Data acquisition strategies in High Energy Physics: insights from the Tracker sub-detector of CMS for HL-LHC

Giacomo Fedi
Imperial College London

IMPERIAL

27 March 2024

Outline

- Introduction: Data acquisition in HEP
- Data acquisition in CMS during LHC upgrade
- Focus on CMS Tracker

Large Hadron Collider (LHC)

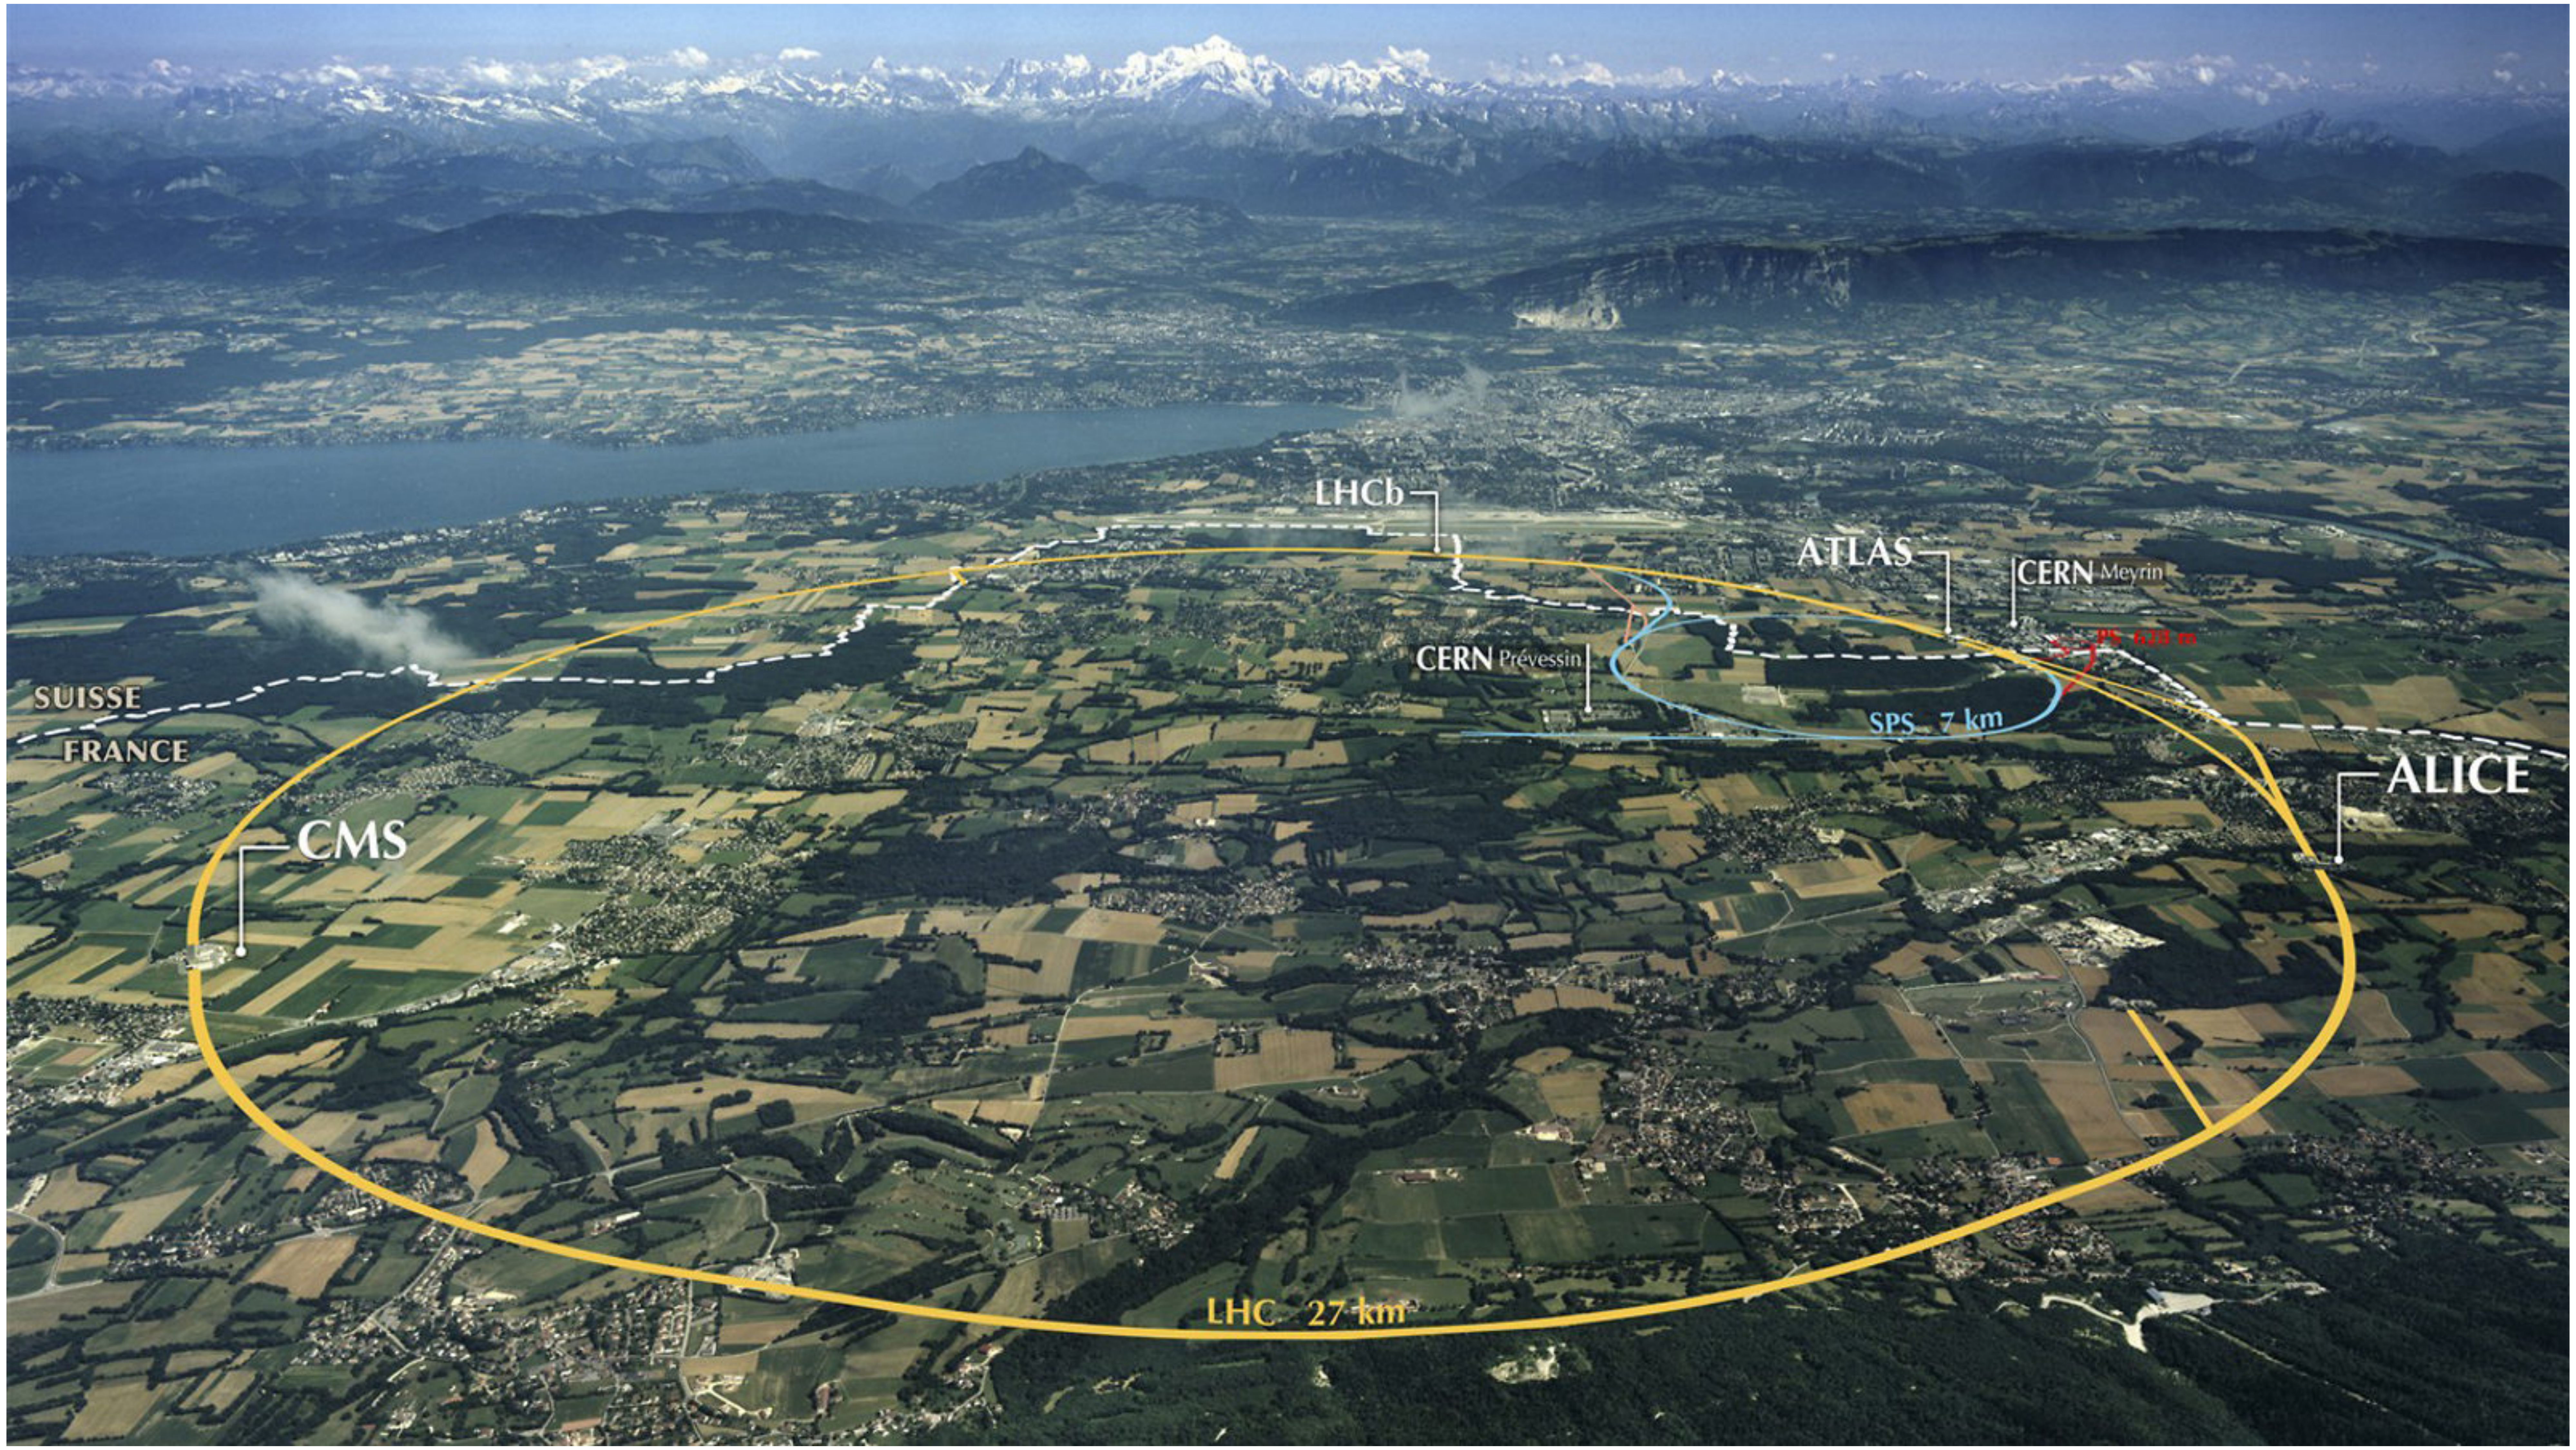

Source: Maximilien Brice, CERN

Compact Muon Solenoid (CMS)

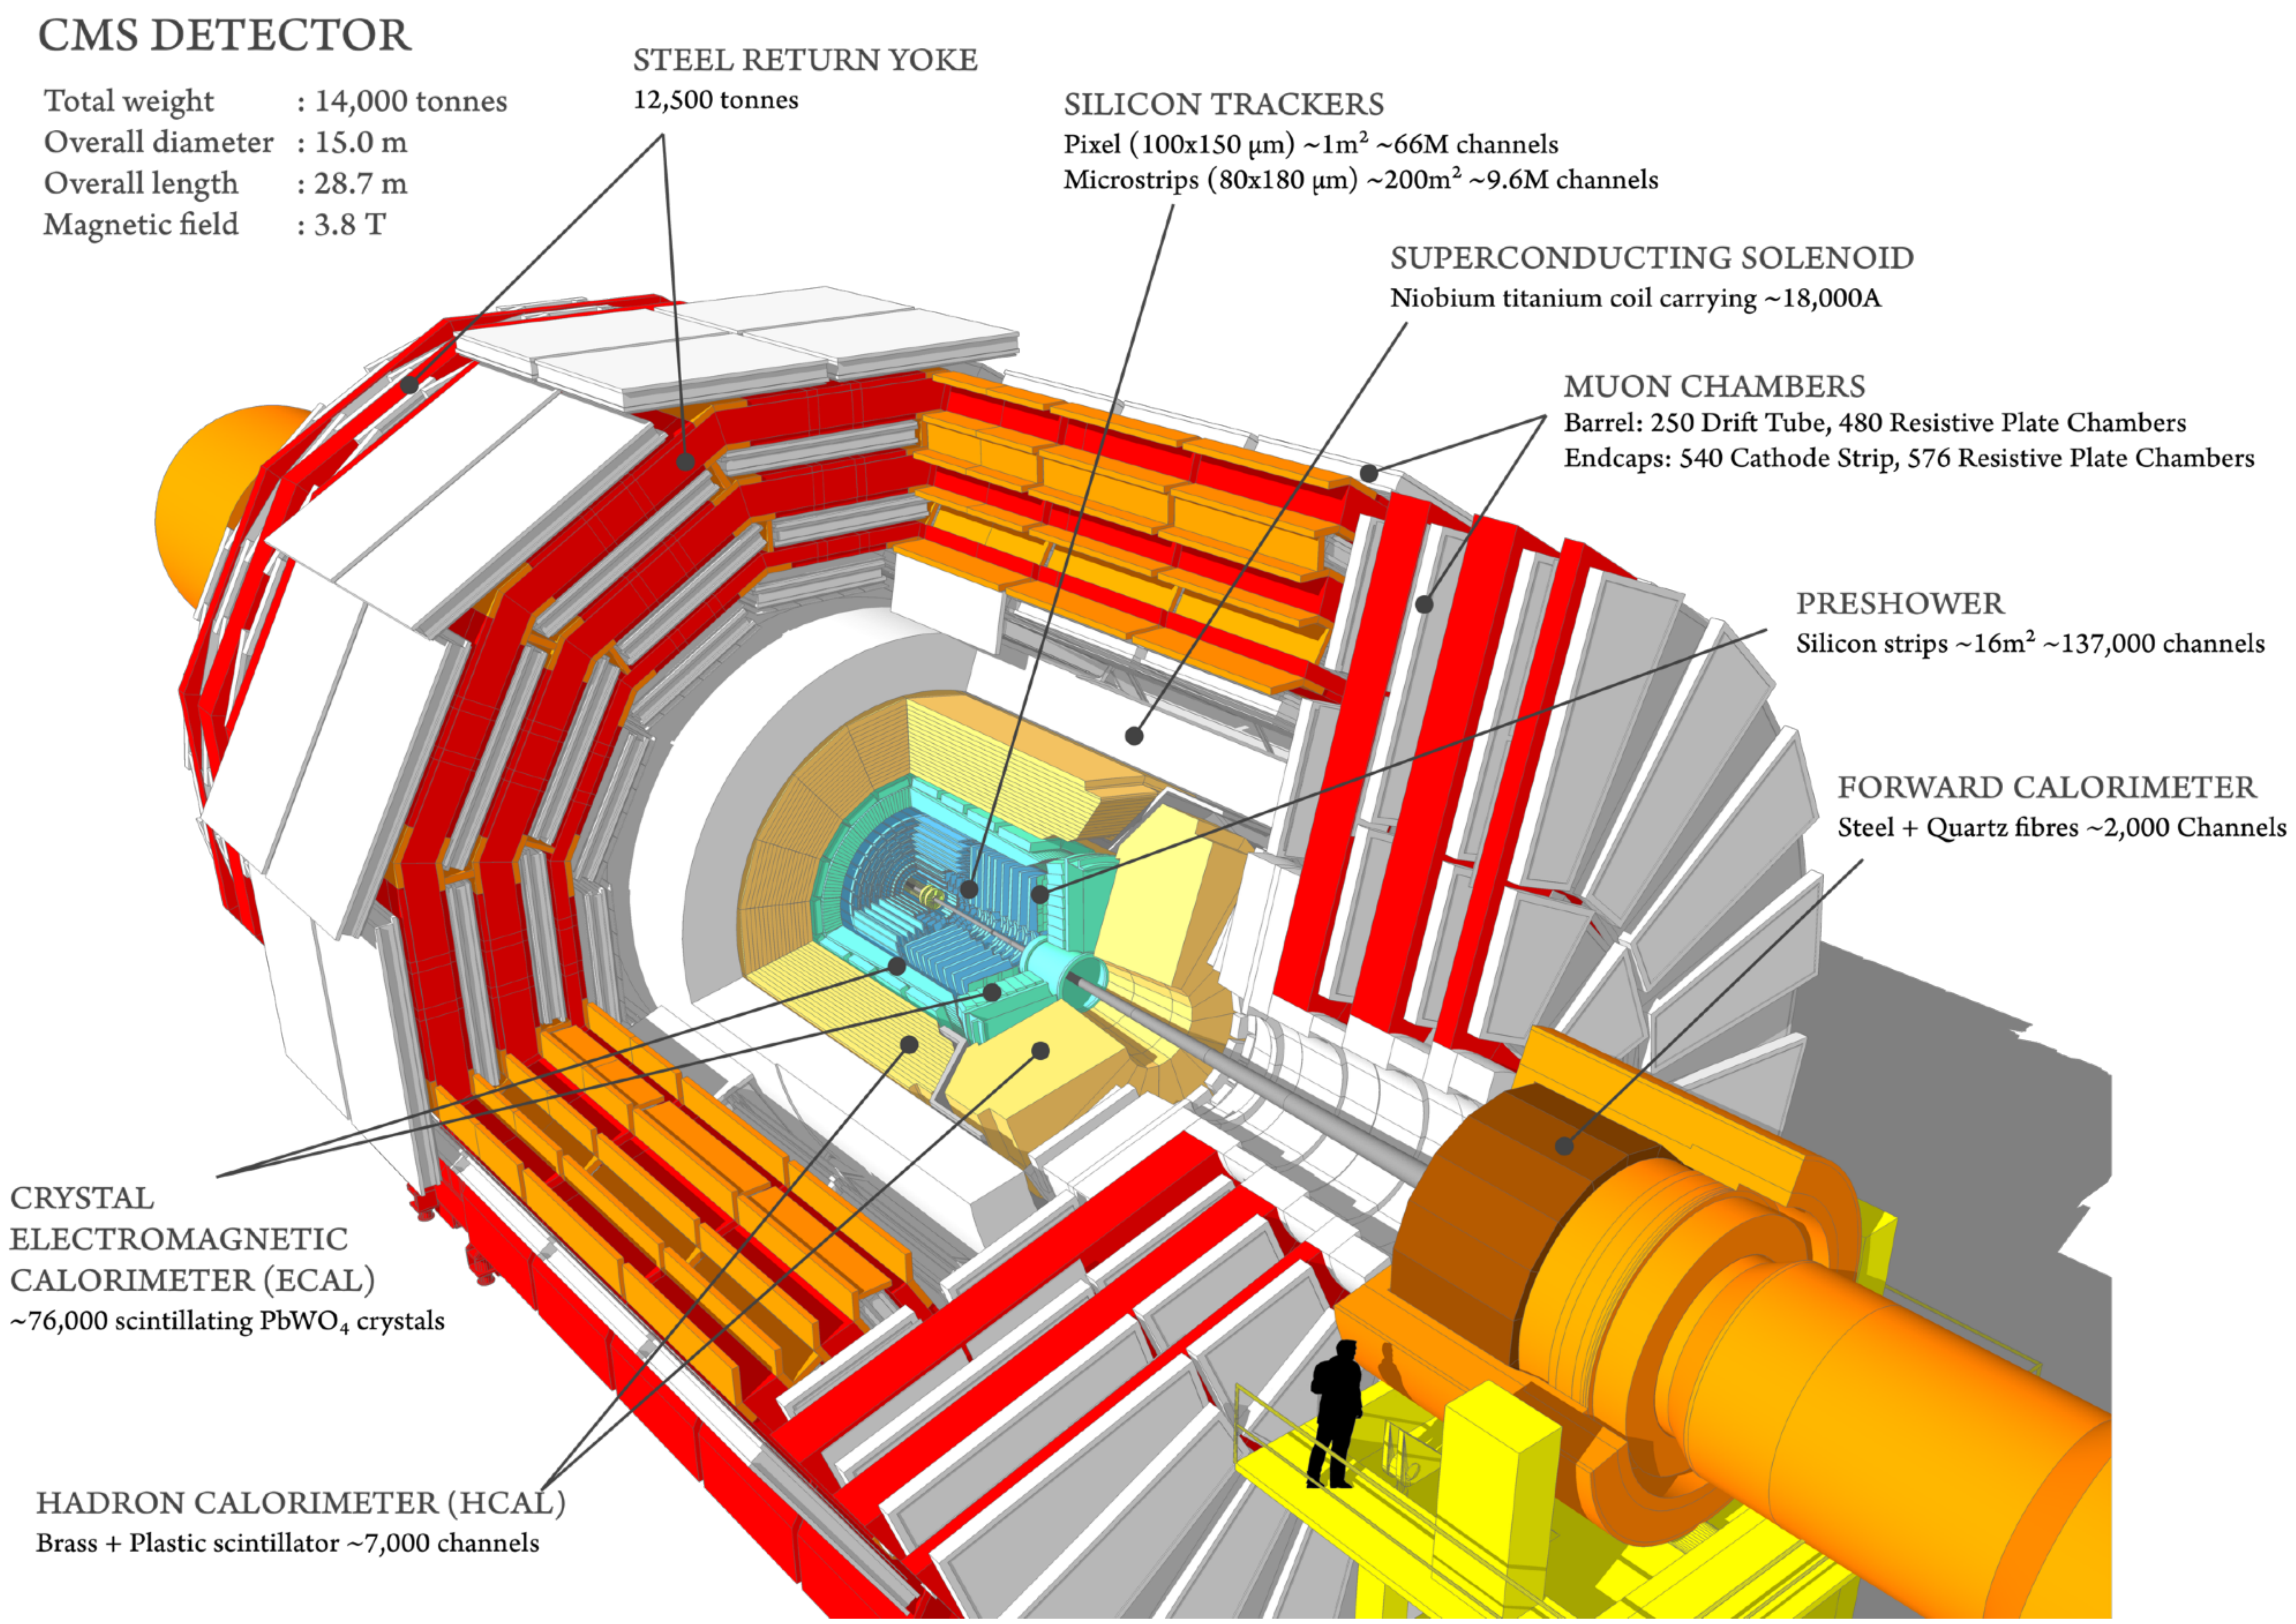

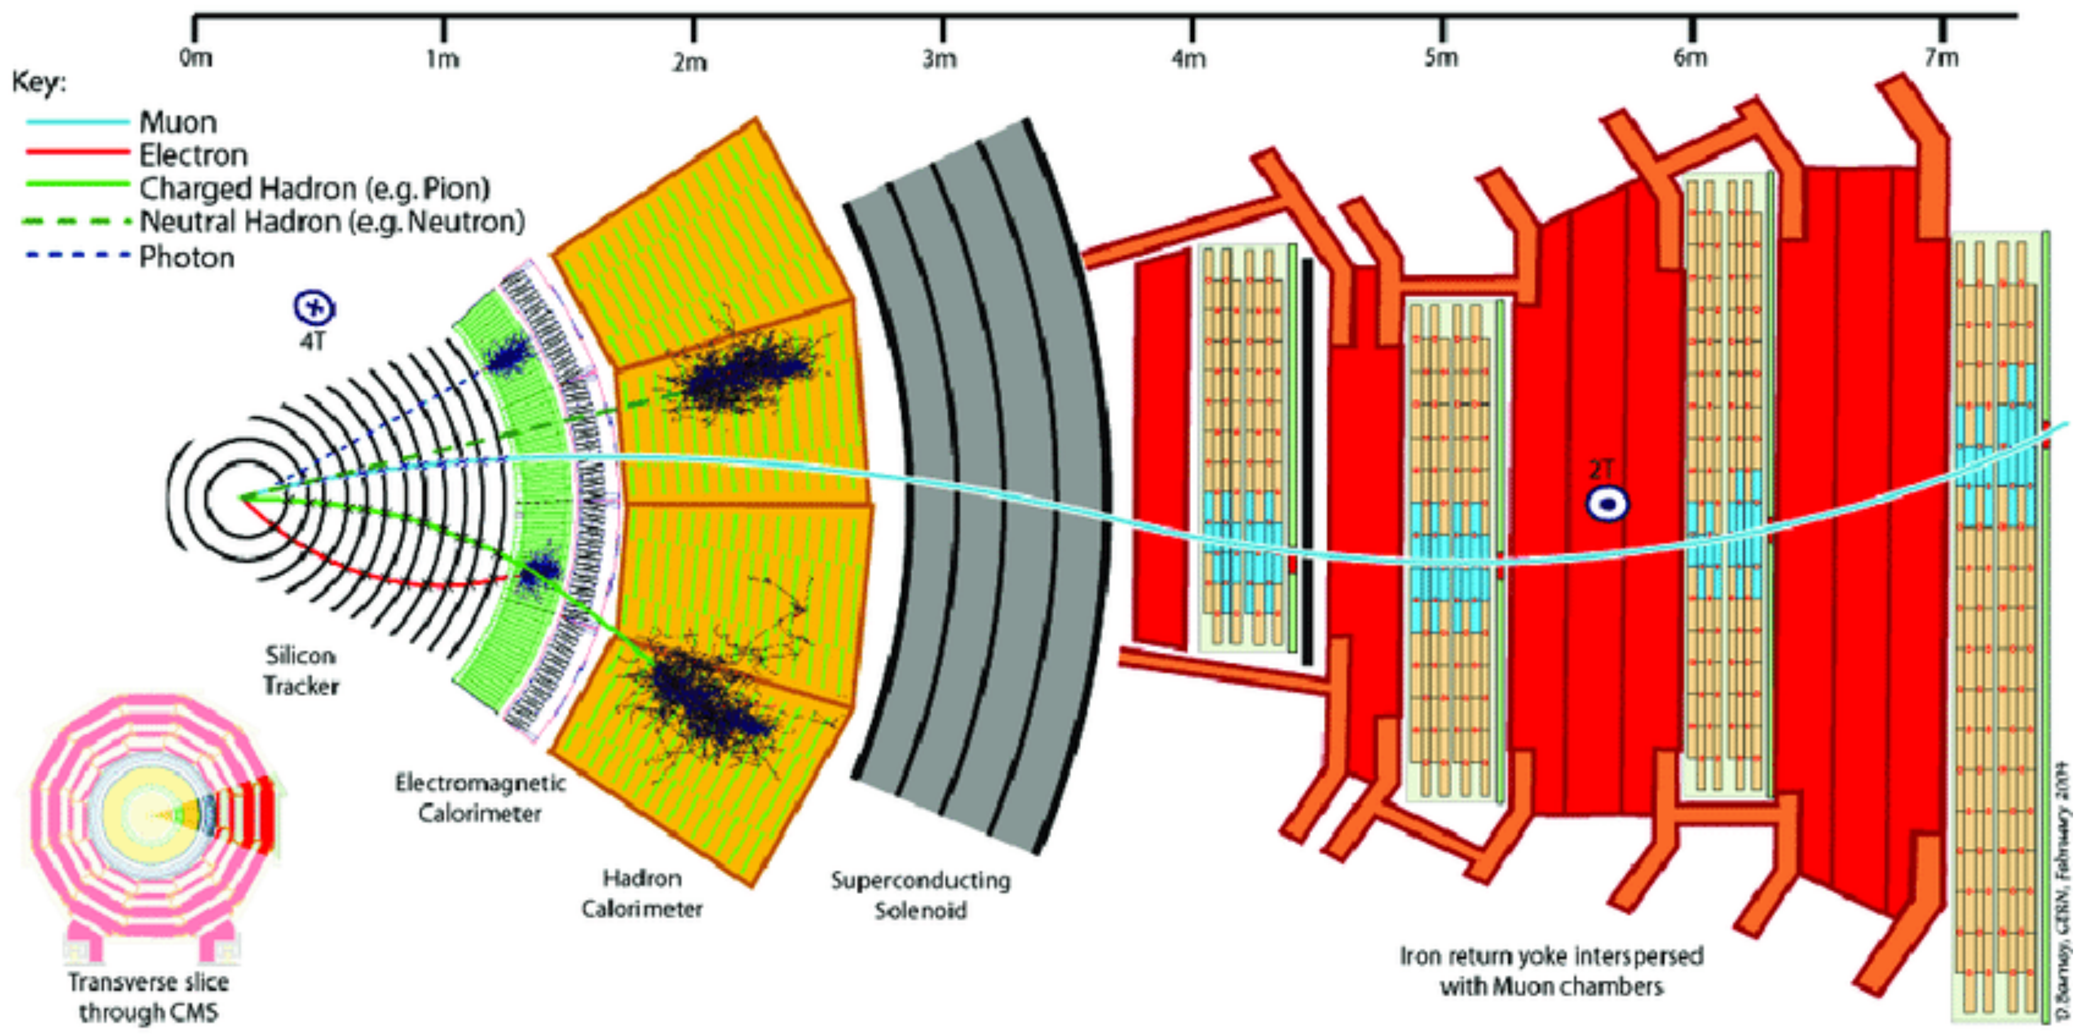

Source: CMS

I p-p collisions and cross sections

Nominal bunch crossing rate: 40 MHz

Avergage raw event size: ~1MB

- Data rate: O(100TB/s)
- Cannot save this amount of data and most of collisions are not interesting -> Solution: Triggering

