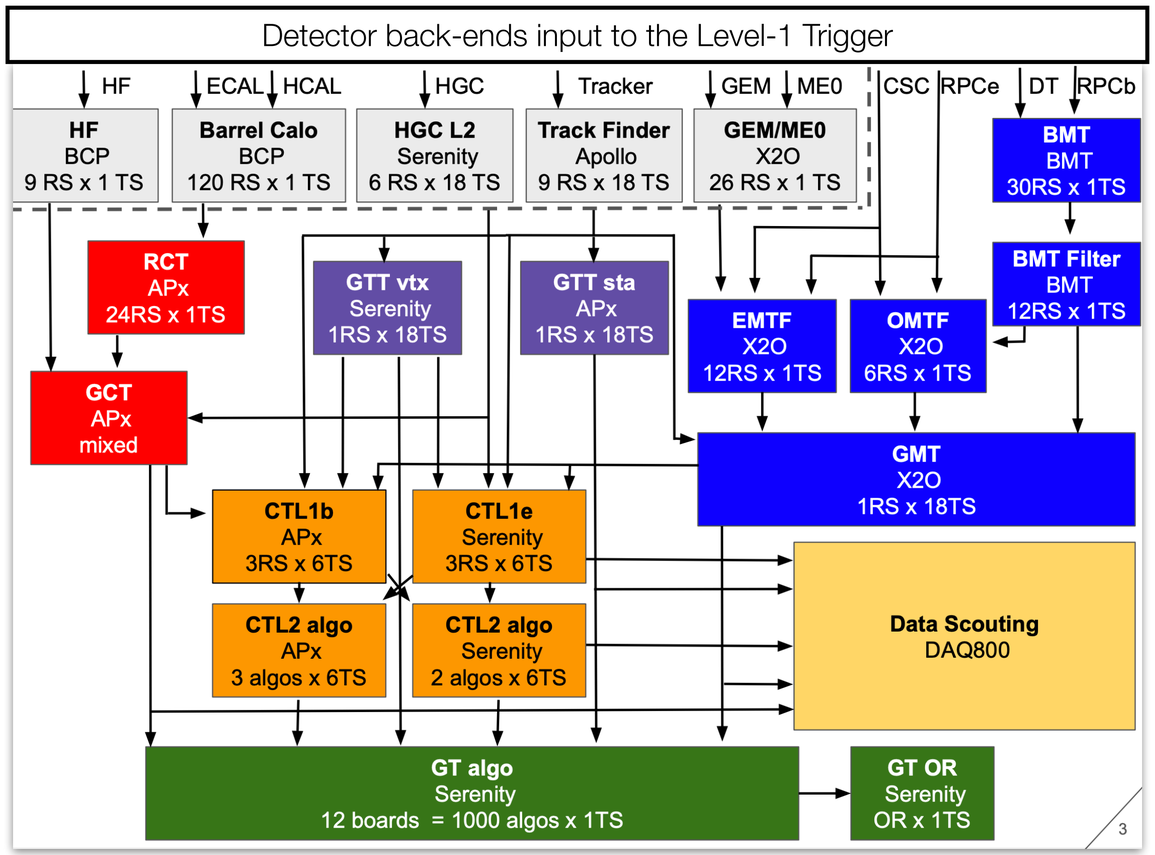

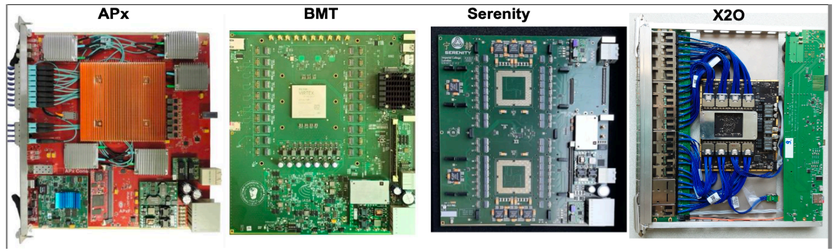

ATCA digital processors
w/ Xilinx VU13P-2 A2577 FPGA
running @ 25 Gbps

The Level-1 Trigger upgrade

- Quite a jump in complexity

Current L1 Trigger

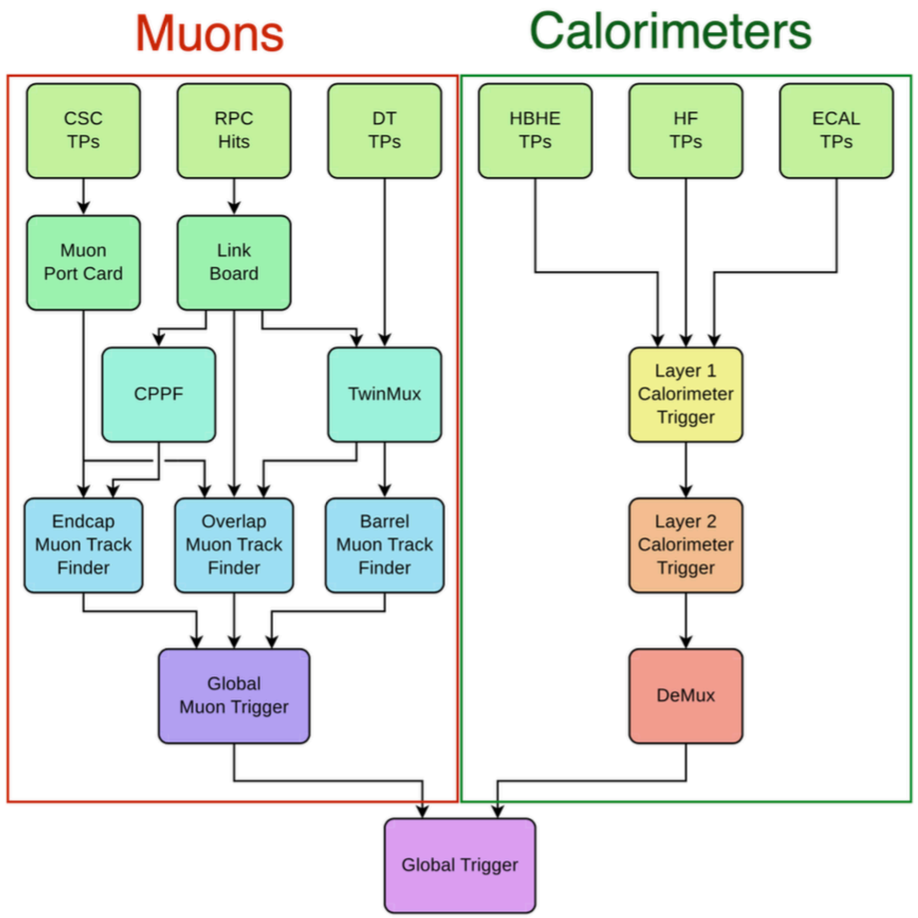

HL-LHC L1 Trigger

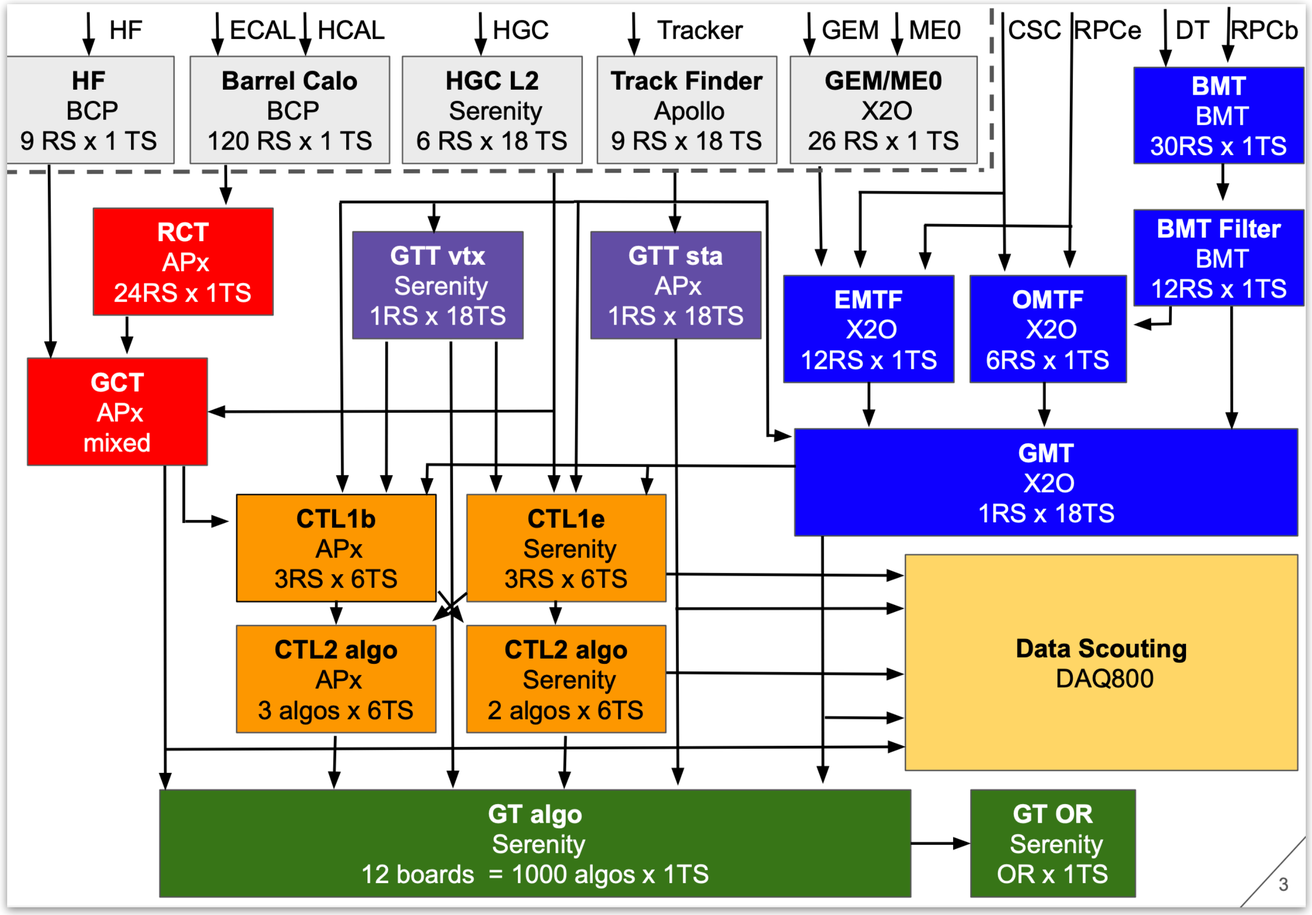

The Muon Trigger upgrade

- Big jump in complexity here as well

Current L1 Muon Trigger

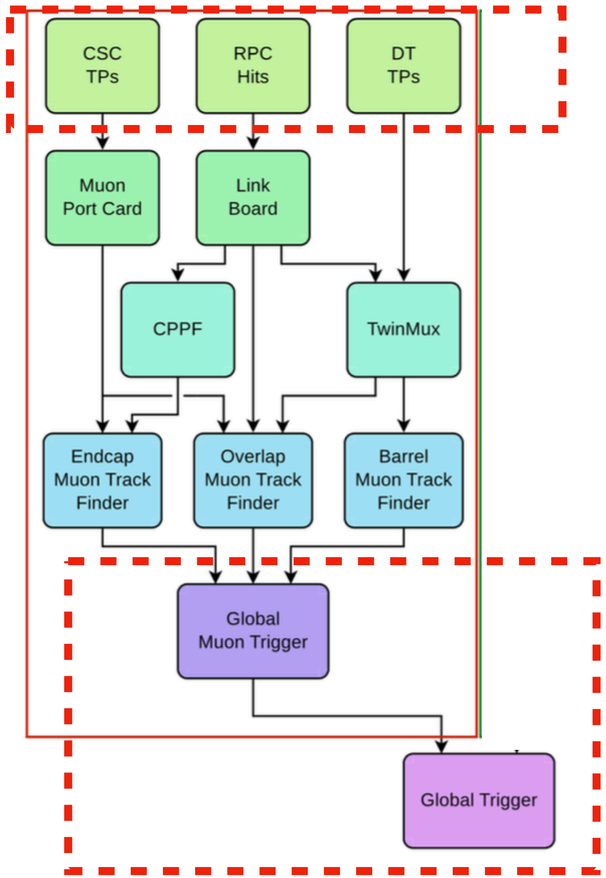

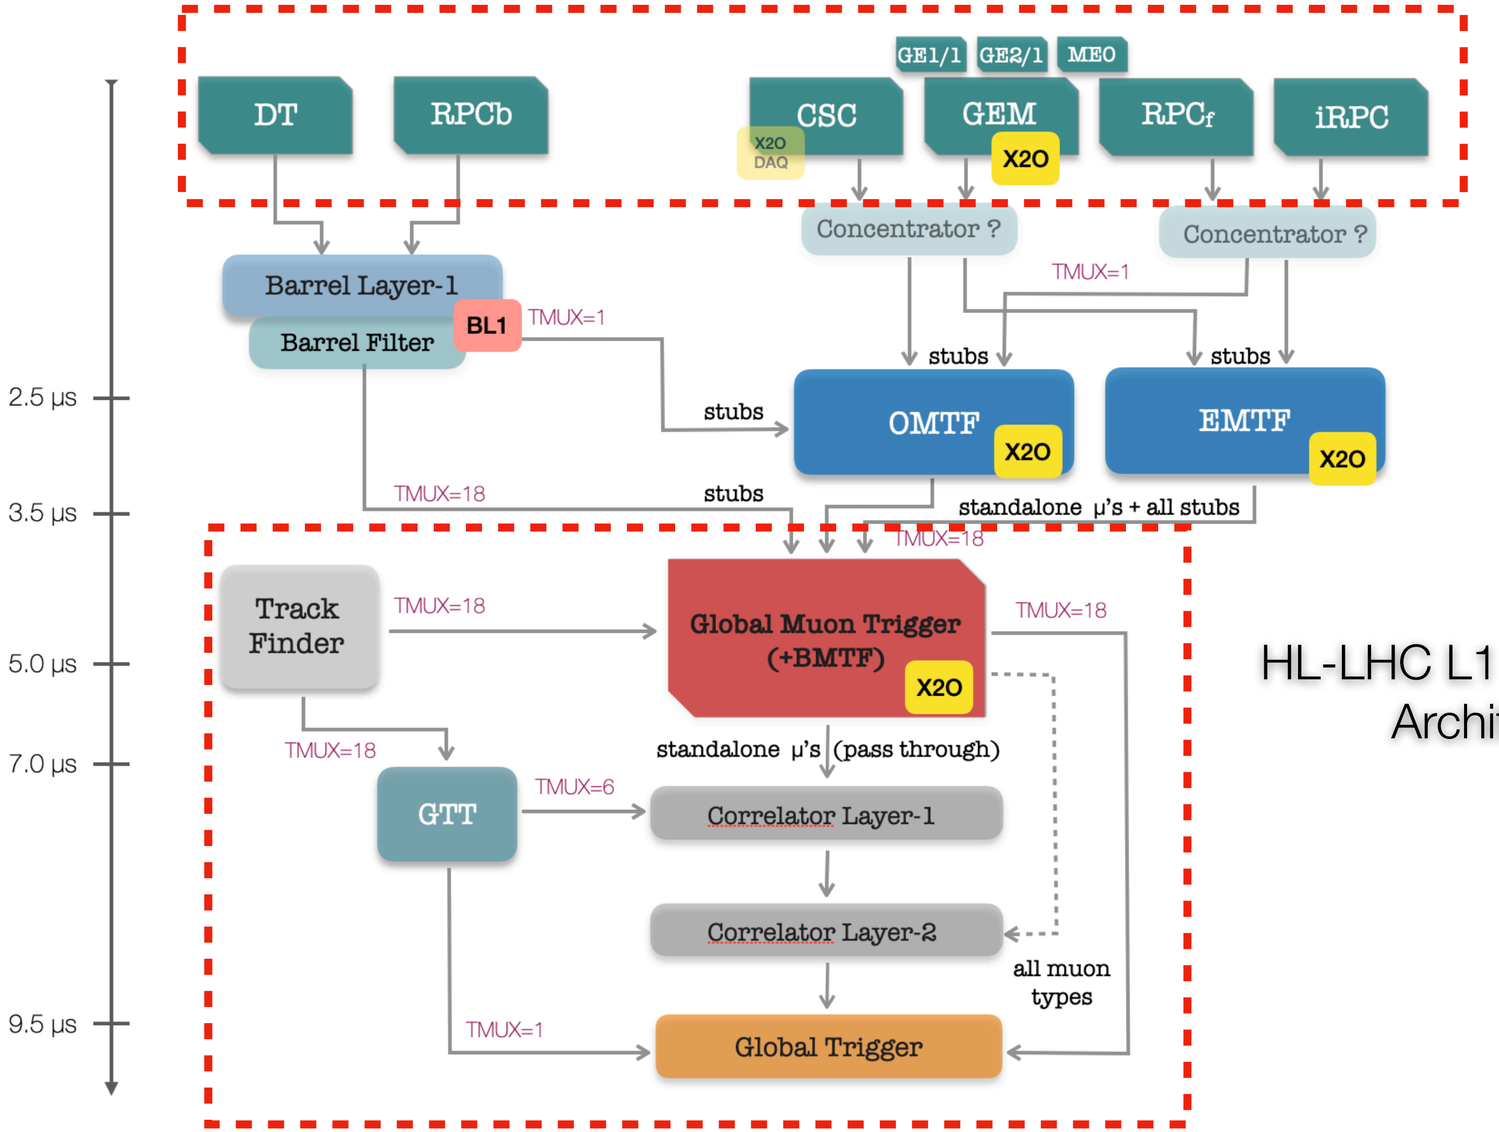

HL-LHC L1 MuonTrigger Architecture

Muon Trigger architecture

- Let's deconstruct this some
- At the core there are 3 systems [the focus here]
  - **EMTF/OMTF** = Endcap/Overlap muon track finders (~2μ sec)
    - Reconstruct standalone muons from μ systems in the corresponding region
      - Prompt and displaced [new] for long lived particles
    - Output to GMT [including all muon stubs] and standalone μ directly to the Correlator and the Global Trigger
  - **GMT** = Global Muon Trigger (~2μs)
    - Operates at 18 time multiplexing factor
    - Reconstructs standalone barrel muons
      - Prompt and displaced
    - Matches tracker tracks with muons and muon stubs from all muon systems
    - Runs topological algorithms e.g.:
      - Track-isolated muons
      - Tau=>3μ BSM decays
      - …
    - Outputs these objects to the Correlator and Global Trigger

