SYSTEM DESIGN AND PROTOTYPING FOR THE CMS LEVEL-1 TRIGGER AT THE HL-LHC

TIPP2021, CANADA, 26TH MAY 2021

A. ZABI, ON BEHALF OF THE CMS COLLABORATION

OUTLINE

- Context of triggering @ HL-LHC: scientific case and system requirements (technological choices)
- Phase-2 L1 trigger conceptual design and instrumentation: System interfaces & Architecture. Key features and hardware prototyping.
- Phase-2 Level-1 trigger algorithm design, firmware developments & testing: selecting physics with sophisticated firmware algorithms and expected performance. System demonstration.

TRIGGERING @ HL-LHC

INTRODUCTION & DESIGN REQUIREMENTS

The scout of the HL-LHC

INTRODUCTION: CONTEXT OF TRIGGERING

- **HL-LHC Upgrade Project: offers an unprecedented opportunity to explore uncharted lands and achieve scientific progress.**
- **A new LHC machine and a new CMS Detector:**

The HL-LHC and the CMS Phase-2 detector

→ ***Set the context of triggering & define system requirements***

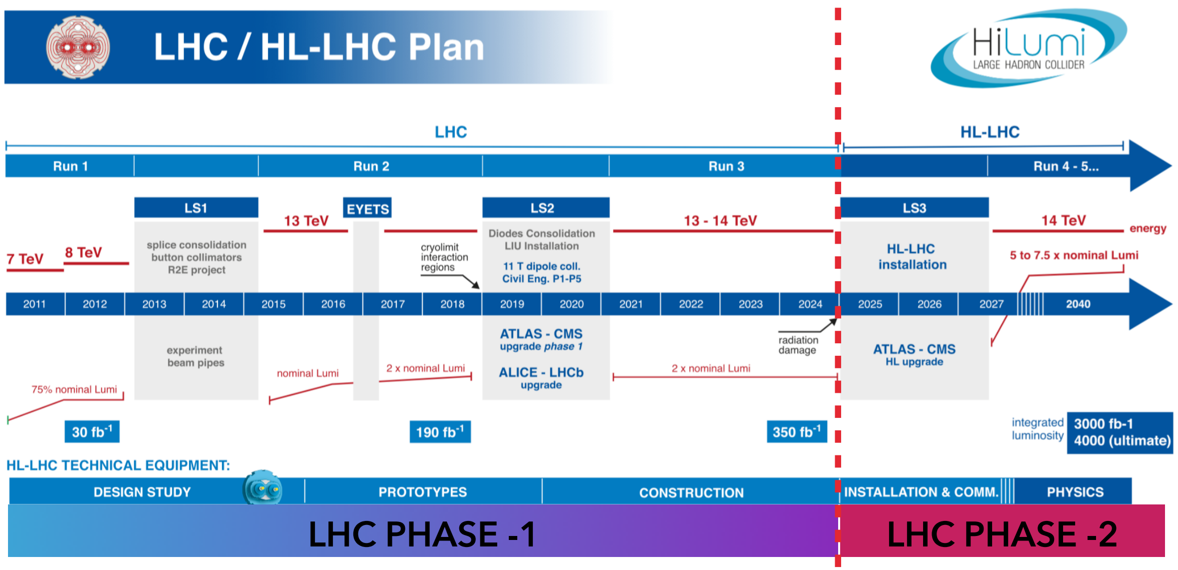

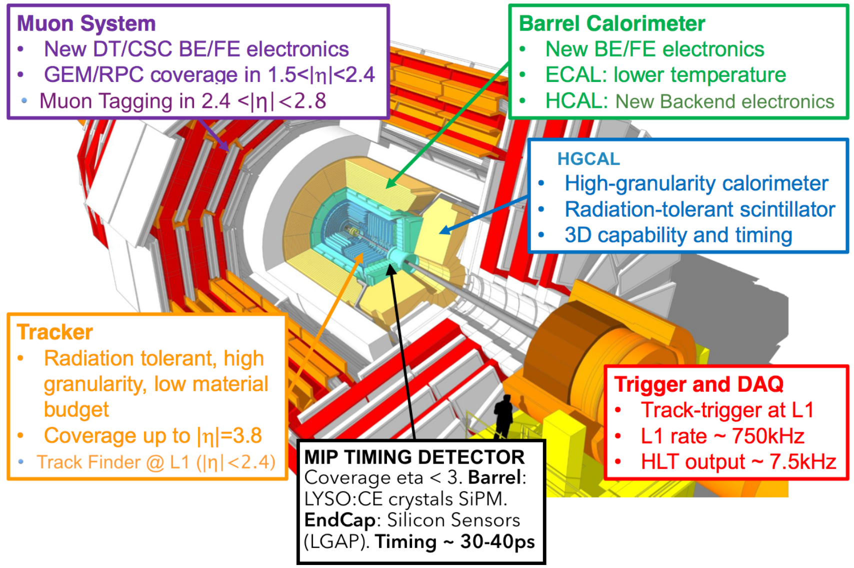

High-Luminosity-LHC: 13 TeV (Nominal : 5×10^{34} & 140 PU, Int Lumi = 3000 fb^{-1})

- **Ultimate: 7.5×10^{34} & 200 PU, Int Lumi = 4000 fb^{-1} (baseline for all TDR studies)**

→ ***unprecedented running conditions, exceeding machine design values 7 fold.***

PHASE-2 TRIGGER UPGRADE: KEY PARAMETERS & STRATEGY

- CMS Phase-2 Trigger:
  - CMS keeps a 2-level triggering approach: **L1 & HLT**
  - **Level-1 (hardware) system**
    - Increase bandwidth 100 kHz → 750 kHz
    - Increase latency 3.8 us → 12.5 us
- Benefiting from upgrade of the CMS detector:
  - Include high-granularity information (calo&μ)
  - Include tracking information (first time!)

  → Manageable object rate (L1 Physics Menu)
- Strategy:
  - Exploit sub-detector back-end electronics
  - Sophisticated reconstructed objects and correlations → Enhanced physics selectivity
  - Expand reach with Scouting System

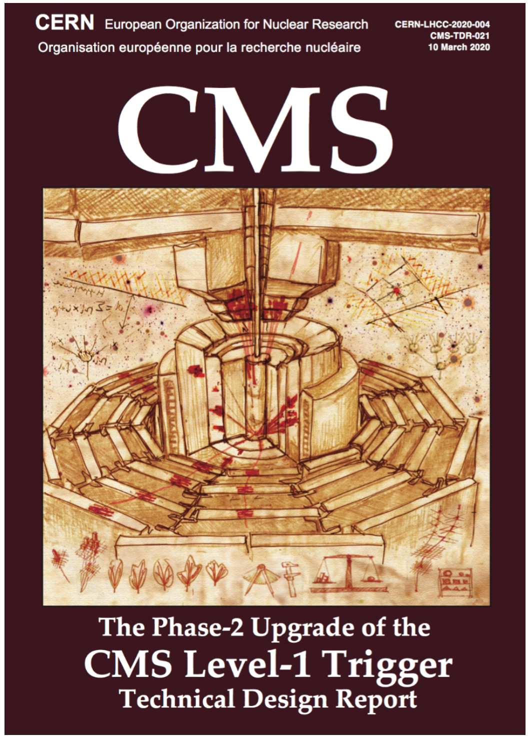

TDR approved in 2020

L1 PHASE II TRIGGER UPGRADE: SCIENTIFIC CASE

Maintaining thresholds is <u>NOT</u> the only motivation for upgrading the L1 trigger. HL-LHC research program opens a door to the unknown → <u>the Phase-2 Level-1 Trigger system is our scout !</u>
The goal is to extend the physics reach by increasing the available phase space

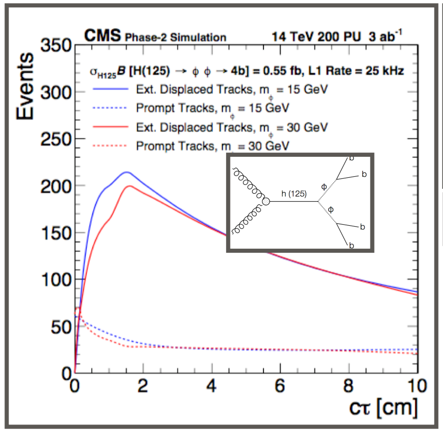

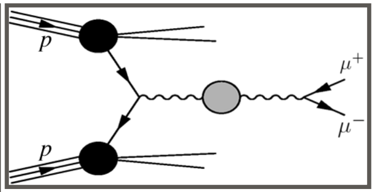

DarkPhoton(displaced mu)

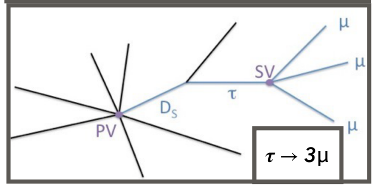

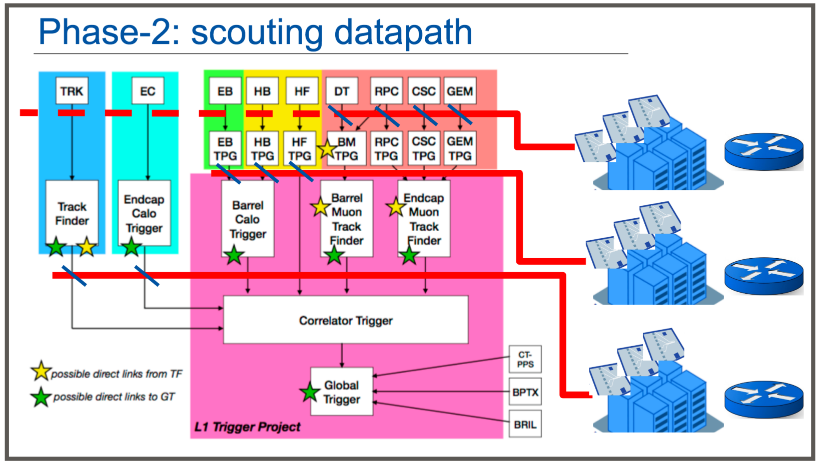

<u>L1 Trigger algorithm requirements</u>:

- Object reconstruction closer to offline performance: higher-level trigger objects **(particle-flow)** w/ optimised response and resilient to pileup (up to 200)
- Sophisticated triggers to select specific topologies: VBF production, rare B-meson decays **(tracking@L1)**, forward muon trigger for τ → μμμ **(muon extended coverage)**, dedicated algos for displaced jets and muons, etc.
- Expand reach: Low mass resonances

<u>Scouting into HL-LHC data @ 40 MHz</u>:

- Physics objects: reconstructed from L1 objects
- Storage: Only high-level information (selected events)
- Specific features: analyse multiple contiguous BX, identify signatures unreachable through standard trigger techniques

L1 PHASE II TRIGGER UPGRADE: TECHNOLOGICAL CASE

The Phase-2 Level-1 Trigger system performs precise physics selection using a global event reconstruction based on enhanced granularity already at hardware level.
Considering: Inputs > 60 Tb/s (vs 2 Tb/s during LHC Phase-1), operate: 7.5e34 along w/ 200 PU events