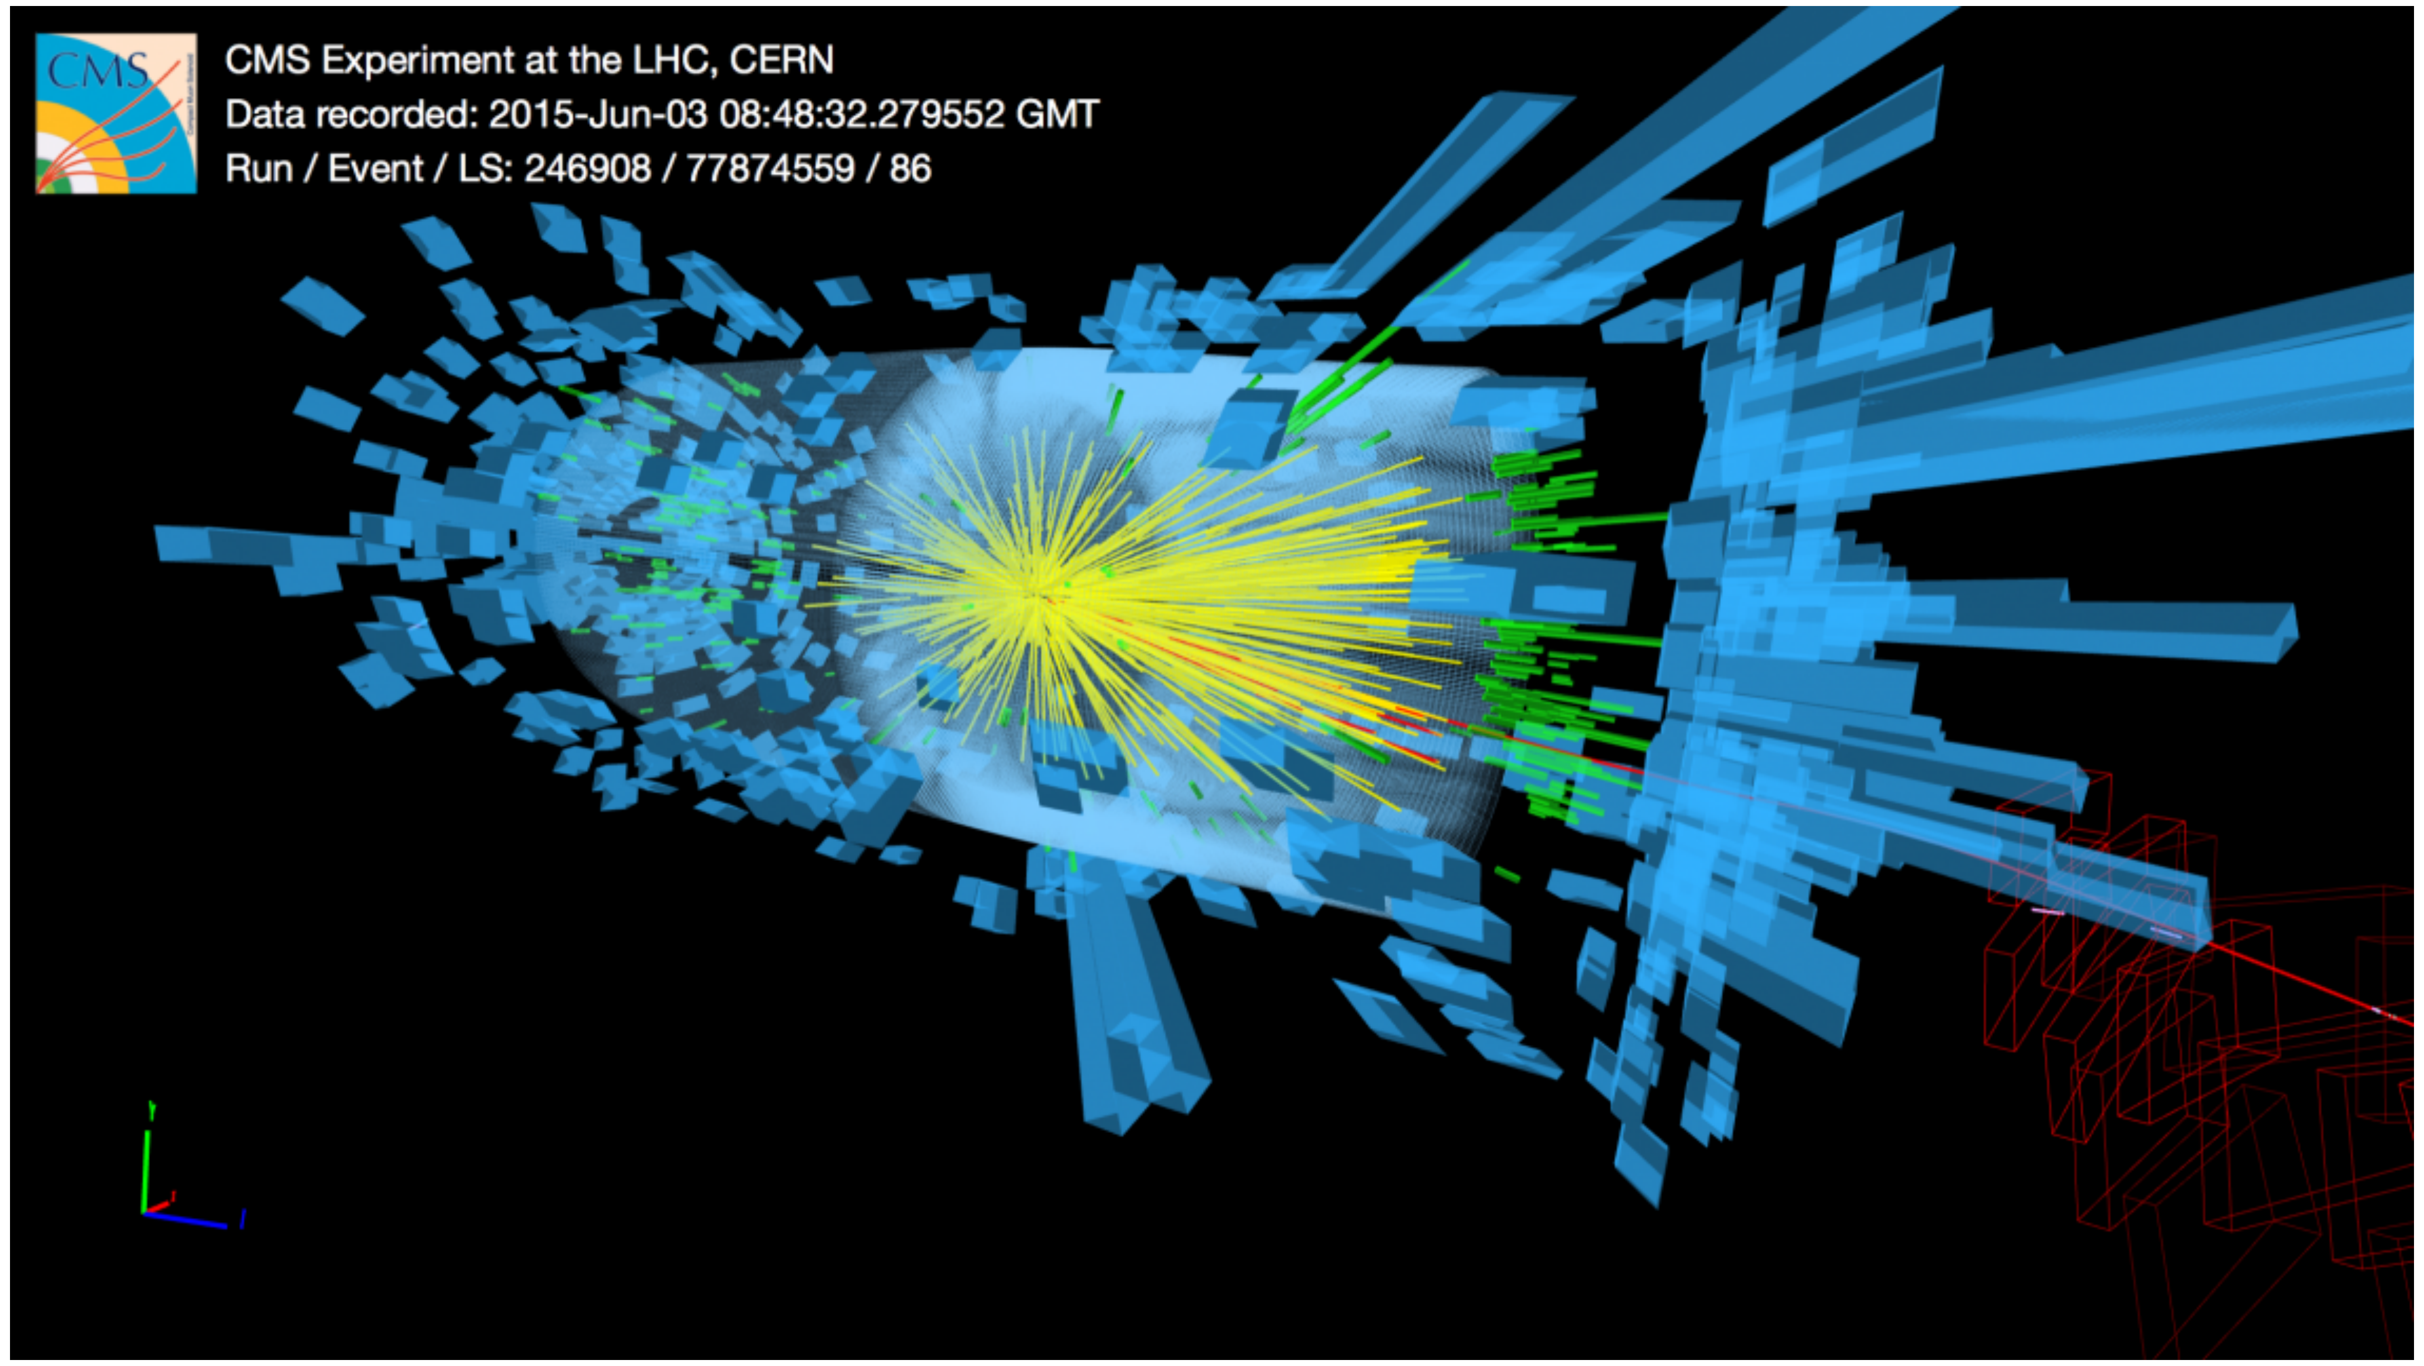

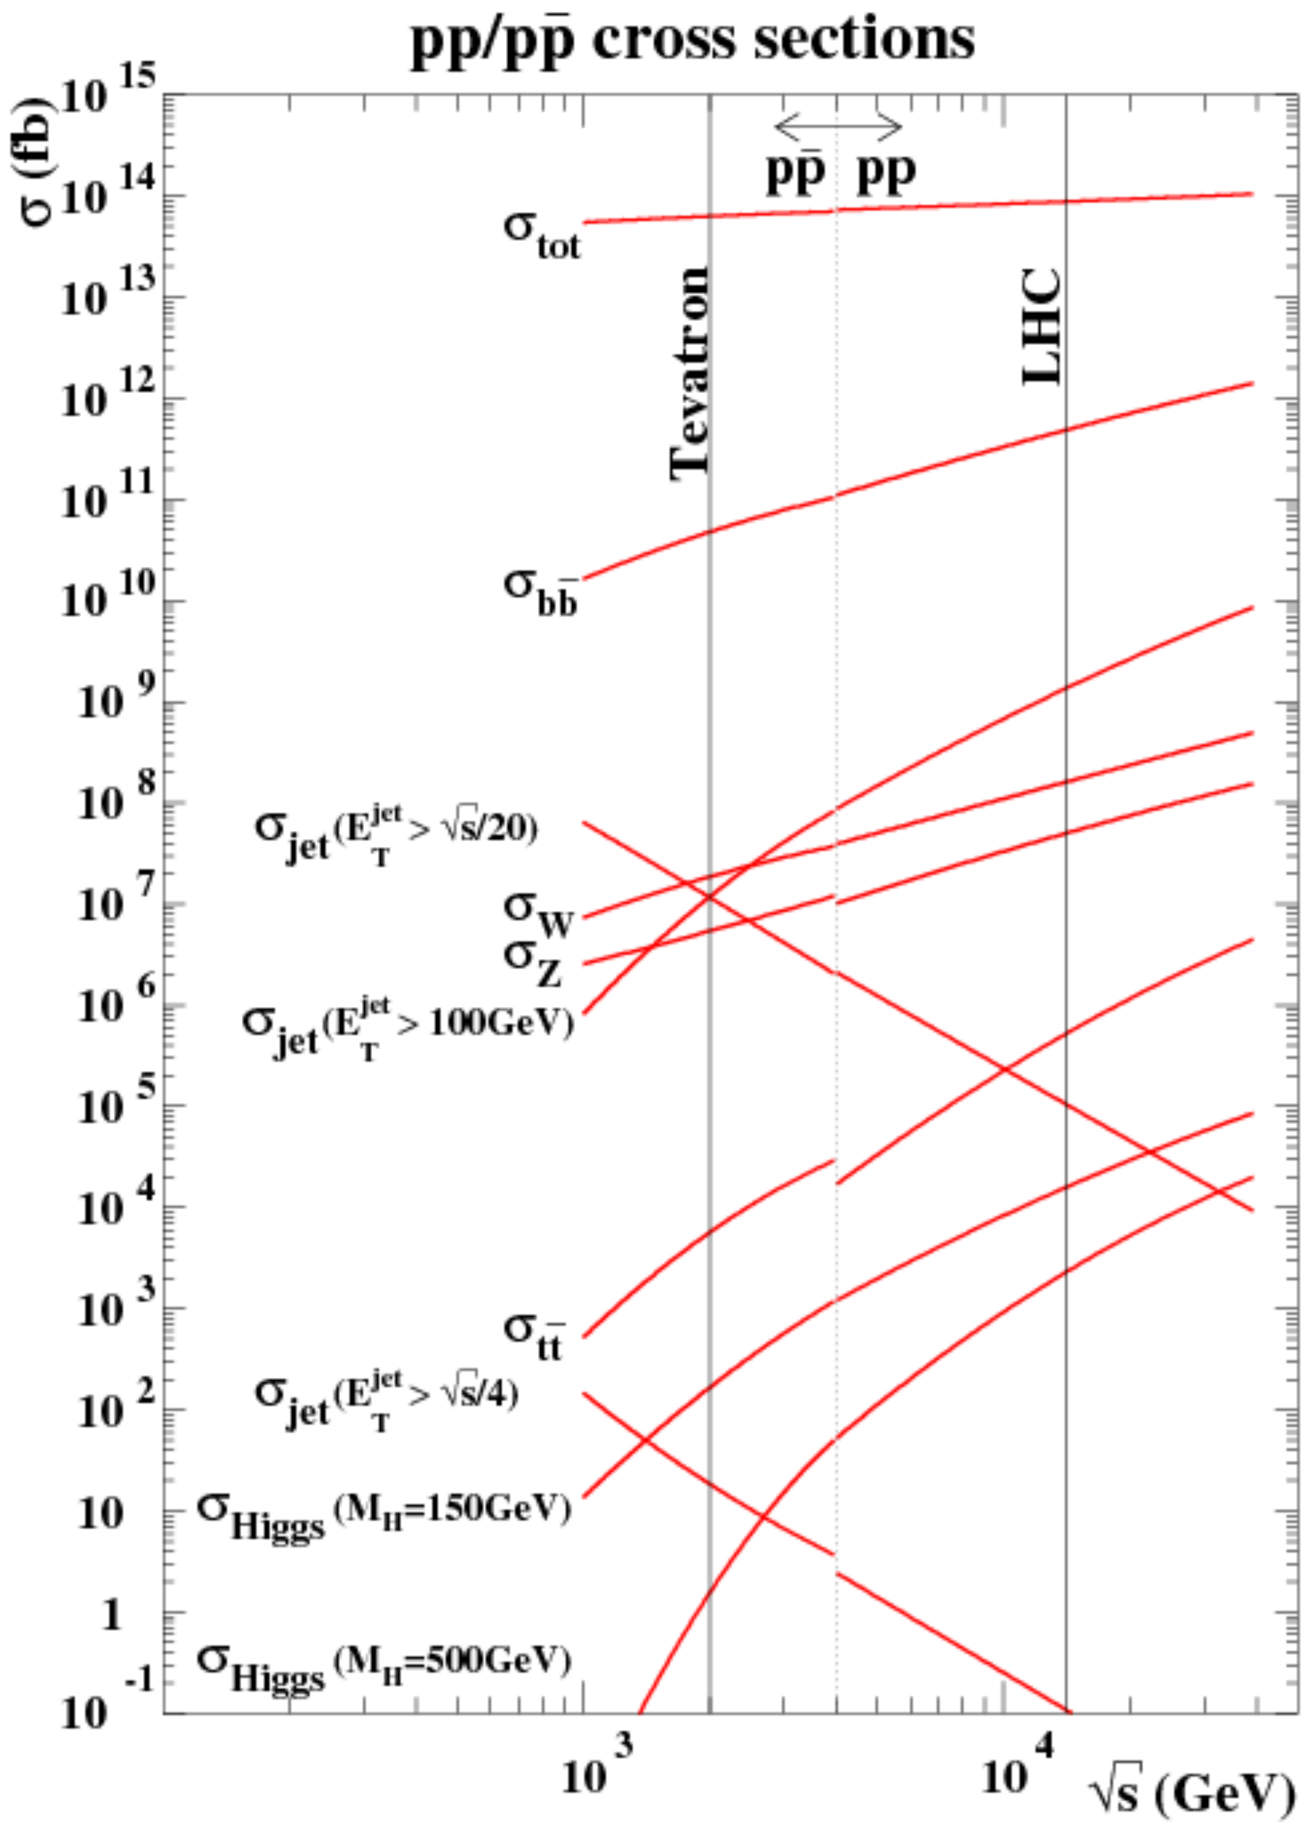

Source: CMS/*Phys. Rep.* 426 (2006) 47-358

Triggering

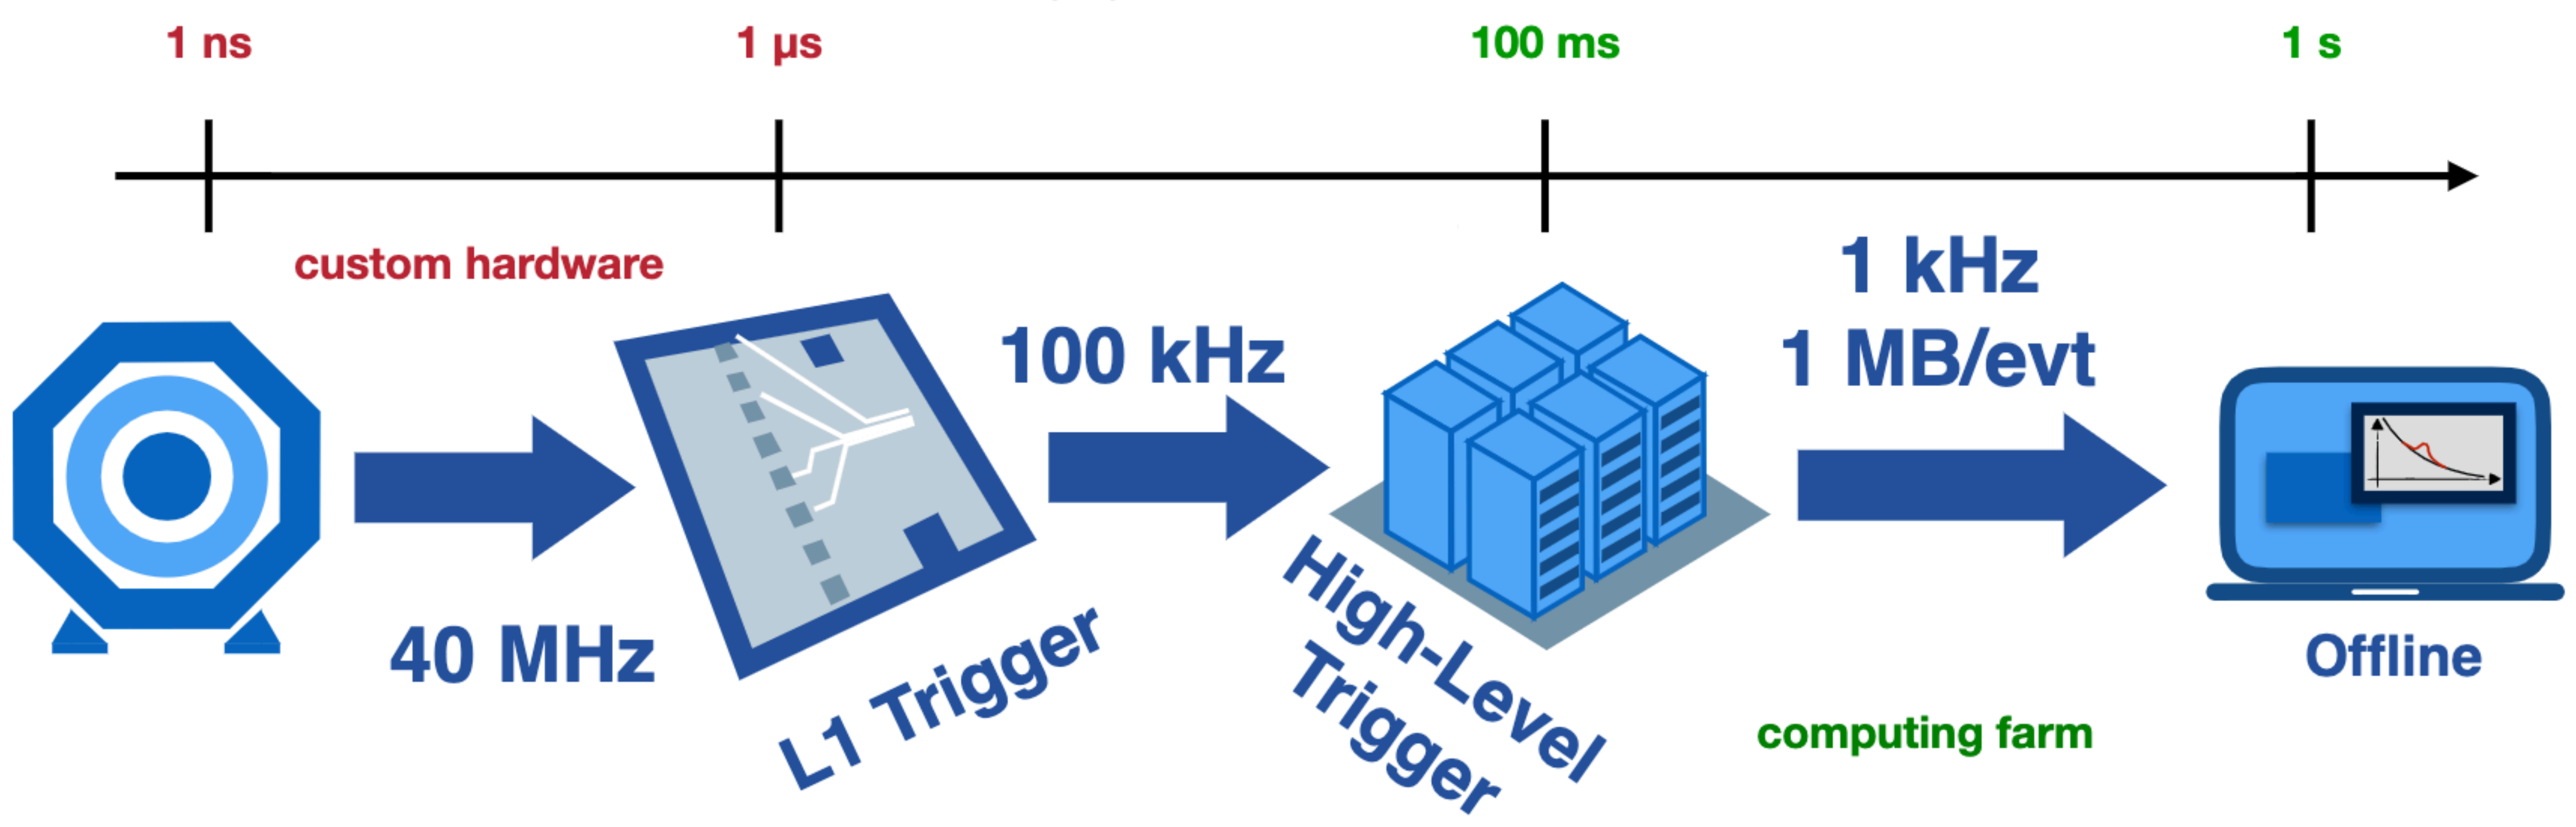

Trigger run very complex physics tests to look for specific signatures, for instance matching tracks to hits in the muon chambers, or spotting photons through their high energy but lack of charge.

- L1 Trigger: electronics + custom chips + FPGAs processing partial and raw data
- High-Level Trigger: PC farm processing the full data obtaining good resolution physics objects
- Offline: fine grained physics object reconstruction

Source: M.Pierini

Trigger and DAQ

- Trigger path: for event selection
- Data path: data sample, data collection, and storage of the selected events

Variations: LHCb recently introduced a trigger less readout; no need for hardware trigger

Subdetectors

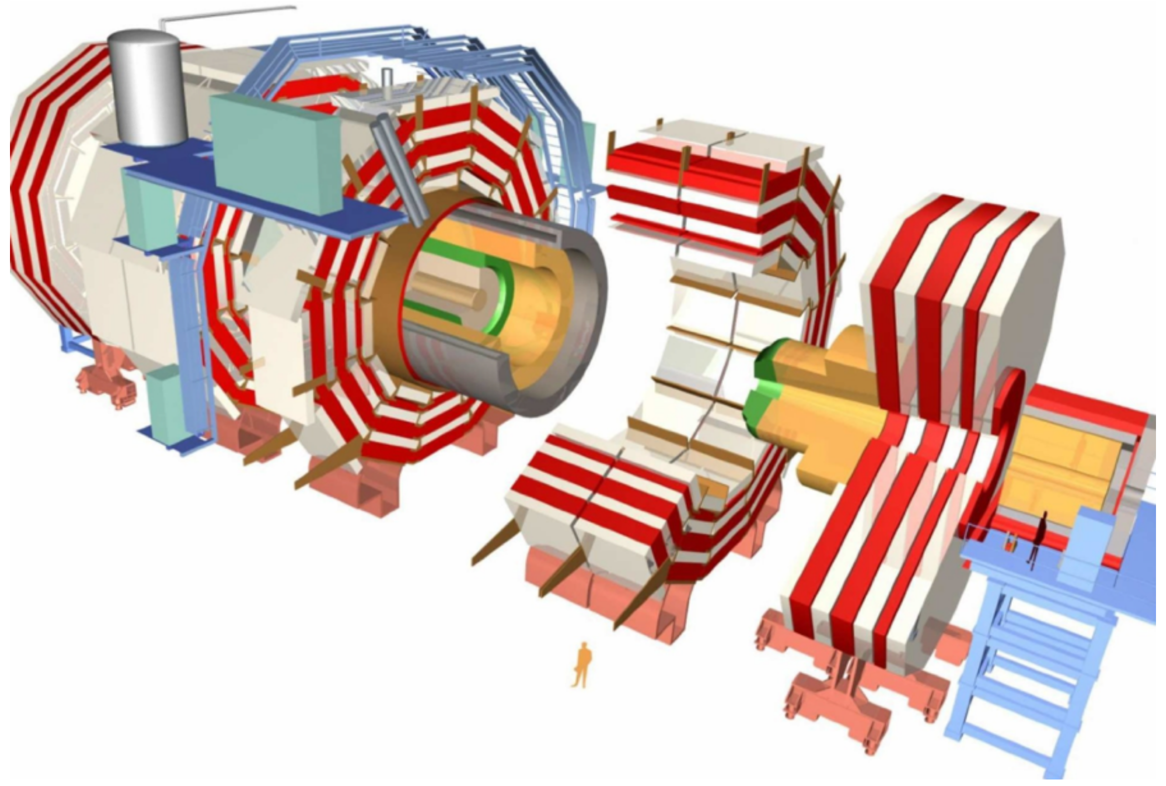

Simplified and generic Trigger and DAQ schematics

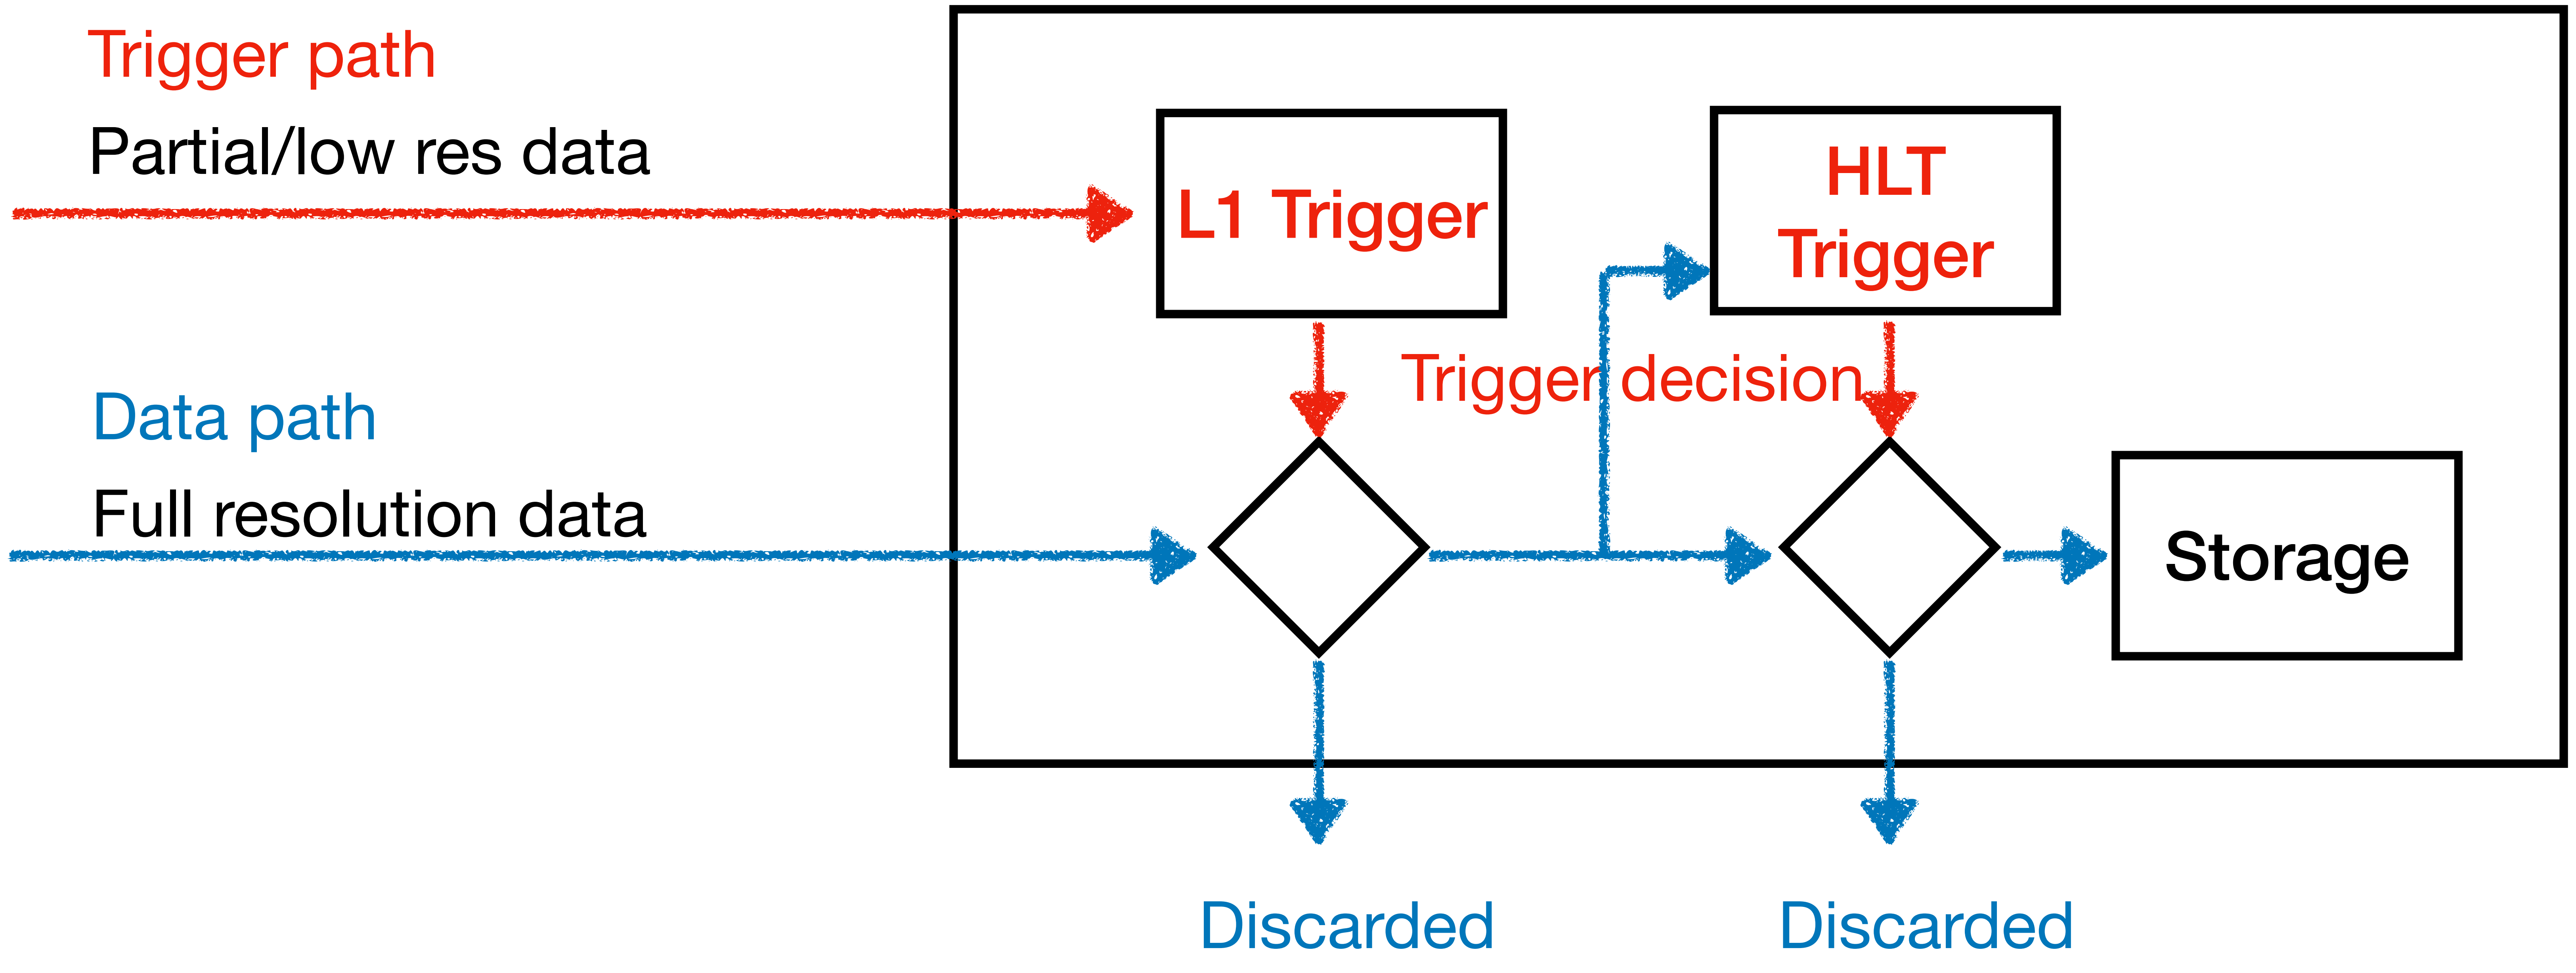

Source: CMS

L1 Trigger

- Selects based on raw data from the sub detectors and a **coarse granularity** information
- Fast, based on **custom electronics** (front-end part) and **FPGAs**
- In the current detector it accepts events at a **rate** of about 100KHz (**rejecting 99.75%** of events)
- Multiple trigger "path" are implemented in order to select **physics "objects"** such as photons, electrons, muons, jets, missing E_T
  - Information based on the muon chamber sub-detector and energy deposits in the calorimeters
  - Each path can be based on multiple objects, e.g. double photons to trigger on Higgs->γγ