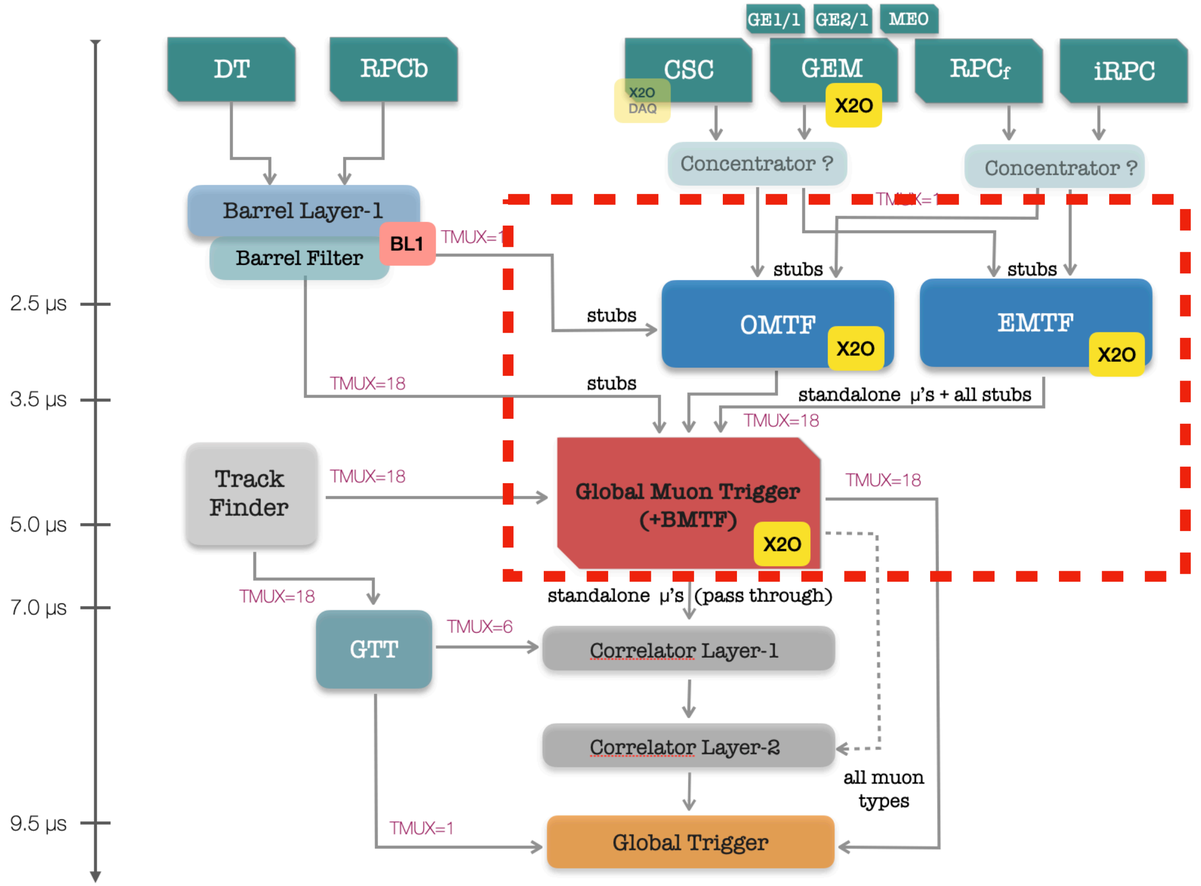

Muon Trigger interfaces

- Up to ~20 interfaces
  - inputs @ 25 / 10 / 3 Gbps
    - 10 with the μ detectors @ TMUX=1
      - Will aim to reduce down to ~6
      - muon stubs
    - 1 with the Track Finder @ TMUX=18
      - Tracker tracks
  - Internal @ 25 Gbps
    - 1 between BMTL1 and OMTF @ TMUX=1
      - muon stubs
    - 1 between BMTL1 and GMT @ TMUX=18
      - muon stubs
    - 2 between OMTF/EMTF and GMT @ TMUX=18
      - standalone (SA) μ's and hybrid stubs
  - outputs @ 25 Gbps
    - 1 to the CL1 @ TMUX=18
      - standalone (SA) μ's
    - 1 to the GT @ TMUX=18
      - all muon types
    - 1 to the CL2?

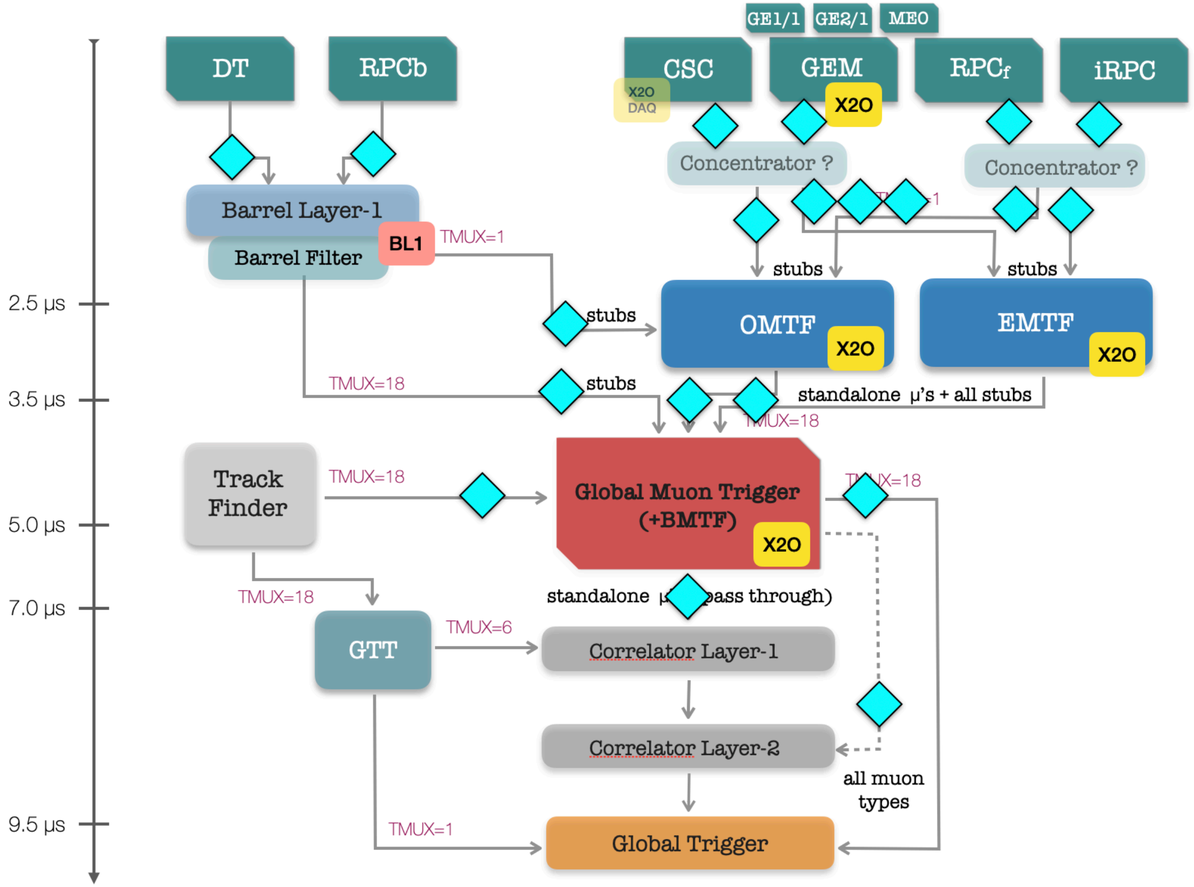

We have defined the data structure/payload for all these

Muon Trigger interfaces examples

EMTF ==> GMT interface data structure

Range	Field	Bits	Format	Notes
Standalone Muons [64 bits]				
12–0	p_T	13	unsigned int, LSB 31.25 MeV	saturates at 256 GeV
25–13	ϕ	13	signed int, LSB $2\pi/2^{13}$	at PCA to beamline
38–26	η	13	signed int, LSB $2\pi/2^{13}$	at PCA to beamline
43–39	z_0	5	signed int, LSB 1.875 cm	range ±30 cm
50–44	d_0	7	signed int, LSB 3.85 cm	range ±246.4 cm
54–51	quality	4	tbd	zero if null muon
55	charge	1	0 = positive, 1 = negative	
59–56	β	4	unsigned int, LSB 6%	linear in 0-1 range
63–60	spare	4		
5 hybrid stubs associated with this muon [64 bits]				
3–0	quality	4	tbd	zero if null muon
6–4	BX	3		bunch crossing
19-7	ϕ_1	13	signed int, LSB $2\pi/2^{13}$	global coordinate
32-20	ϕ_2	13	signed int, LSB $2\pi/2^{13}$	global coordinate
44-33	η_1	12	signed int, LSB $2\pi/2^{12}$	global coordinate
56-45	η_2	12	signed int, LSB $2\pi/2^{12}$	global coordinate
61–57	time	5	signed int, LSB 25/32ns	
63–62	spare	2		

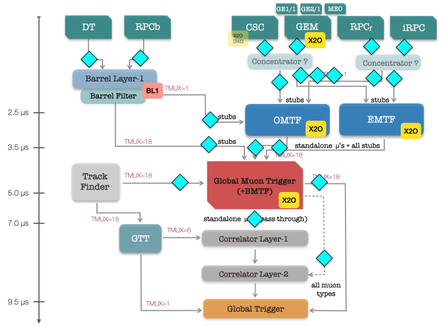

GMT will send 12 track-matched muons, sorted by pT.
Isolation and beta included ("TrackerMuons")
Each track-matched muon involves 96 bits.