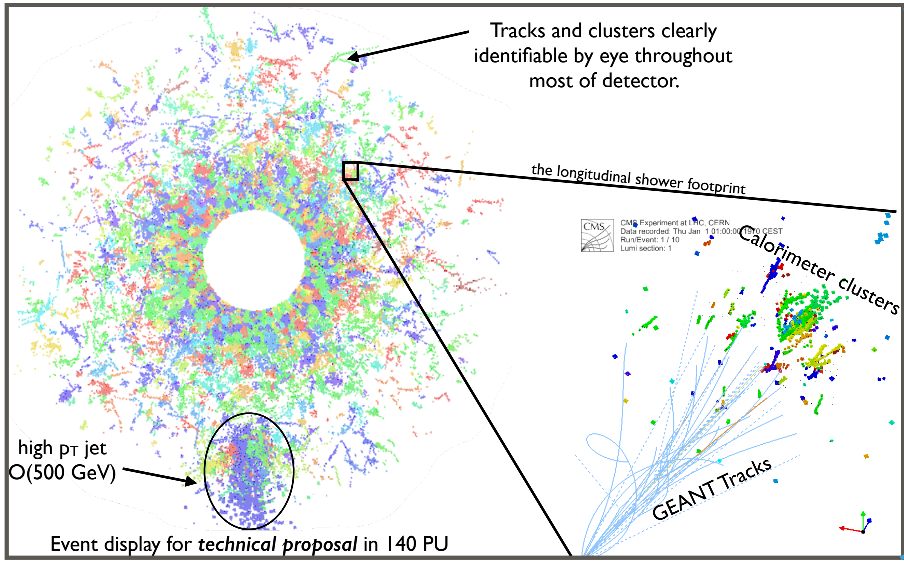

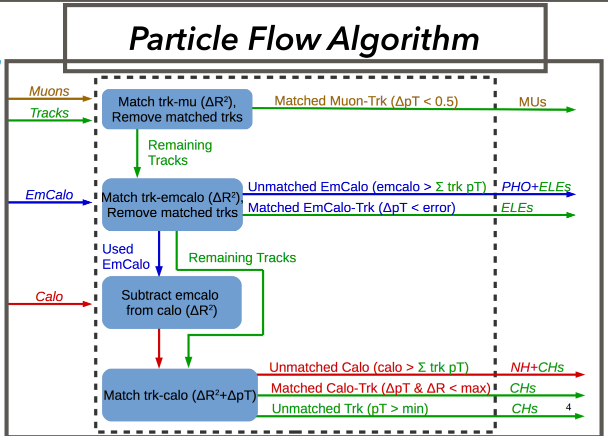

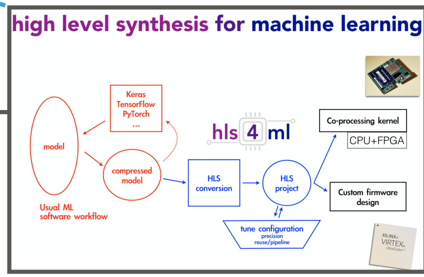

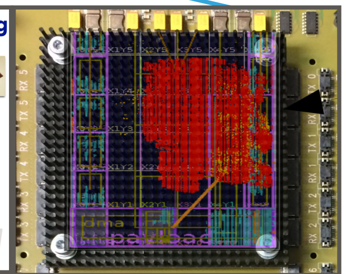

L1 Trigger requirements:

- Cutting-edge hardware: modern technology
  → FPGA VU9P x 8 resources of Virtex 7 (Phase-1), 28 Gb/s links
- High-Level-Synthesis: used successfully, much faster turn-around, novel techniques based on machine learning → *The Phase-2 L1 Trigger can do much more!*
- Advanced Architecture: platform and interconnections (ATCA) → robust, flexible & modular design
- Handling all technical issues: integration, commissioning, etc.