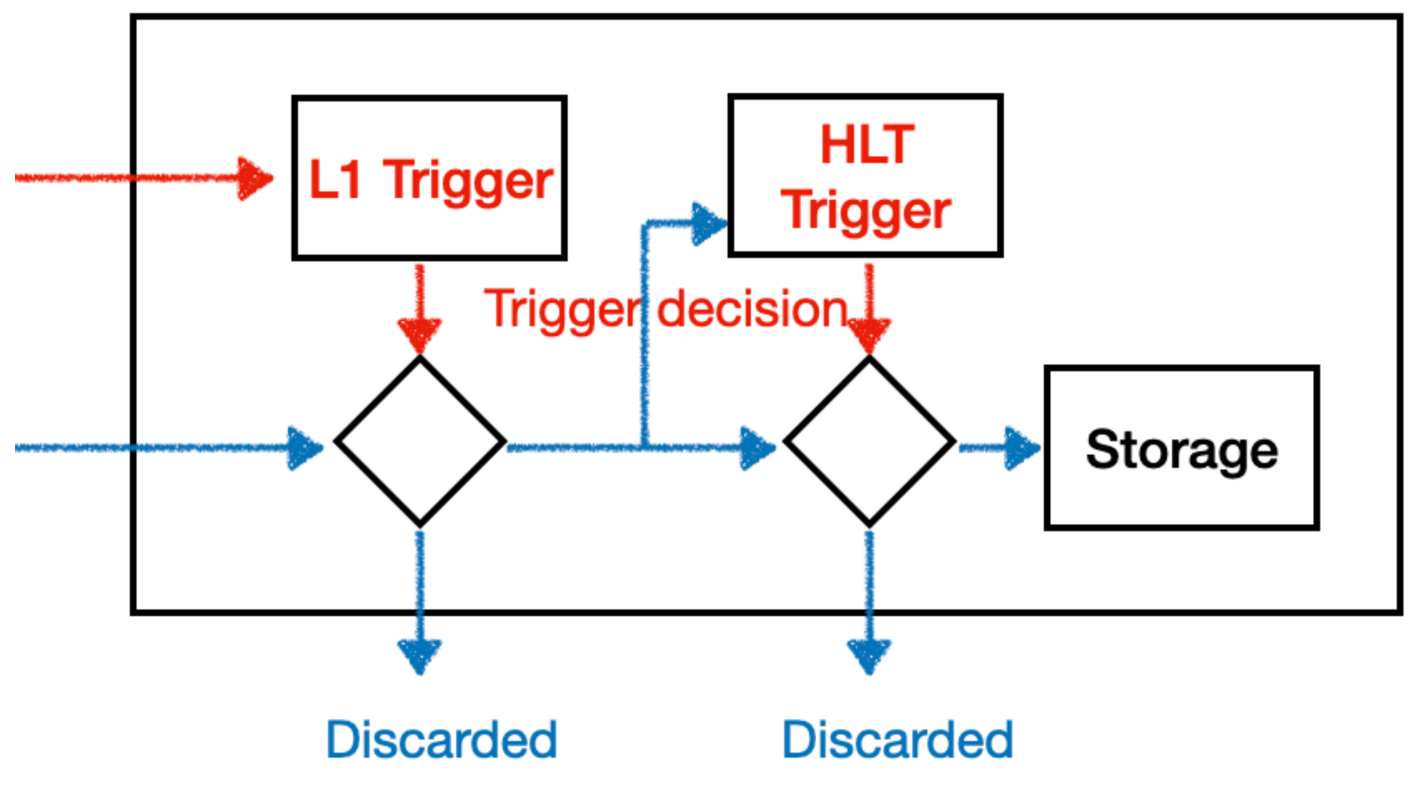

High Level Trigger (HLT)

- Selects based on the full detector data with a **good granularity**, but often less granular than the offline object reconstruction
- Relatively slow O(100) ms process done in PC farms
- In the current detector it accepts events at a **rate** of about 1KHz (rejecting 99% of events)
- Multiple trigger "path" are implemented in order to select **physics "objects"** such as photons, electrons, muons, jets, missing E_T and their correlations
  - Try to find the signature of interesting **physics processes**, e.g. Higgs decays, top decays, beyond the Standard Model particles, etc

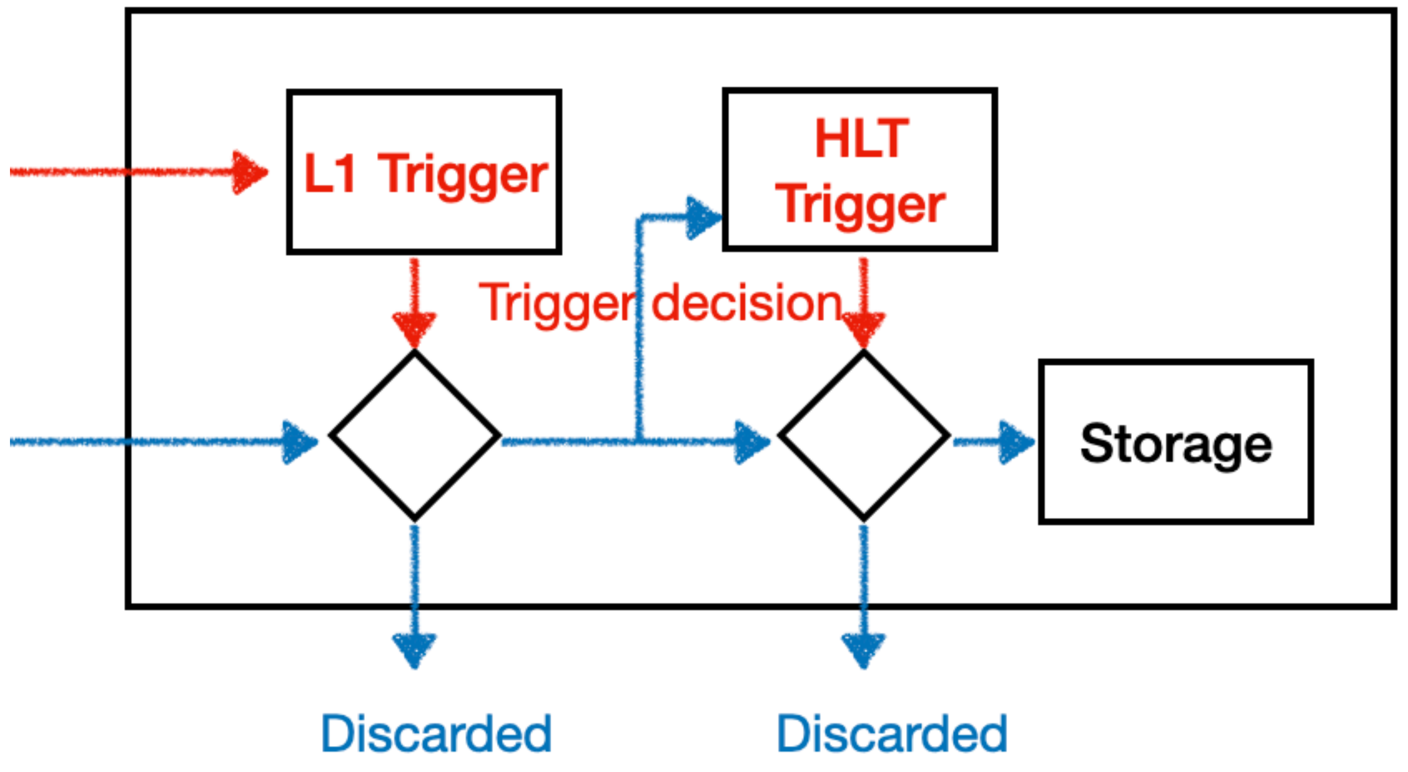

Data Path

- Data collection starts in the electronics of the sub detectors
  - Signals from **sensitive material** (calorimeter crystals + PMT, tracker silicons, etc) is **readout**
- Data is sent to the “back-end” electronics in case the L1 trigger gives the green light
- Data from each sub detector is routed by a commercial **switched network** to combine all the sub detector info
- Data is used by **HLT** to make the last trigger decision
- Data with HLT green light is saved into **local storage**
- Data is used for the offline event reconstruction and distributed to multiple data centres - Grid - for data analysis

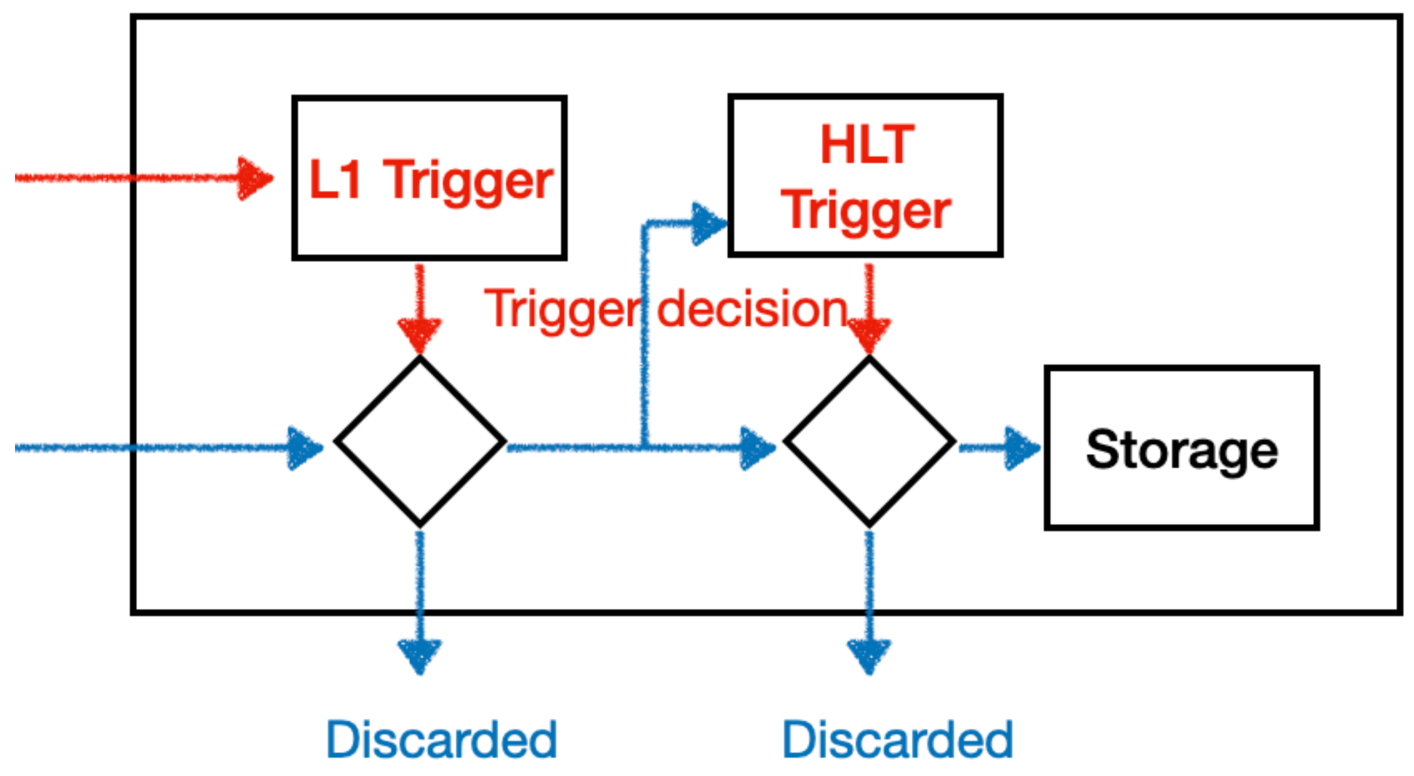

Alternative ways of DAQ

- **Data scouting**
  - Reduced the **event size** by saying HLT information
  - Avoids the HLT data buffering bottleneck
- **Data parking**
  - Exploits the computational margin arising from the LHC fill luminosity decay
  - Event are **reconstructed later** in time (normally “prompt” reconstruction is done after 48h) and saved on **tape**
  - Avoids the reconstruction resources bottleneck

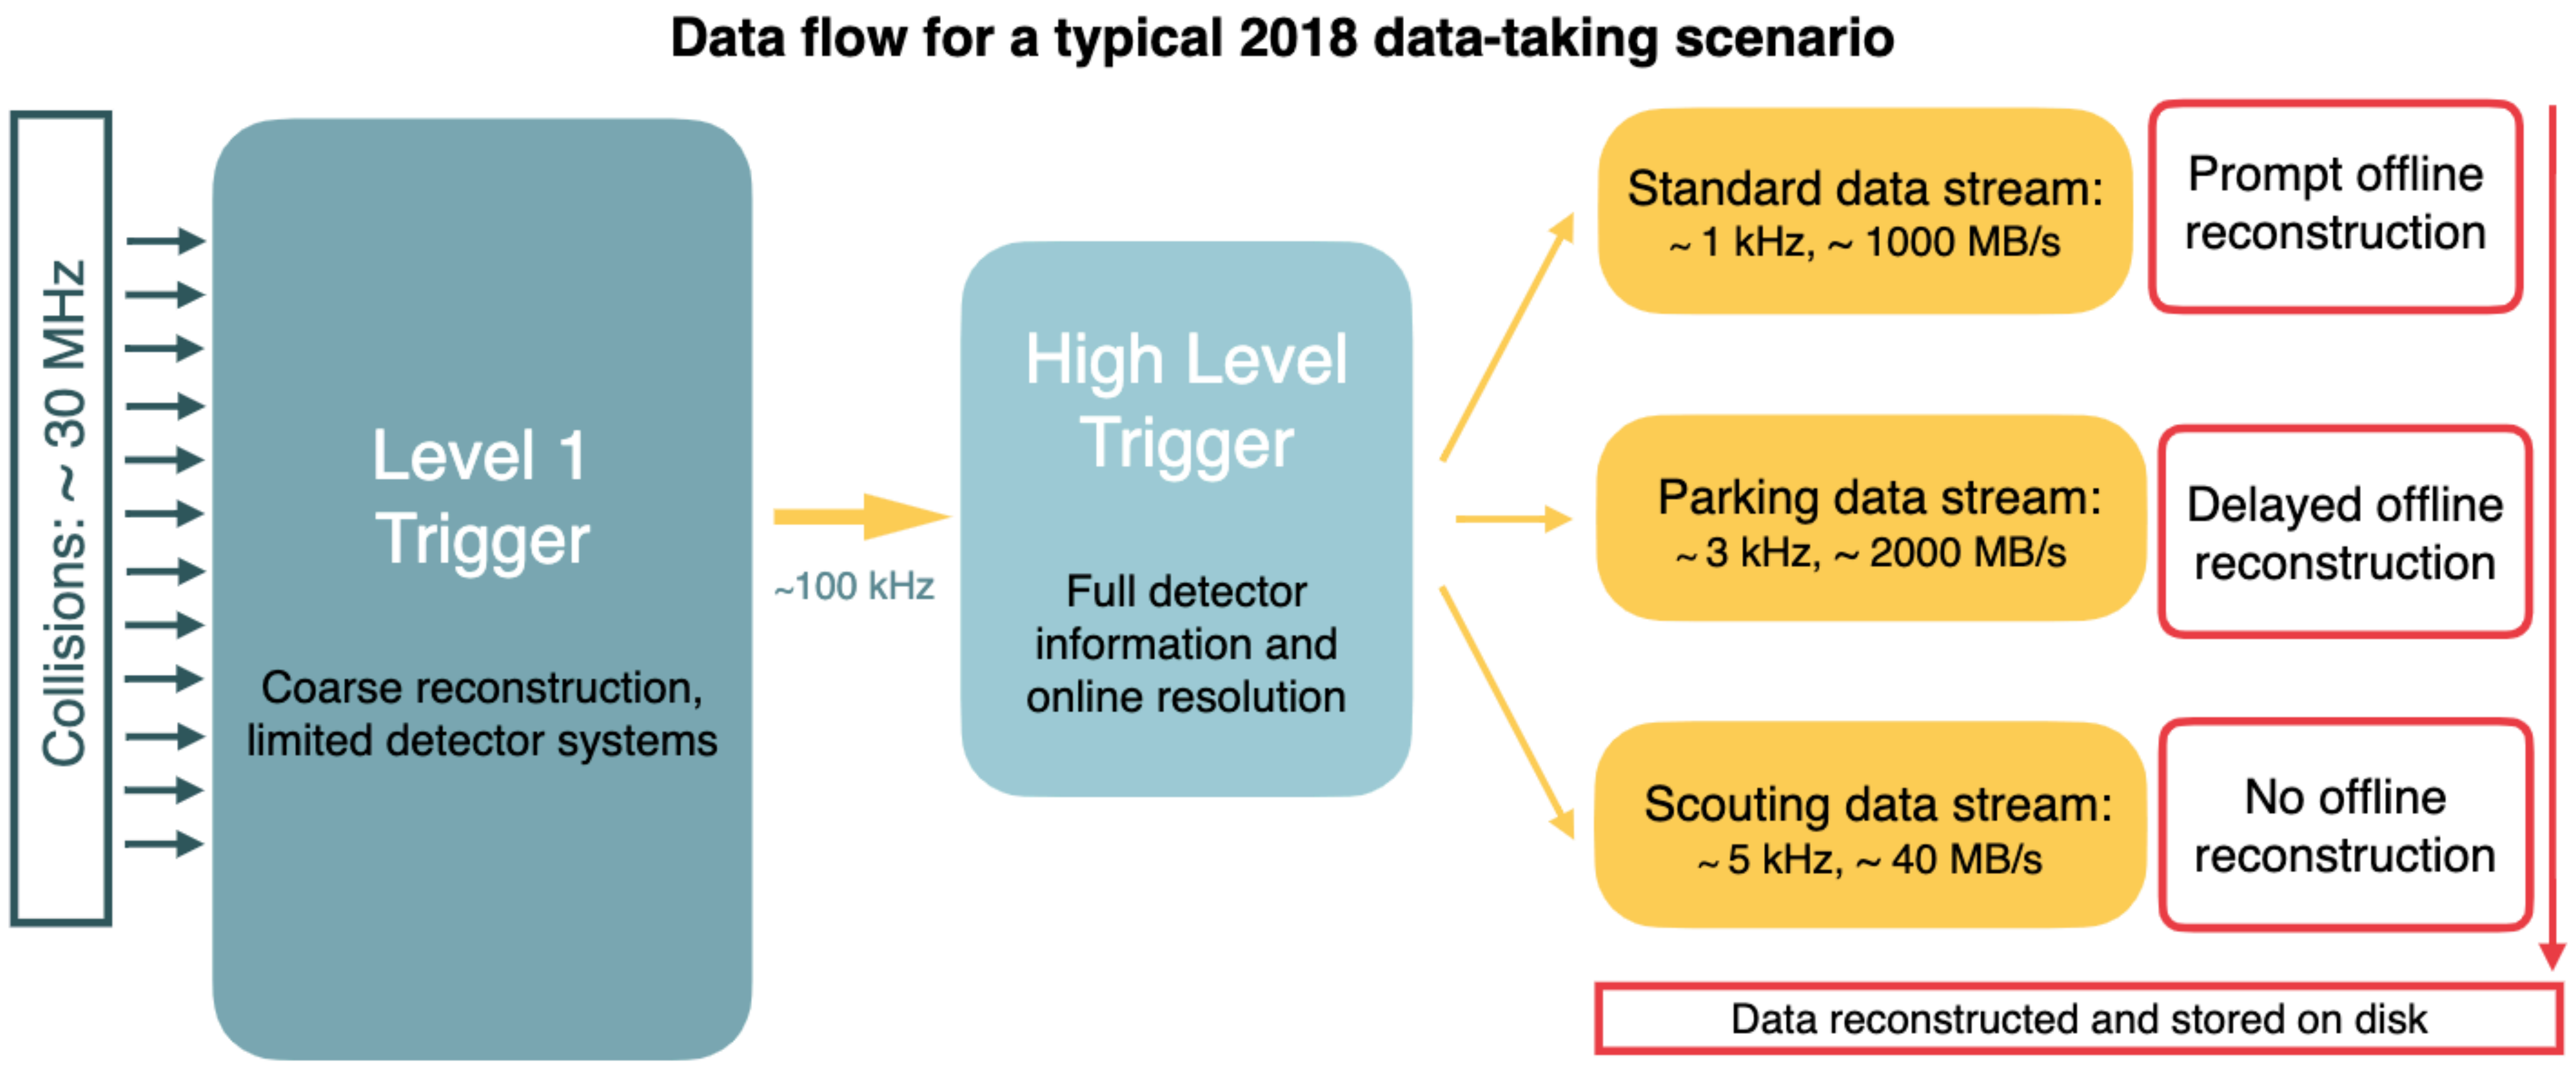

arxiv2403.16134

Data acquisition in CMS during LHC upgrade

Source: CMS

HL-LHC Plan

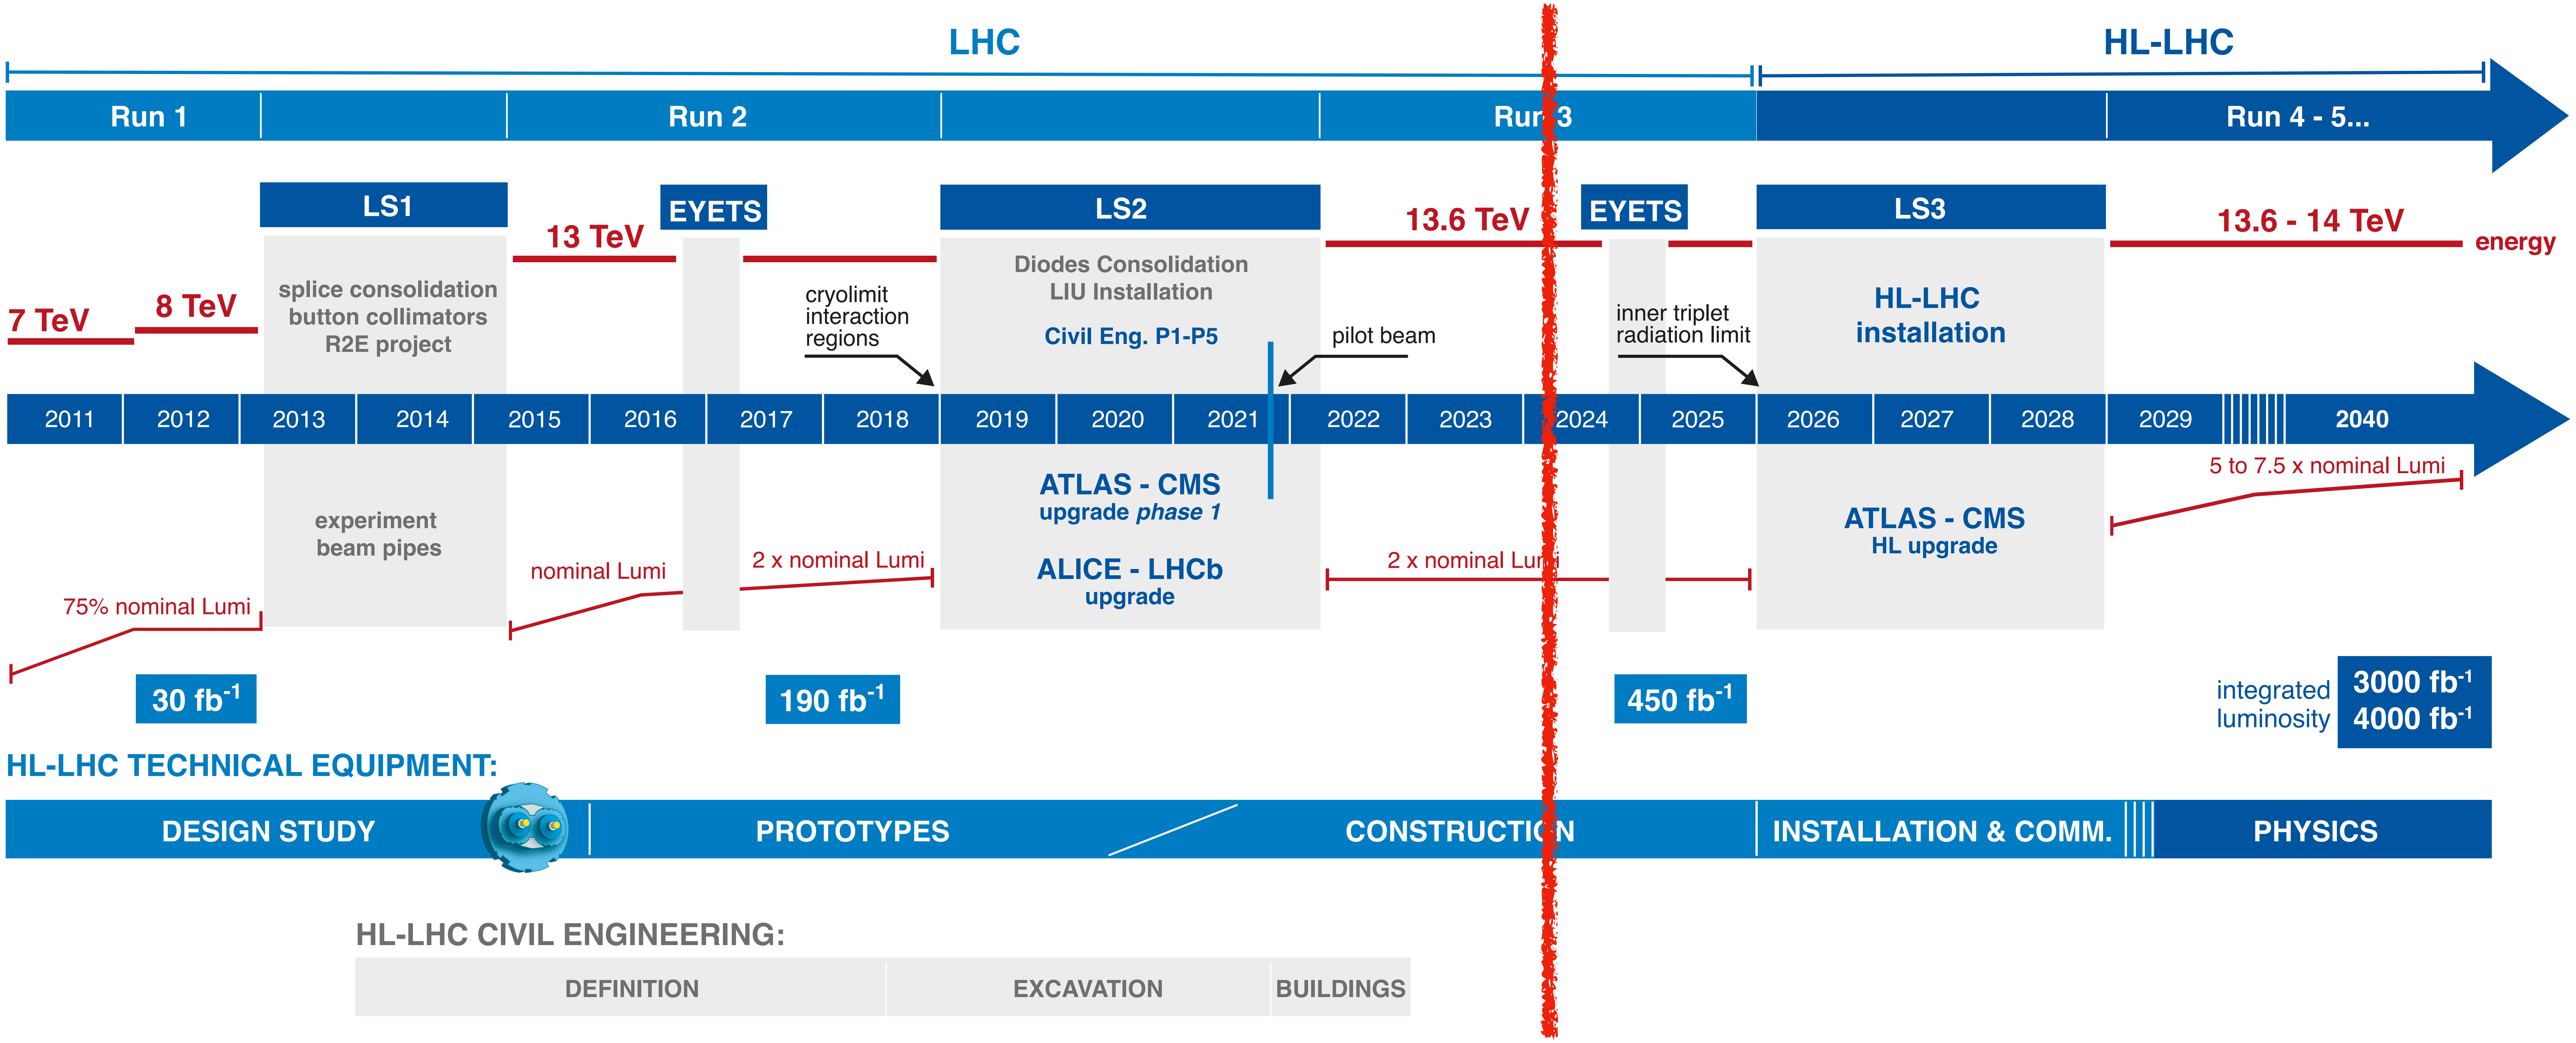

Increase of number of collisions -> increase of event complexity and radiation -> need new sub-detectors

Source: LHC

HL-LHC CMS detector overview

L1 Trigger HLT/DAQ
- Tracks in L1-Trigger at 40 MHz
- PFlow selection 750 kHz L1 output
- HLT output 7.5 kHz
- Latency within 12.5 μs
- 40 MHz data scouting

Muon Systems
- DT & CSC new FE/BE readout
- RPC back-end electronics
- New GEM/RPC $1.6 < \eta < 2.4$
- Extended coverage to $\eta \simeq 3$

Endcap Calorimeter (HGCAL)
- 3D showers and precise timing
- Si, Scint+SiPM in Pb/W-SS

Barrel Calorimeter
- ECAL crystal granularity readout at 40 MHz
- with precise timing for e/γ at 30 GeV
- ECAL and HCAL new Back-End boards

MIP Timing Detector
- Barrel layer: Crystals + SiPMs
- Endcap layer: Low Gain Avalanche Diodes

Tracker
- Increased granularity
- Design for tracking in L1-Trigger
- Extended coverage to $\eta \simeq 4$