- **Send to GMT a fixed number of SA μ's**
  - 64 bit info for **prompt muons** (up to **4/sector**)
  - 64 bit info for **displaced muons** (up to **4/sector**)
  - SA info contains p_T, eta, phi, z0, d0, beta
  - With each SA muon include up to **5 of its hybrid stubs** [64 bits each]
  - Deserialize, sort, and prioritize at the GMT
- **Send to GMT up to 32 loose hybrid stubs**
  - right behind the SA muon information
  - 64 bit info for each [=> will add geometrical address]
- **Payload** [estimate in backup]
  - **1 link per TF board to each GMT board @50% occupancy**
  - **Each GMT board gets 12(EMTF)+6(OMTF) = 18 input links**

Range	Field	Bits	Format	Notes
Four-vector frame [64 bits]				
0	valid	1	0 = null object, 1 = valid object	
16–1	p_T	16	linear, LSB = 31.25 MeV	Saturates at 256 GeV
29–17	ϕ	13	two's complement, LSB = $2\pi/(2^{13})$	
43–30	η	14	two's complement, LSB = $2\pi/(2^{13})$	
53–44	z_0	10	two's complement, LSB = 0.05859 cm on ±30 cm	Same as track trigger
63–54	d_0	10	two's complement, LSB = 0.03 cm on ±15.4 cm	Reduced to 10 bits
Next frame [32 bits]				
0	q	1	0 = positive, 1 = negative	
8–1	Q	8		Track-match related
12–9	Isolation	4	undefined, ignore for now	3 thresholds for absolute, 3 thresholds for relative
16–13	β	4	16 bins, LSB = 6%	Linear increments in [0,1]
31–17	spare	15		Reserved

EMTF/OMTF/GMT Hardware

- Unified home-made ATCA platform (“X2O”)
- Modular design w/ three separate (halogen-free) modules
  - Power/Control/Monitoring module
    - AMD/Xilinx Kria Zynq FPGA
  - Optical Module
    - QSFP optics 30 cages 120 Rx + 120 Tx links @25Gbps
  - Processing module
    - With Xilinx VU13P-2 FPGA, running all μ algorithms
      - 3.7M LogicCells, 1.7M LUTs, 94Mb BRAM, 12k DSP
- System size
  - EMTF/OMTF/GMT = 12/6/18 boards, 36 in total
- Status
  - Versions of each module are at, or very close to, final production specs
  - Extensive tests performed on single boards
    - Signal integrity
    - Temperature
    - Optics BER

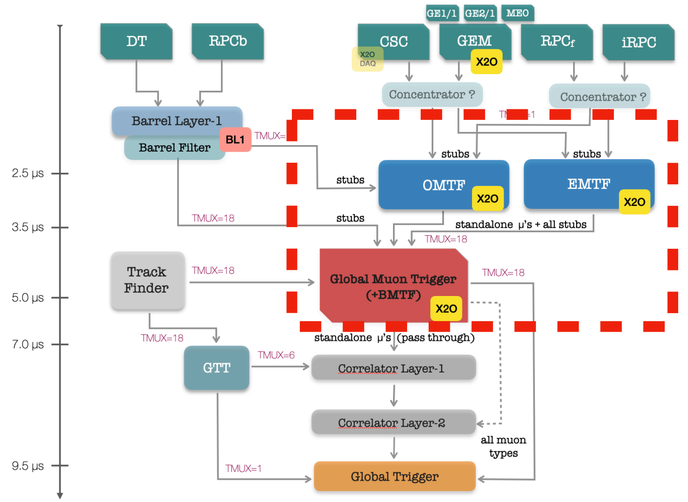

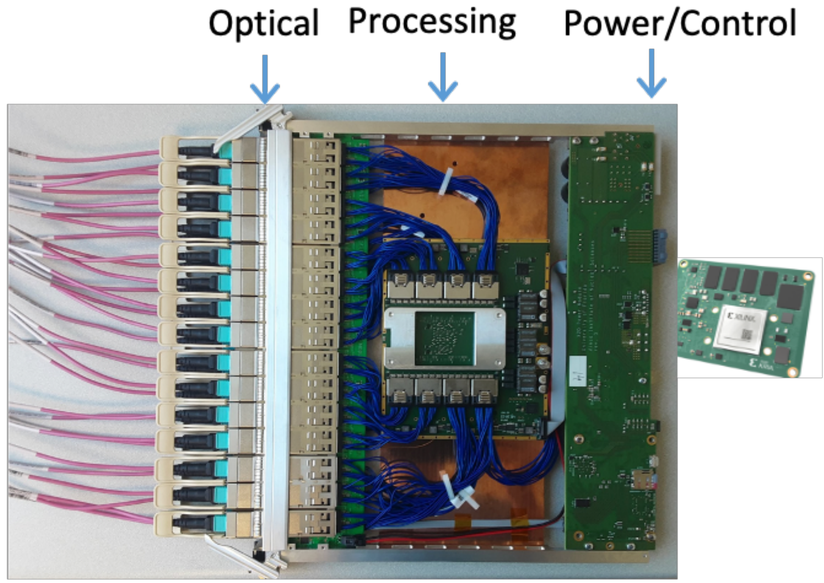

X2O ATCA modular blade

Muon trigger board tests

X2O QSFP signal integrity (electrical)

All positions tested with latest optical module and QSFP

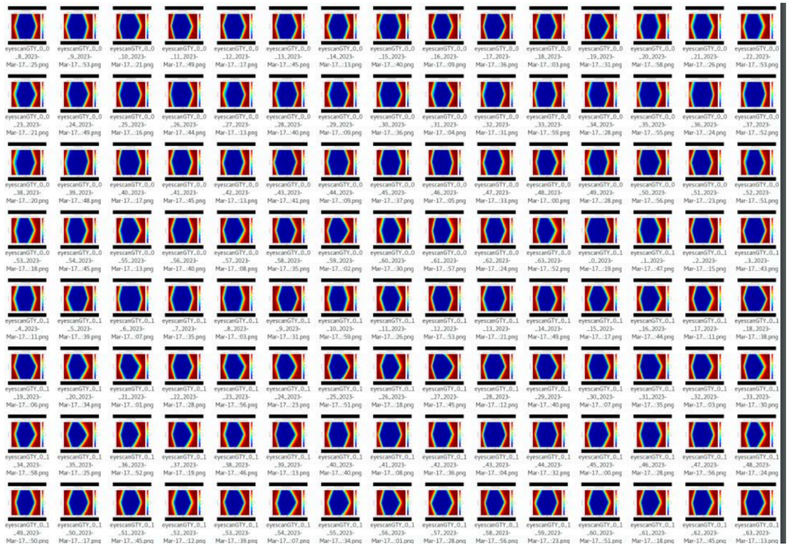

X2O thermal performance

- FPGA temperature of 85C with 200W at fan level 10/15
- QSFP temperature < 50C with 40% airflow on all cages
- Three board side-by-side test with all optics on shows optics T< 50C with 25C input air and airflow at 40%
  - QSFP optics specified up to 70C

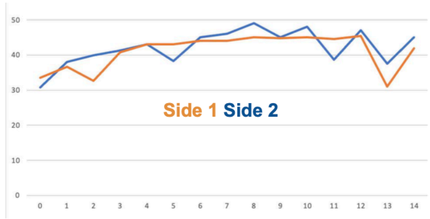

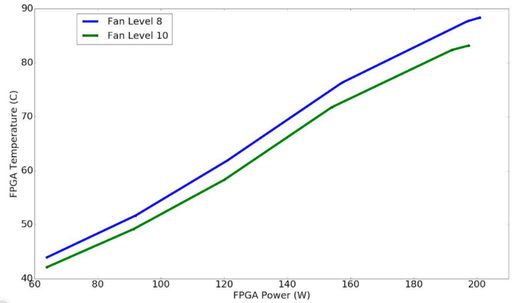

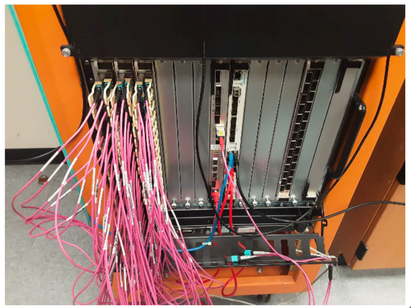

X2O QSFP optical performance

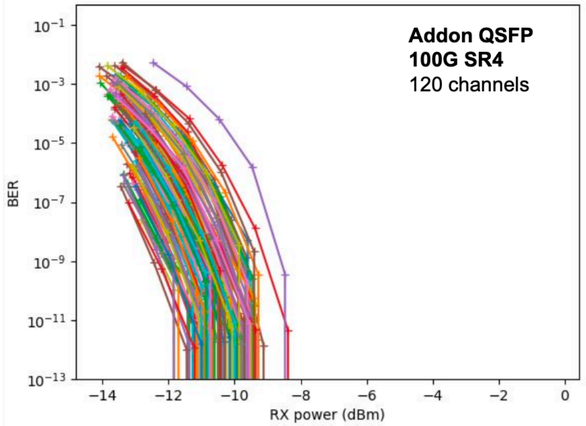

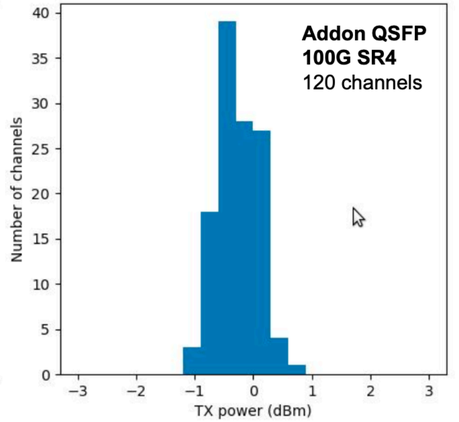

- New results with proper precision attenuator for 120 optical channels
- BER measured as a function of average optical power for different attenuation
- No errors for attenuation levels of 8 dB or more for all channels.
- Lower TX power - higher RX sensitivity → ~7.5 dBm margin

All looks very good.