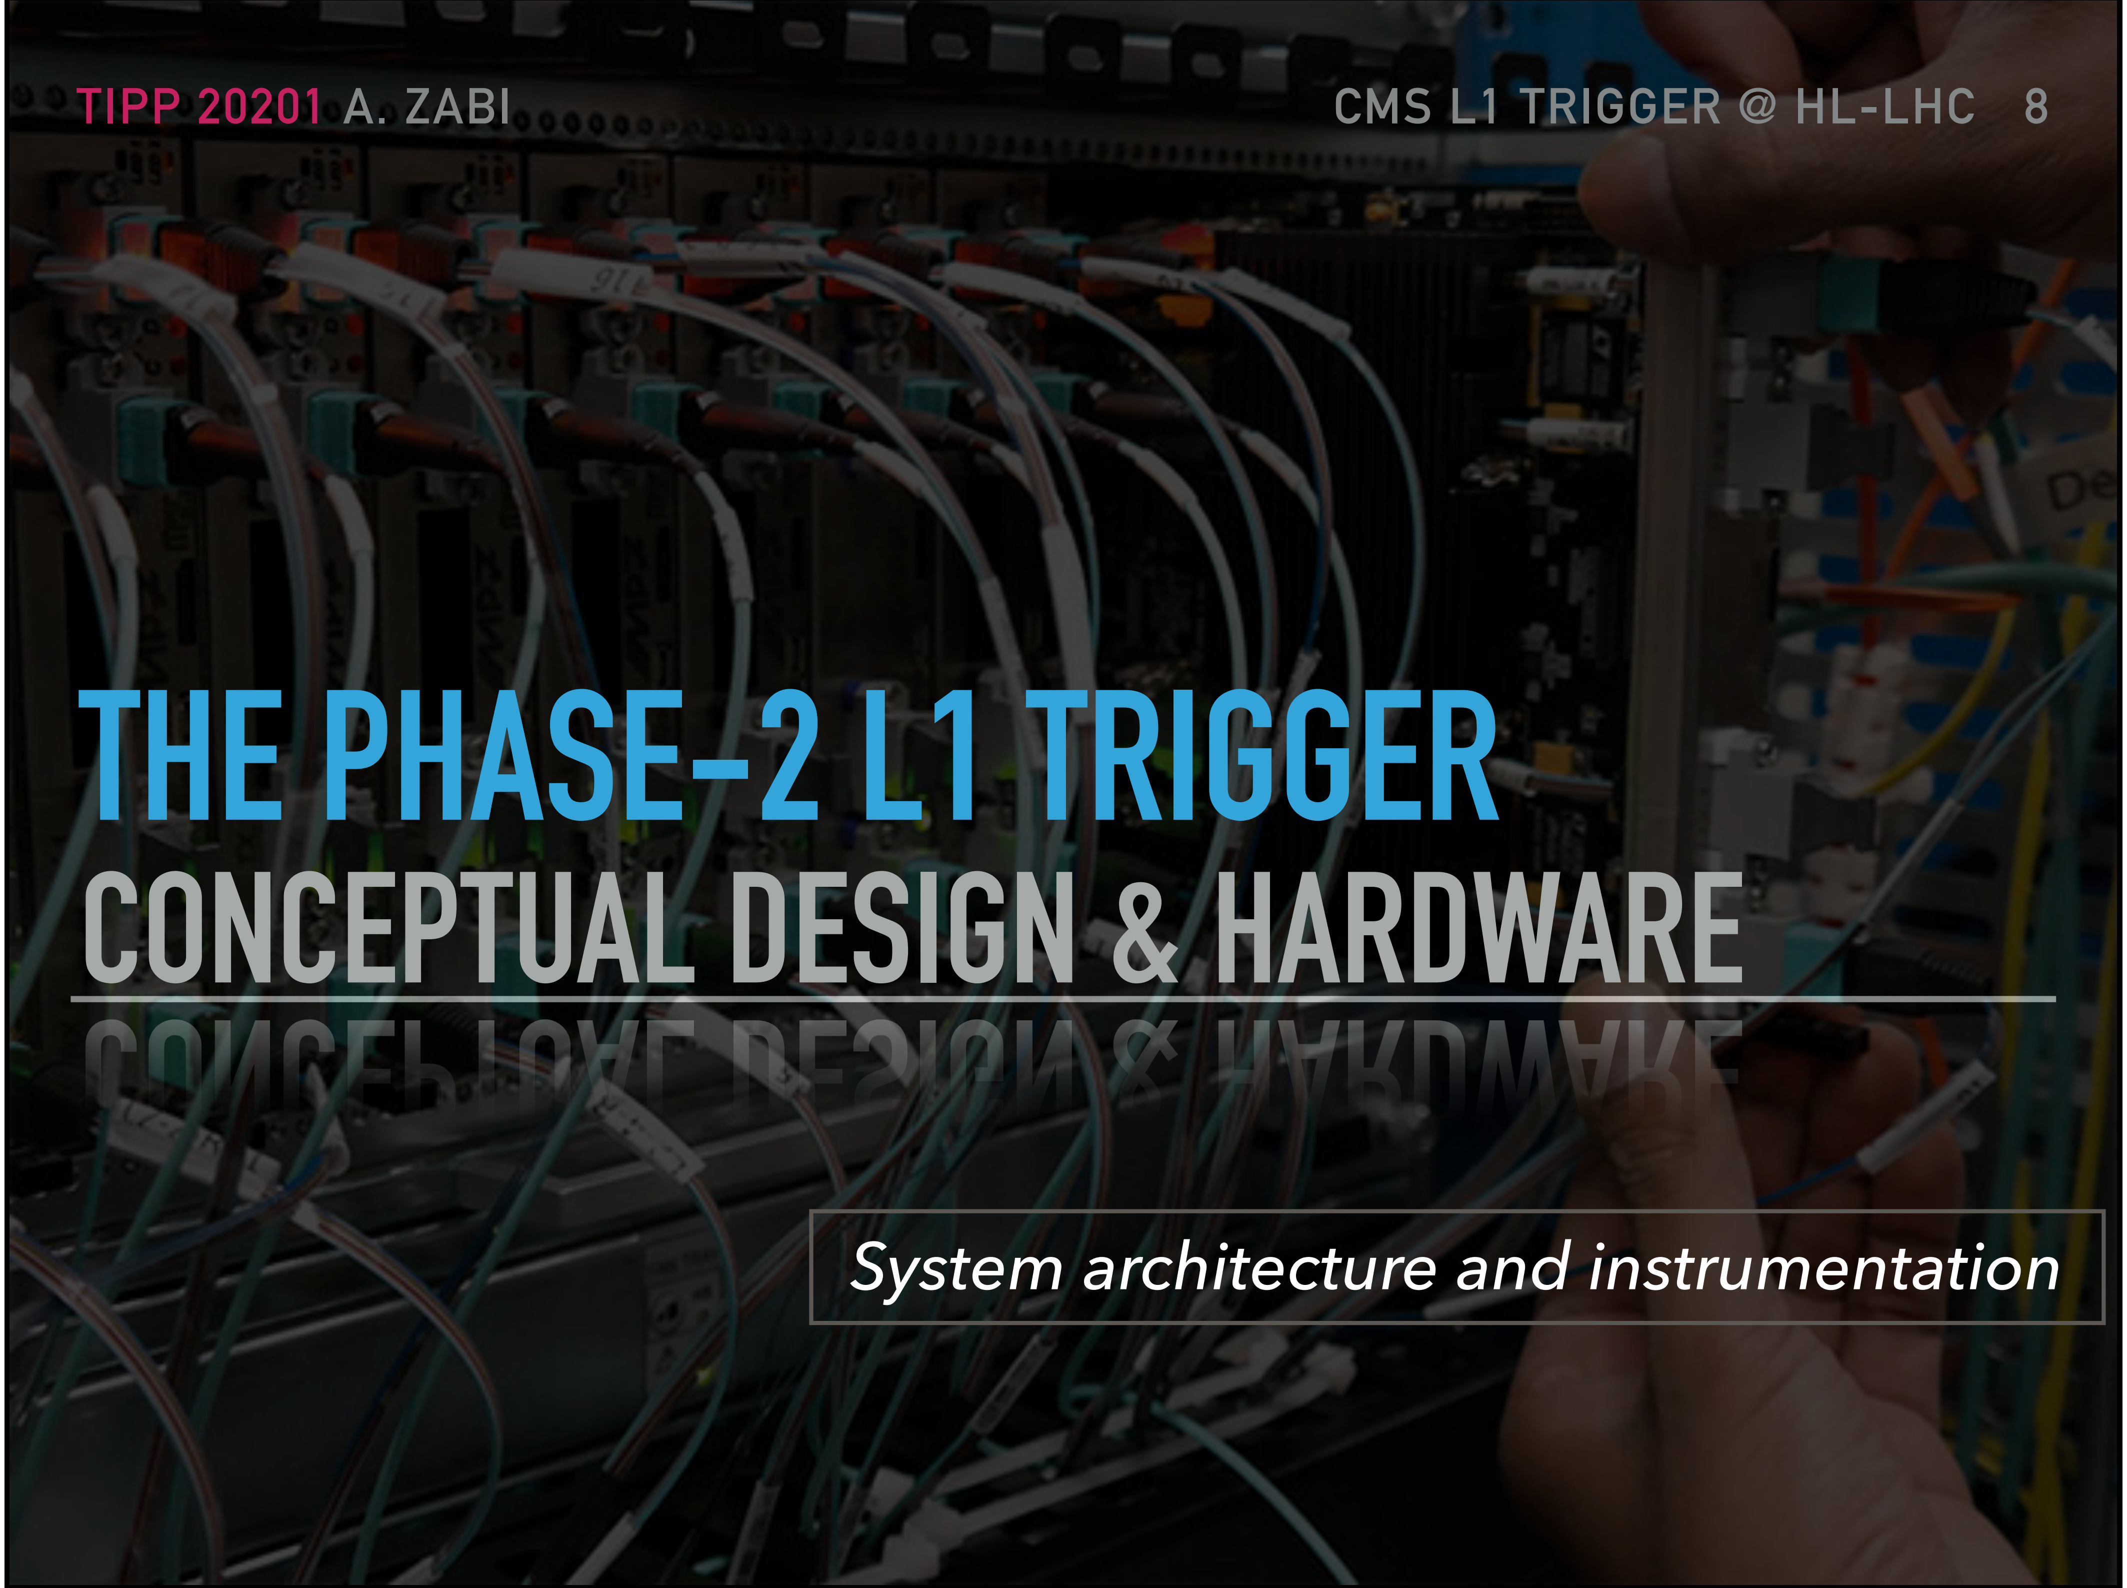

THE PHASE-2 L1 TRIGGER

CONCEPTUAL DESIGN & HARDWARE

System architecture and instrumentation

LEVEL-1 PHASE II TRIGGER UPGRADE SYSTEM

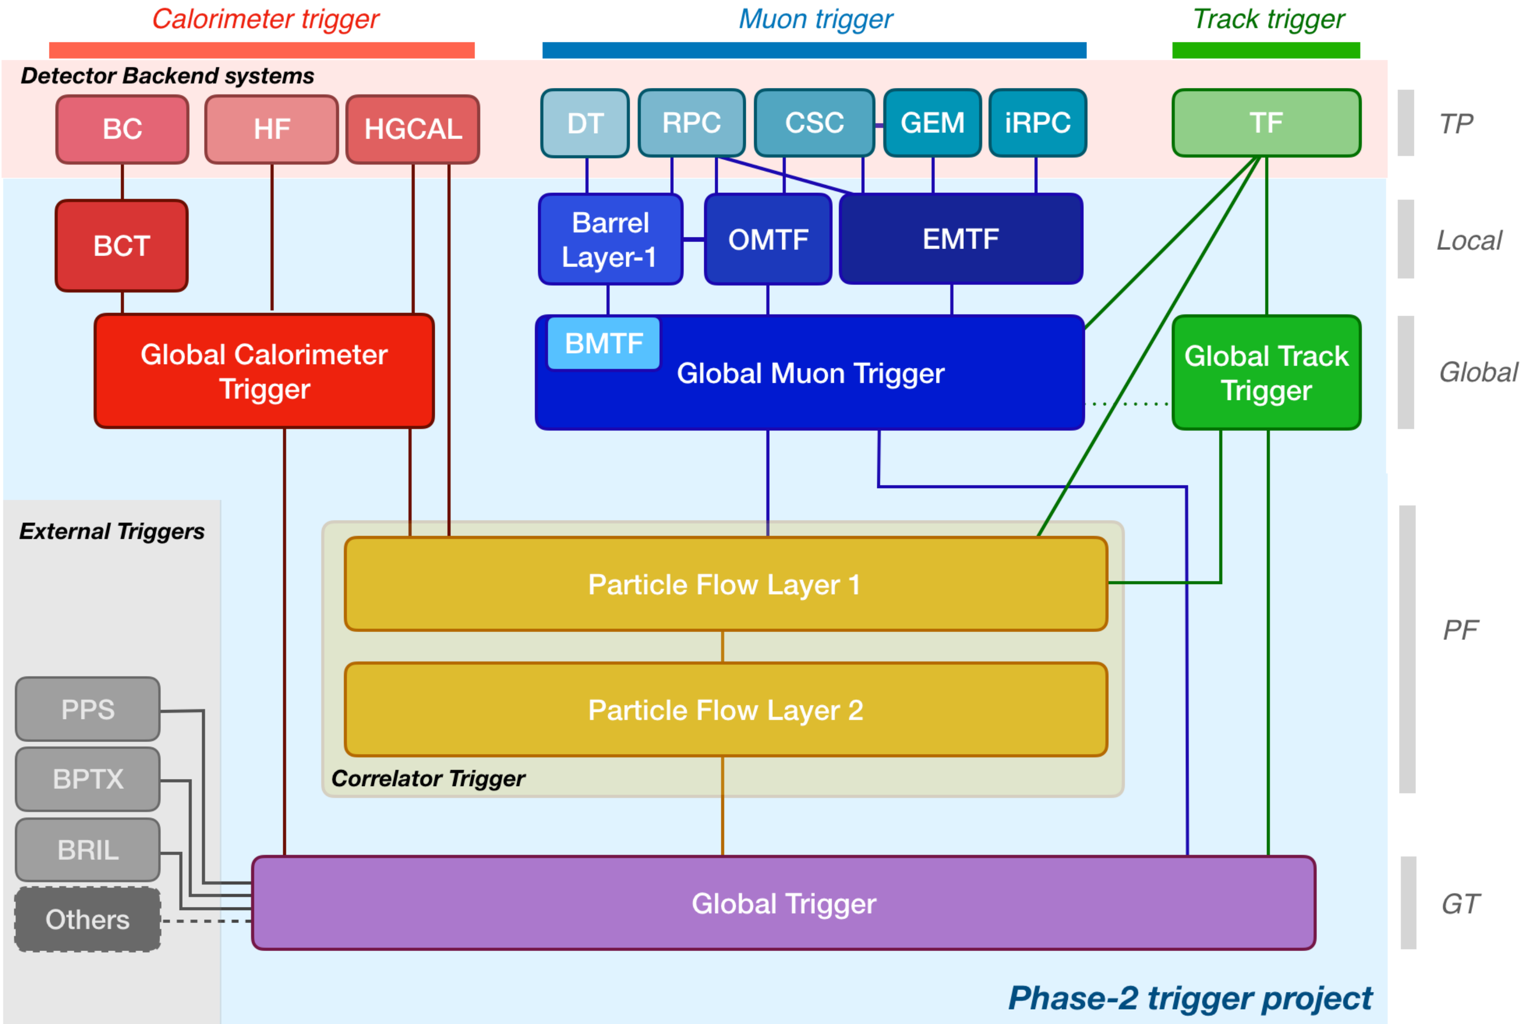

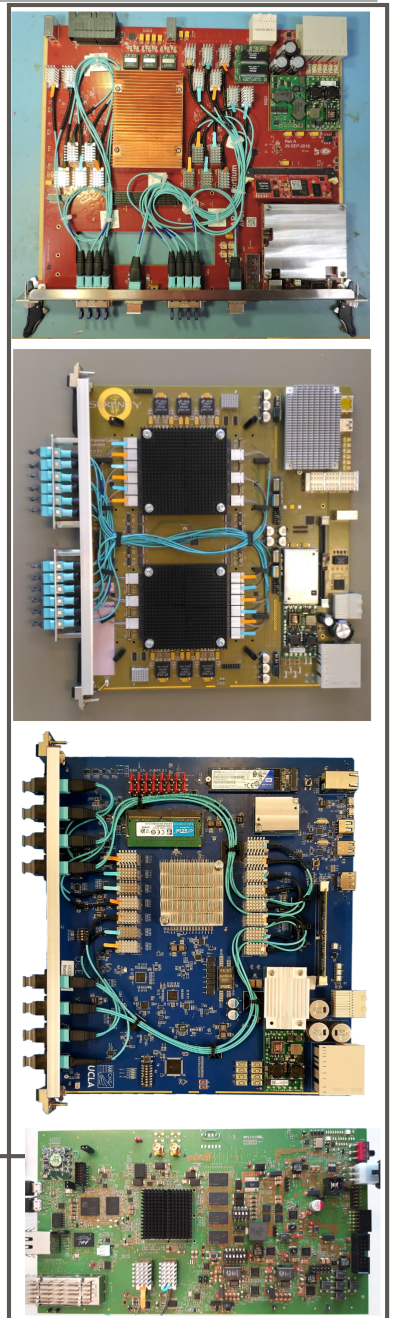

- Level-1 Architecture: Efficient distribution and processing of trigger primitives, provision appropriate resources and interconnections, retain enough headroom future flexibility & Robustness
- Level-1 technological choices: generic processing engines (inspired from Phase-1 upgrade)
- Key design feature: Correlator Trigger. Collects all inputs and feed sophisticated algorithms
- Design Constraints : HW processors ~ 100 links , FPGA resources < 50 %, Latency (< 9.5 us (keep 20%) while HGCAL/TF~5us)

HARDWARE PROTOTYPES

Design philosophy: Generic Processing Engines → I/O, FPGA

- FPGA : Xilinx Virtex Ultrascale / Ultrascale+
- Optics : Samtec Firefly x4 flyover
  - 32 modules for 120 RX + 120 TX links
  - A single optical module has 4 RX + 4 TX
- Processors on board running commercial linux for flexible configuration and monitoring

25 Gb/s

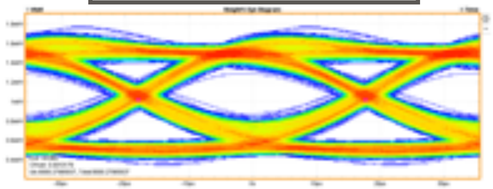

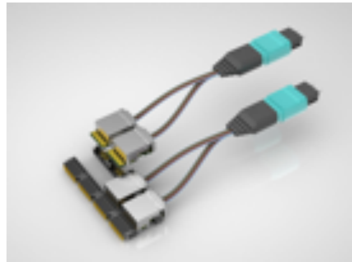

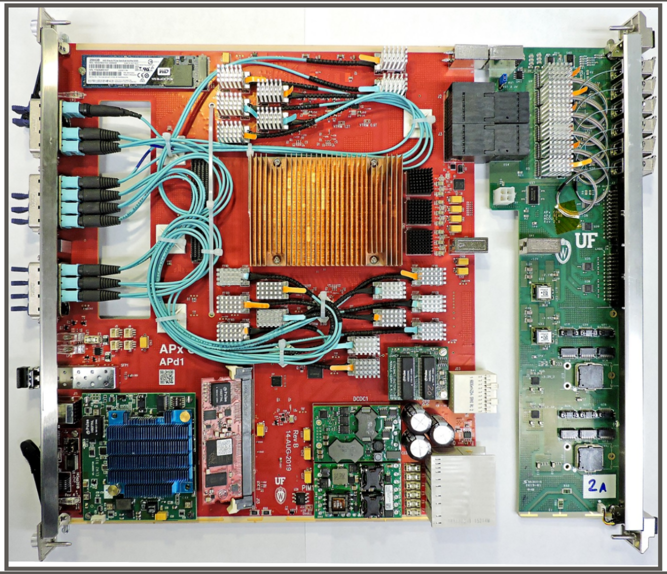

Serenity:

- Carrier board w/ 2 sites hosting daughter cards (any combination of FPGA)
- Up to 144 bidirectional links (extendable to 192)
- Control & Monitoring: COM express (x86 processor)
- IPMI management through CERN IPMC

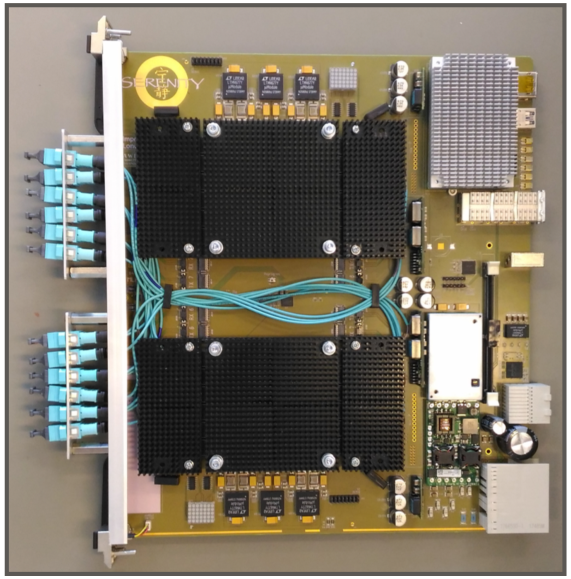

APX:

- Powered by a VU9P FPGA with 2.5M logic cells
- 100 bidirectional links up to 28 Gb/s
- Control, management, and monitoring by an embedded linux mezzanine (ELM) (ZYNQ SoC)
- Shelf management via custom IPMI mezzanine (OS)

HARDWARE CAN DO MORE: EVOLUTION

Design evolution (since TDR): increased I/O and computing power → sophisticated algo, flexibility

- FPGA : larger A2577 pin package FPGA from Xilinx (Virtex UltraScale 13P)
- Optics : New denser version of on-board flyover Samtec Firefly, alternatives (ex: QSFP)