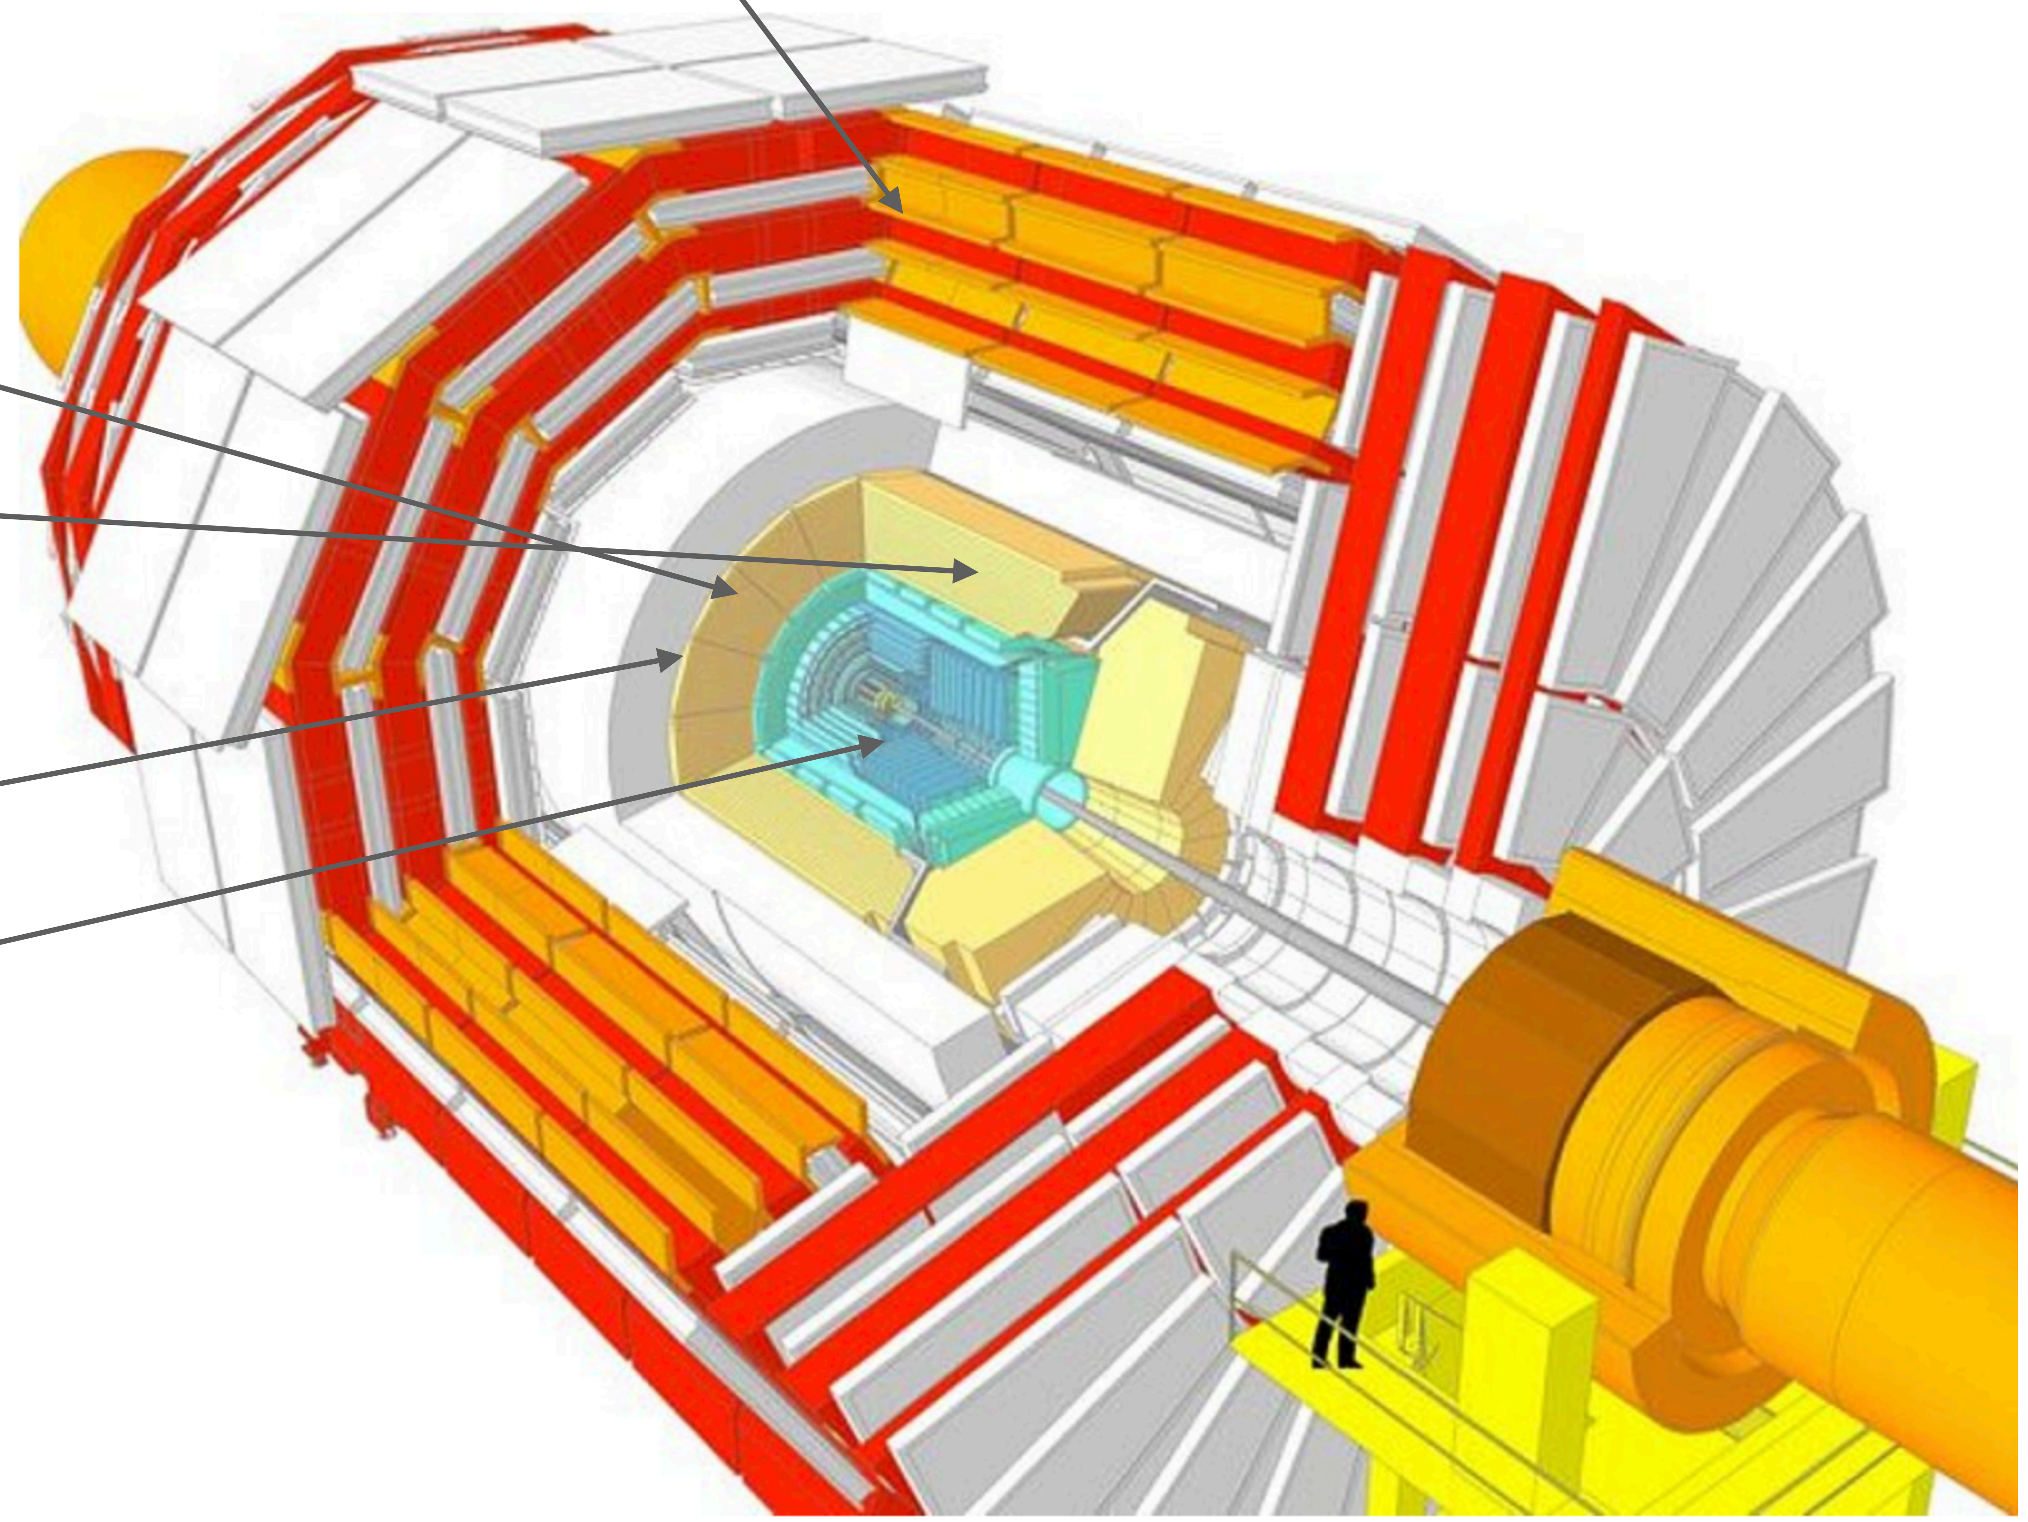

Source: CMS

CMS facility - P5

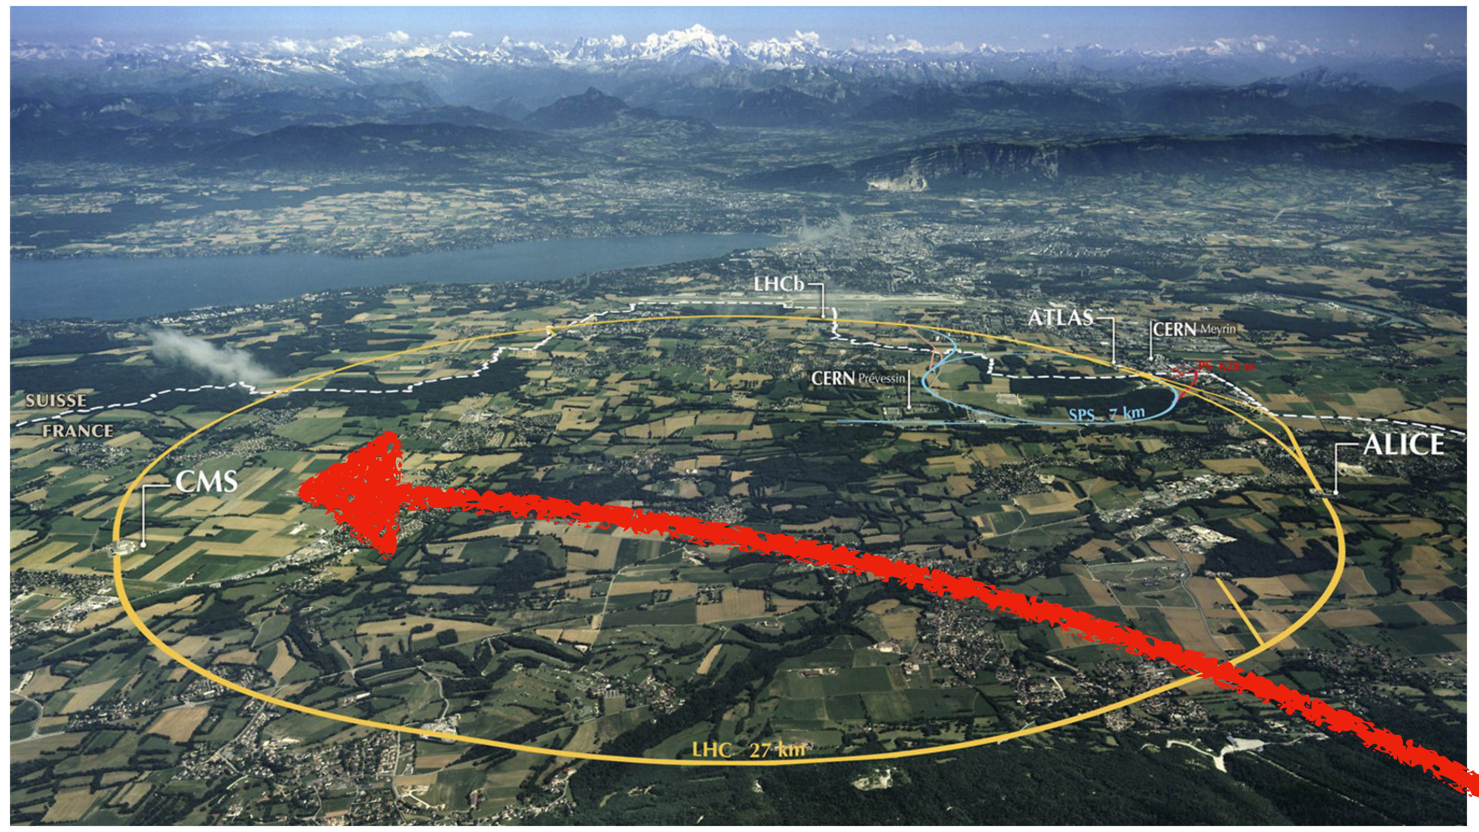

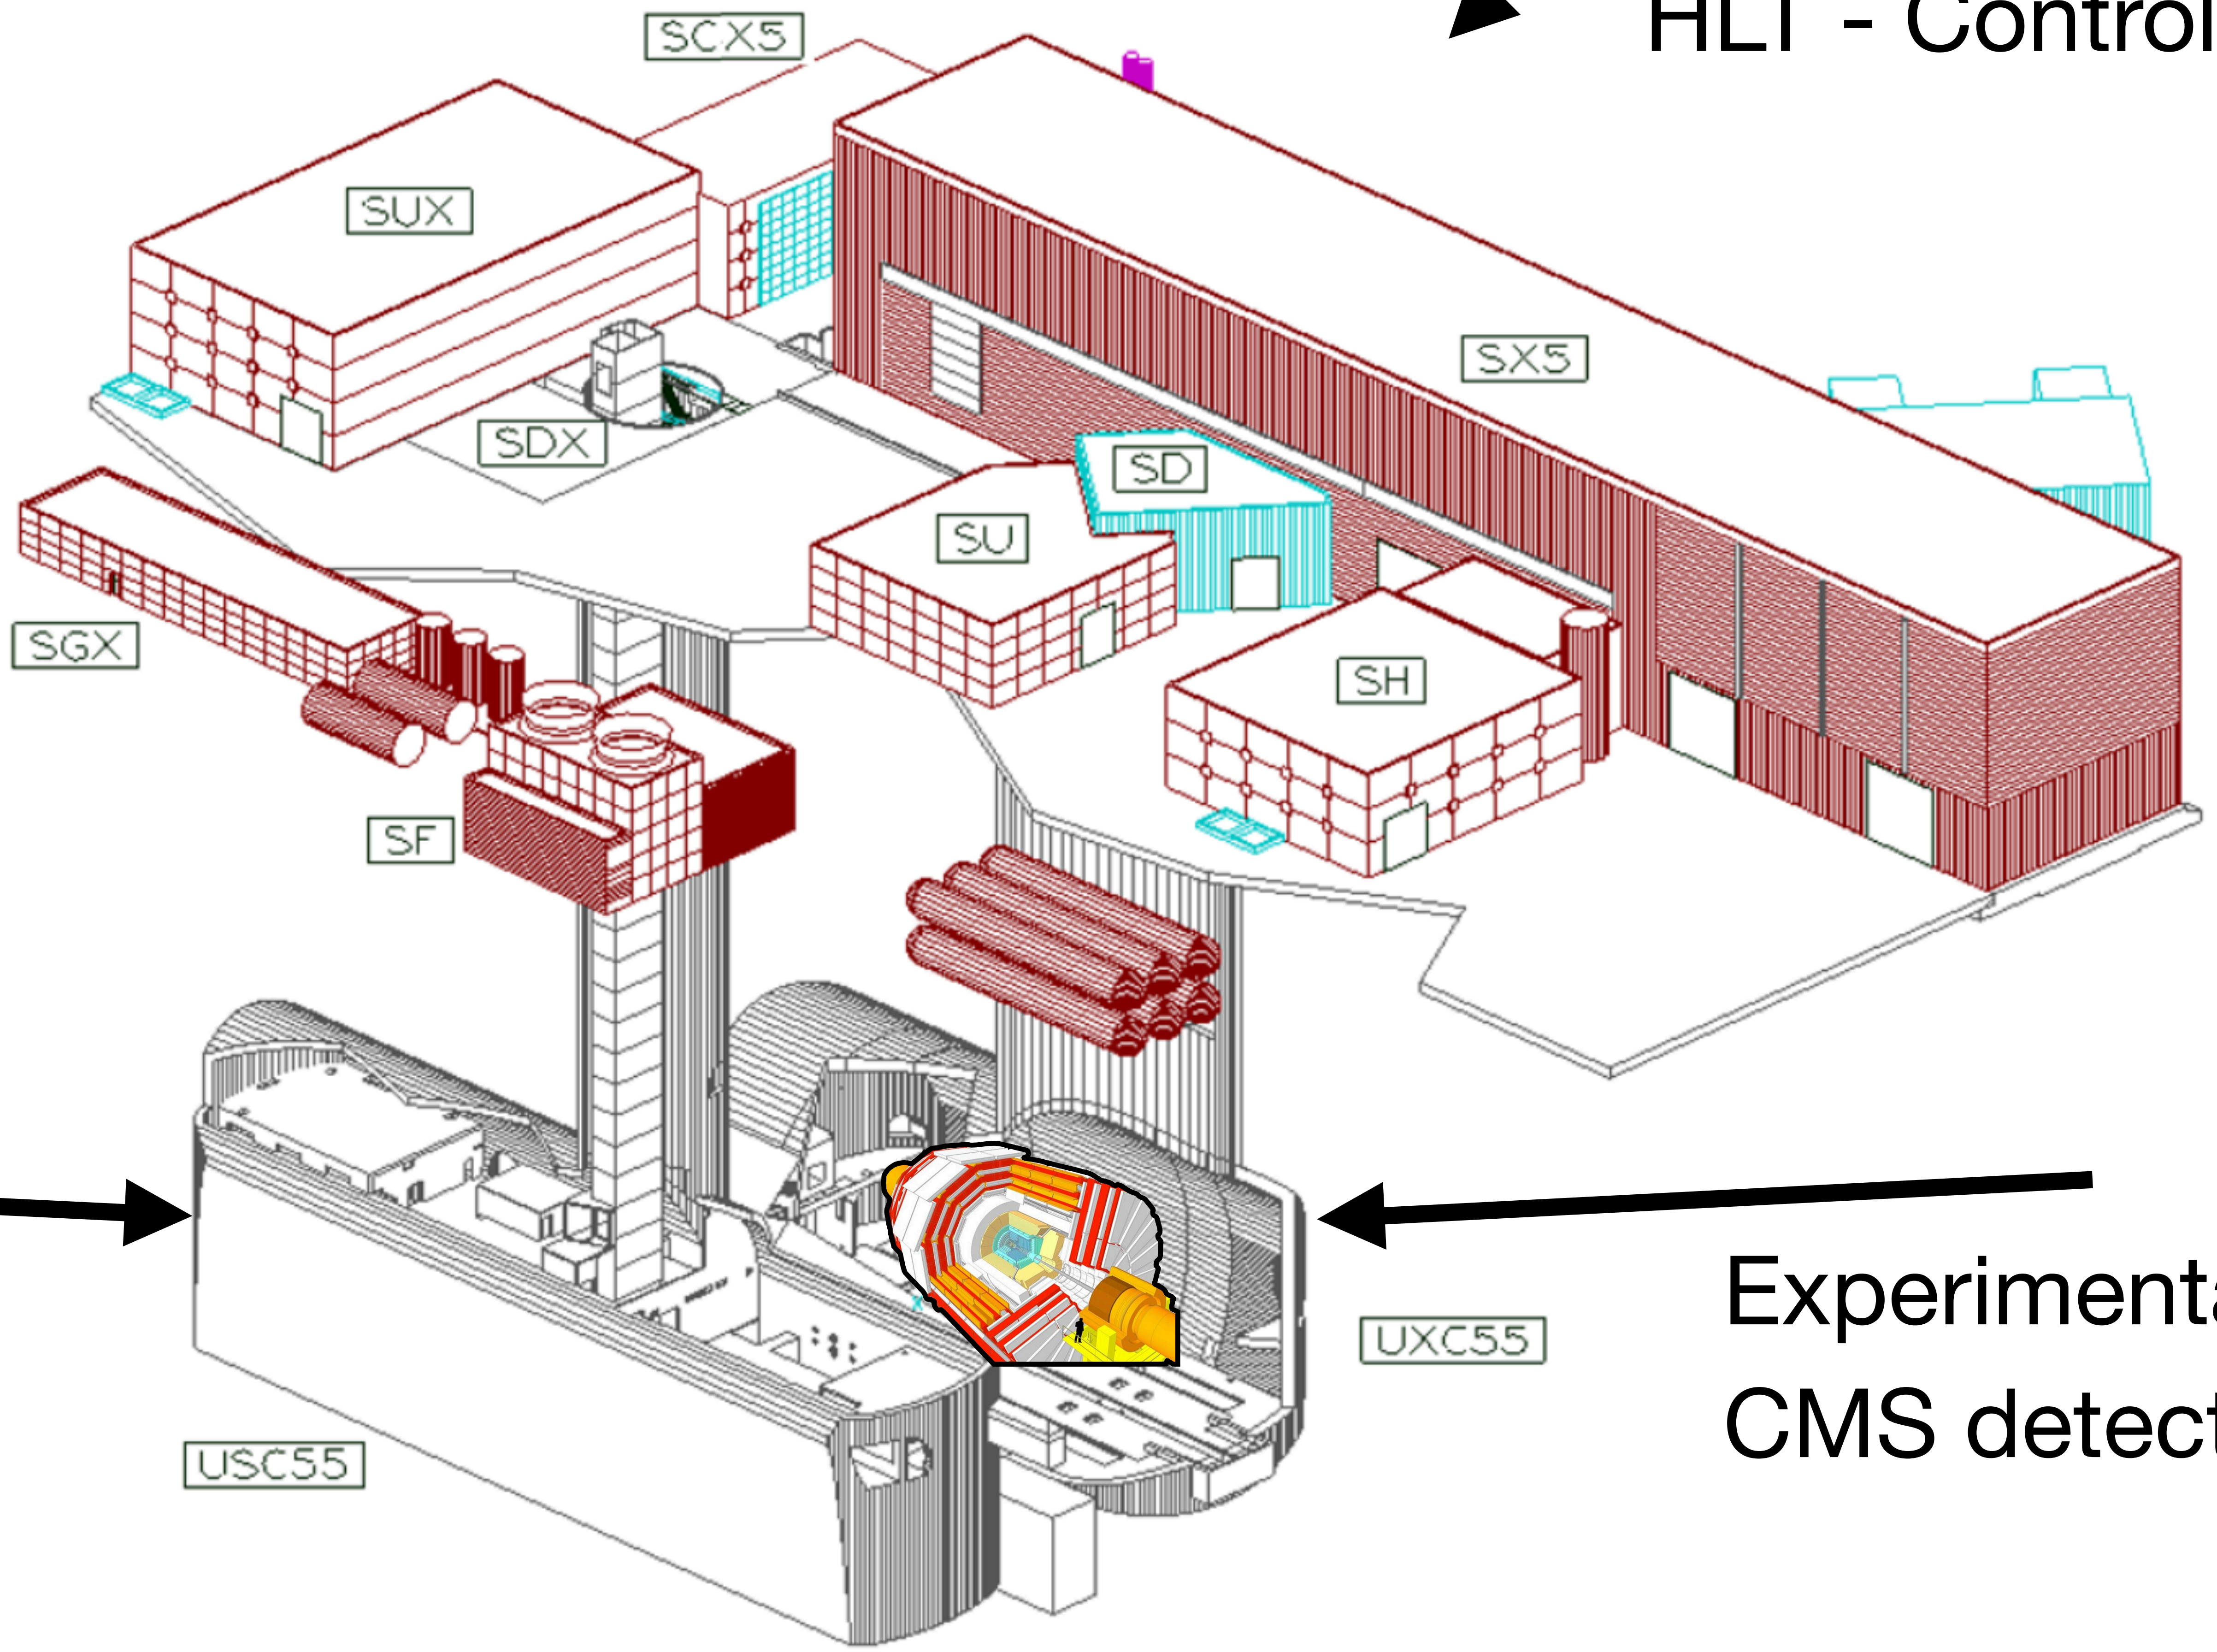

Surface
HLT - Control room

Service cavern
L1 Trigger - DAQ electronics

Experimental cavern
CMS detector

Source: CMS-TDR-022

HL-LHC - CMS T-DAQ

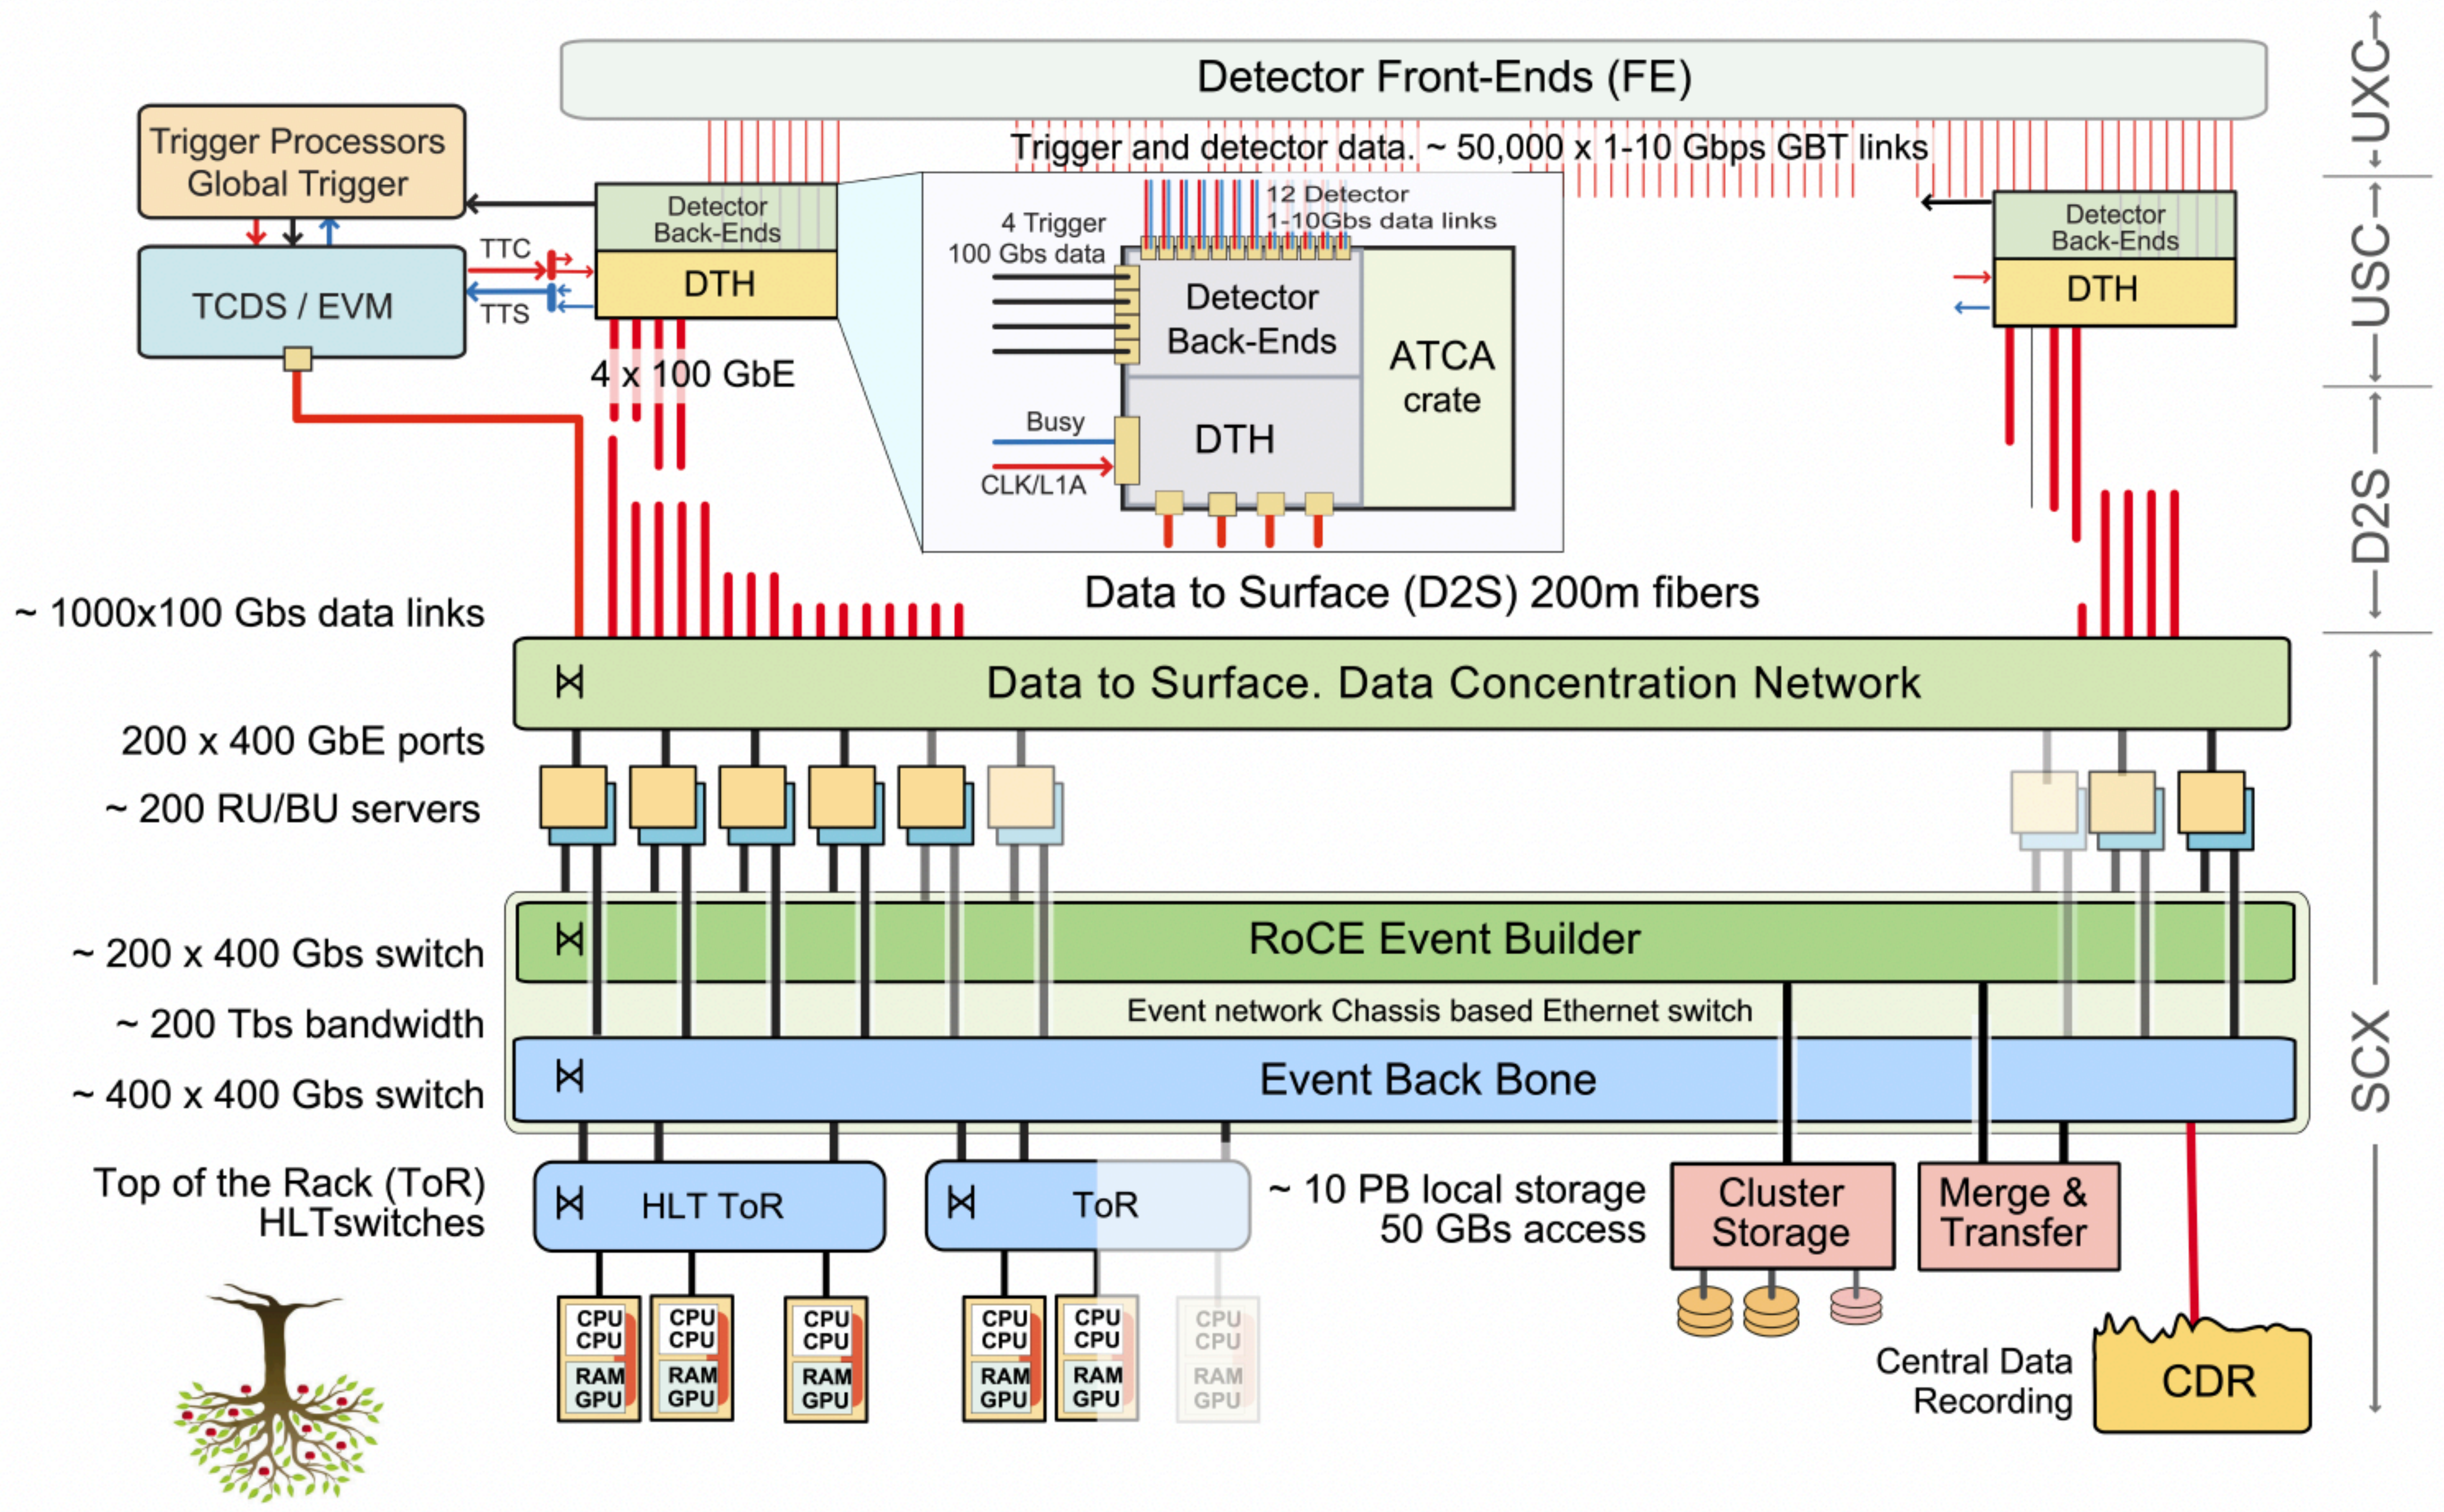

Source: CMS

HL-LHC - L1 Trigger

- Output trigger **rate 750KHz**
- Total maximum **latency 12.5 μs**
- Particle identification "particle flow" now implemented offline will be implemented in the L1 trigger
- 40 MHz data scouting
- Implemented in **ATCA boards** using a commercial **FPGA** (VU13P)
- Some algorithm based on machine learning models developed for FPGAs
- **Track reconstruction** at L1 (not available at present)