Muon Trigger algorithms in a nutshell

Standalone Barrel Muons: BMTF

- Runs on the GMT board
- Uses DT and RPC hits
- Outside in Kalman Filter track building
- Allows for displaced muons [not beam constrained]

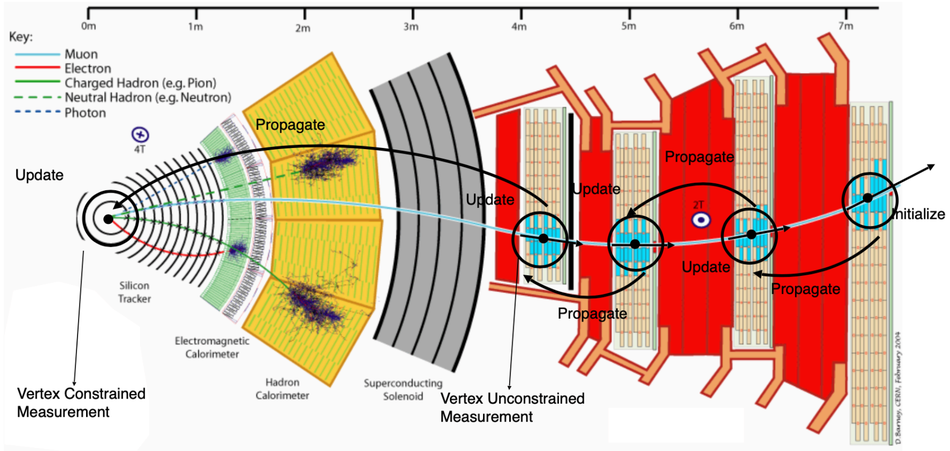

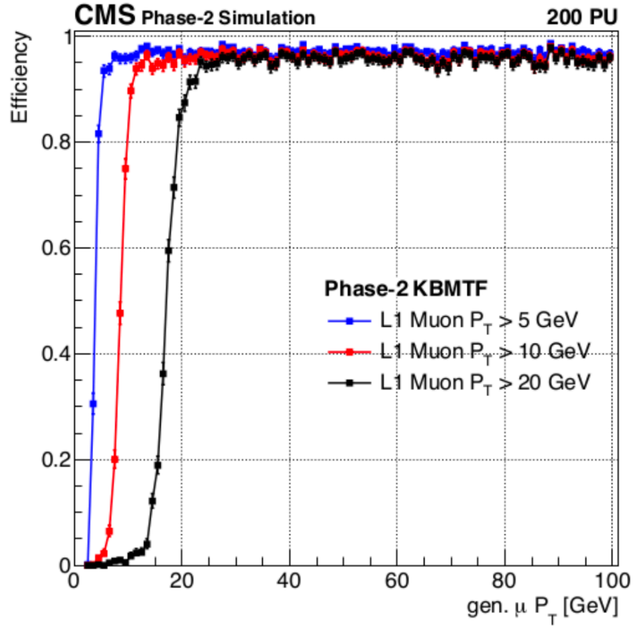

Efficiency vs pT

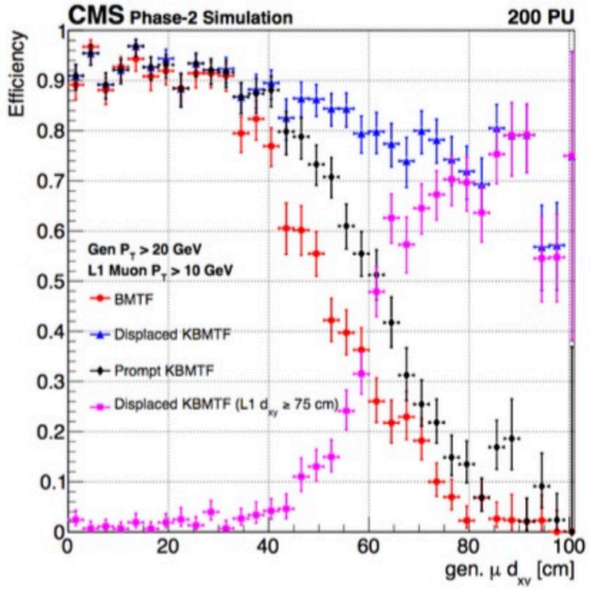

Efficiency vs dxy

Good eff up to ~100 cm

Muon Trigger algorithms in a nutshell

Standalone Overlap Muons: OMTF

- Naive Bayes Classifier
- Uses pre-determined p_T binned hit patterns from DT/RPC/CSC
- Probabilistic comparison of actual hits selects best muon track

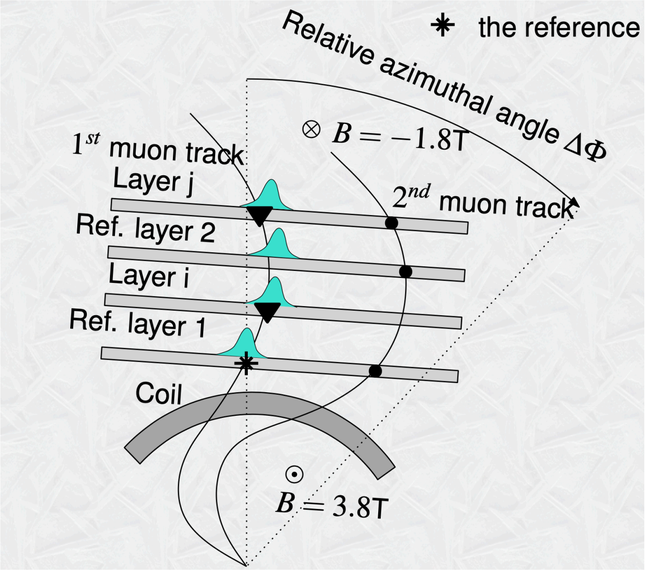

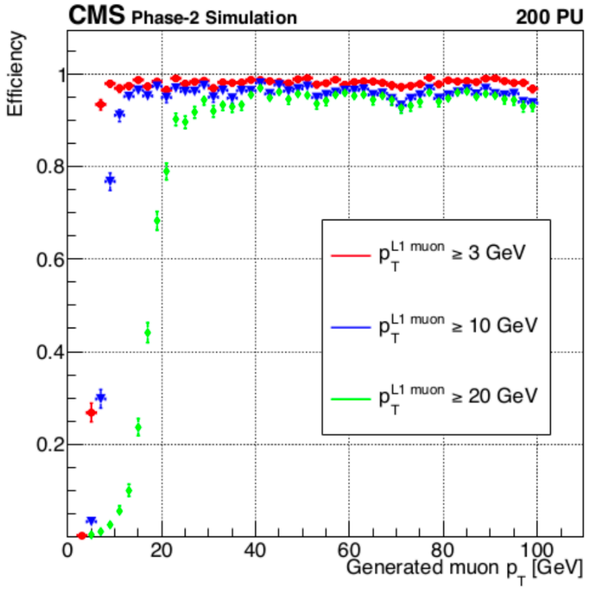

Efficiency vs pT

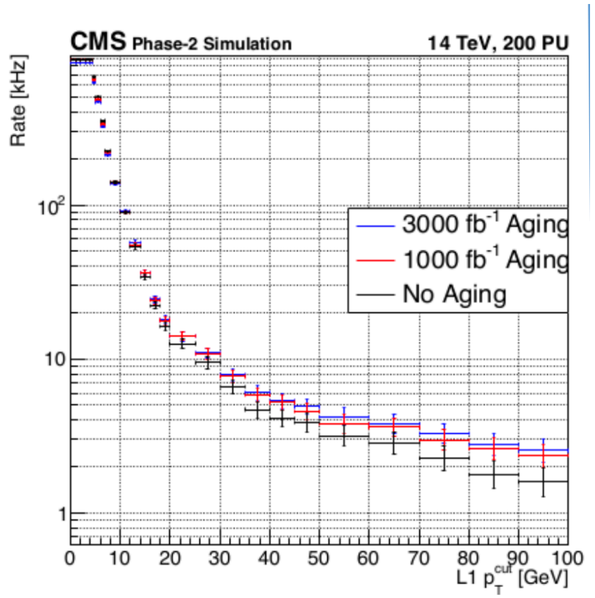

Rate vs pT

~10kHz @ 20 GeV threshold [target for these track finders]

[+ageing simulation]

Muon Trigger algorithms in a nutshell

Standalone Overlap Muons: EMTF

- Fiducial pattern segmentation {eta, $1/p_T$} of the detector layer consistent with μ trajectories
- Uses CSC/GEM/RPC/iROC/ME0 hits
- A pattern fires if enough hits on it
- Hit information {eta, phi, phi-bend, time, quality} is input to a Neural Network to assign p_T and displacement
- Patterns can also be tailored for displaced μs