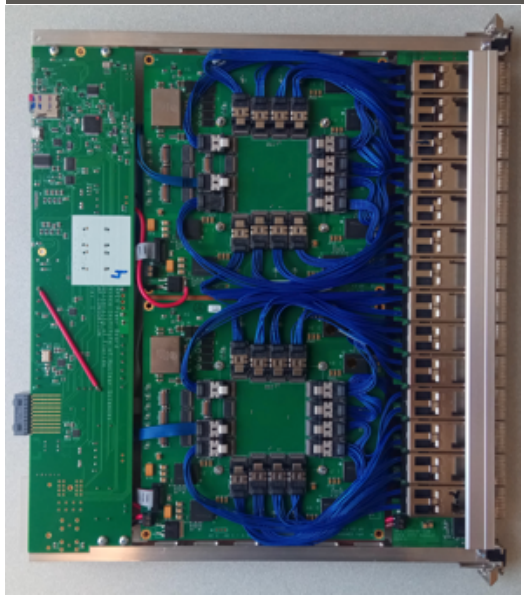

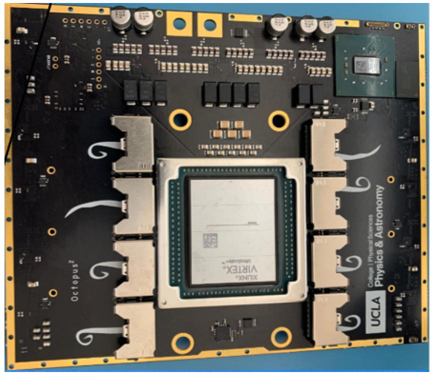

Variant: Intensive processing in muon trigger. Octopus Mezzanine (A2577-VU13P)

<u>Samtec Firefly x12</u>

- One module: x12 RX or x12 TX
- Note Module in alpha-stage
- Ongoing test:

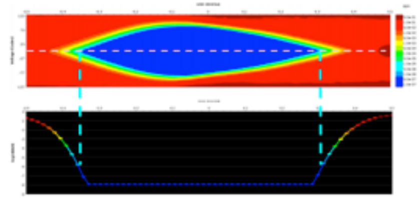

X2O: Evolution from OCEAN Prototype

- Modular design (x2 FPGA)
- Optical Module
  - Up to 28 QSFP cages (112 links)
  - Compatible with 25G and 10G transceivers
- Power Module: Off-the-shelf ZYNQ mezzanine, DC-DC converters, IPMC running on the ZYNQ
- Inter-module connections with cables

<u>Alternatives: QSFP</u>

- widely used in industry
- x4 TX / x4 RX (x8 TX / x8 RX QSFP DD)
- Under qualification (BER etc)

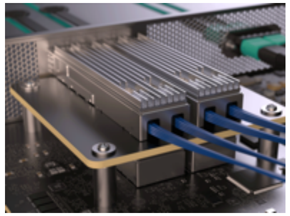

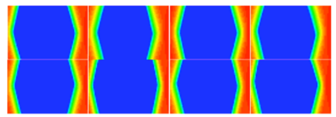

QSFP on X2O

SCOUTING @ 40MHZ: SCRUTINISING THE DATA

- Enables many features: real-time diagnostics (even at lower level systems), monitoring, testing new algorithms and developing menus, selecting an reconstructing physics objects w/o rate limitation.
- Analyses conducted trough queries (from storage)
- Demonstrated during LHC Run-2 with Level-1 Phase-1 muon output, now being prepared for Run-3 data taking
- Uses DTH board (DAQ800) designed for large readout detectors

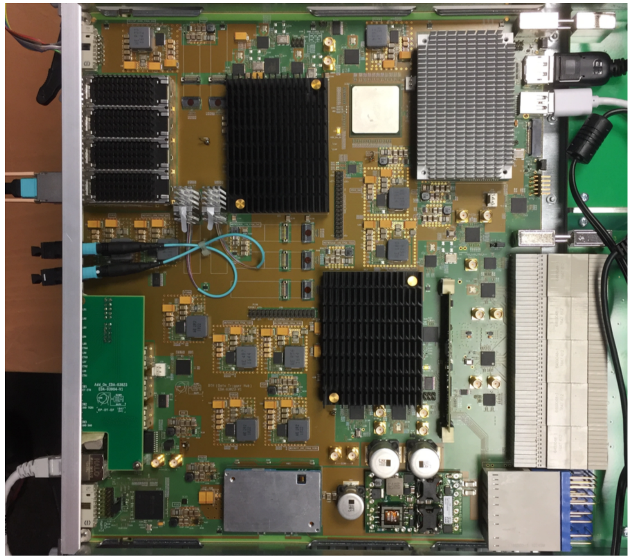

DTH

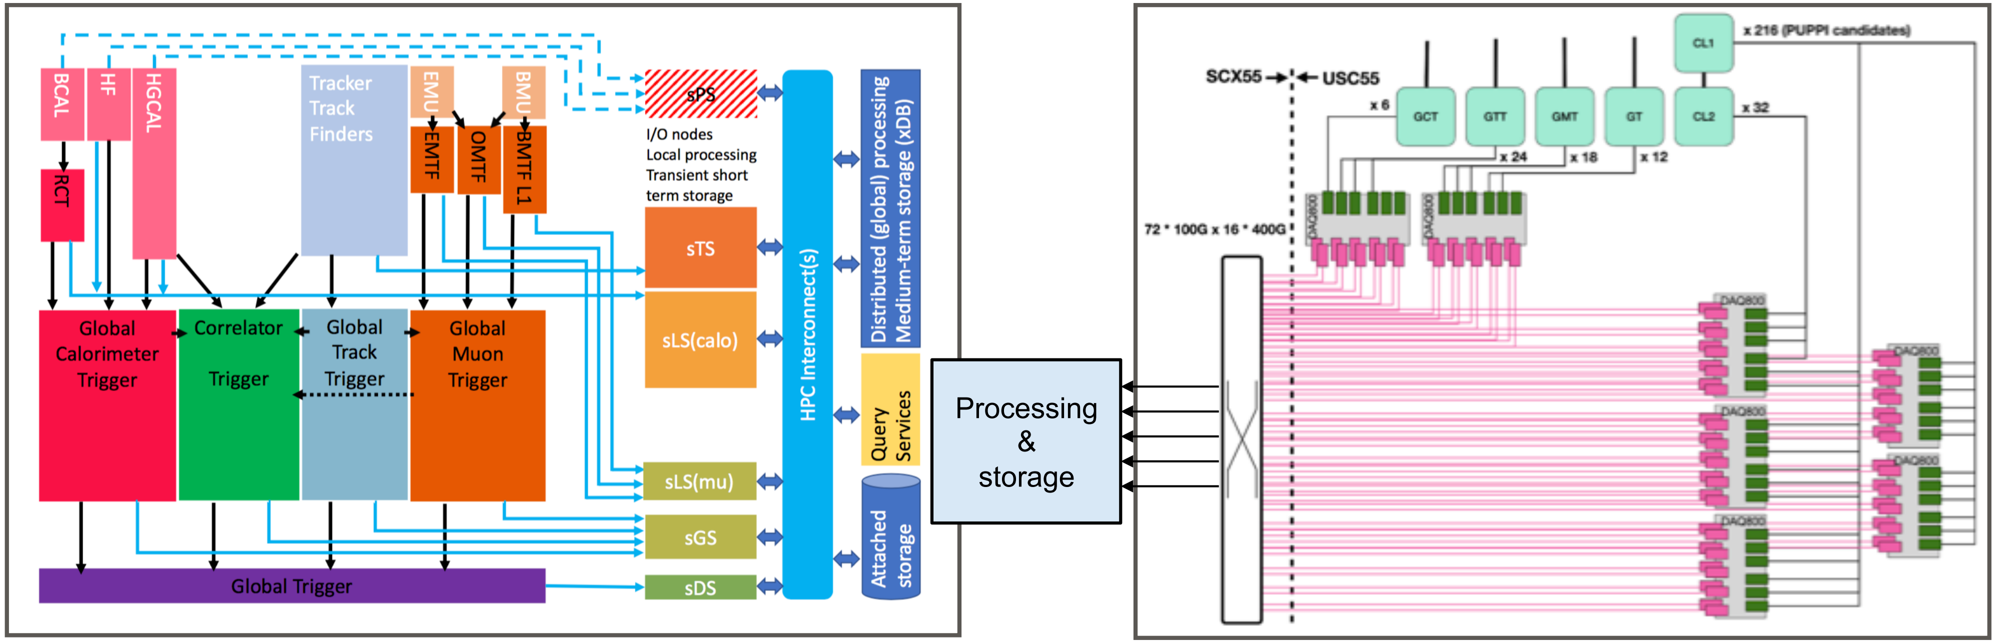

THE PHASE-2 L1 TRIGGER

ALGORITHMS, FIRMWARE & TESTING

selecting physics

PHASE II LEVEL-1 TRIGGER: ALGORITHMS & MENU

Algorithms for the Level-1 trigger:

- Extensive use of tracking to reach near offline performance (sharper efficiency turn-on curves) + reconstruction of **Primary Vertex.**
- Exploit complementarity of different object flavour:
  - **Standalone objects:** robust triggers based on independent sub-detectors
  - **Track-matched objects:** tracking used to confirm standalone Muon and Calo objects, significant improvement with simple design
  - **Particle-flow objects**: ultimate performance improvement, combine all information to match offline algorithms, require most processing time and resources for calculation

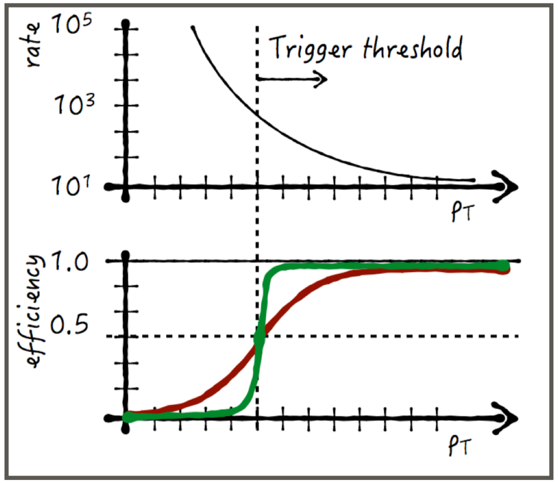

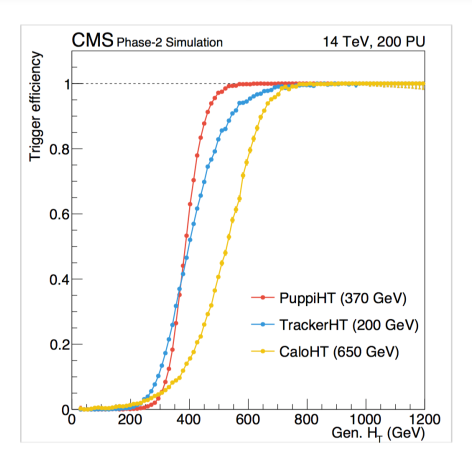

Particle-flow/PUPPI HT

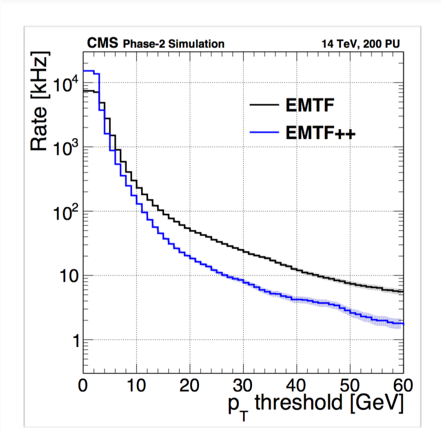

Standalone forward Muon reconstruction Endcap

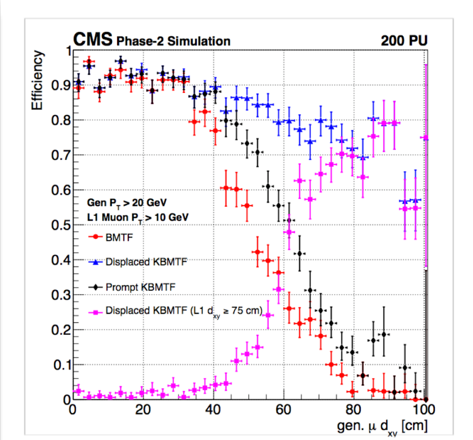

Displaced Muon trigger Barrel (Kalman Filter) →*Run-3*

Level-1 Menu:

- Simplified: Phase-1 physics built from Run-2 L1 Menu (346 kHz)
- Extended: new triggering strategy to expand physics reach (+110 kHz)

GLOBAL EVENT RECONSTRUCTION @ LEVEL-1

- Availability of tracks & high-granularity calos
- Implement global event reco @ L1 (like PF)
- Additionally it makes sense to mitigate pileup
- *Challenge : can we run full PF+PUPPI in time w/ hardware for L1?*

PF takes in everything | **PF is local** | **PF Links** | **Can we run a local PU Algo?**

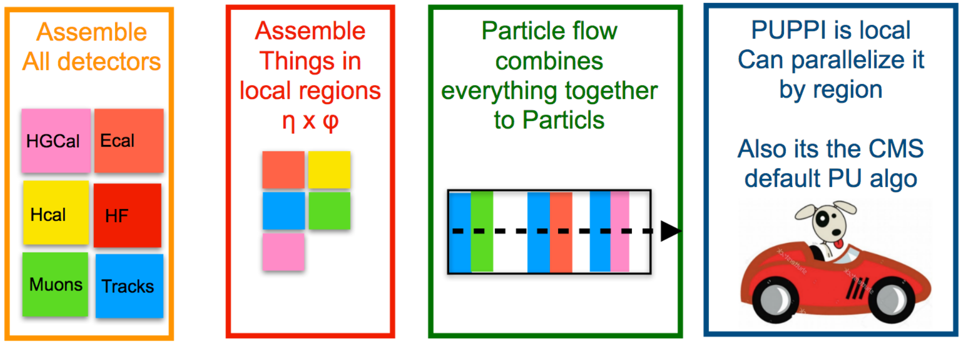

- **Demonstrated a working PF+PUPPI algorithm**
- PF+PUPPI hugely reduces the event complexity
- Allows for a lot of flexibility in downstream design
- L1 Algorithms looks like offline reconstruction
- PF+PUPPI developed with Vivado HLS (a lot of written by physicists along with engineers)

Sang Eon Park's poster: The Particle Flow Algorithm in the Phase II Upgrade of the CMS Level-1 Trigger

CMS Collaboration. Particle Flow CMS *JINST* **12** (2017) P10003, arXiv:1706.04965.
D. Bertolini, P. Harris, M. Low, and N. Tran, PUPPI, *JHEP* **10** (2014) 059, arXiv:1407.6013.