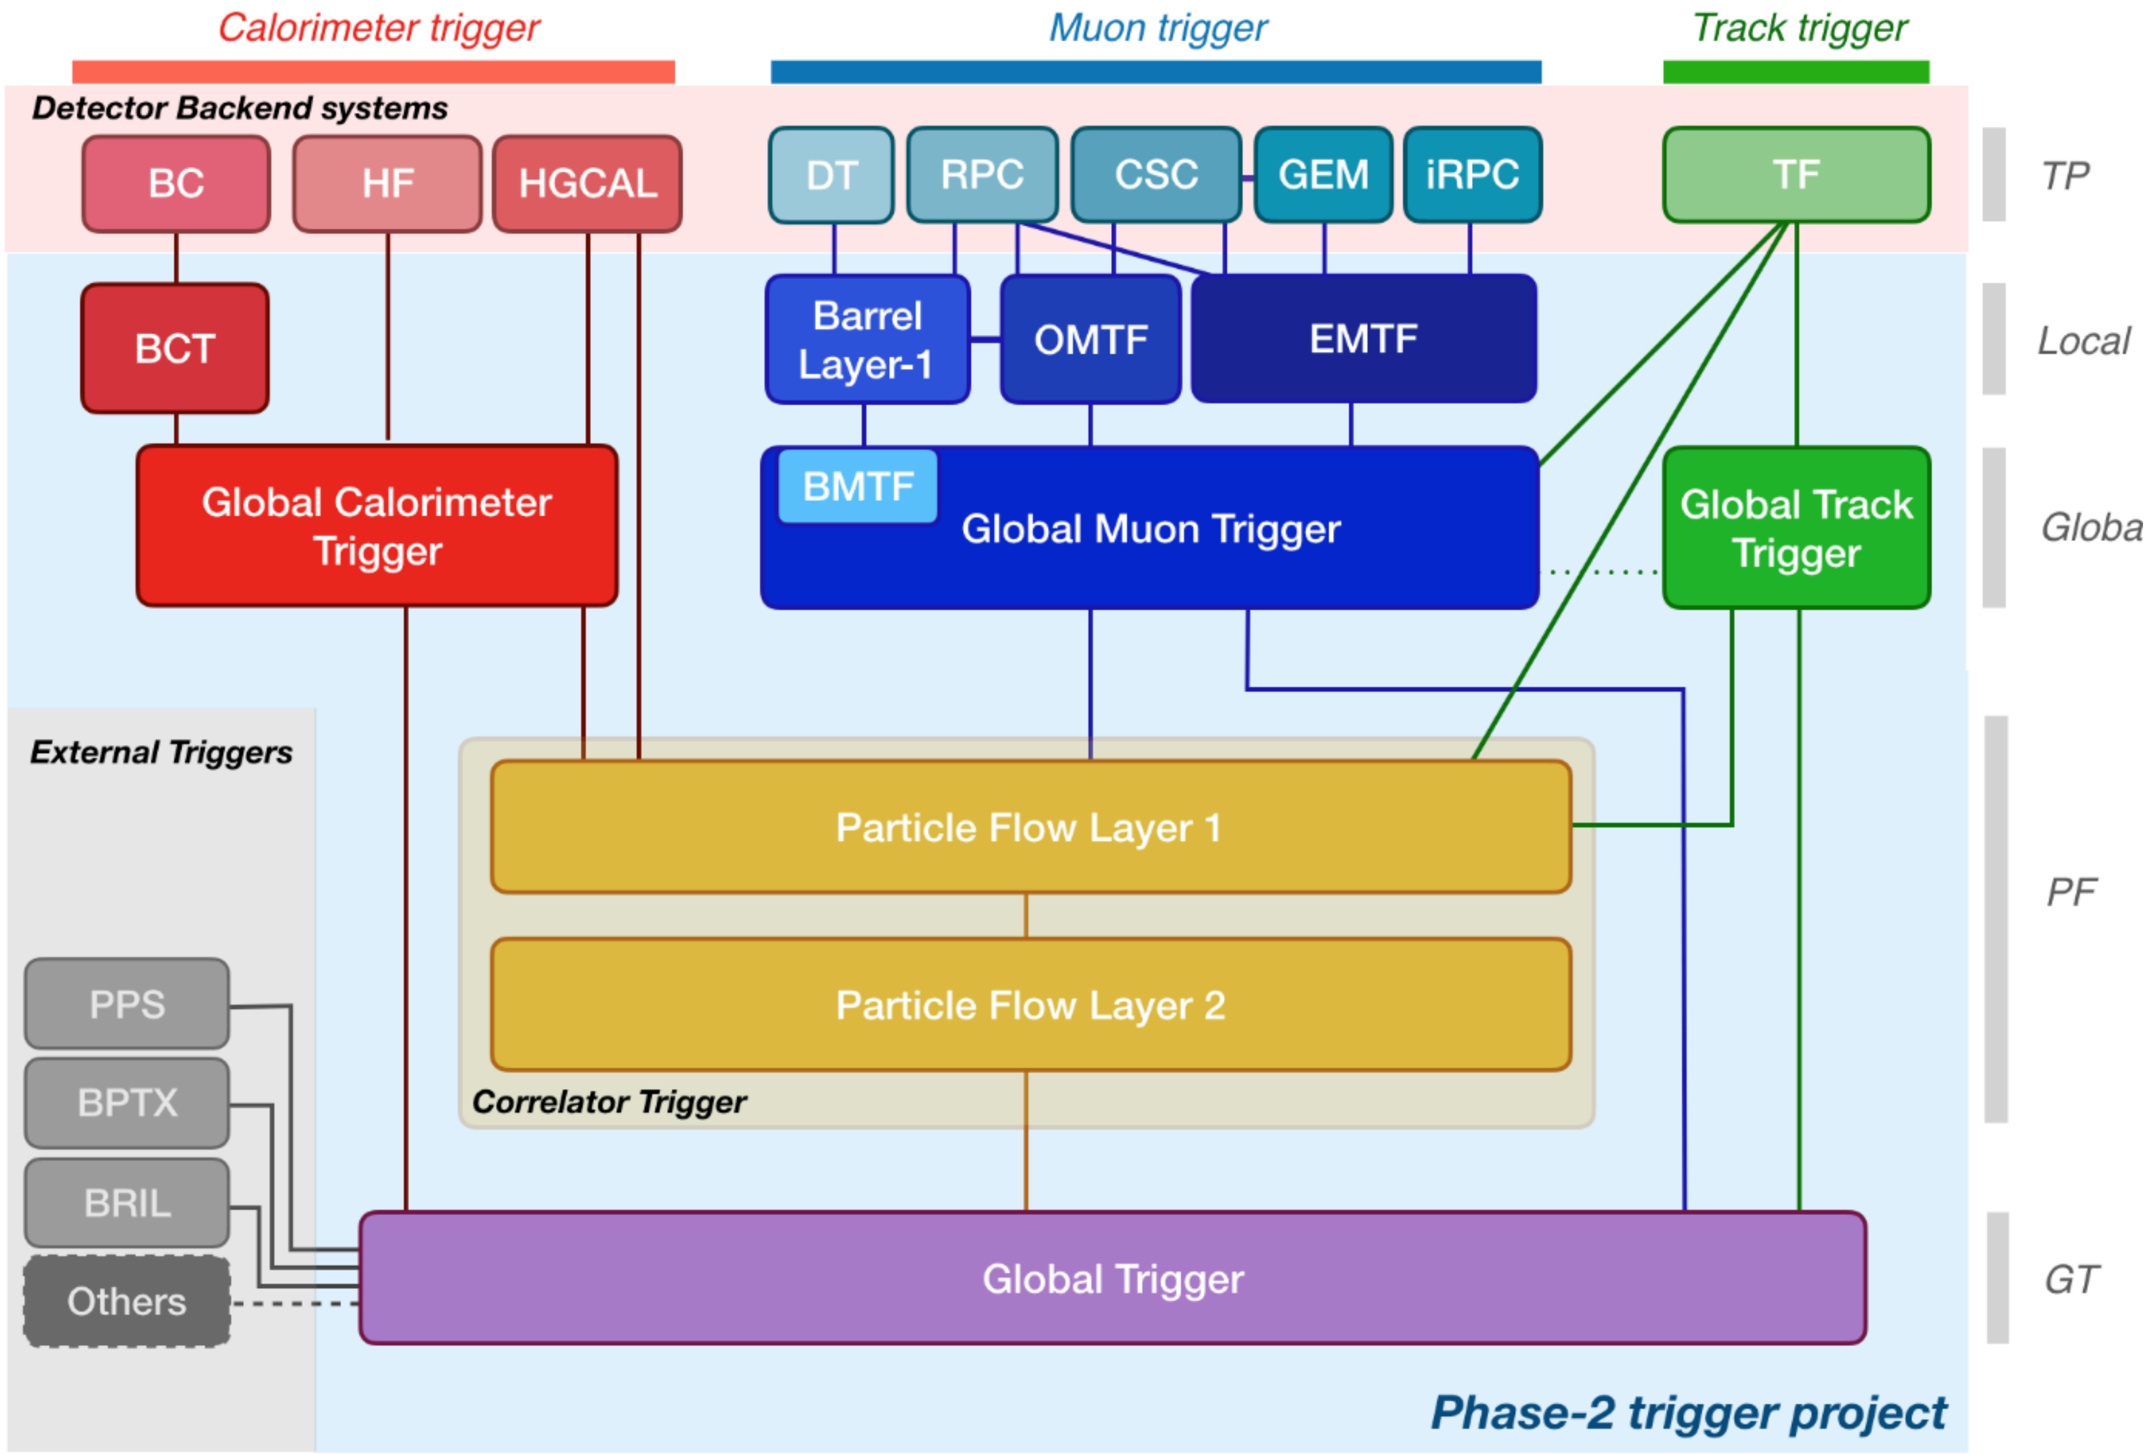

Source: CMS

FPGA 1/2

- FPGA: Field Programmable Gate Array
- FPGA is an integrated circuit which contains **basic elements** of digital (mostly) electronics: memories, lookup tables, flip-flops, etc… which can be combined to create complex digital circuits
- It's **programmable**, you can always change your implementation
- The circuit is implemented into the FPGA trough a circuit description file (in HEP often called "**firmware**")

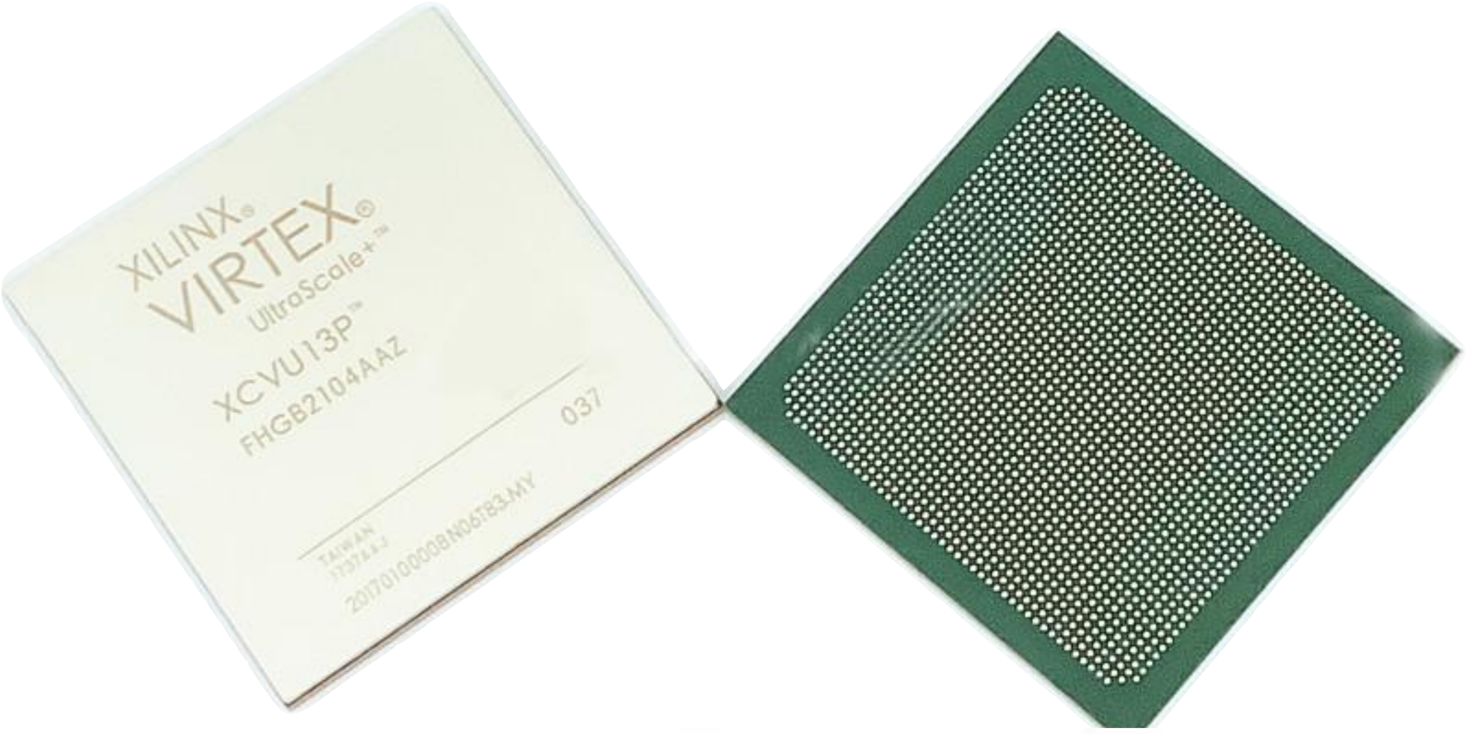

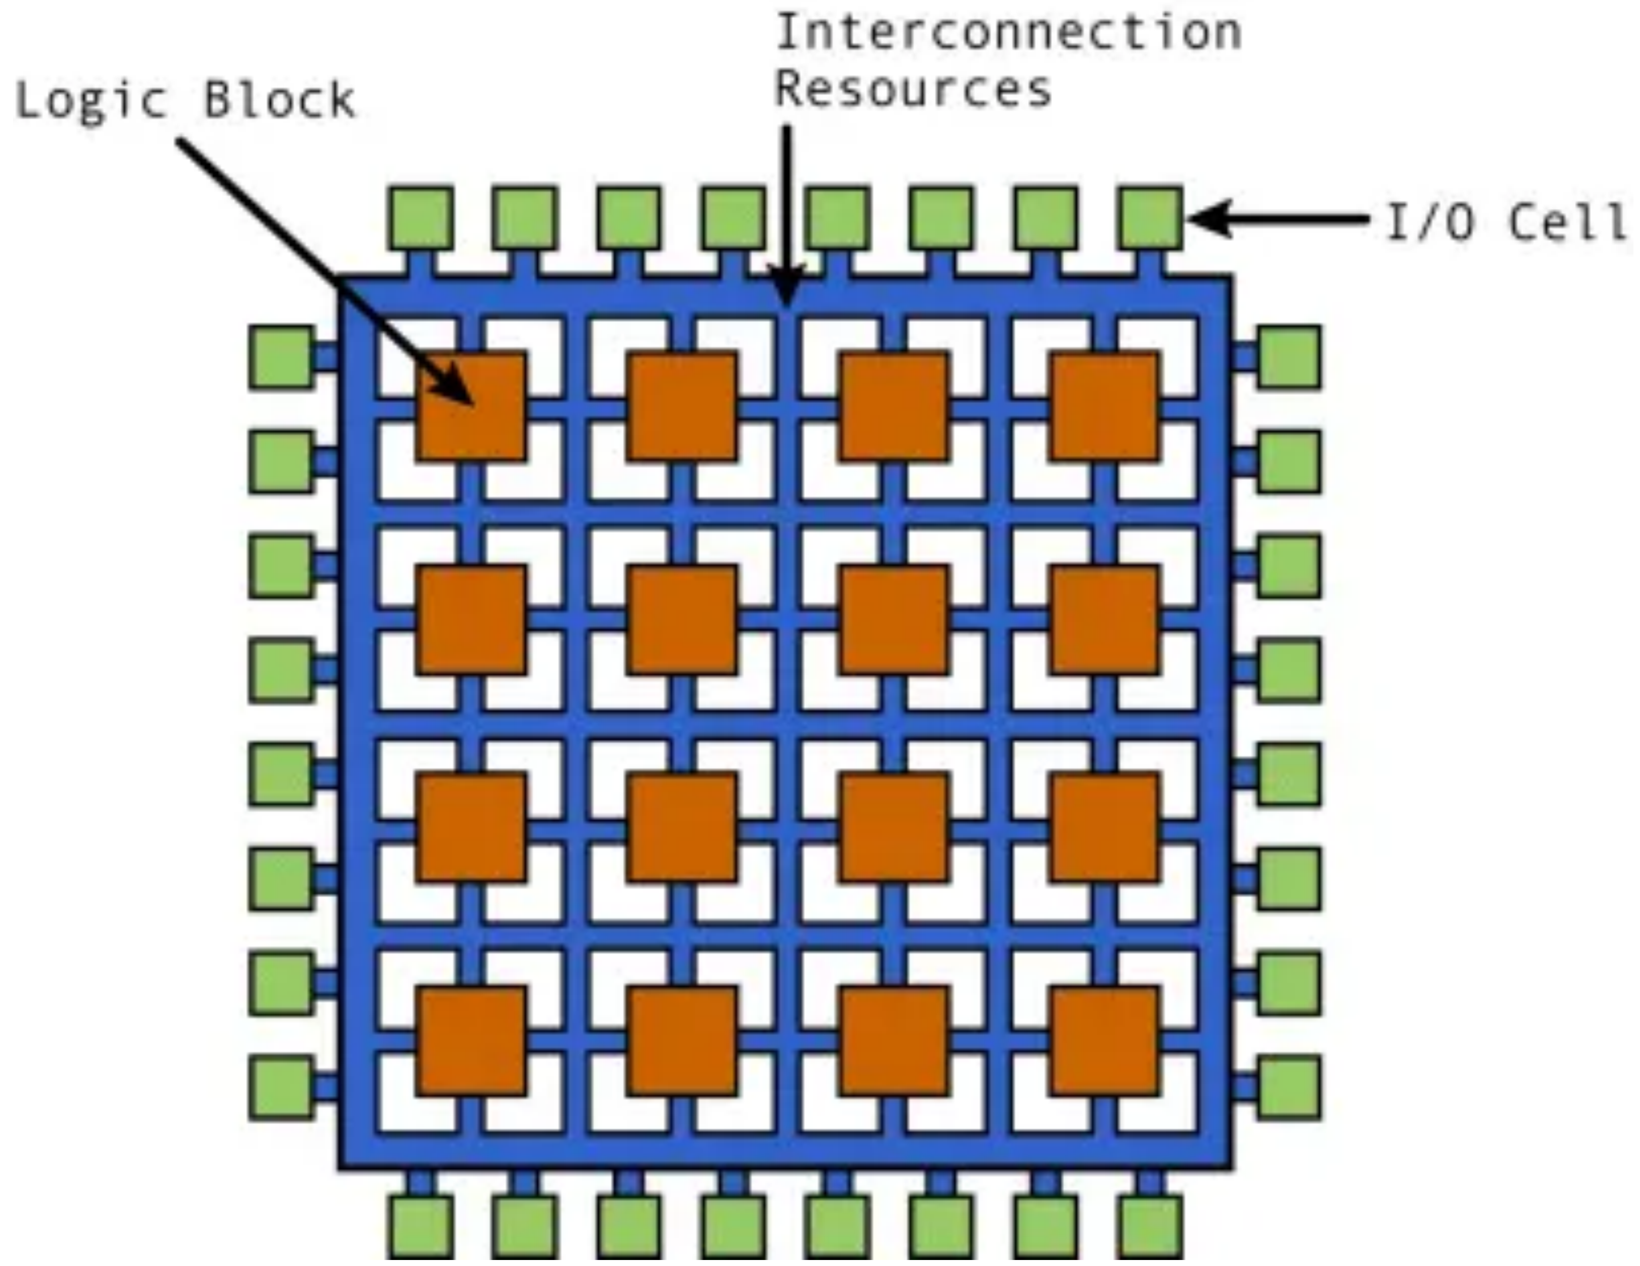

Source: Link

FPGA 2/2

- Firmware is generated from a hardware description language (e.g. VHDL)
  - **Algorithm -> Hardware circuit**
- Recently **high level languages** (C, java) have been interfaced to HDLs to allow a wider audience to program FPGAs, e.g. Xilinx HLS

- Benefits from FPGAs usage in HEP
  - **Flexibility**: the implemented circuit can be changed reprogramming the FPGA
  - **Parallelisation**: multiple data processing at the same time
  - **High data throughput**: each FPGA can connect more than 100 high speed links (up to 25 Gb/s per link in our case)

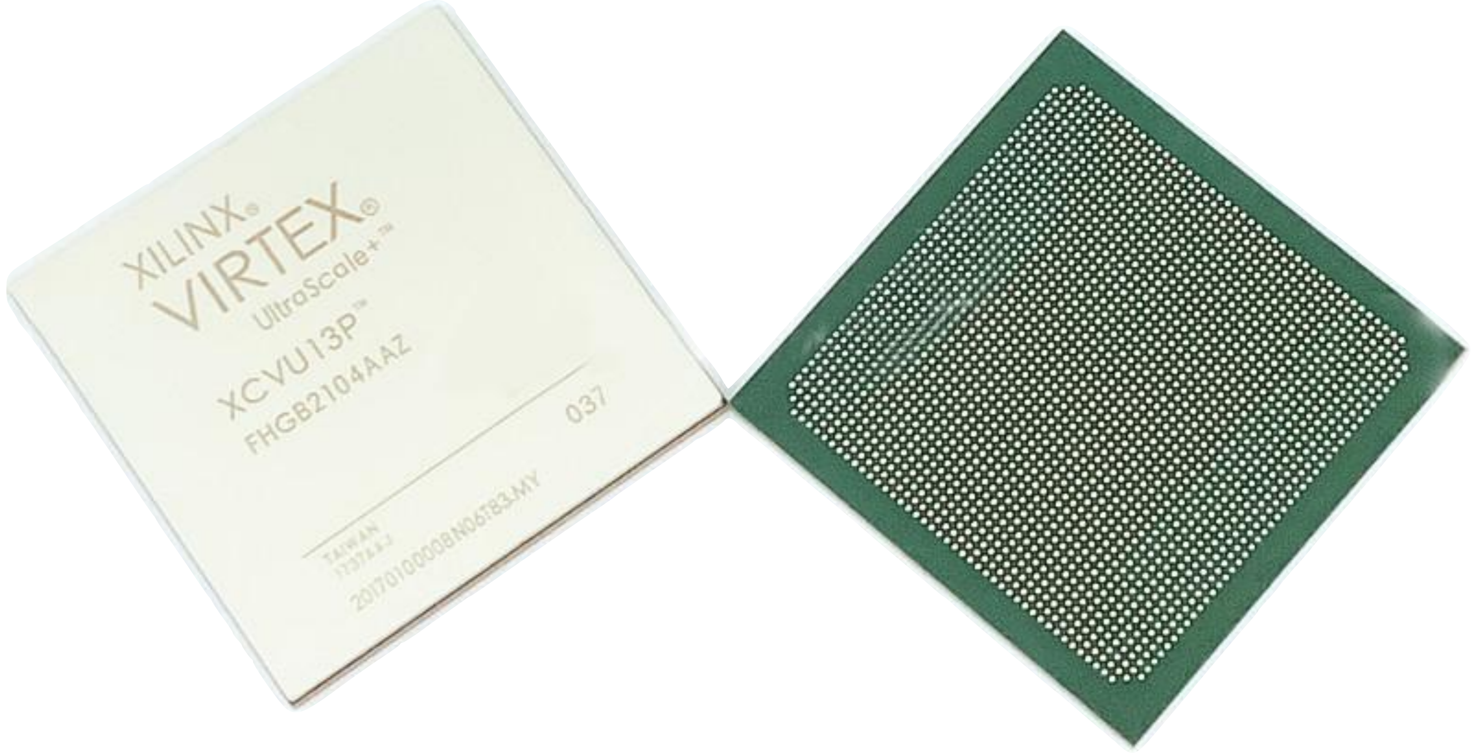


```vhdl
library ieee;
use ieee.std_logic_1164.all;
use ieee.numeric_std.all;

entity signed_adder is
  port
  (
    aclr : in   std_logic;
    clk  : in   std_logic;
    a    : in   std_logic_vector;
    b    : in   std_logic_vector;
    q    : out  std_logic_vector
  );
end signed_adder;

architecture signed_adder_arch of signed_adder is
  signal q_s : signed(a'high+1 downto 0); -- extra bit wide

begin -- architecture
  assert(a'length >= b'length)
    report "Port A must be the longer vector if different sizes!"
    severity FAILURE;
  q <= std_logic_vector(q_s);

  adding_proc:
  process (aclr, clk)
    begin
      if (aclr = '1') then
        q_s <= (others => '0');
      elsif rising_edge(clk) then
        q_s <= ('0'&signed(a)) + ('0'&signed(b));
      end if; -- clk'd
    end process;

end signed_adder_arch;
```

Example of VHDL code

Source: Wikipedia

Machine learning models on FPGAs

- L1 Trigger group is already prototyping the **future algorithms**
  - Strong use of high level language (**HLS**) to program the FPGA
  - Implementation of machine learning models into the FPGA: **HLS4ML** and **Conifer** (Decision Forests)