Patterns for one eta region: layers (y axis), $1/p_T$ squares

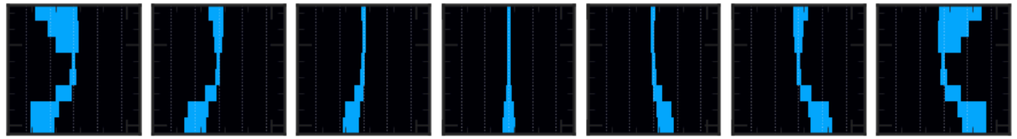

	ME1/1	ME1/2	ME2	ME3	ME4	RE1	RE2	RE3	RE4	GE1/1	GE2/1	ME0
ϕ	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔
θ	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔
bend	✔	✔	✔	✔	✔							✔
quality	✔	✔	✔	✔	✔							✔
time												

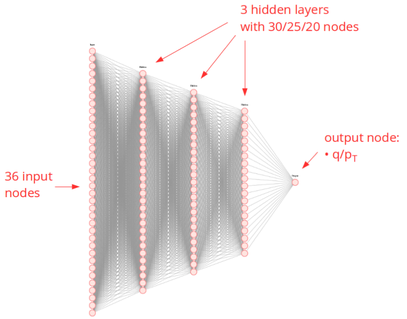

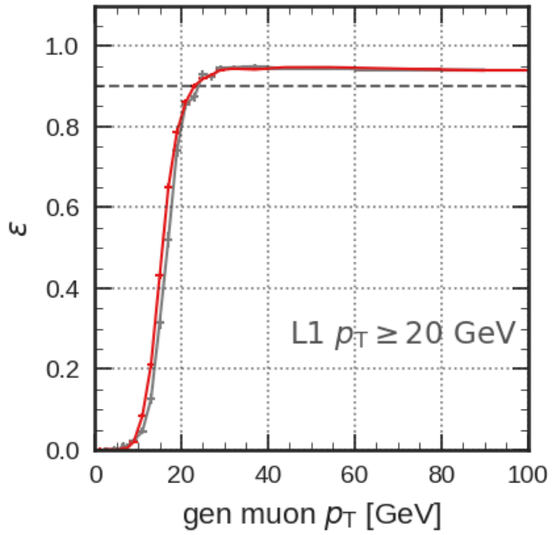

Efficiency vs pT

Rate vs pileup

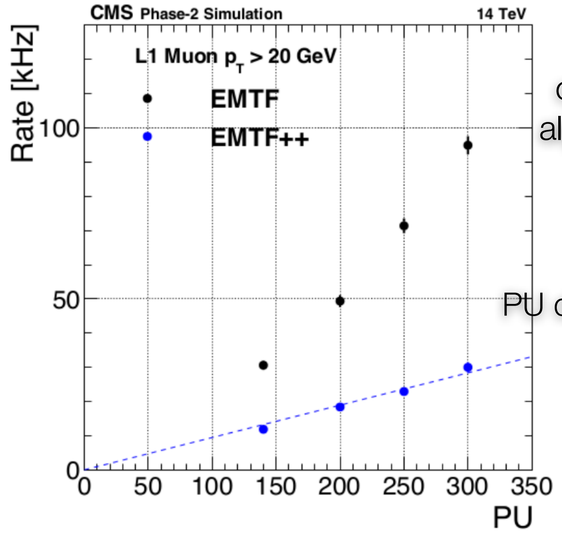

current algorithm

new algorithm @HL-LHC PU dependance reduction

Muon trigger algorithms in a nutshell

GMT Tracker Track matching

- Everything in one box!
  - Input all Level-1 tracks from outer tracker TMUX=18
  - Input all reconstructed muons, and muon stubs, from all muon detectors
- Propagate tracks to match muons/stubs in eta/ p_T windows at different muon detector stations
  - Magnetic field / multiple scattering / detector resolutions
- Create track-matched muons with different qualities
  - Many choices on this
  - Can fill in the gaps from muon track finders
- Use the pT of the tracker track instead of muon
  - Much better resolution, sharp turn on, reduced rates

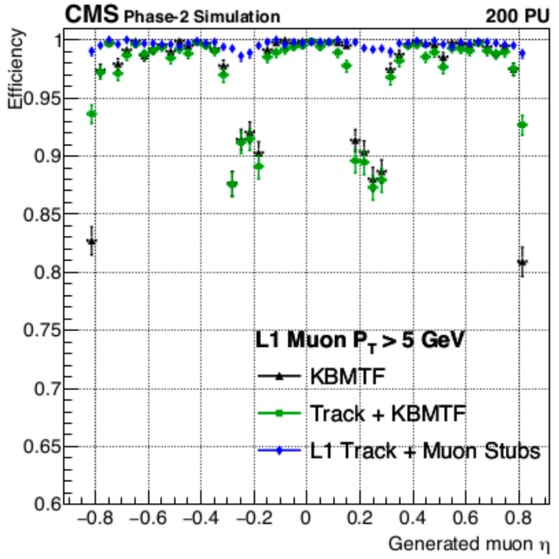

Efficiency vs rapidity

Track matching to stubs fills in the gaps

Rate vs pT threshold

Track matching reduces rate significantly

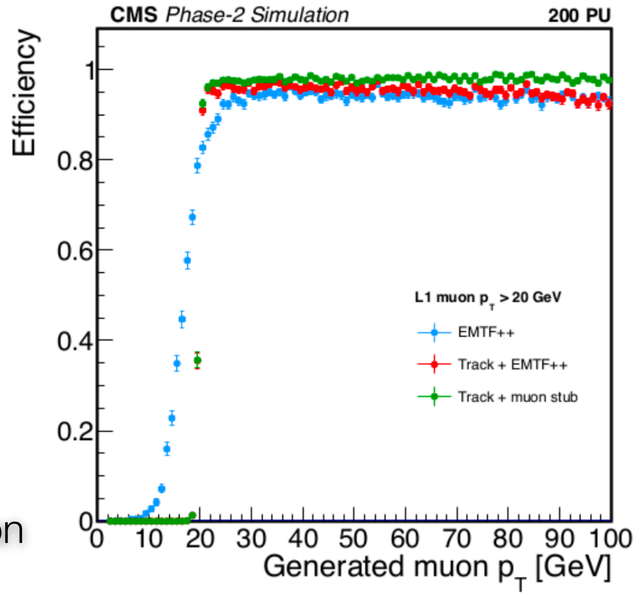

Efficiency vs pT

Track matching sharpens the turn-on

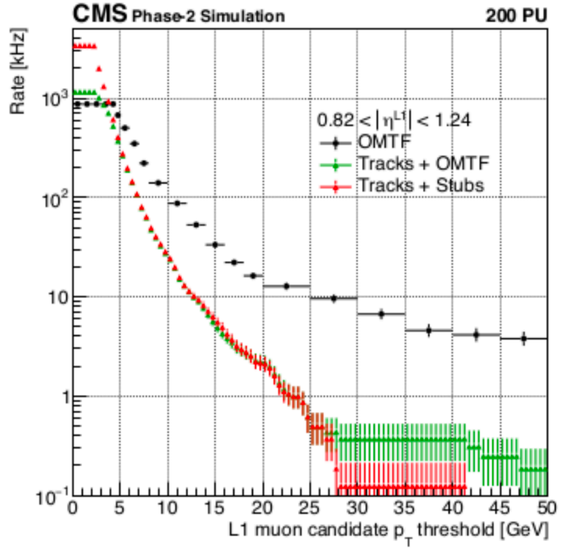

Muon trigger algorithms in a nutshell

- GMT topological algorithms
  - Exploit that info from all μ detectors in one place
  - Exploit that all tracker tracks available as well
  - Can see if a muon is isolated from tracks
    - Relative or absolute isolation @ different pt thresholds
  - Can gather several muons into a single object such as a BSM tau=>3μ decay
    - Find suitable triplets, reconstruct invariant mass