- What are the advantages of particle flow?

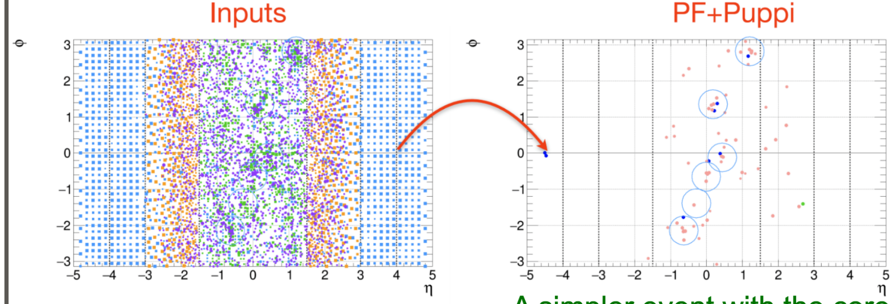

A simpler event with the core physics preserved

- Better resolution
- Ease of use → build particle level algorithms
- Reduction in bandwidth and reduction of resources

Testing this on firmware

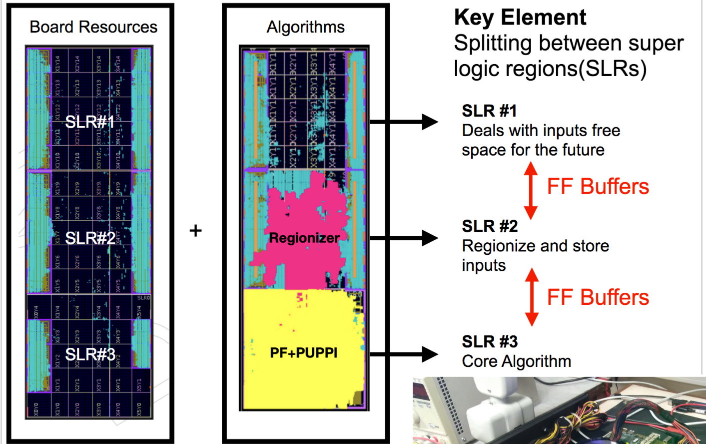

Hardware demonstrator

ALGORITHM INTO FIRMWARE

Firmware design:

- Algorithm developed mostly in C → High Level Synthesis (HLS)
- New fixed point arithmetic in C++ [taken from Xilinx libraries] → emulator firmware
- Continuous integration of the firmware in repository

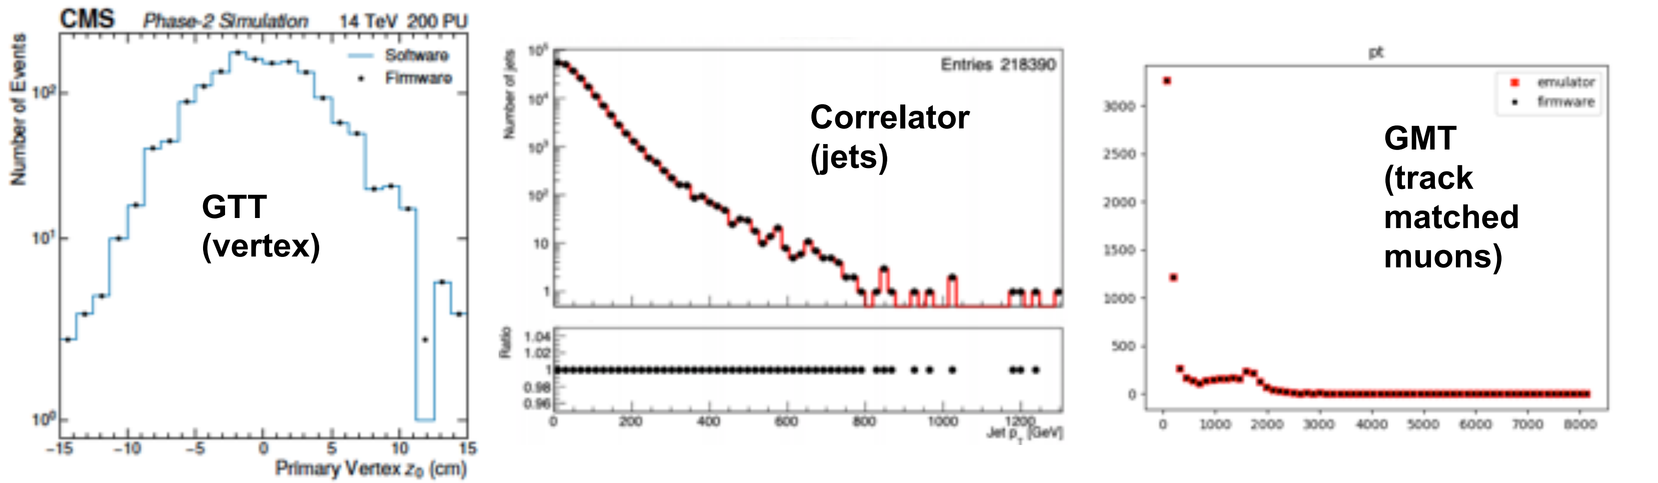

Excellent correspondance firmware & emulator

Firmware integration:

- All algo & manage I/O
- Verify timing, resources utilisation & latency.
- Common framework wrapper → firmware implementation board agnostic

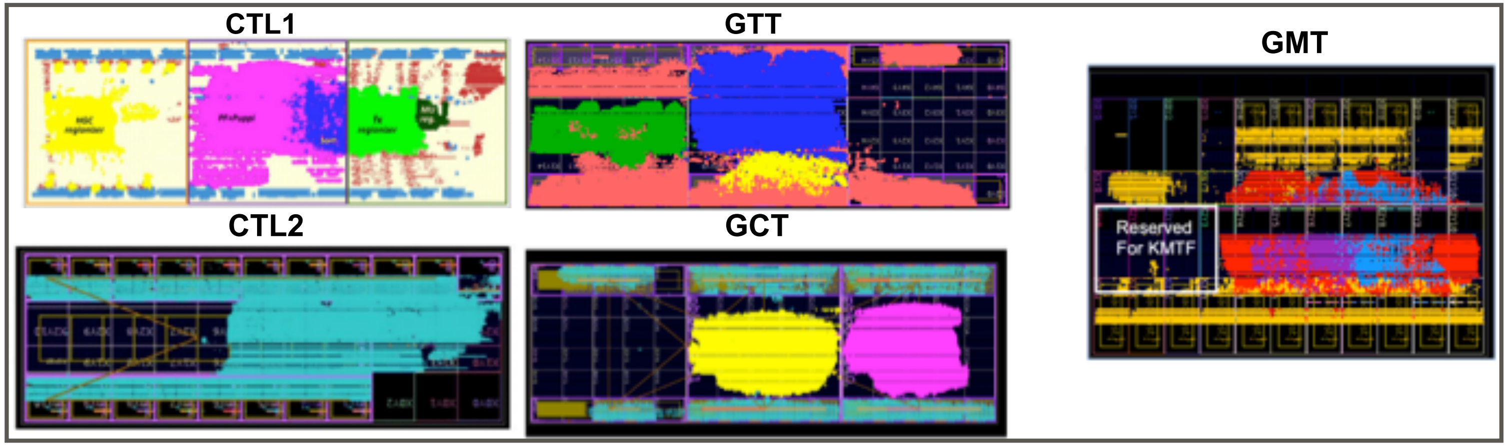

TESTING AND SYSTEM DEMONSTRATION

Phase-2 Level-1 Trigger system demonstration

- Single-board and multiple board tests
- Integration centers across the globe: larger scale integration planed @ CERN (904)
- **Board connection: protocol**
  - Links (asynchronous) operation @ 25.78 Gb/s
  - L1 Trigger boards sending packets only once (no retransmission) → error proof
  - Protocols (64/66b or 64/67b) encoding achieved low error rate, validated recovery mechanism etc.

APx prototype (Colorado)

Serenity in GT test crate (B40, CERN)

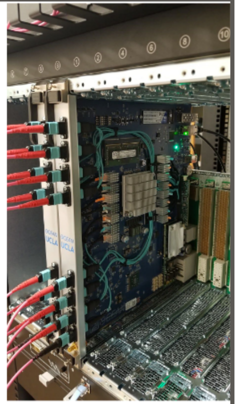

Ocean (UCLA)

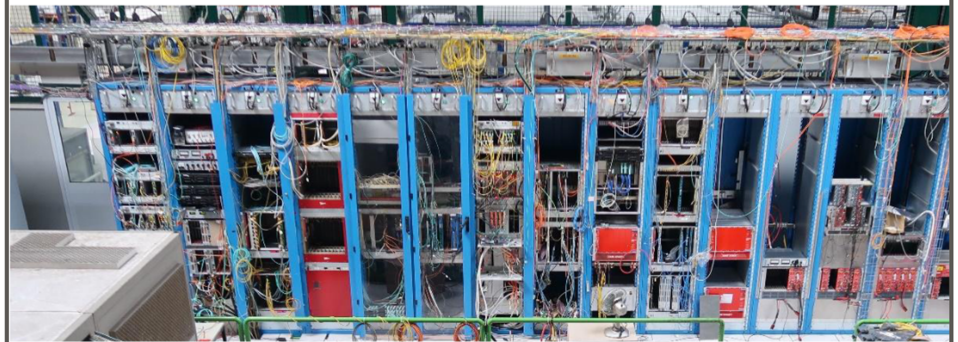

Electronics racks in building 904 (CMS electronics integration centre)