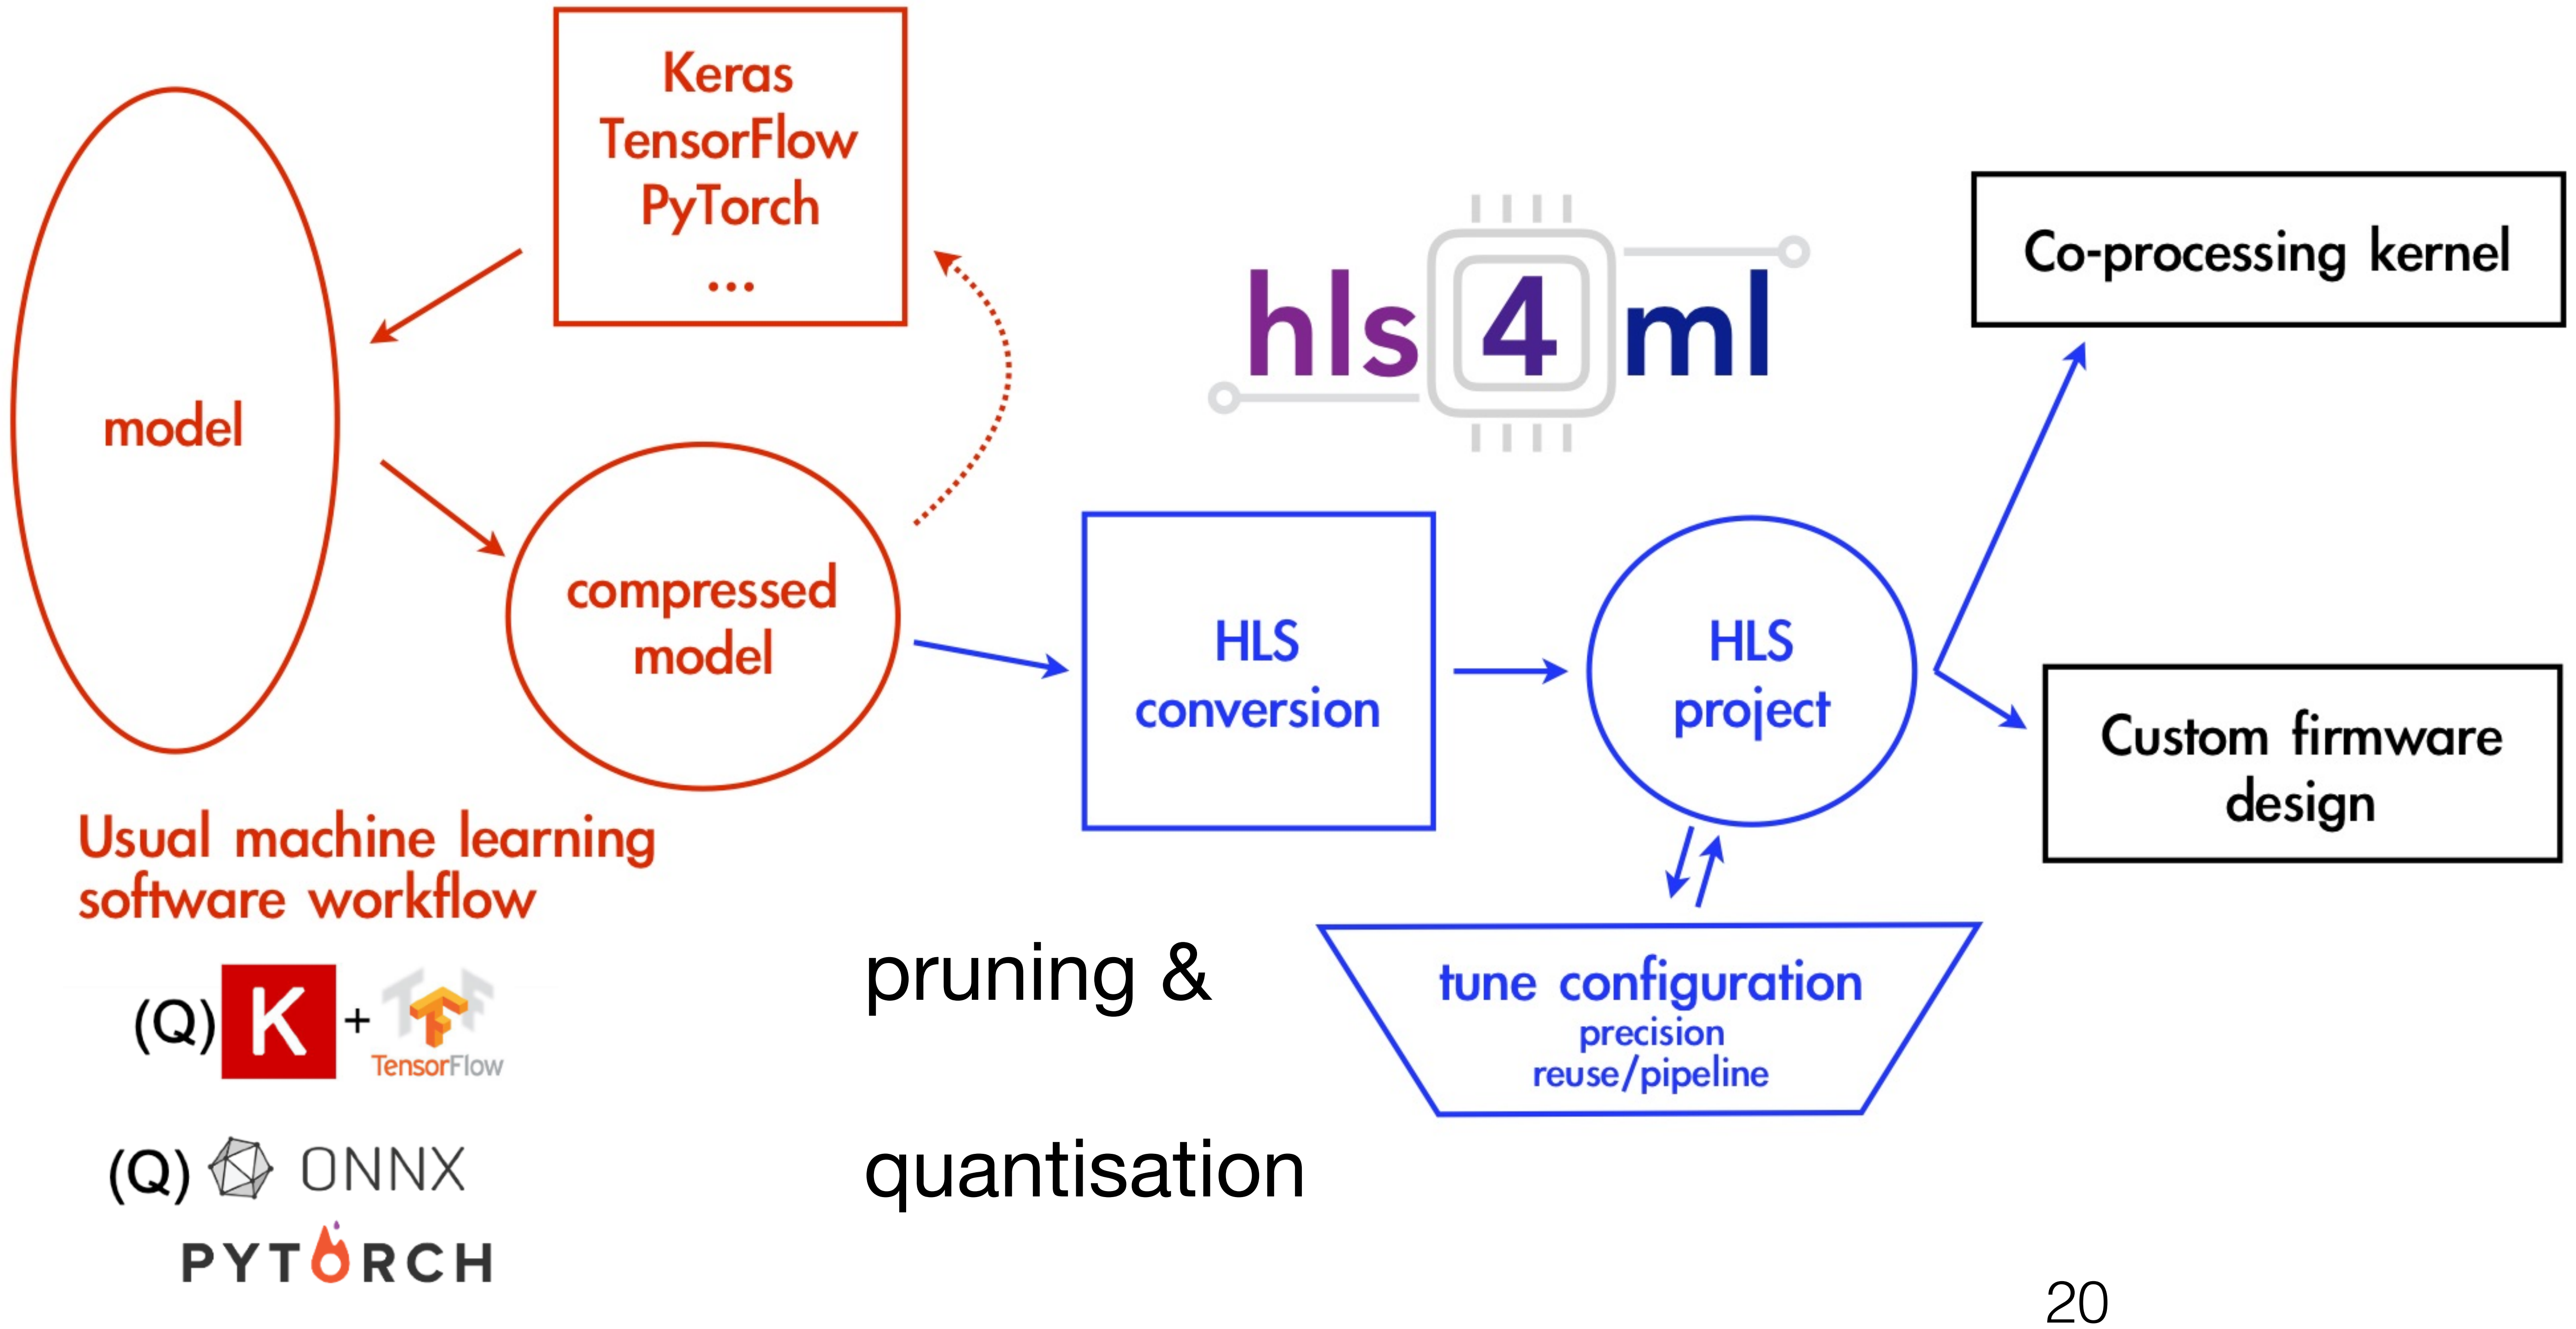

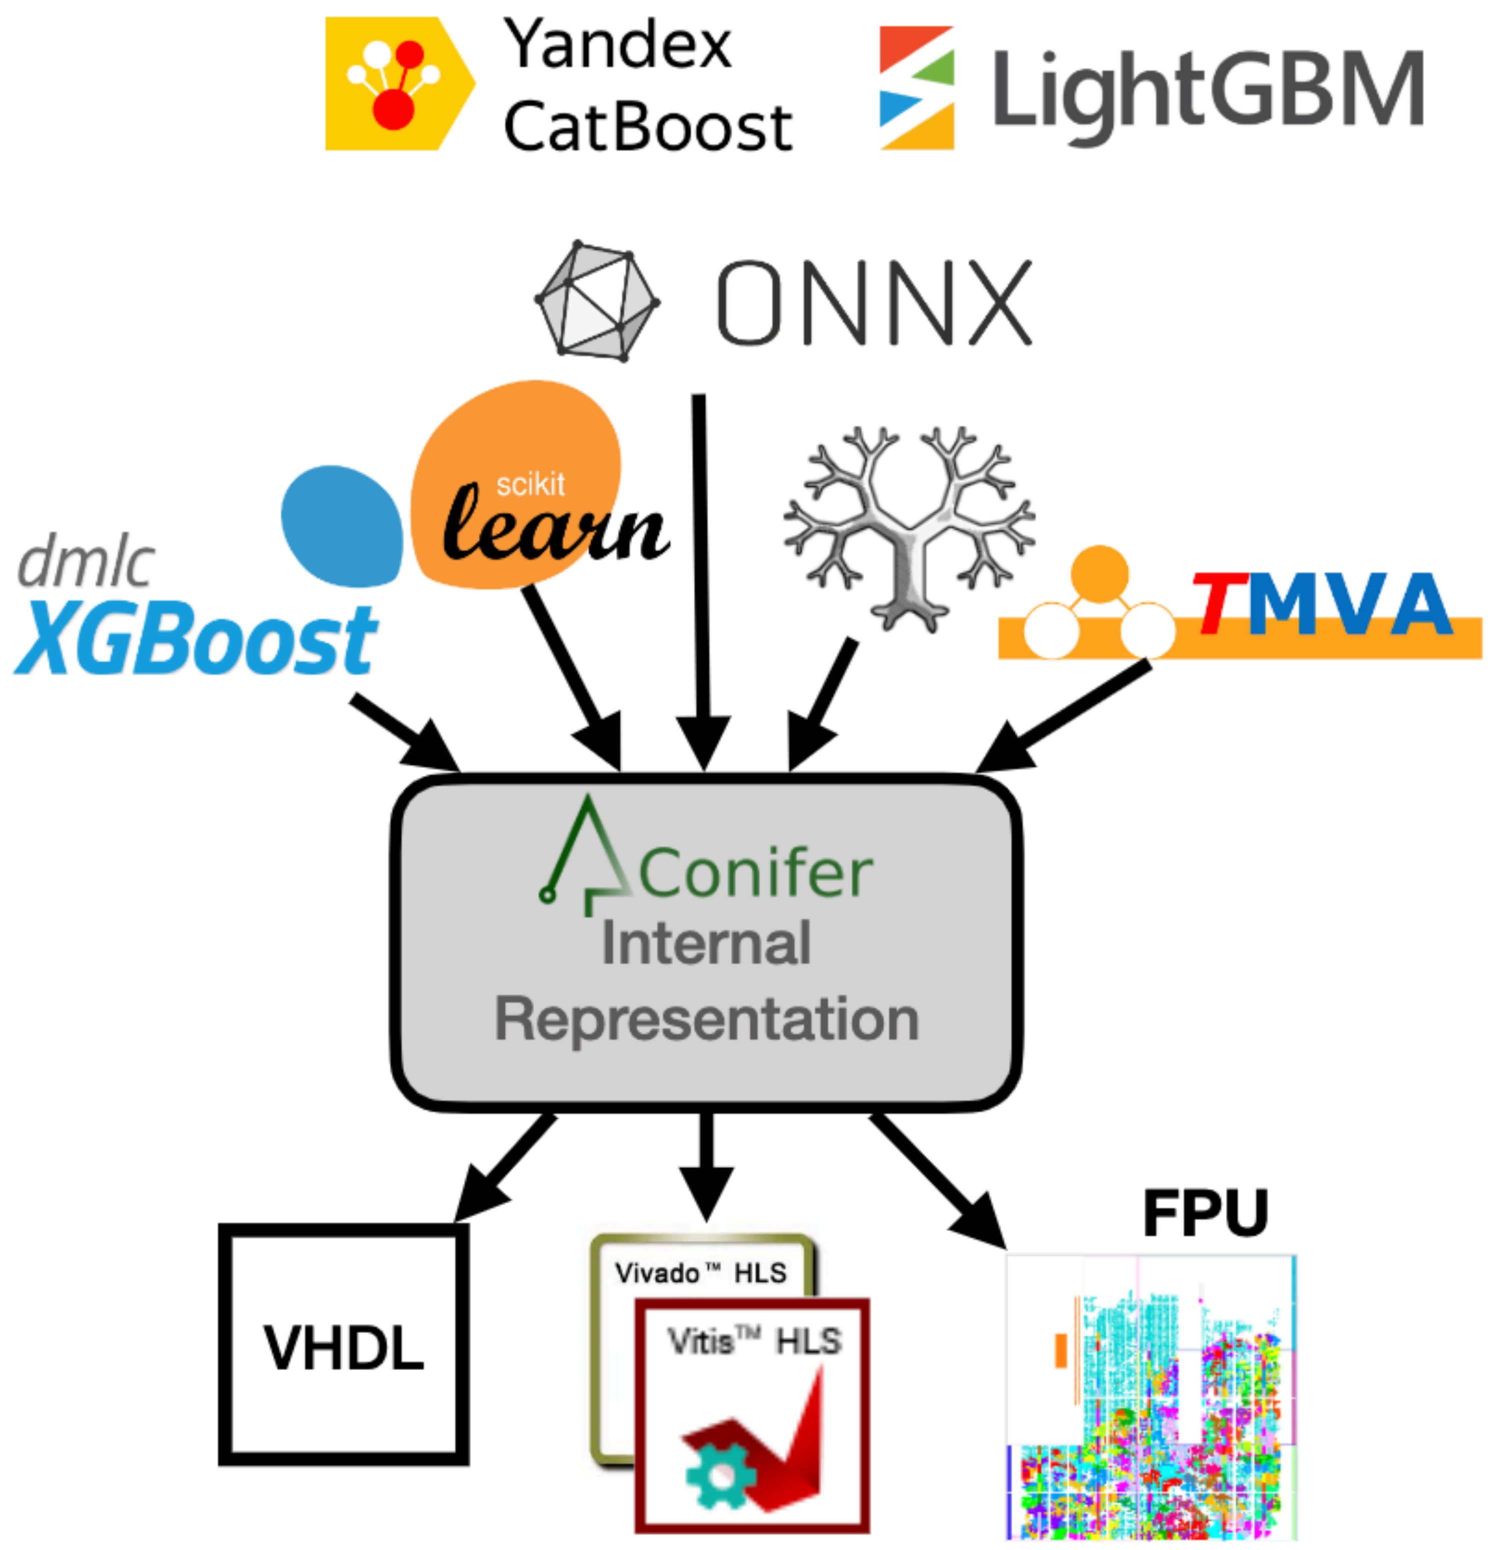

Example: Jet identification

- ML model to identify the **Jet flavour**: i.e. light quark/gluon Vs b quark
- Implemented using **HLS4ML** toolkit
- 2x 1D convolution layers, 3x dense layers, one feature output
- Latency: in conjunction with jet finding the total latency is less than 1 μs

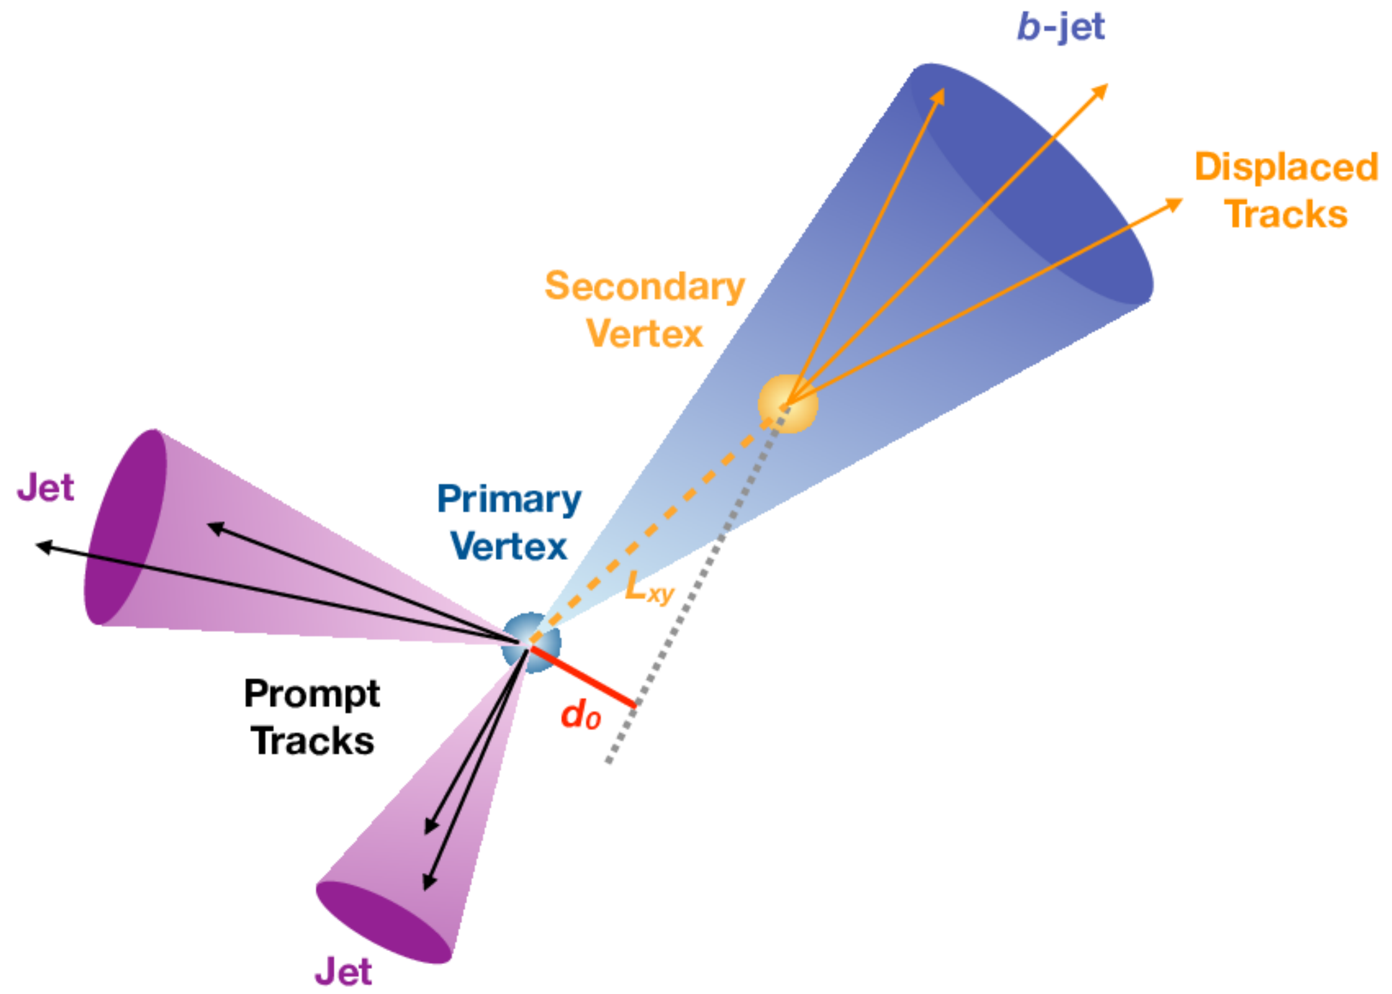

Source: CERN

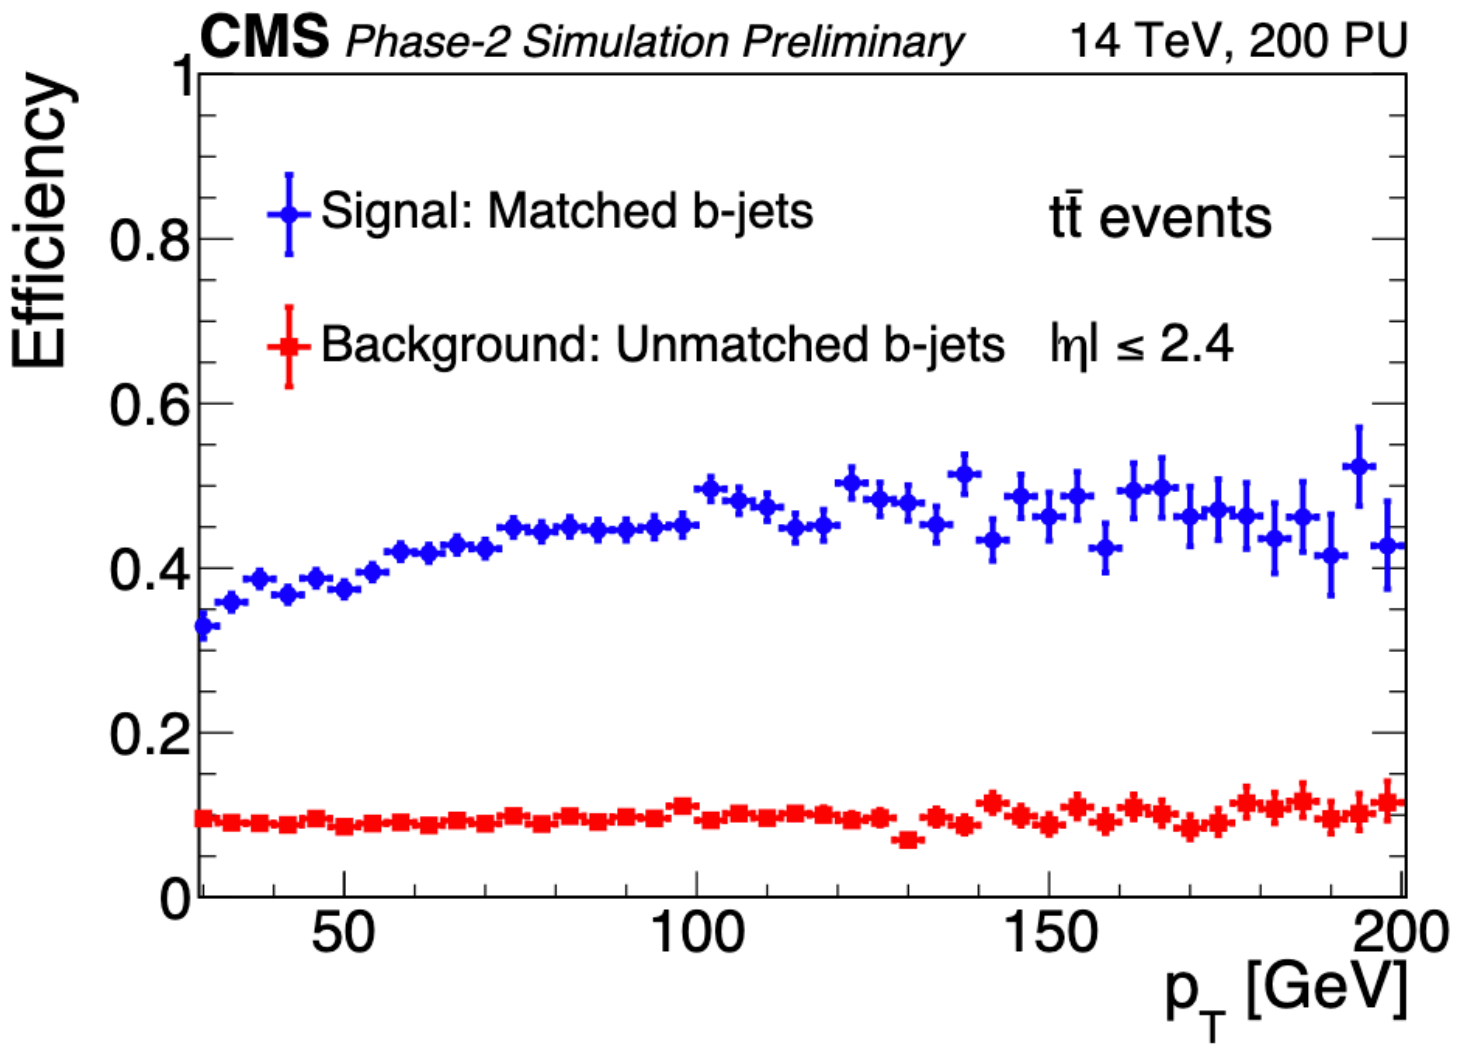

DP-2022-021

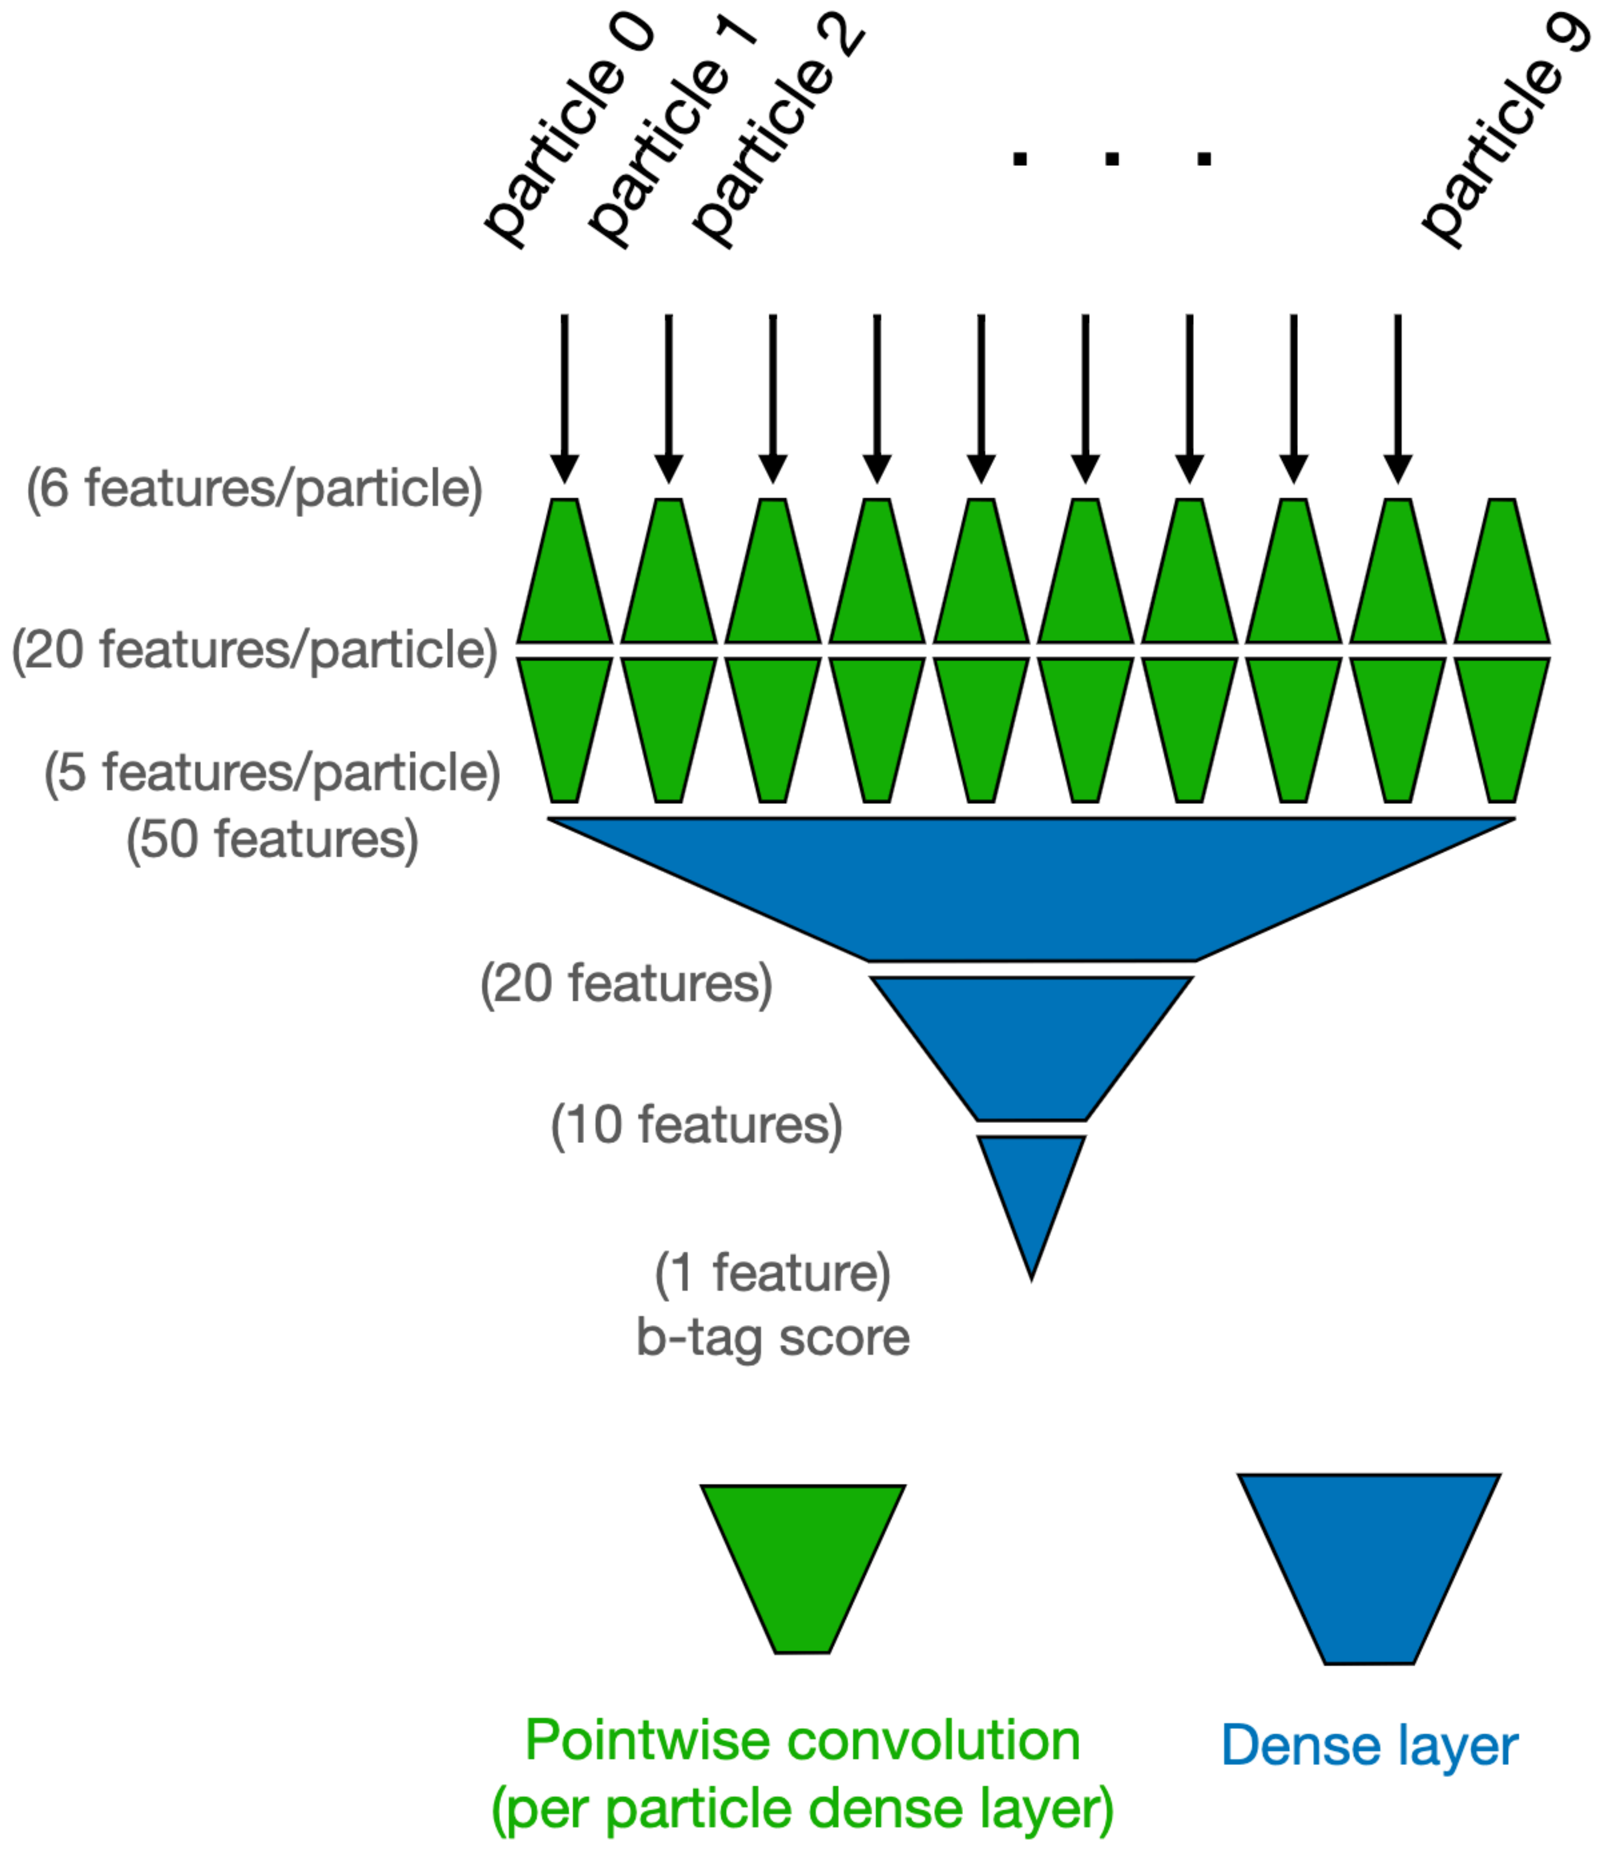

HL-LHC - HLT Trigger

- Event processing 3-4 seconds
- Output trigger **rate 7.5 kHz**
- Implementation of **heterogeneous computing**, e.g. **GPU** (introduced during current data taking), FPGA(?), TPU(?)
  - Testing abstraction libraries, e.g. Alpaka
- About 1000 PC in the HLT farm