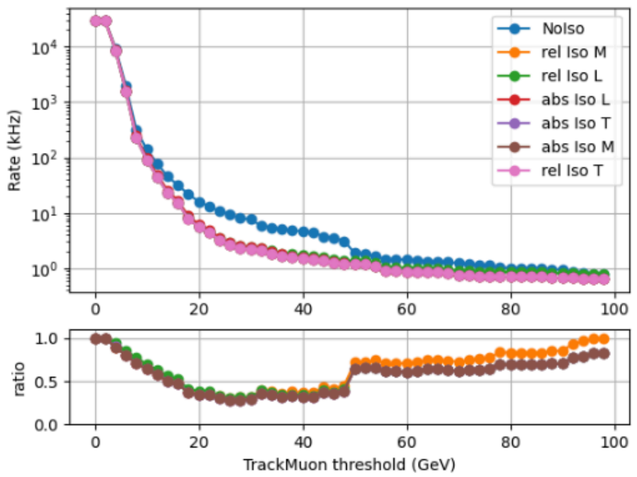

Rate vs pT threshold
Isolated muons vs all muons

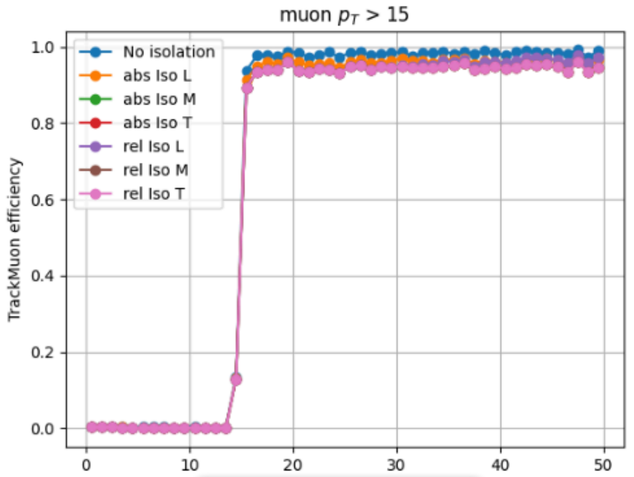

Efficiency vs pT
Isolated muons

28

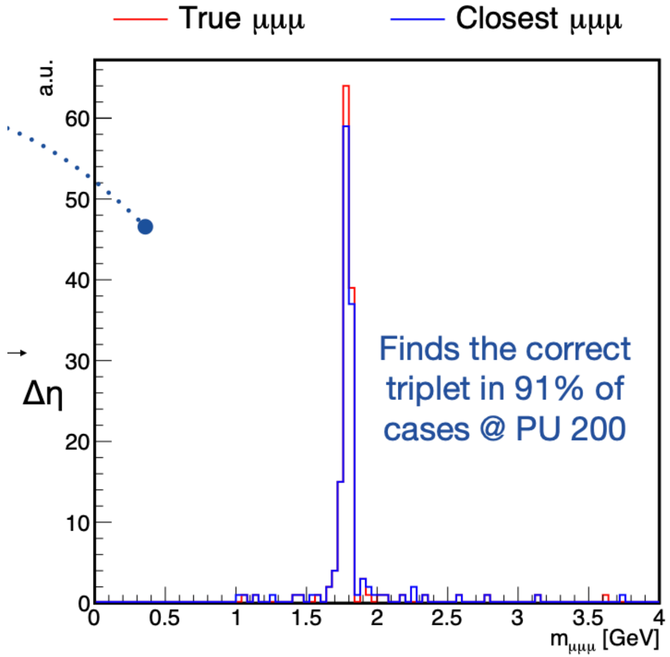

Finds the correct triplet in 91% of cases @ PU 200

Reconstructed $\mu\mu\mu$ mass as a cluster of 3 closest TkMu.

Muon Trigger firmware overview

- Advanced firmware development on all fronts
  - A corresponding software emulator to compare with
- All development on VU13P FPGA

Example: Firmware pipeline for EMTF algorithms

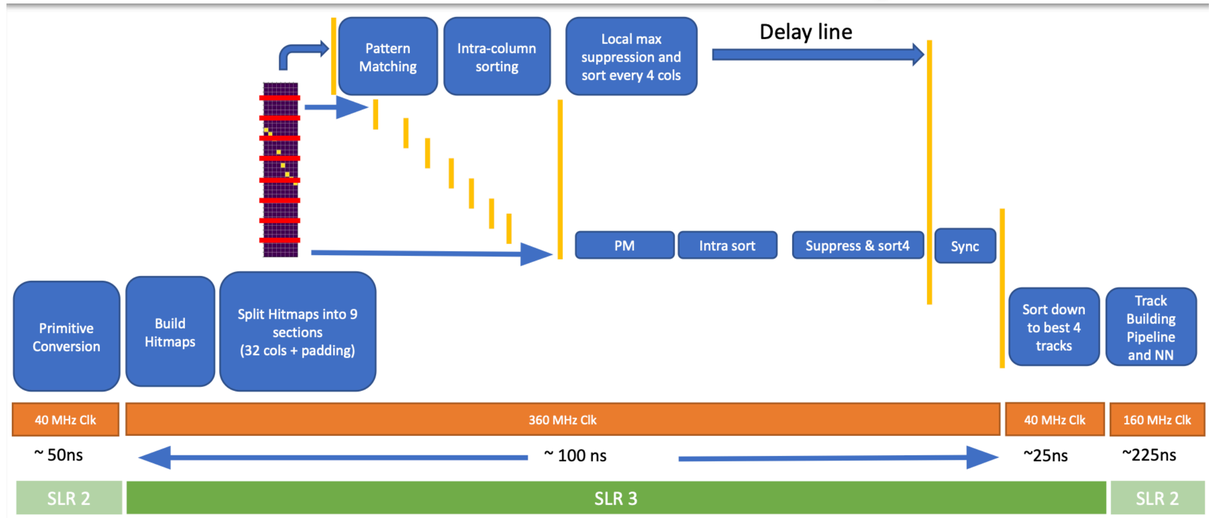

Split tasks by SLR (Super Logic regions). Total ~400 ns latency !

Endcap region firmware

Site Type	SLR0	SLR1	SLR2	SLR3	SLR0 %	SLR1 %	SLR2 %	SLR3 %
CLB	0	0	22873	16188	0.00	0.00	42.36	29.98
CLBL	0	0	11871	8677	0.00	0.00	40.54	29.63
CLBM	0	0	11002	7511	0.00	0.00	44.51	30.38
CLB LUTs	0	0	117166	82652	0.00	0.00	27.12	19.13
LUT as Logic	0	0	107691	81048	0.00	0.00	24.93	18.76
LUT as Memory	0	0	9475	1604	0.00	0.00	4.79	0.81
LUT as Distributed RAM	0	0	1261	0	0.00	0.00	0.64	0.00
LUT as Shift Register	0	0	8214	1604	0.00	0.00	4.15	0.81
CLB Registers	0	0	90825	89984	0.00	0.00	10.51	10.41
CARRY8	0	0	10755	0	0.00	0.00	19.92	0.00
F7 Muxes	0	0	328	12592	0.00	0.00	0.15	5.83
F8 Muxes	0	0	33	5950	0.00	0.04	0.03	5.51
F9 Muxes	0	0	0	0	0.00	0.00	0.00	0.00
Block RAM Tile	0	0	316	0	0.00	0.00	47.02	0.00
RAMB36/FIFO	0	0	299	0	0.00	0.00	44.49	0.00
RAMB18	0	0	34	0	0.00	0.00	2.53	0.00
URAM	0	0	0	0	0.00	0.00	0.00	0.00
DSPs	0	0	1454	0	0.00	0.00	47.33	0.00
Unique Control Sets	0	0	605	1005	0.00	0.00	0.56	0.93

Resource usage is **less than 25% of chip**
- Less than 12% of LUTs, DSPs, and BRAM
- Missing: Hybrid Stub Logic and Algorithm wrapper

Main algorithm is split across 2 SLRs, both less than 50% full
- SLR3 - Pattern matching and sorting roads
- SLR2 - Primitive Conversion, Track Building, and NNs

Muon Trigger firmware overview

OMTF

Overlap region firmware

- Also moving towards a Neural Network, as Endocarp

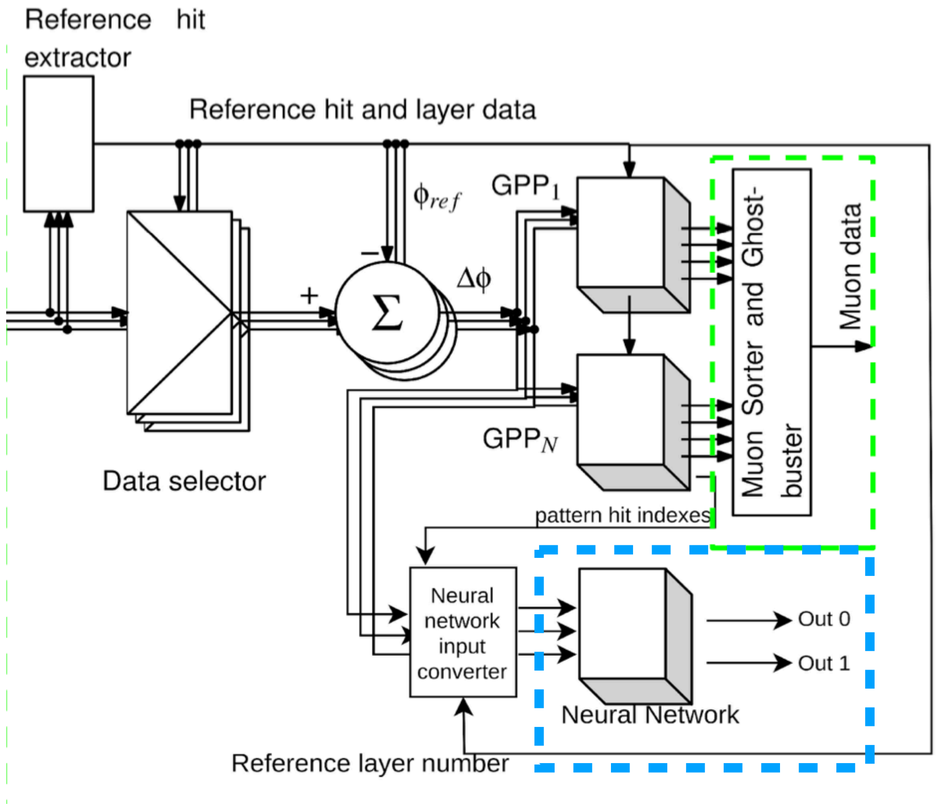

Resources

Module	LUT	LUTRAM	FF	BRAM	DSP
Neural network	164690 (9.53%)	47 (<0.01%)	19015 (0.55%)	0 (0%)	441 (3.59%)
OMTF	184718 (10.69%)	18127 (2.29%)	108223 (3.13%)	252 (9.38%)	0 (0%)
OMTF+NN	**349408 (20.22%)**	**18174 (2.30%)**	**127238 (3.68%)**	**252 (9.38%)**	**441 (3.59%)**