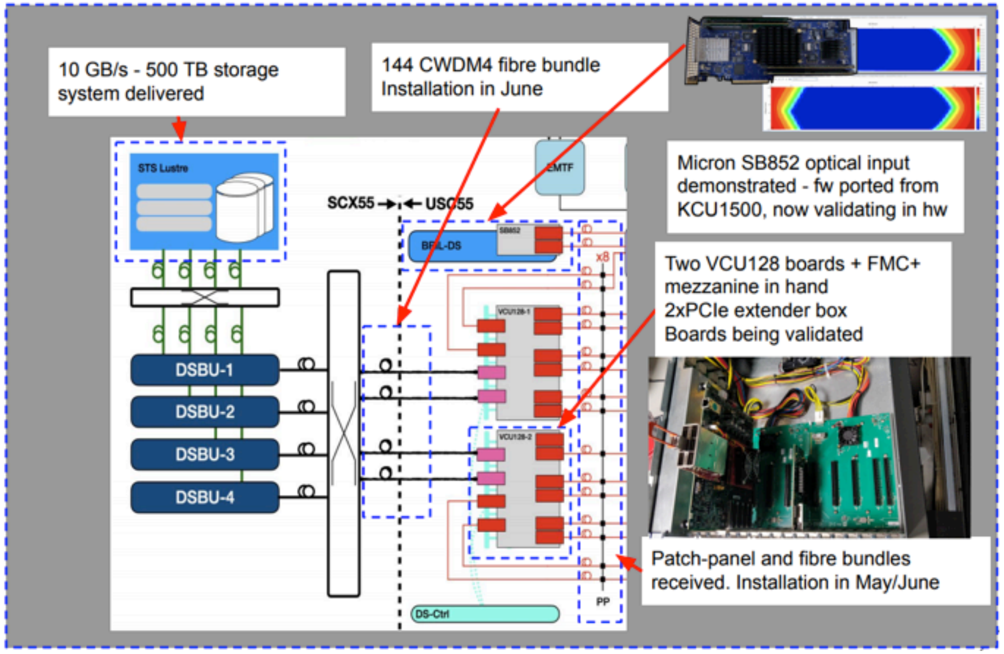

Scouting System for LHC Run-3

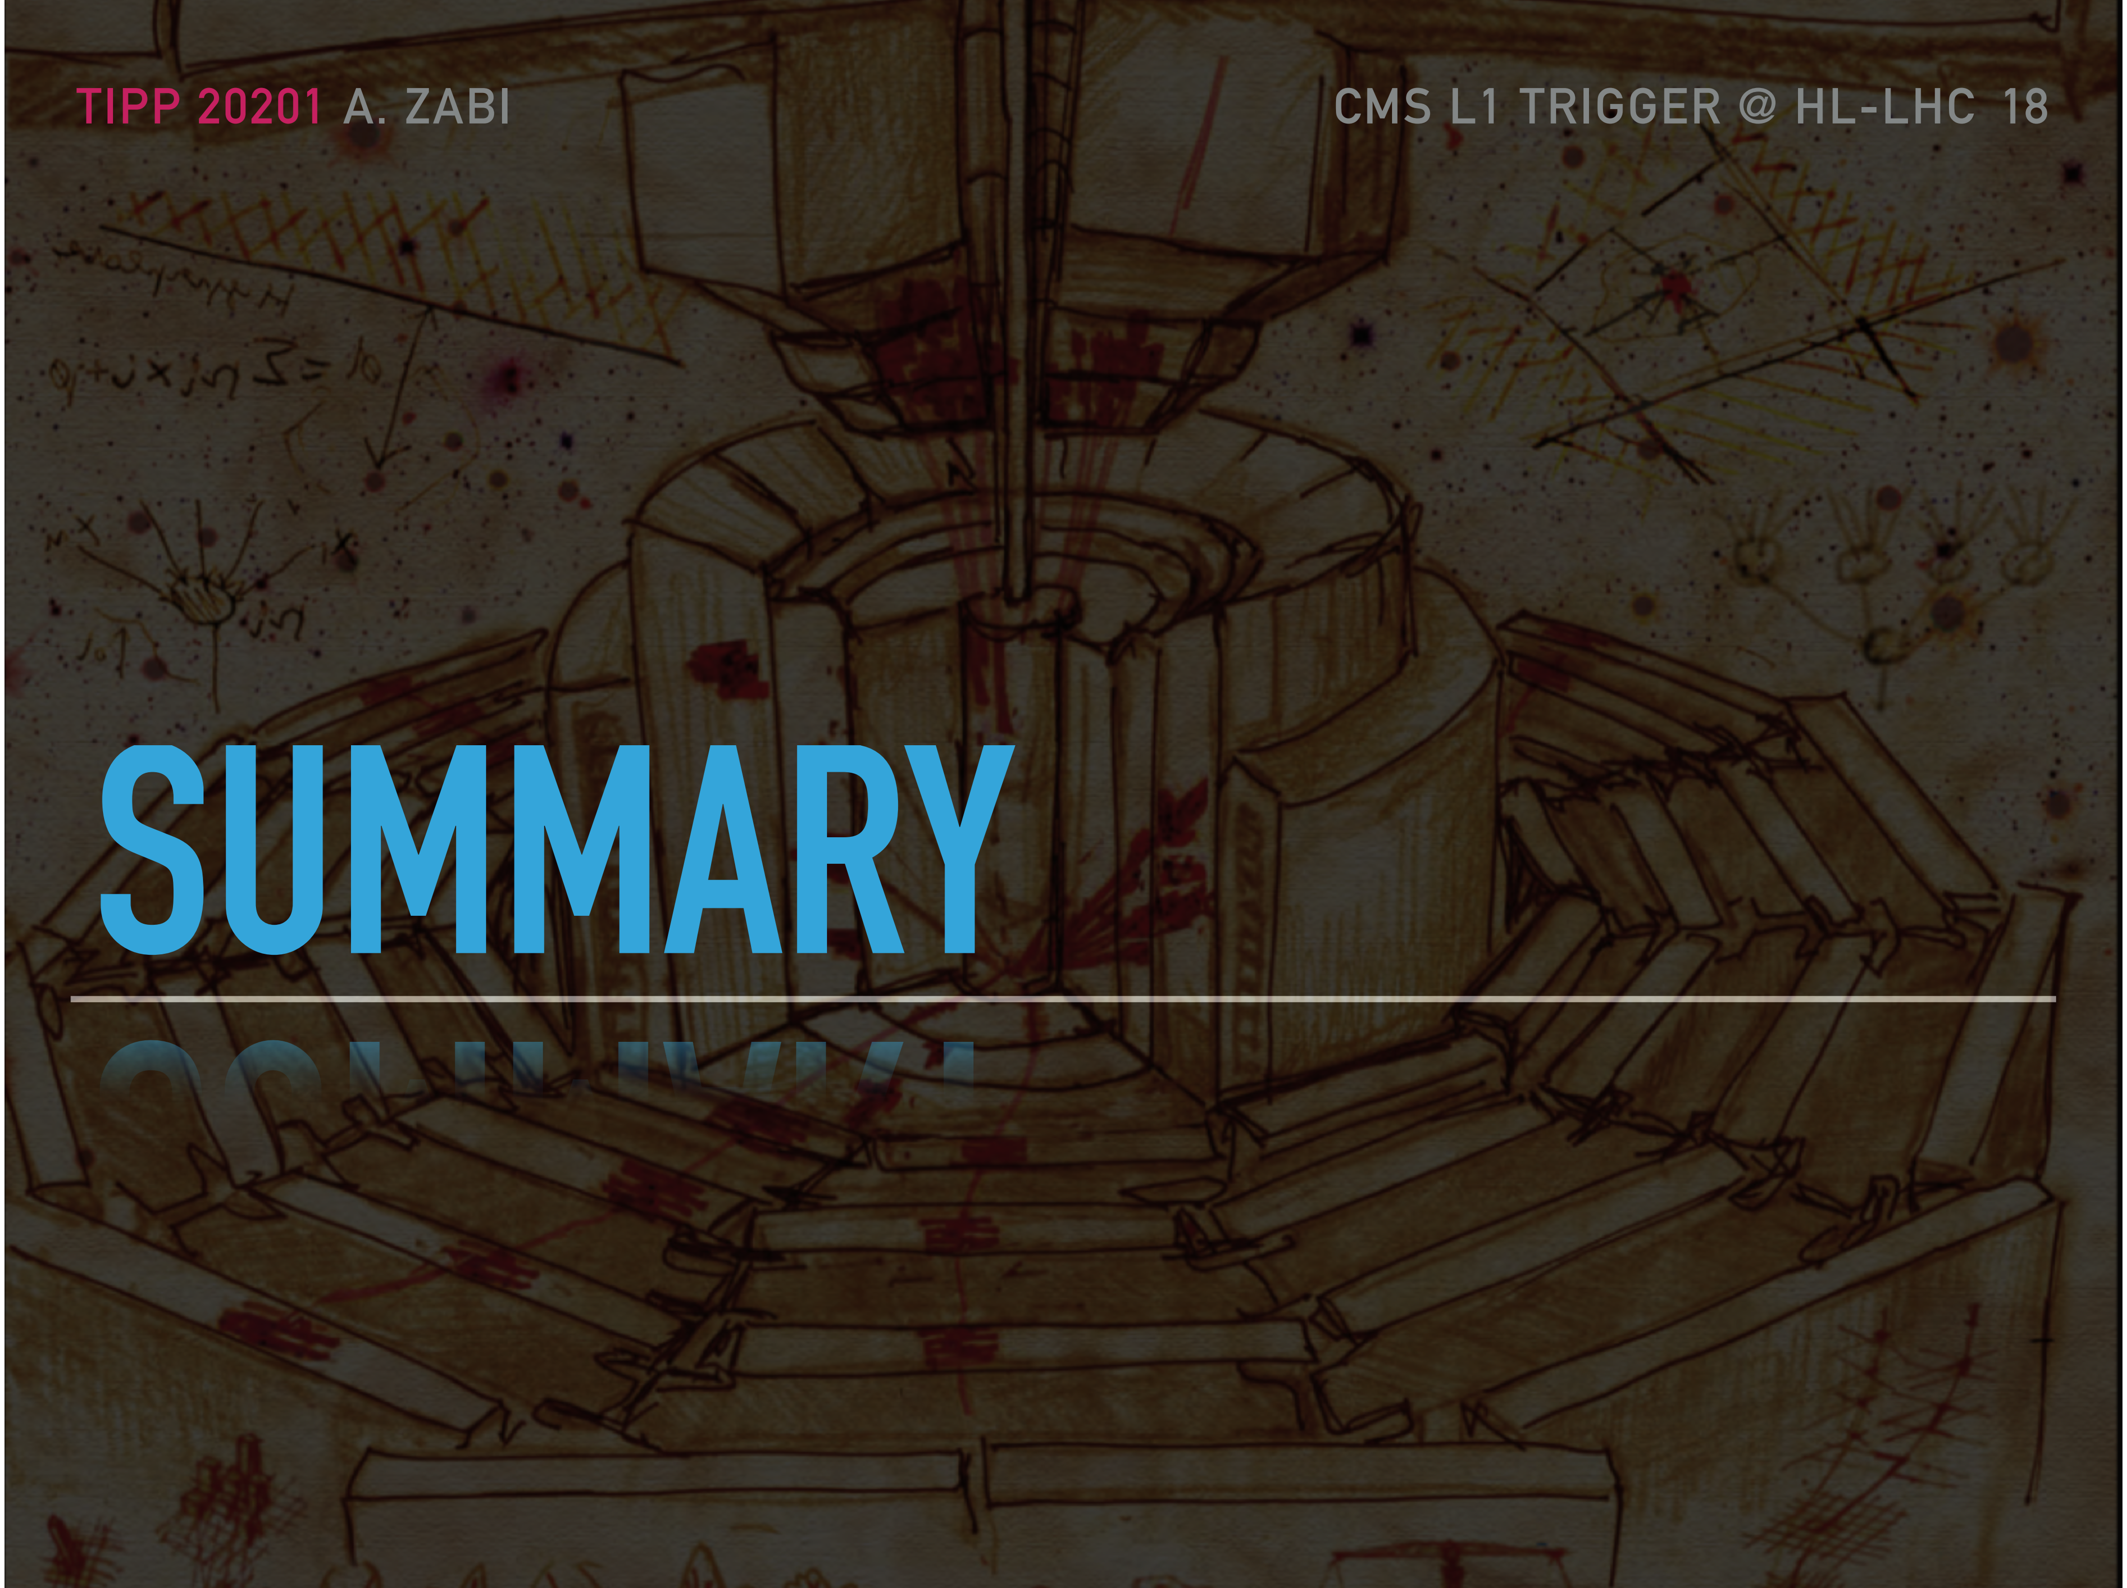

SUMMARY

CMS PHASE II L1 TRIGGER UPGRADE

- CMS proposing solid solutions to triggering and data acquisition challenge @ HL-LHC
- Phase-2 Level-1 Trigger Upgrade project: TDR approved in 2020 (https://cds.cern.ch/record/2714892?ln=en), steady progress with construction
- Level-1 Hardware trigger with enhanced capabilities complying with physics requirements. Sophisticated algorithms (particle-flow) are prototyped in FPGAs and exploit target hardware (VU9P/28Gb/s links)
- Modular and flexible architecture
- Hardware development lines pursuing 4 flavours of ATCA boards meeting the requirements of the project.
- Hardware demonstration ongoing and planned for testing with live data during LHC Run-3

BACK UP

PHYSICS @ HL-LHC

CMS Phase-2 physics drivers

- **Exploring the unknown :** Searches for new physics beyond the Standard Model (SM) DM, LLP,etc.
- **Standard Model as tool for discovery :** Precise knowledge of SM processes, probe anomalous couplings, 4 tops, VBS, VBF, etc. Higgs Sector: couplings (Hcc, Hμμ), differential xc, self-coupling HH
- **Understanding the Standard Model:** parton shower, underlying event, differential measurements