Example of current time reduction by using GPUs

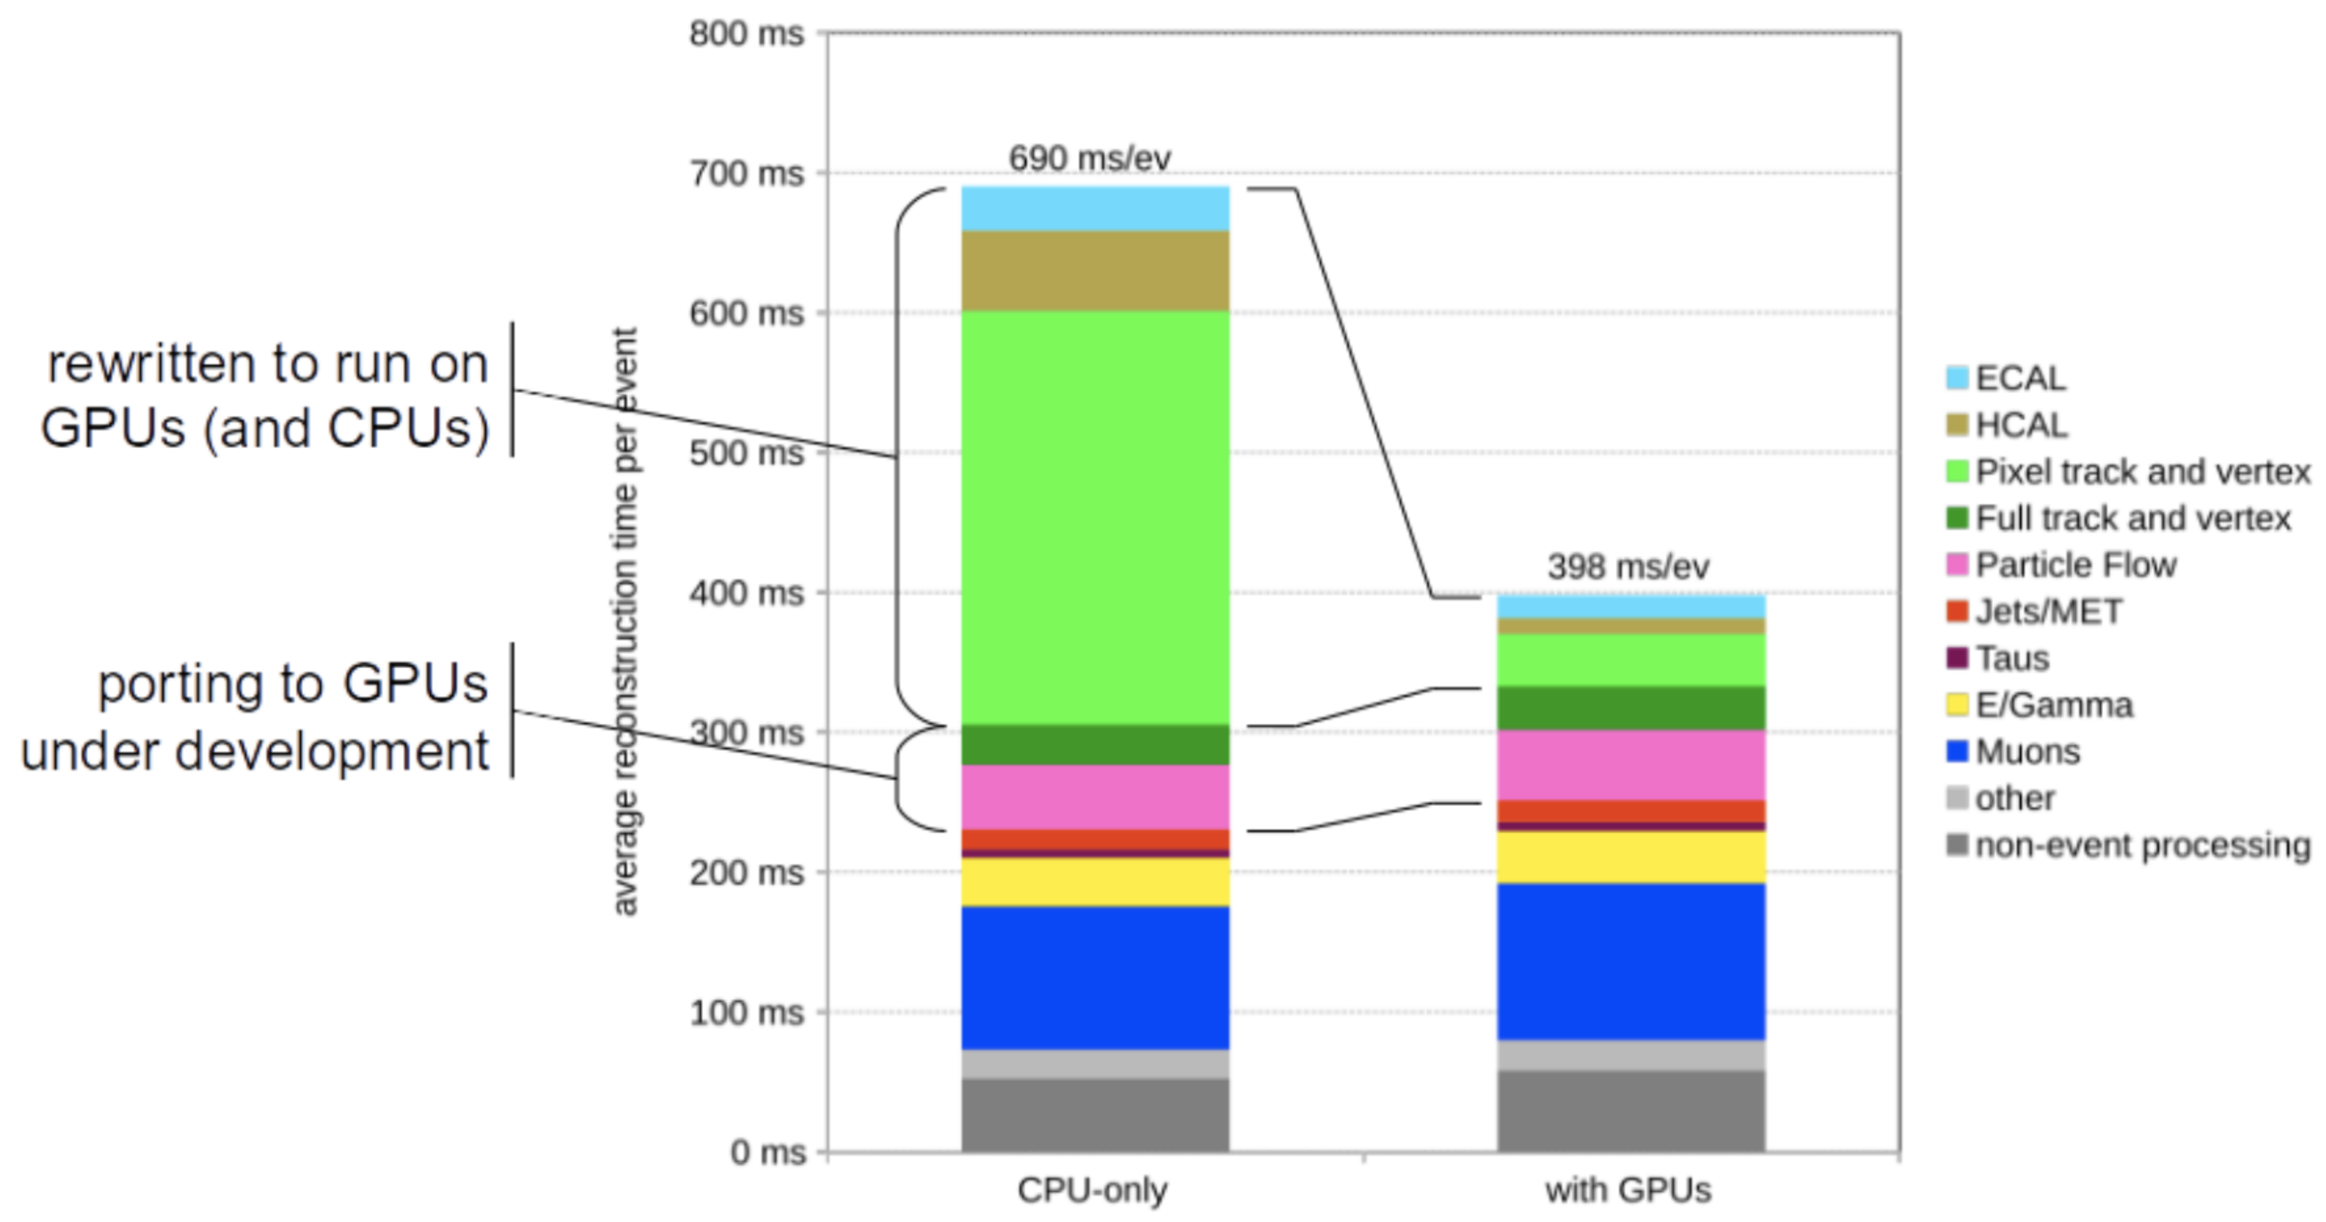

Source: CERN

Focus on CMS Tracker

Source: CMS

CMS Tracker subdetector

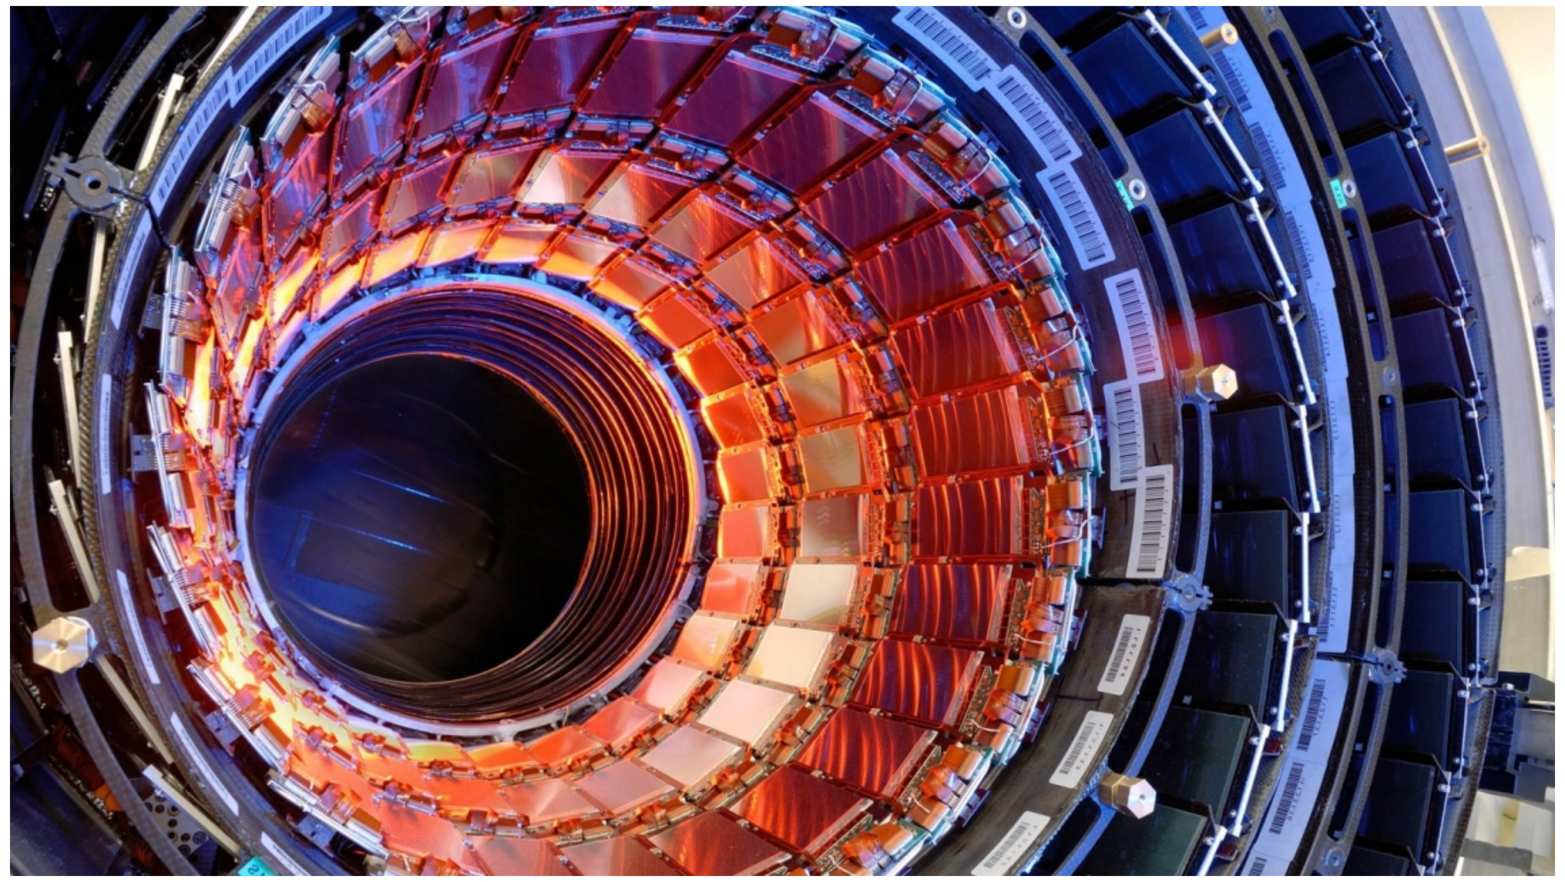

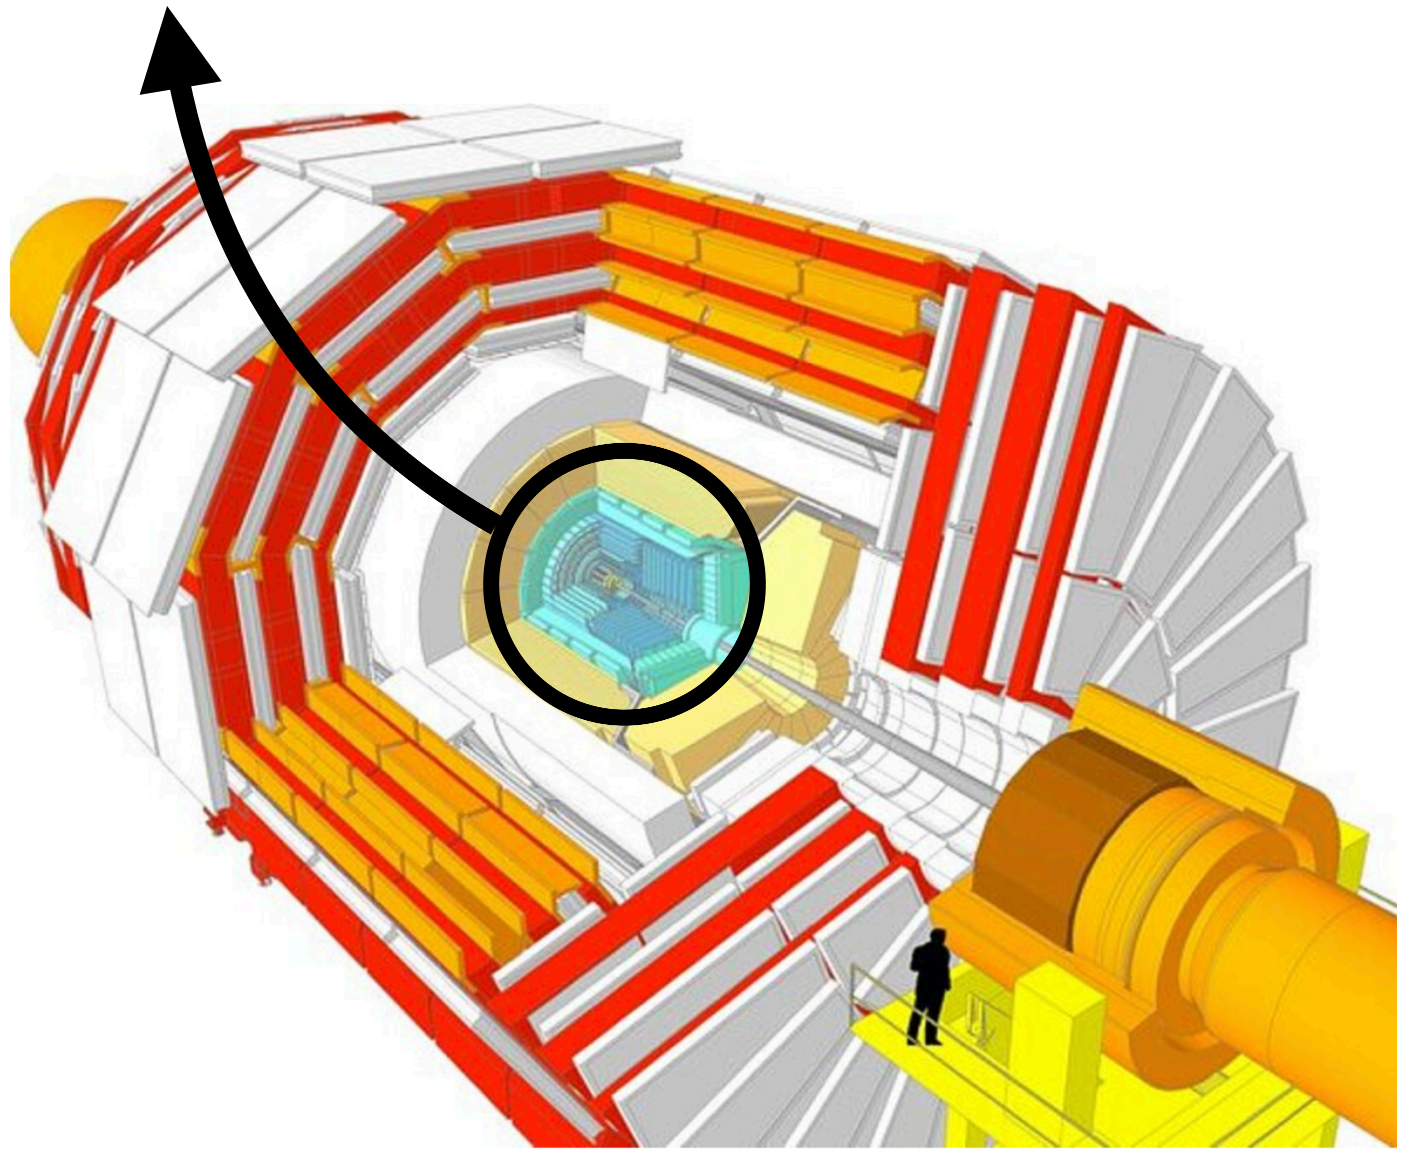

- Detects the passage of **charged particles**
- With pattern finding and fitting it enables the reconstruction of the **trajectory** of charged particles
- Thanks to the strong and uniform magnetic filed (3.8T) we can measure their **momentum**

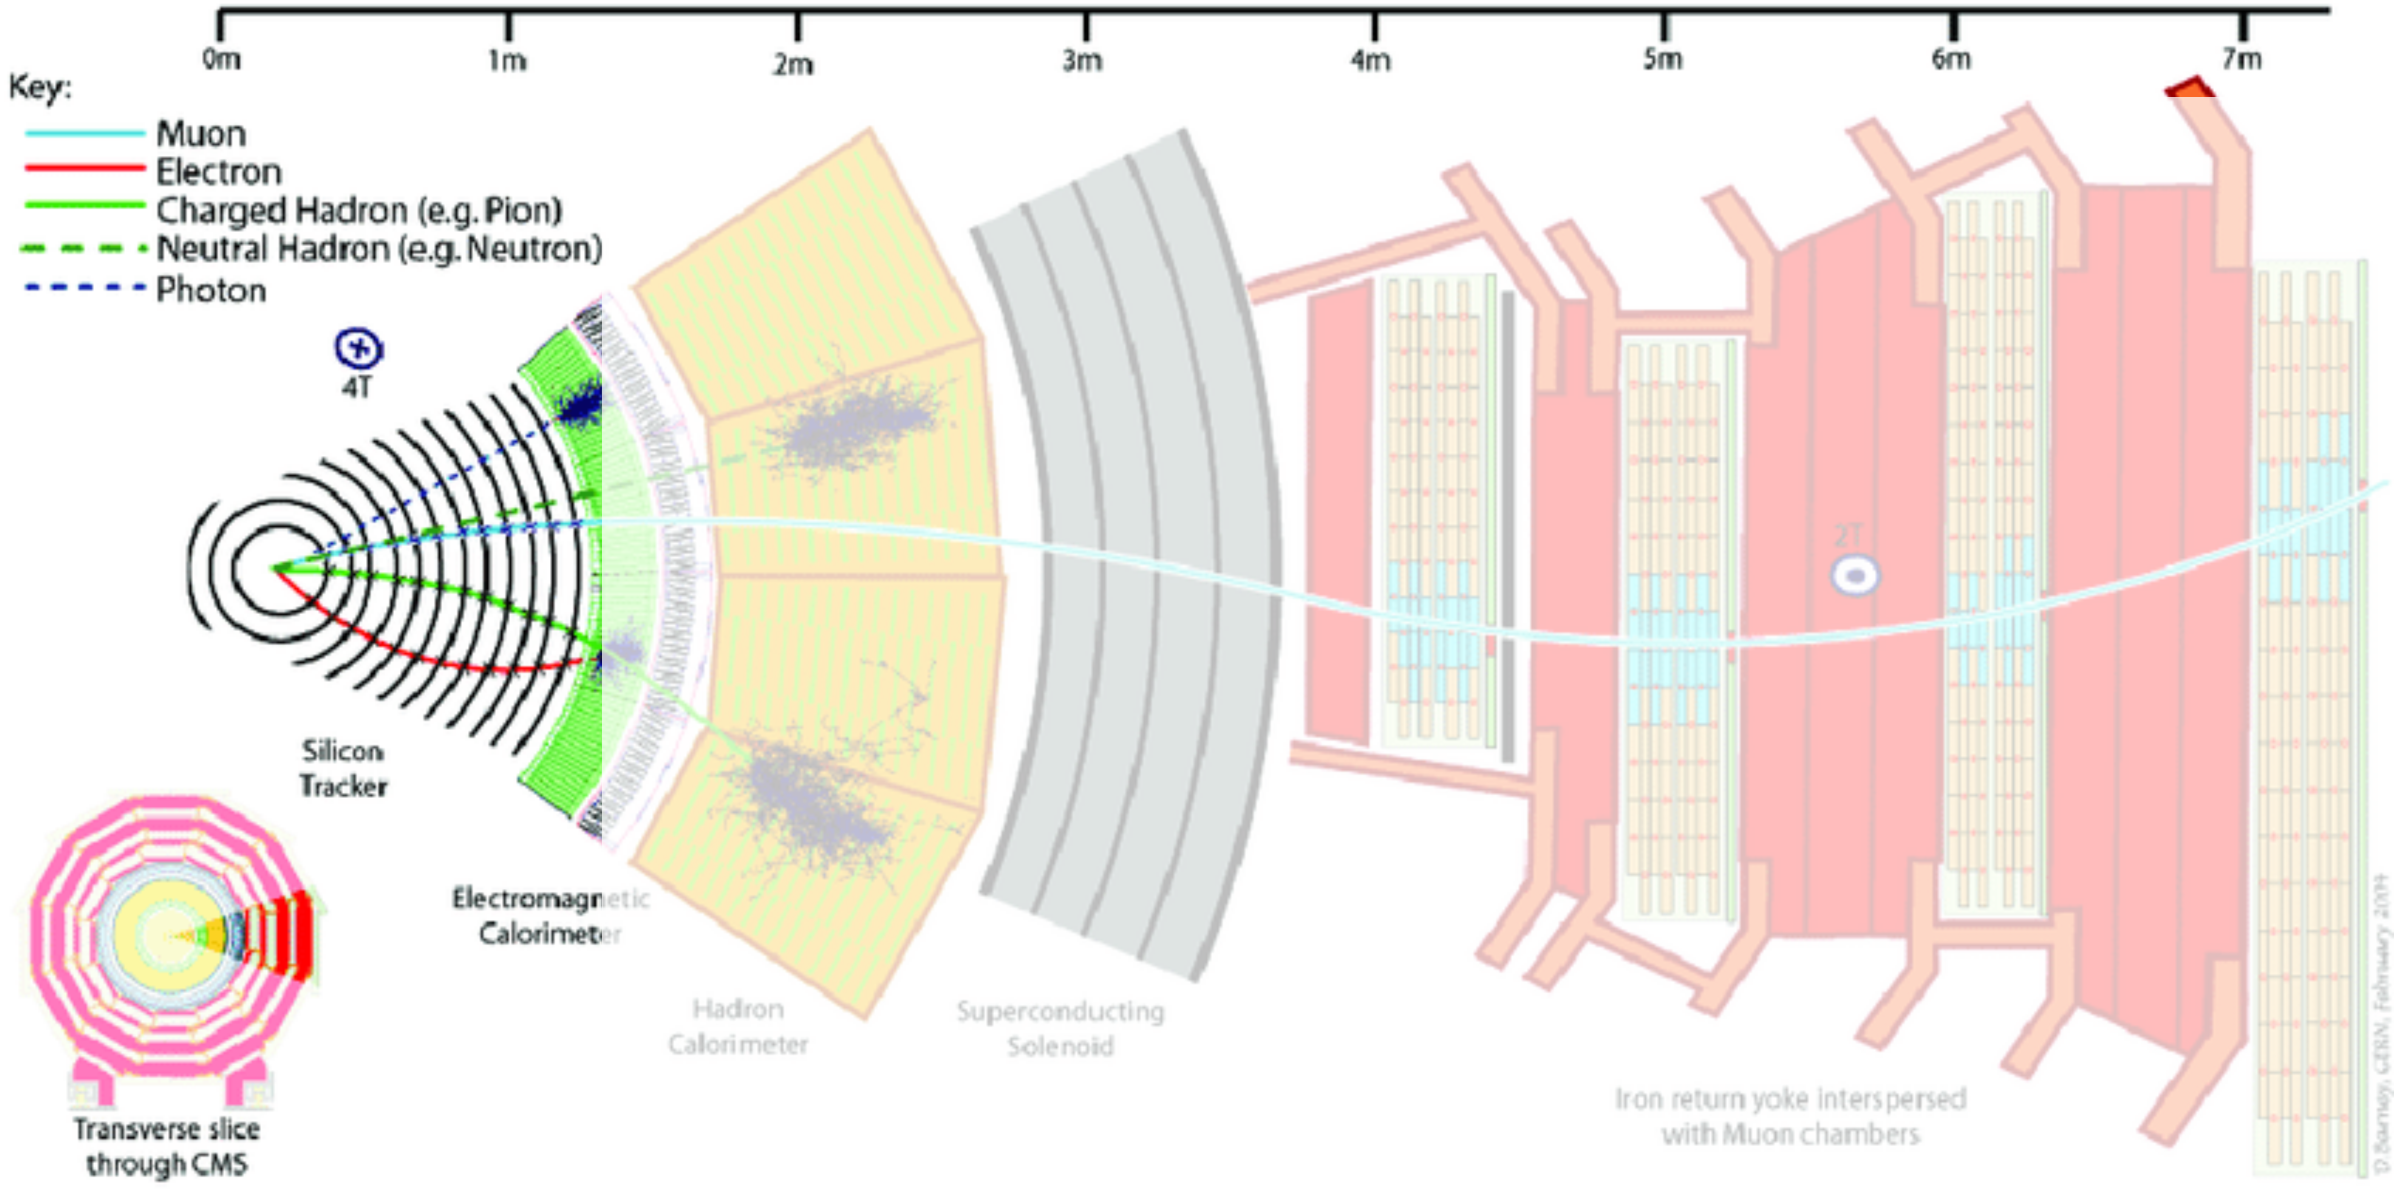

Source: CMS

The HL-LHC challenge

Each bunch crossing generates **multiple proton-proton collisions**, one collision at most is an interesting event, other collisions constitute “noise” called “**pile-up**”

During HL-HLC data taking we expect event with **pile-up 200** and we want to generate L1 trigger information from the tracker: L1 tracks to be reconstructed within **5 μs**

Event display of a high pile-up event

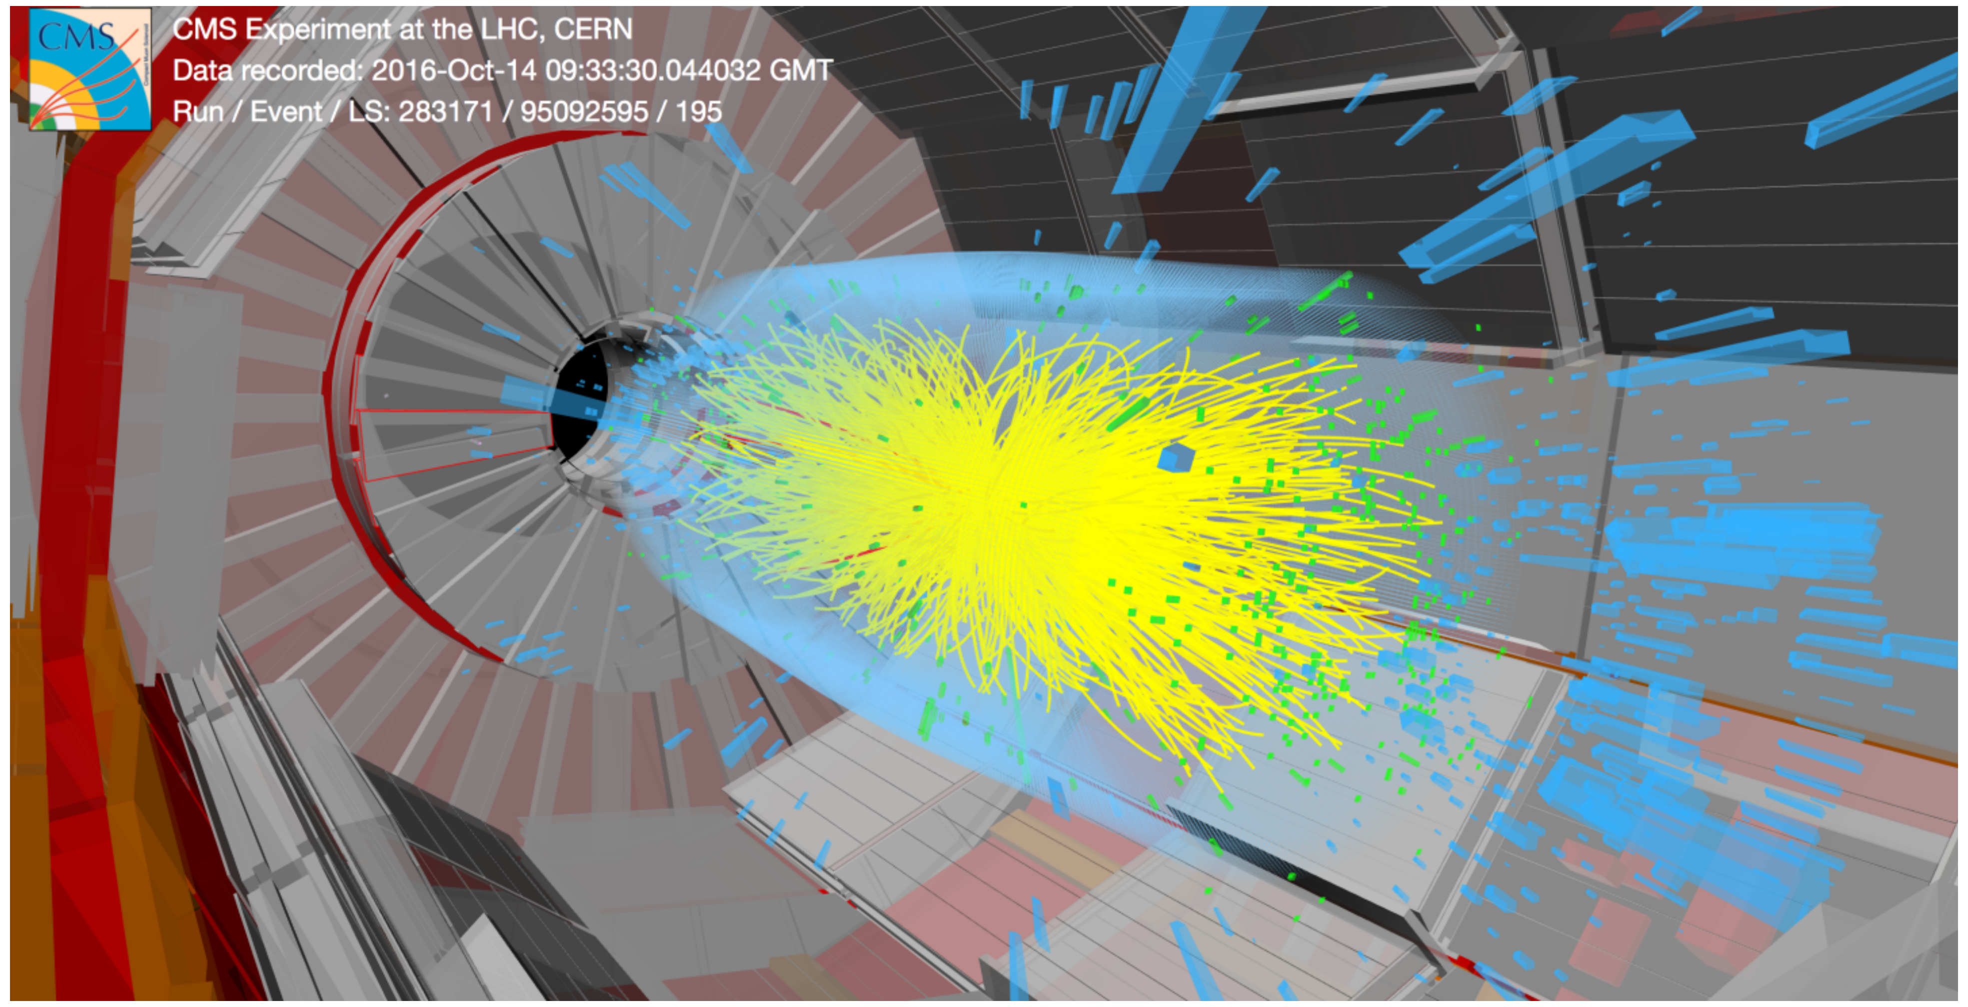

<pile-up> ~ 100

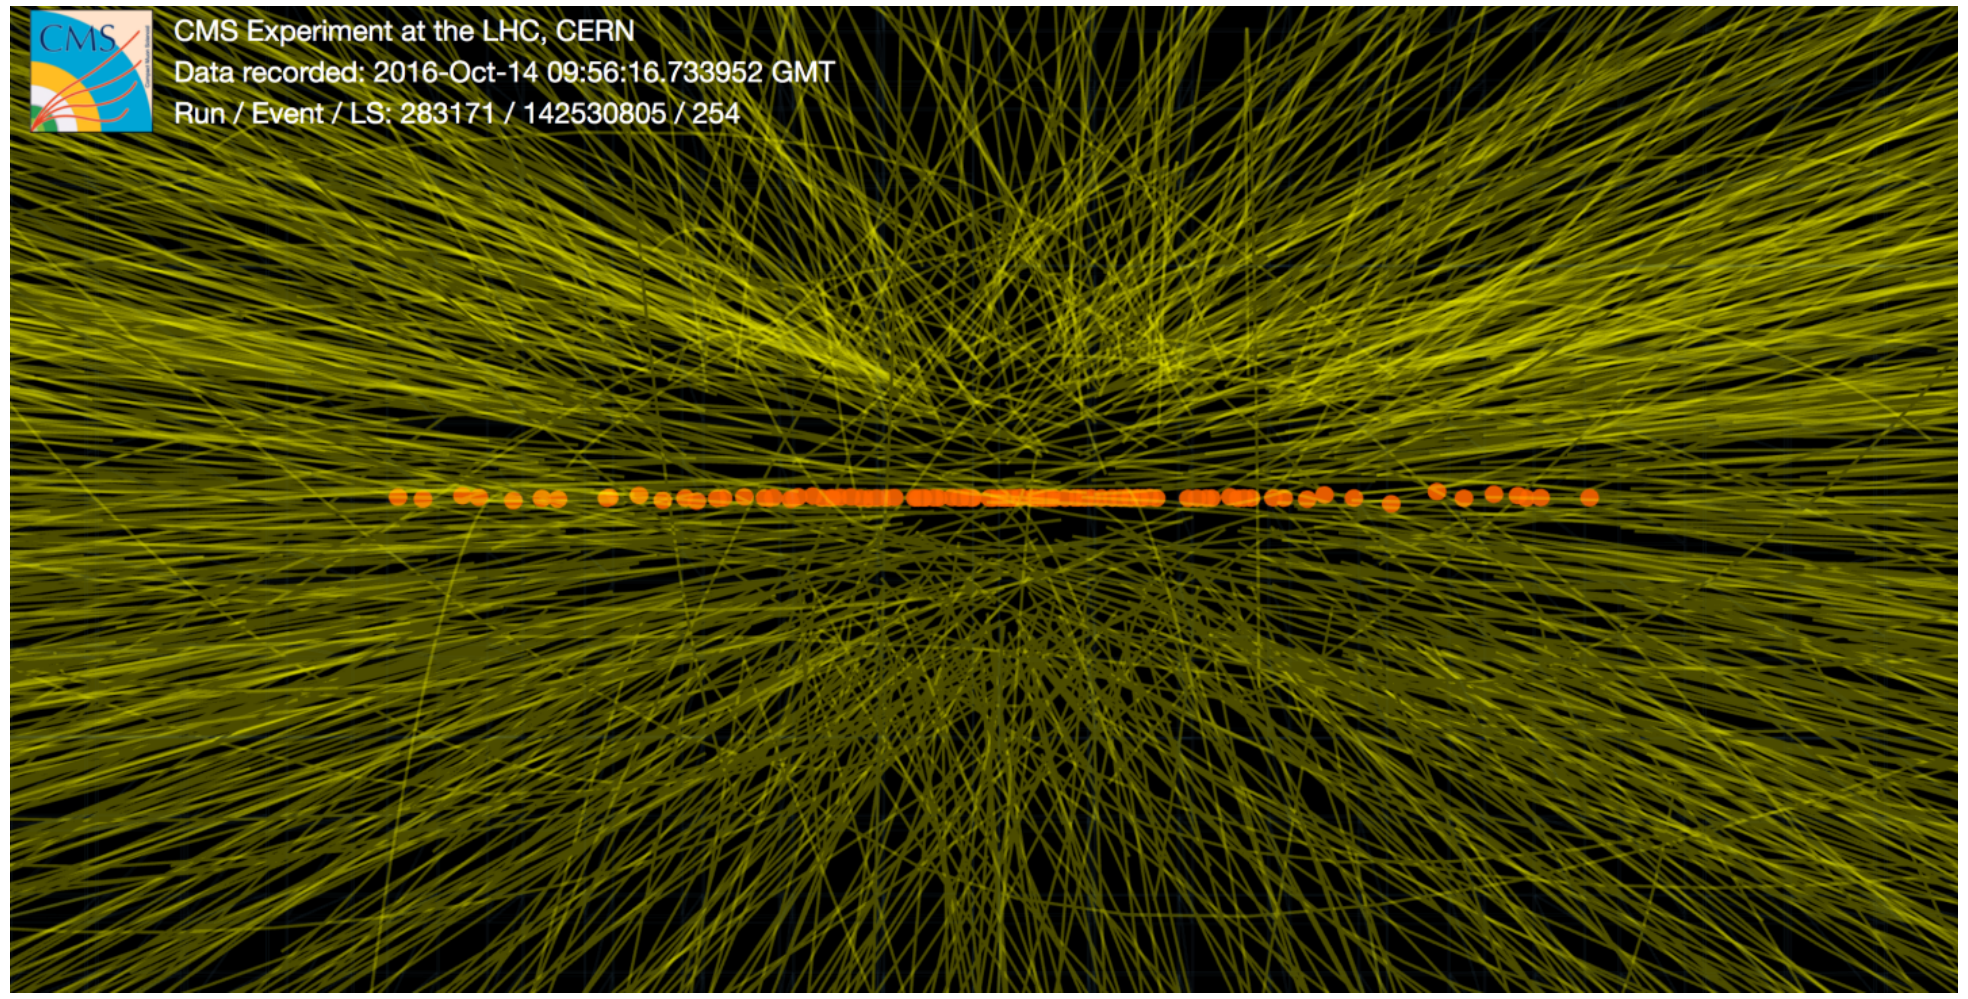

Source: CMS

HL-LHC CMS Tracker

General increase of granularity and radiation hardness

Some key features

- **Tilted geometry** of part of the tracker
- Reduced front-end data rate via in-situ trigger data filtering (**pt modules**)
- Reconstruction of the **charged particle trajectory** at trigger **Level 1** (hardware trigger) by using the outer tracker data
- **Material budget** reduction