- % of usage w.r.t XCVU13P-FSGA-2577

Timing

Module	CLK	LATENCY	ii
Neural network	160 MHz	16 cycles	1
OMTF+NN	**160 MHz**	**45 cycles**	**1**

Modest resource usage and small latency of ~280 ns

Muon Trigger firmware overview

GMT firmware

- Different SLRs for different tasks
- Also fitting well on the FPGA

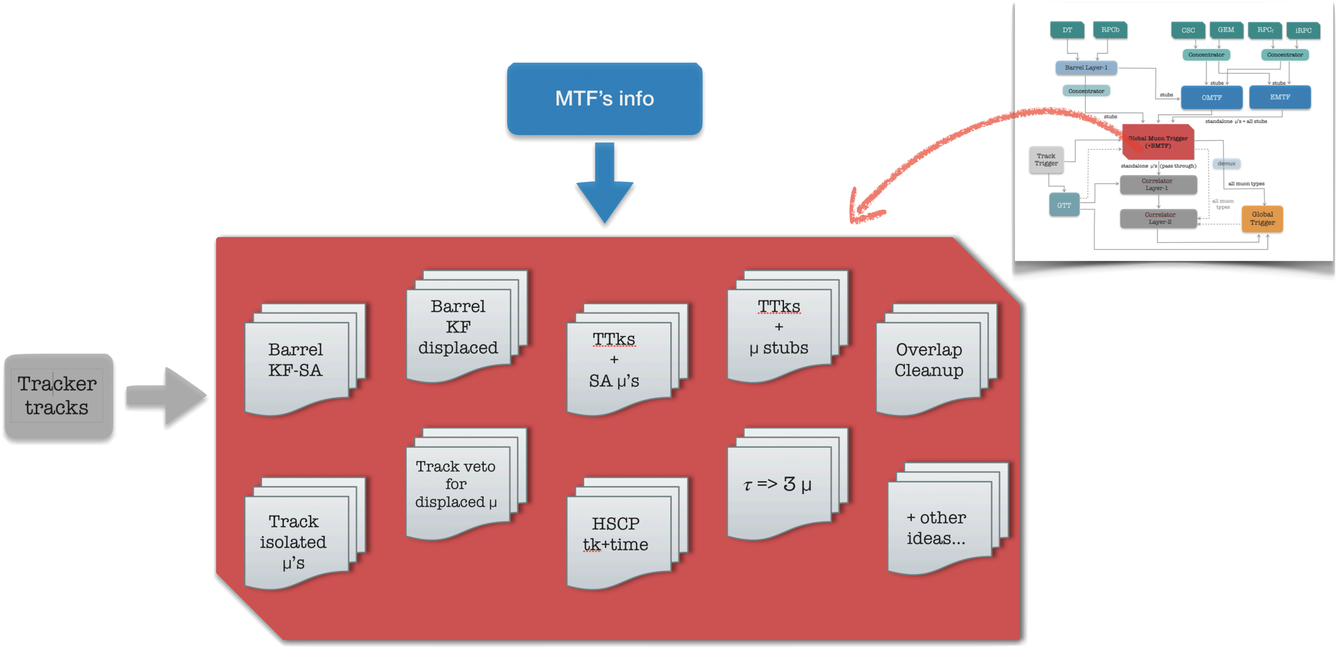

Global Muon Trigger "GMT"

Global Muon Trigger firmware

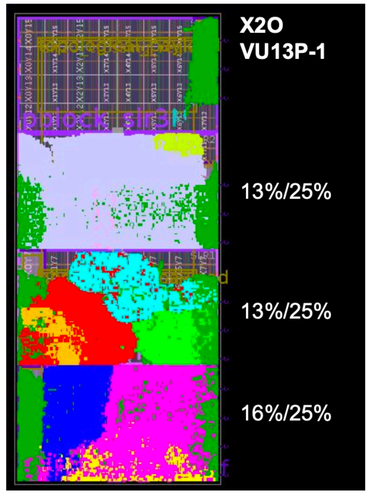

- framework+links
- KMTF regionizer
- KMTF
- KMTF ghost cleaner
- EMTF ghost cleaner
- OMTF ghost cleaner
- STA muon sorter
- Tracker Muon reco
- Muon Sorter

Integration slice tests in preparation

- At CERN b904
- Board-to-board integration
- With other trigger systems

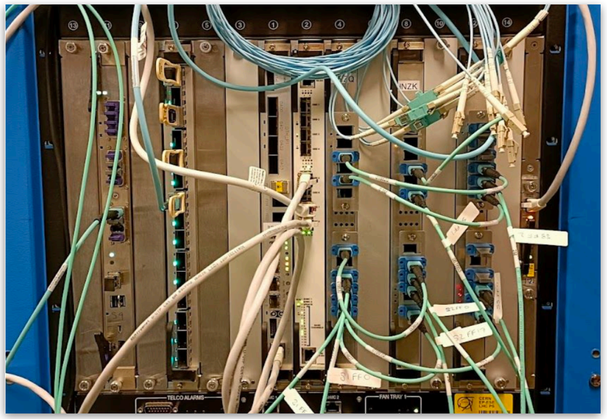

Summary

- CMS has developed an advanced Muon Trigger @L1 for the HL-LHC era
  - Will preserve/expand the CMS physics program
  - Fast, modular, flexible
- Much progress towards implementation
  - Algorithms, emulators, firmware, hardware
  - No show stoppers foreseen
- Integration slice tests picking up
- Expect to produce all boards next year
- Installation of the final system…to follow
- Then commission, take data and… discover something !

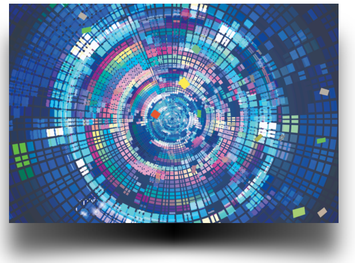

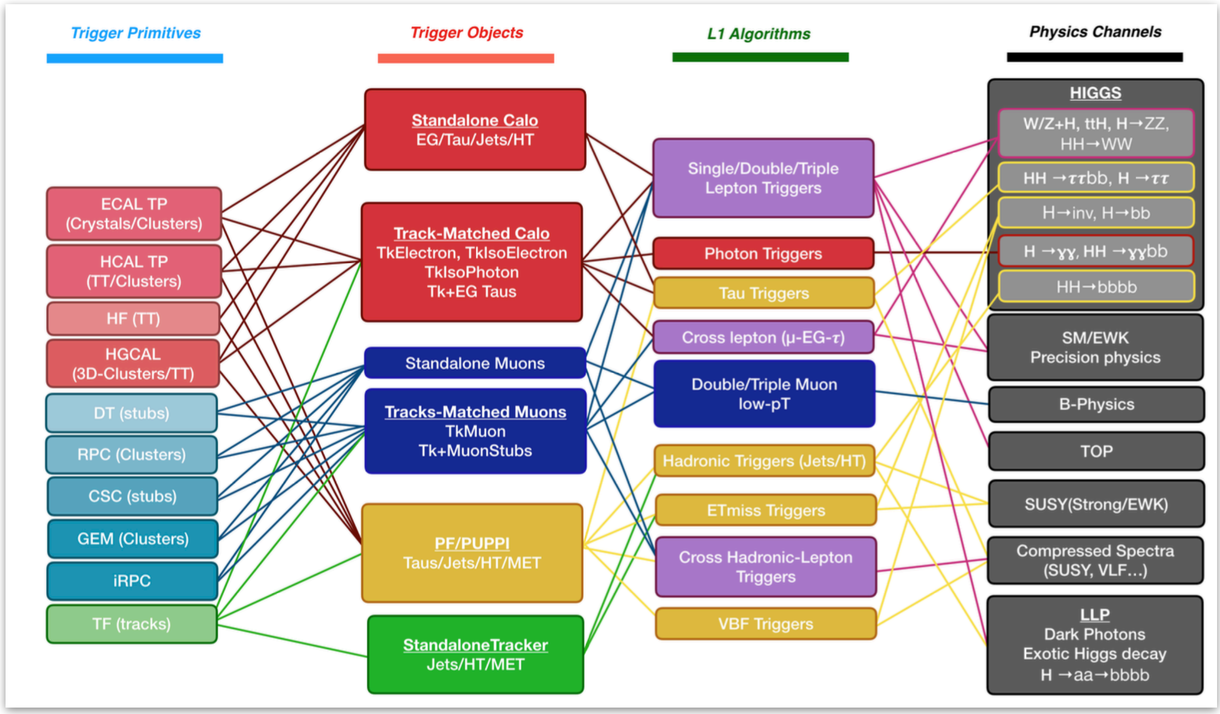

Backup etc

L1 Data Scouting: overview

- What: **acquire & analyze the L1 information for all events** (40 MHz)
- Why: look for physics signatures identifiable with just L1 information but that would evade the L1 → HLT → Offline chain, e.g.:
  - Too large “irreducible” backgrounds, e.g. narrow resonances of unknow mass, as already done using HLT scouting for muons & hadronic resonances.
  - Signal identification requires an algorithm that can’t fit the L1 fixed latency and resource budget, e.g. has too complex combinatorics on some events.
  - Signal from time-correlation across several BXs, e.g. slow or long-lived BSM.
- How: FPGA-equipped boards that receive L1 data via optical links and transfer it to PCs and the software world via TCP/IP or PCIexpress
- When: HL-LHC, but is being demonstrated already in 2018 (muons), Run 3 (muons + calo) limited by L1 reconstruction quality