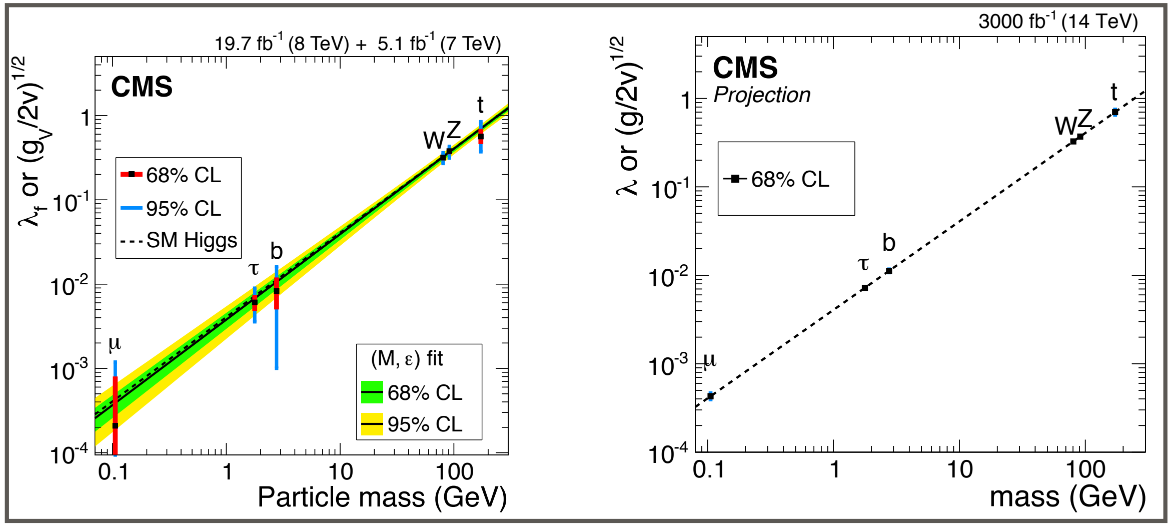

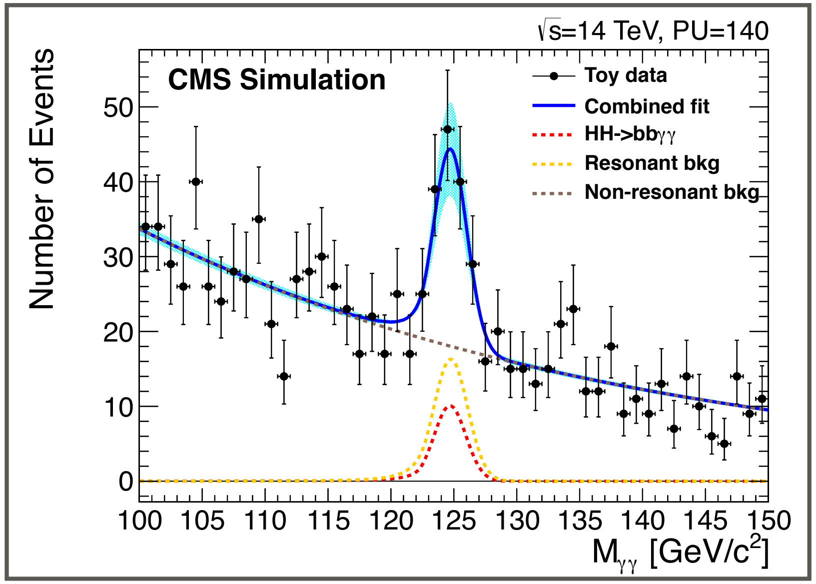

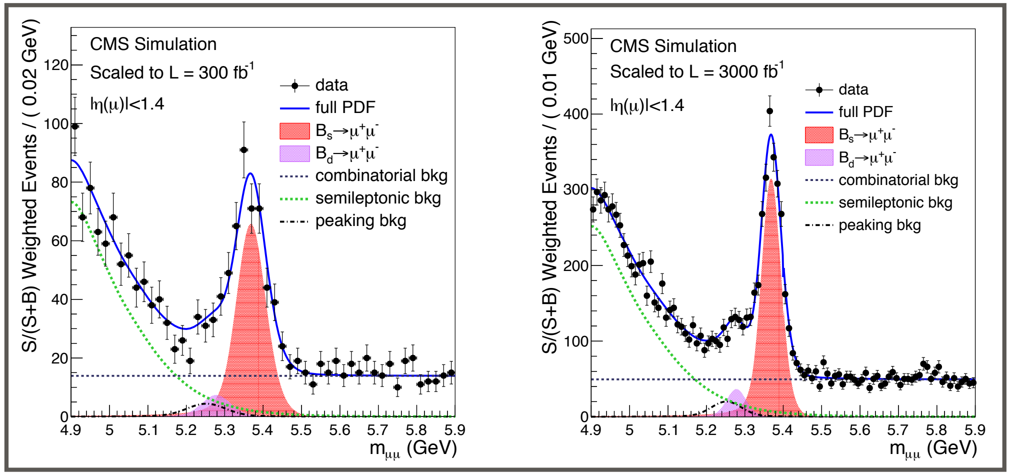

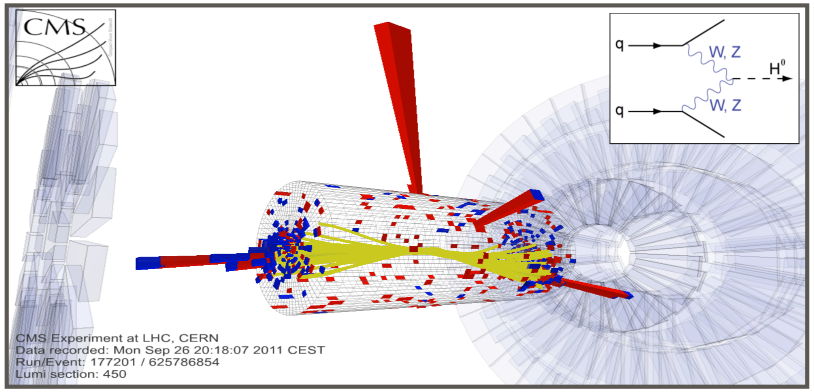

LEVEL-1 PHASE II TRIGGER UPGRADE SYSTEM

Phase-2 L1 trigger: latency

Latency budget = 9.5 us (20% margin to get to 12.5 us)

- **5 us** on region processing
- **4.5 us** on producing triggerable objects (including correlations) & final decision

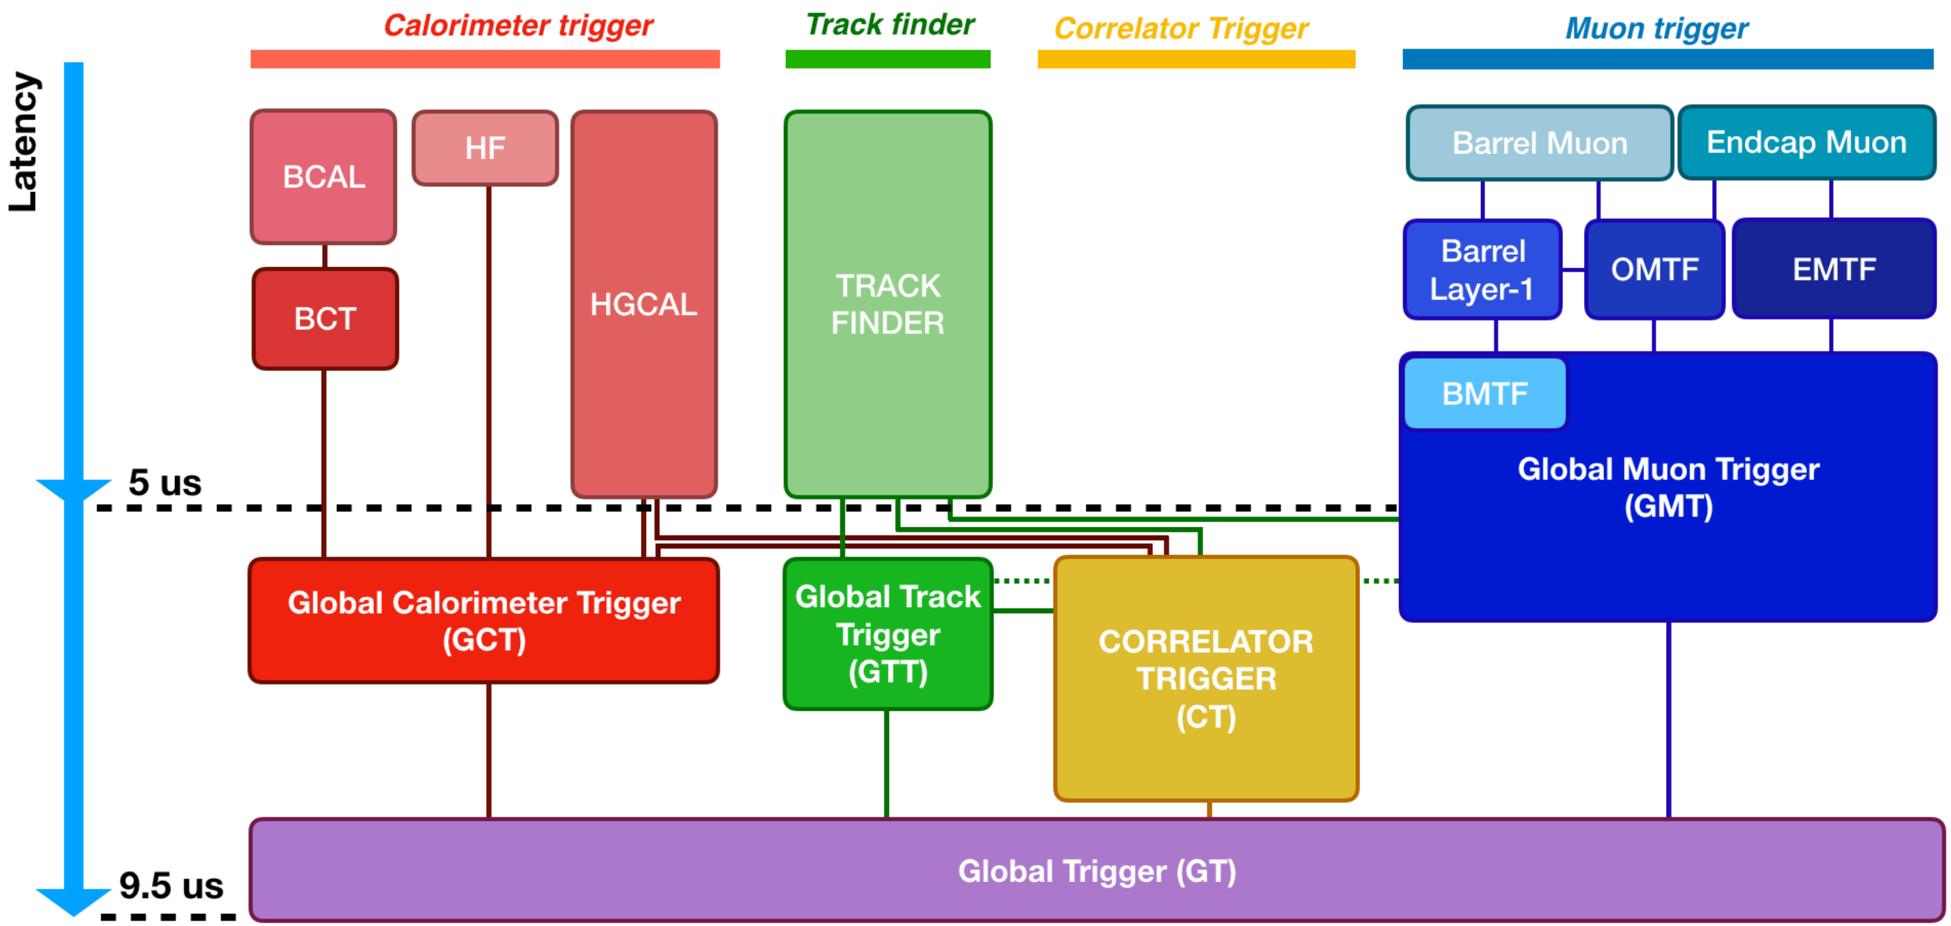

PHASE II TRIGGER UPGRADE INTERFACES

TRACK FINDER INPUT

MAIN FEATURES: Outer Tracker

- Trigger primitives |η|<2.4 & 9 phi sectors
- TMUX = 18
- Track finding algorithm: hybrid (fully reco tracks)
- Upper limit of 110 Tracks/sector (+40% extended tracking for displaced tracks)
- L1 Track word = 96 bits
- OUTPUT: 3 exact copies Correlator Trigger, GTT and GMT
- Latency = 5us

L1 Tracking System Overview

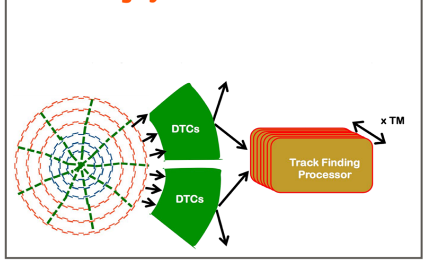

Track parameter	Number of bits
q/R	15
ϕ	12
$\tan(\lambda)$	16
z_0	12
d_0	13
χ^2/dof	4
bend-χ^2	3
hit mask	7
track quality MVA	3
other quality MVAs	6
spare	5
total	96

Table 2.1: Content of the L1 Track word.

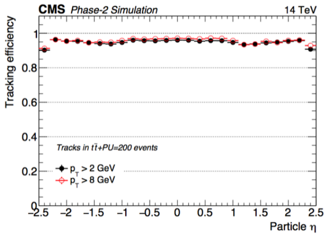

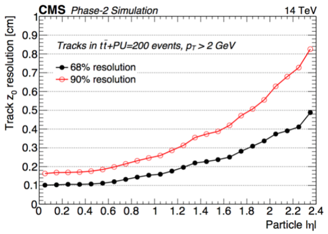

Excellent tracking efficiency

BARREL CALORIMETER INPUT

MAIN FEATURES: Barrel Calorimeter

- Provide ECAL crystal level information (optimal matching with tracks)
- HCAL allows to access depth info (pile-up removal in jets under study)
- TMUX = 1
- *Objects and format:*
  - *ECAL Crystals = 16 bits*
  - *ECAL Clusters (3x5) = 40 bits*
  - *HCAL Towers = 16 bits*
  - *HF Towers = 10 bits*
- *OUPUTS: GCT → Correlator Layer-1*
- *Latency = 1.5us (ECAL) & (HCAL)*

Both system provide timing information ECAL ~ 30 ps pT>50 GeV, HCAL~0.5 ns pT>50 GeV

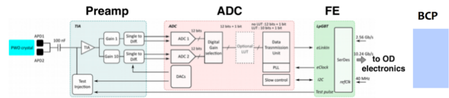

Moving Trigger Primitive logic in back end : BCP

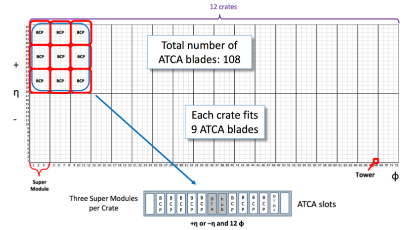

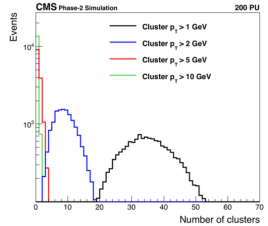

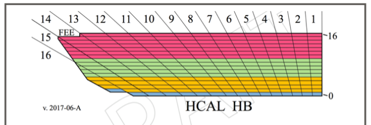

Figure 2.9: The grouping of scintillating layer segments into electronics channels in one azimuthal slice of an HCAL half-barrel. The scintillating tiles are numbered 0-16, the pseudorapidity segments 1-16, and the readout depths 1(blue), 2(orange), 3(green), and 4(magenta). FEE indicates the location of the front-end electronics.

ENDCAP CALORIMETER INPUT

MAIN FEATURES: HGCAL

- High granularity calorimeter (50 Layers)
- TP Backend system to build 3D-Clusters & shower shapes.
- TMUX = 18
- Objects and format:
  - 3D Clusters = 128 - 416 bits
  - Towers = 16 bits
- OUTPUT: 2 exact copies to Correlator Trigger Layer-1 & GCT.
- Latency = 5 us

No timing information available.

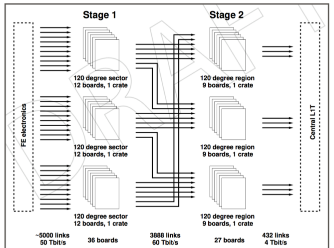

Data type	Content	Minimal size / bx (bits)	Maximum size
Header	Total energy, BX, $N_{clusters}$	32 × 6	-
Cluster	Position	32	-
	Energies	32	-
	Shape variables	64	352
Total clusters		128 ×(≈ 900 clusters)	416 ×(≈ 300 clusters)
Tower	Energy	12	-
	EM fraction	4	-
Total 2600 towers		43 k	-

Table 2.4: Data sent to the central L1T

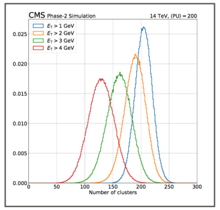

MUON BARREL INPUT

MAIN FEATURES: DT&RPC

- DT: new electronics, including TDC w/ ~2ns resolution (better BX ID, spatial resolution, fake rejection)
- DT TP Algorithms : Majority Mean-Timer and Analytical Method both exploiting fine detector information (DT cells & correlation among superlayers)
- RPC : new readout w/ timing (1/16 BX granularity = 1.5 ns) + faster links
- DT+RPC Super Primitives: combines RPC clusters (timing) and DT TP (position)
  → better performance (robust against ageing)
- TMUX = 1
- Objects and format
  - DT+RPC = 64 bits
  - Expecting 28 segments (long + transverse)
- OUTPUT: GMT(BMTF) & OMTF
- Latency = 2-3 us

Data field	Position	Bending angle	BX in-orbit	Sub-BX time	Chi2	Quality	Superlayer	RPC info
Bits	17	12	12	5	4	4	2	8

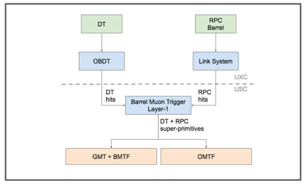

Muon Barrel Super primitives

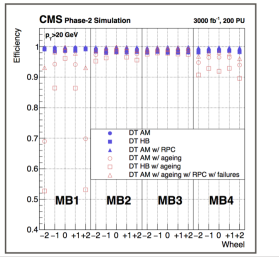

MUON ENDCAP INPUT

MAIN FEATURES: CSC/RPC/iRPC/GEM

- CSC: new electronics, finer granularity of TP (improved bend direction and angle resolution)
- RPC : same as barrel upgrade
- iRPC : 1.1 ns time resolution
- GEM/ME0: precise measurement of muon bending angle (better reach and rate control)
- CSC-GEM super primitives can improve bend direction resolution (under study). Could also CSC + (i)RPC
- TMUX = 1
- Objects and format
  - CSC segment = 32 bits
  - RPC Clusters = 15 bits
  - iRPC Clusters = 41 bits
  - GEM Clusters = 14 bits
  - GEM ME0 segment = 24 bits
- OUTPUT: OMTF & BMTF
- Latency = 1 us - 1.75 us (CSC)

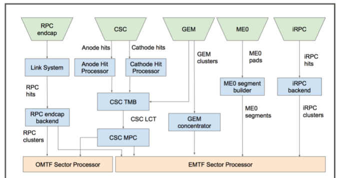

Figure 2.17: Muon endcap trigger primitive architecture.

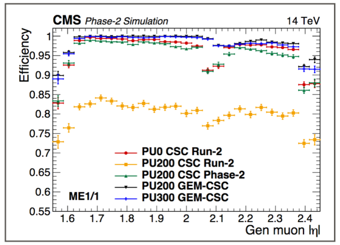

TECHNOLOGICAL CHOICES

Phase-2 L1 Trigger Design: Key technological features (inspired from the L1 Phase-1 upgrade)

- **FPGA:** The extensive use of state-of-the-art FPGAs → optimised reconstruction, identification, isolation and energy calibration of trigger objects using high-granularity detector information.
- **High-speed optical links:** facilitate the aggregation of data from across the entire detector → A complete view of the detector (evaluation of global quantities MET, pileup, specific VBF)
- **Flexible and modular architecture:** Reconfigured to adapt to HL-LHC running conditions and physics needs. Extra resources → Compute sophisticated quantities → richer menu and increased selectivity

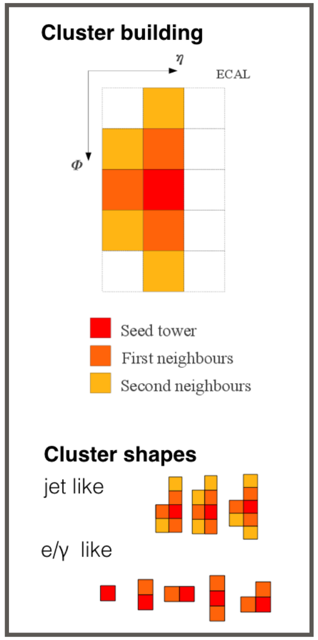

Phase-1 EG algo

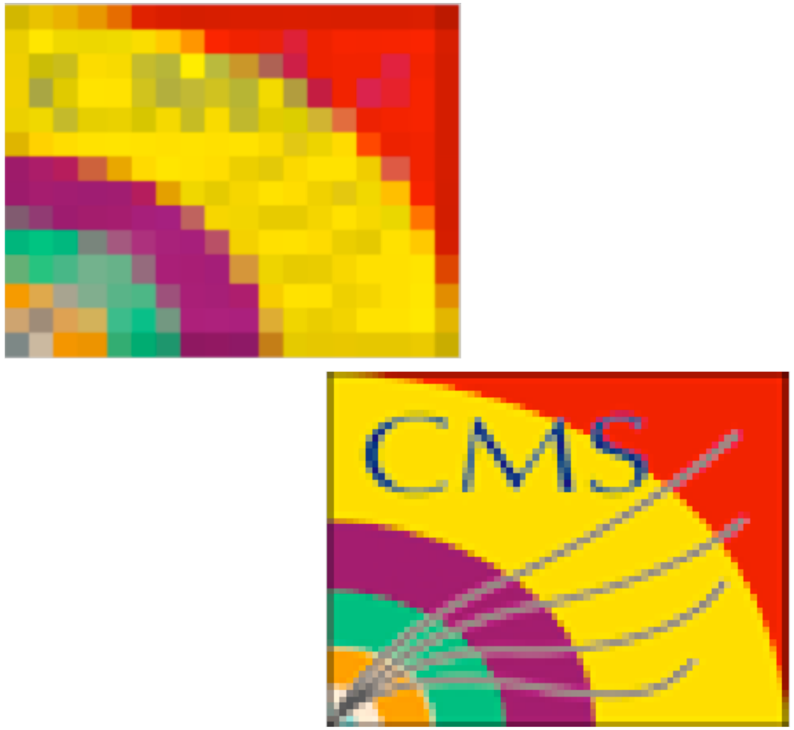

Phase-1 Upgrade: increased calorimeter granularity

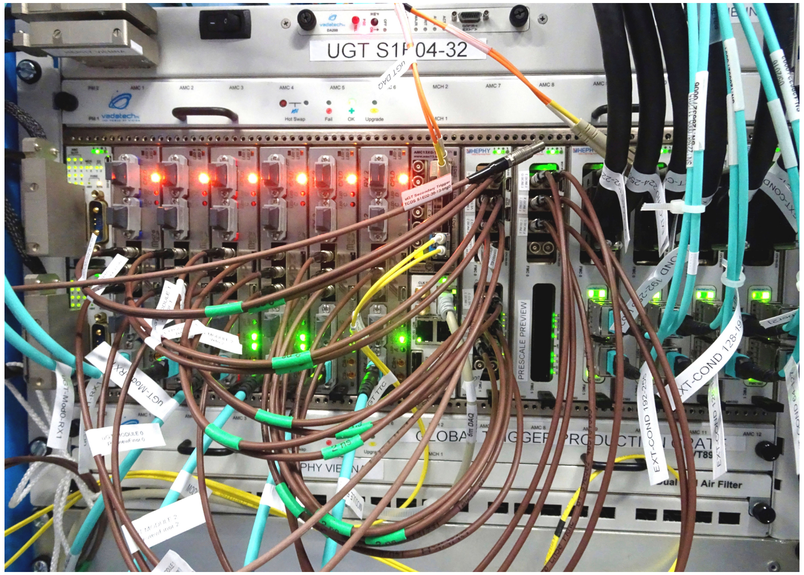

Phase-1 Upgrade: Expanded architecture