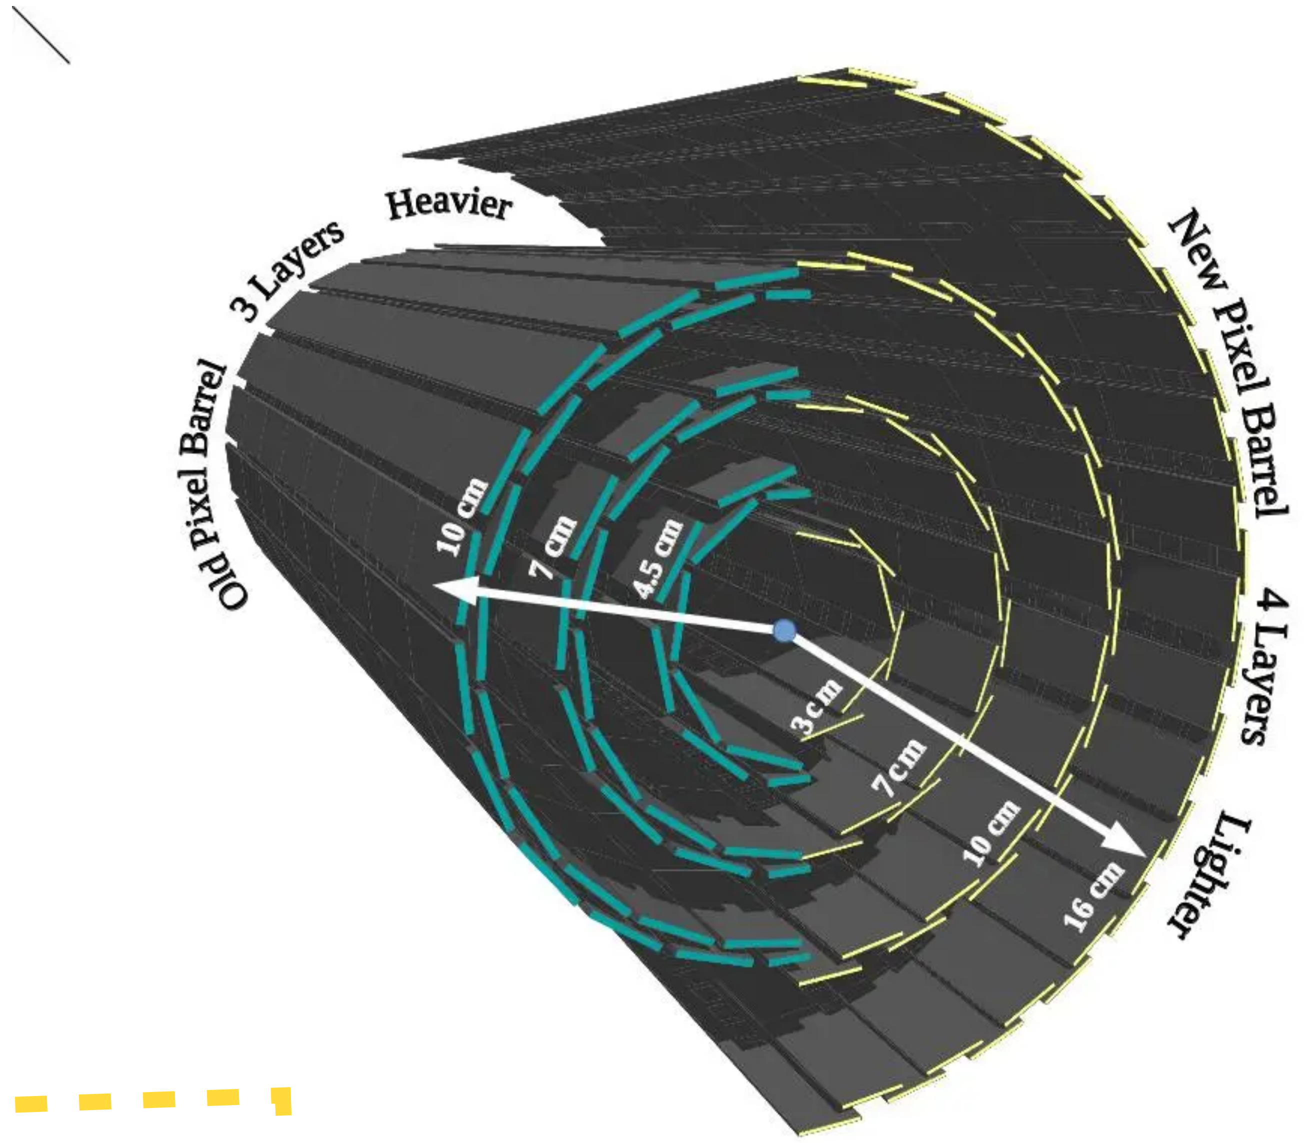

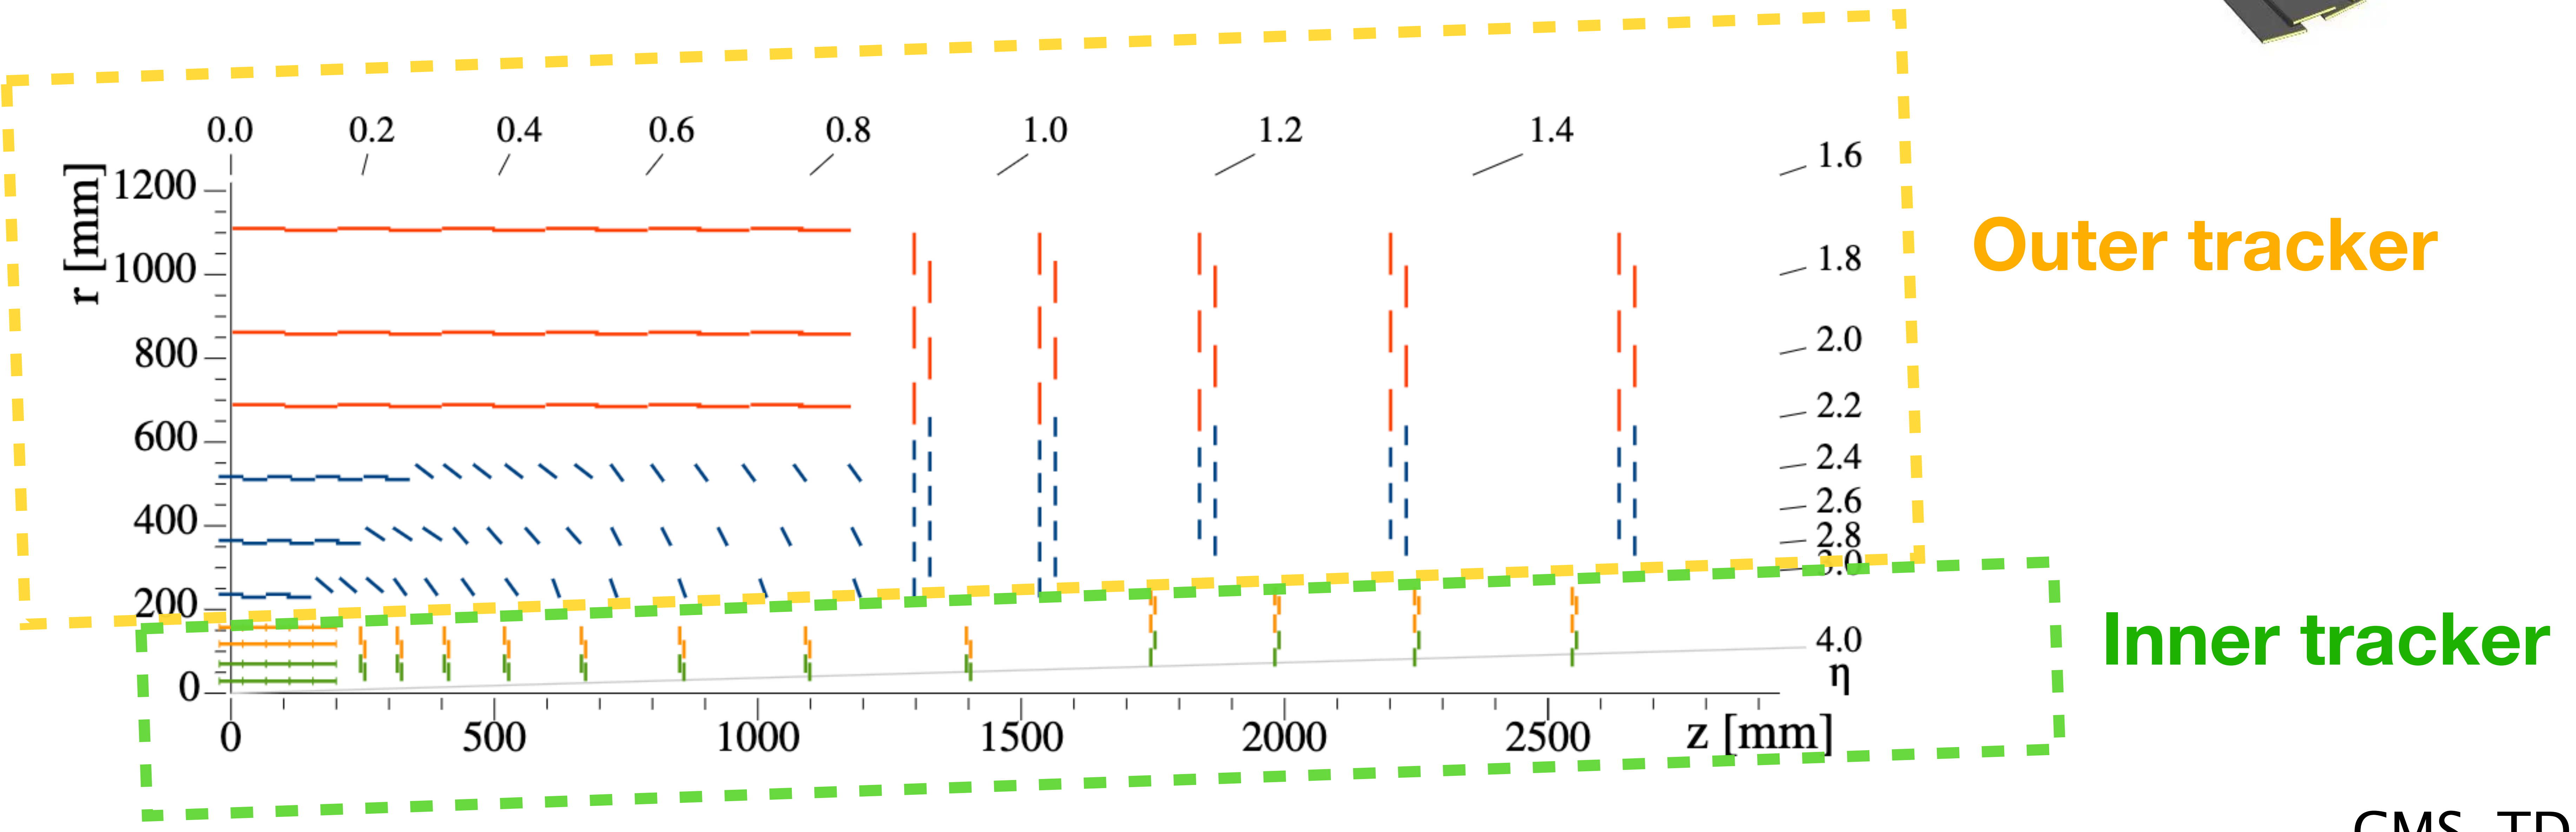

CMS-TDR-014

HL-LHC Inner Tracker

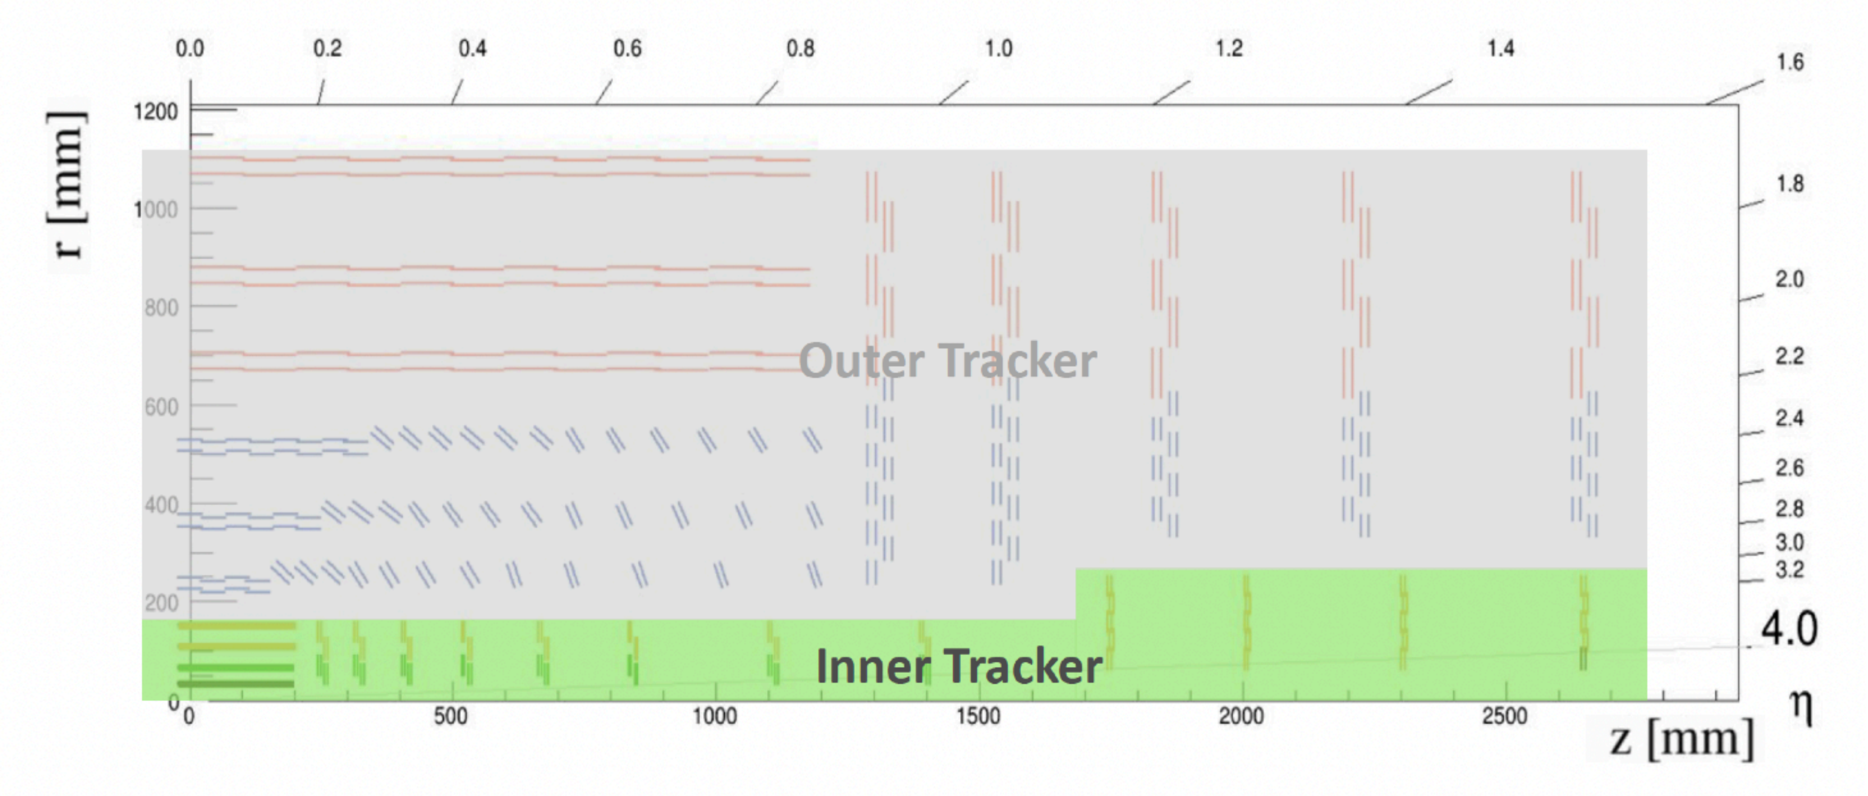

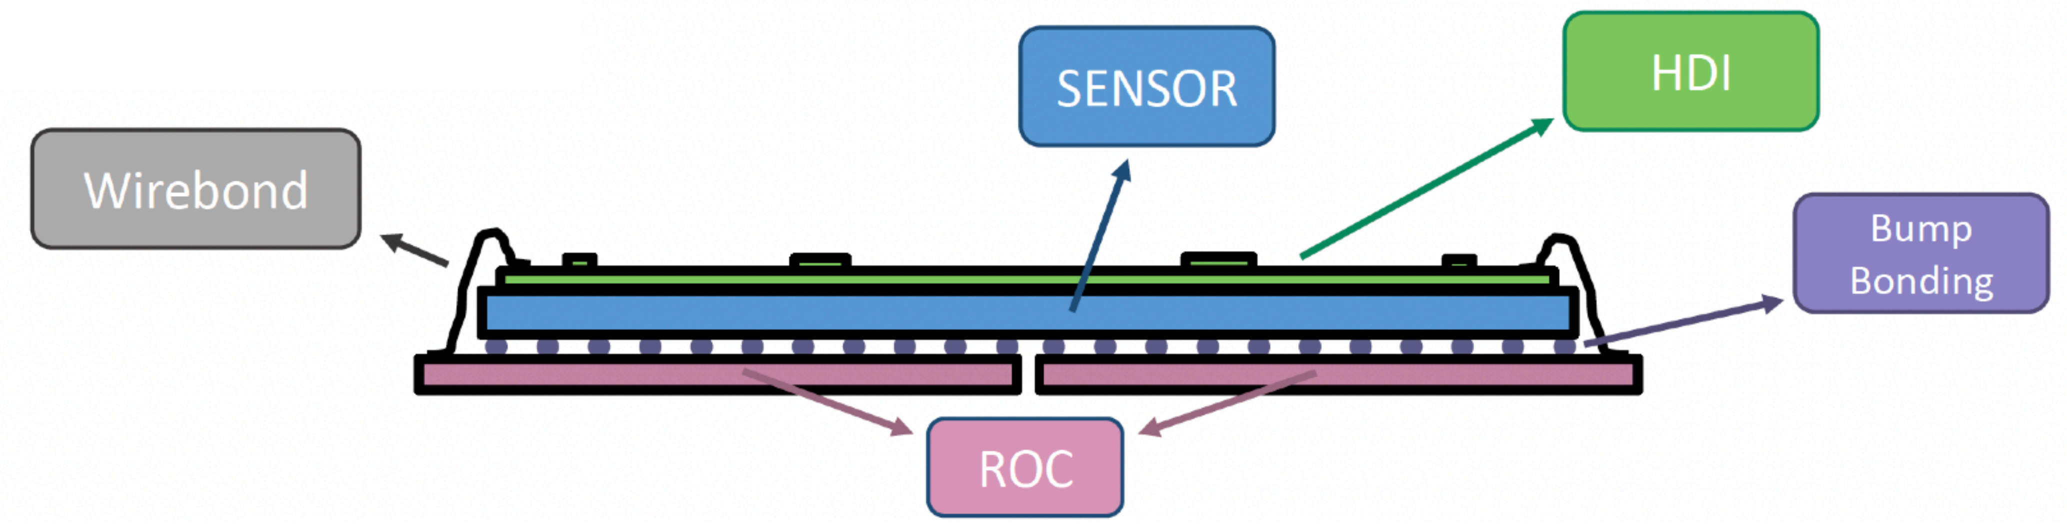

Based on silicon pixel sensors

Modules

- Two types of Pixel Modules: 1x2 and 2x2 readout chip
- Read Out Chip (ROC) bump bonded on sensor
- Serial powering scheme with up to 11 modules per chain
- 25x100μm² pixel cells with 150μm active thickness

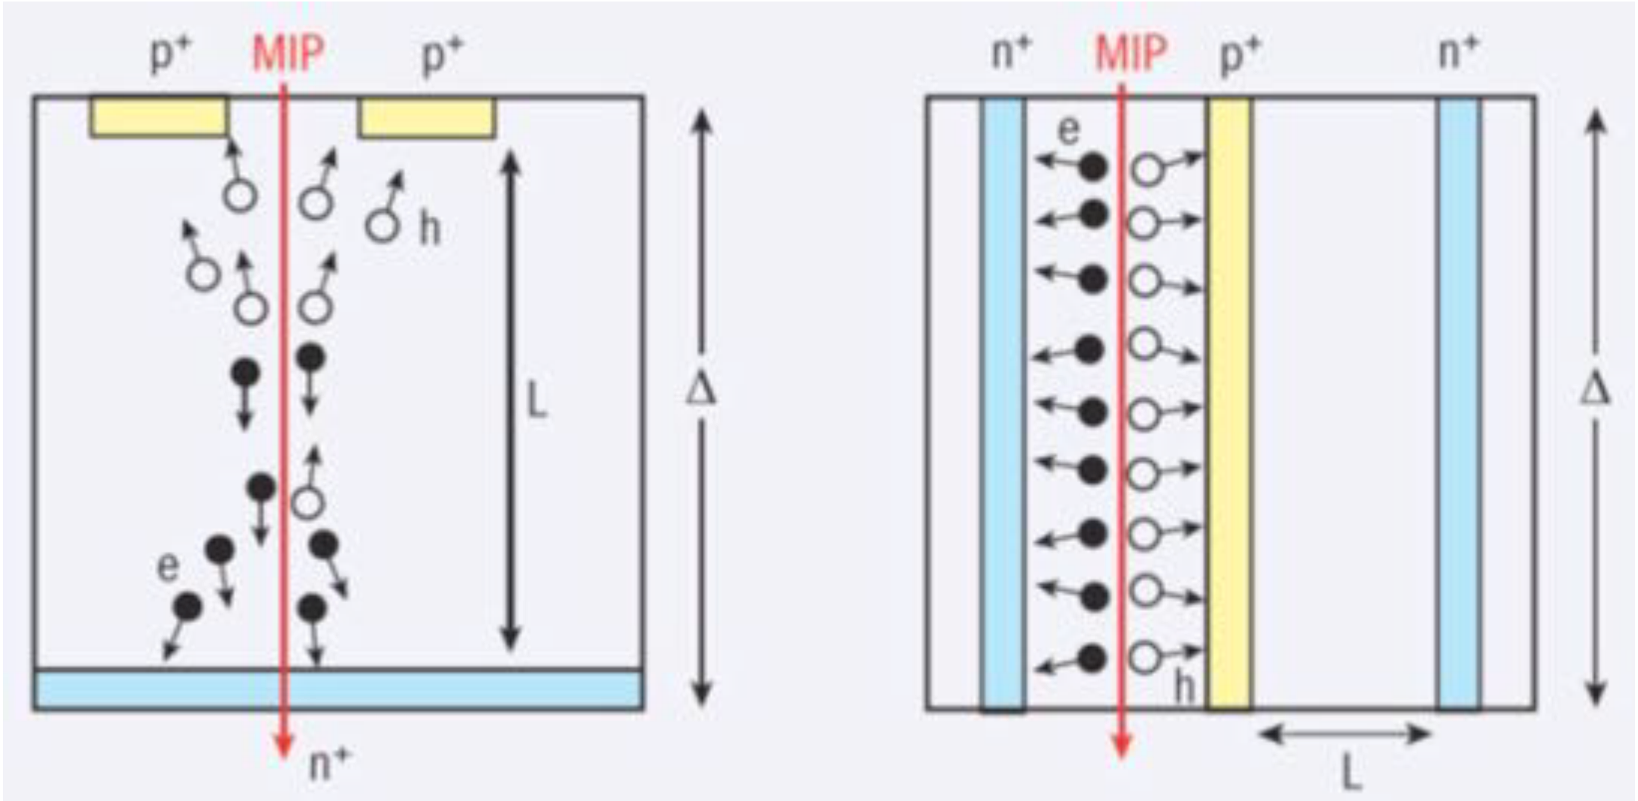

Source: CMS

HL-LHC Inner Tracker - Read-out

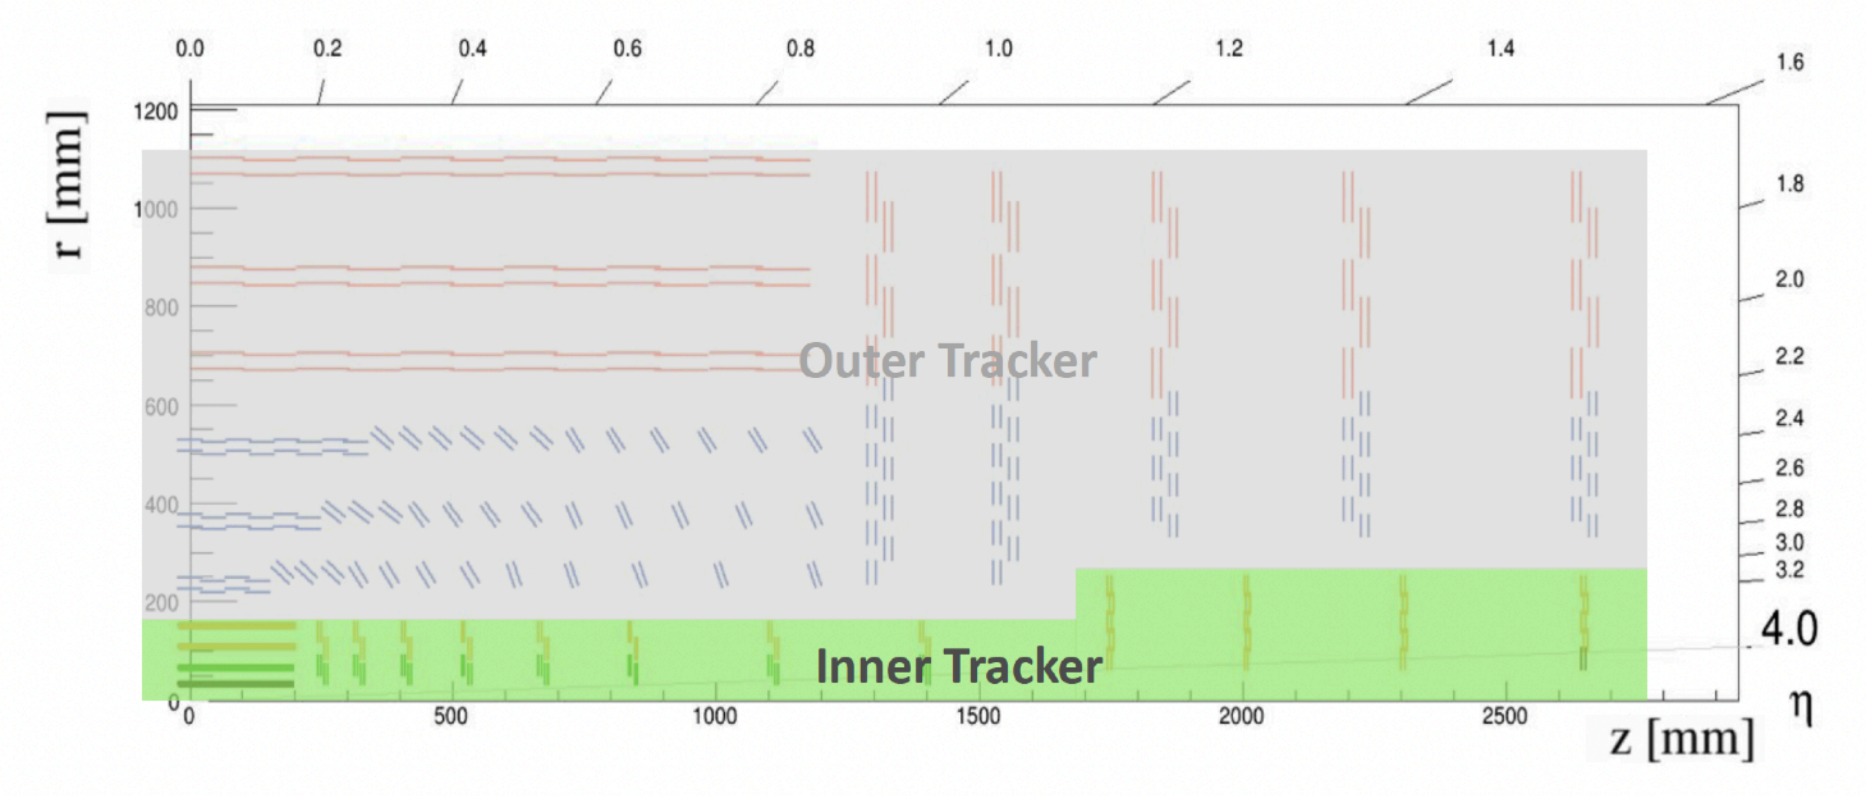

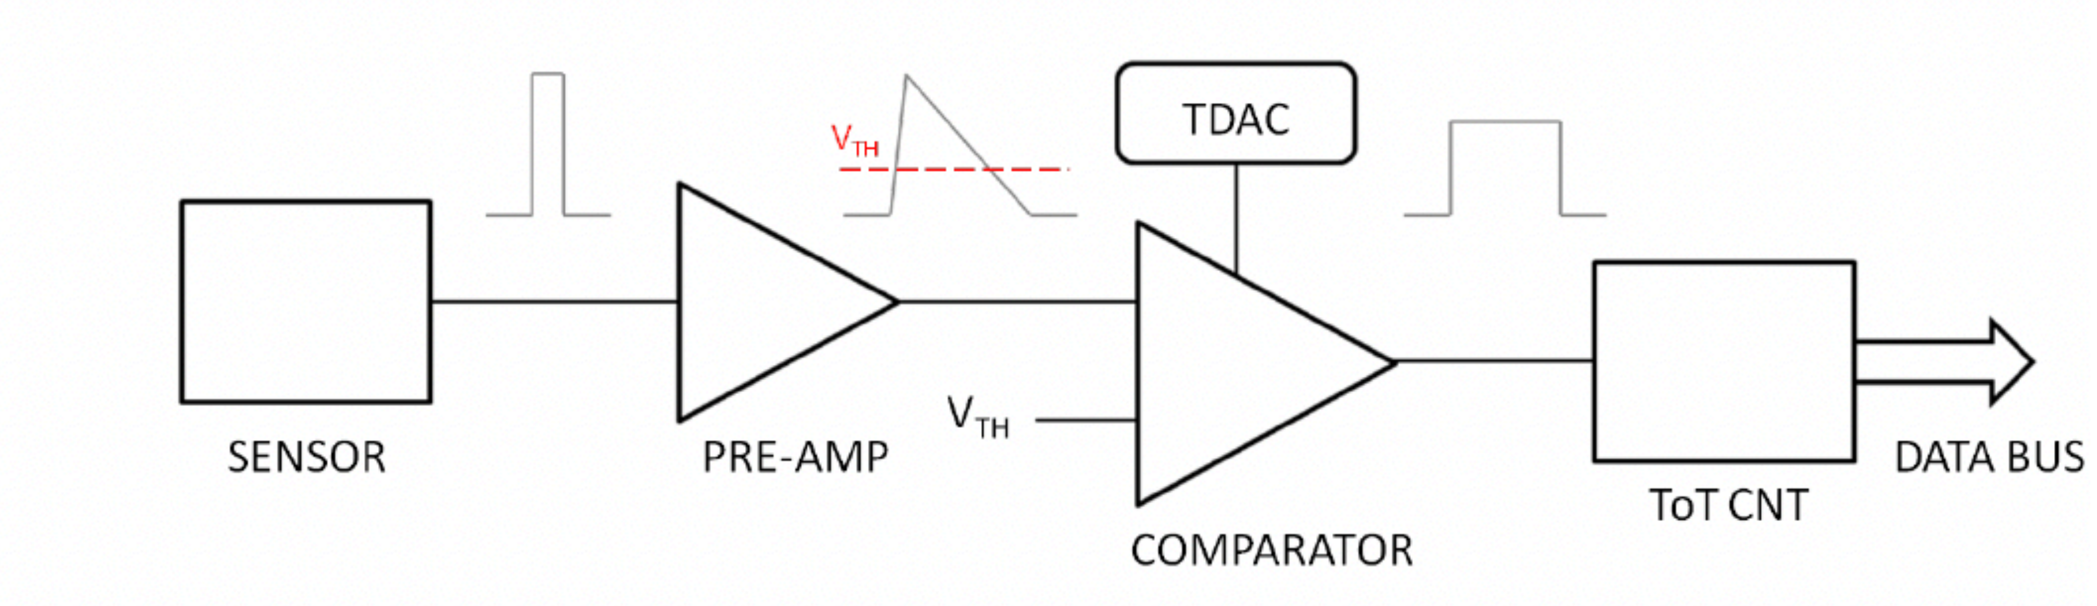

Based on CERN RD53 ASIC project:

- Based on CMOS 65nm technology
- Radiation tolerant up to 1 Grad
- Low power consumption < 1 W/cm^2
- Serial powering via on-chip shunt-LDO regulators
- Full size custom chip (ASIC): 432x336 channels

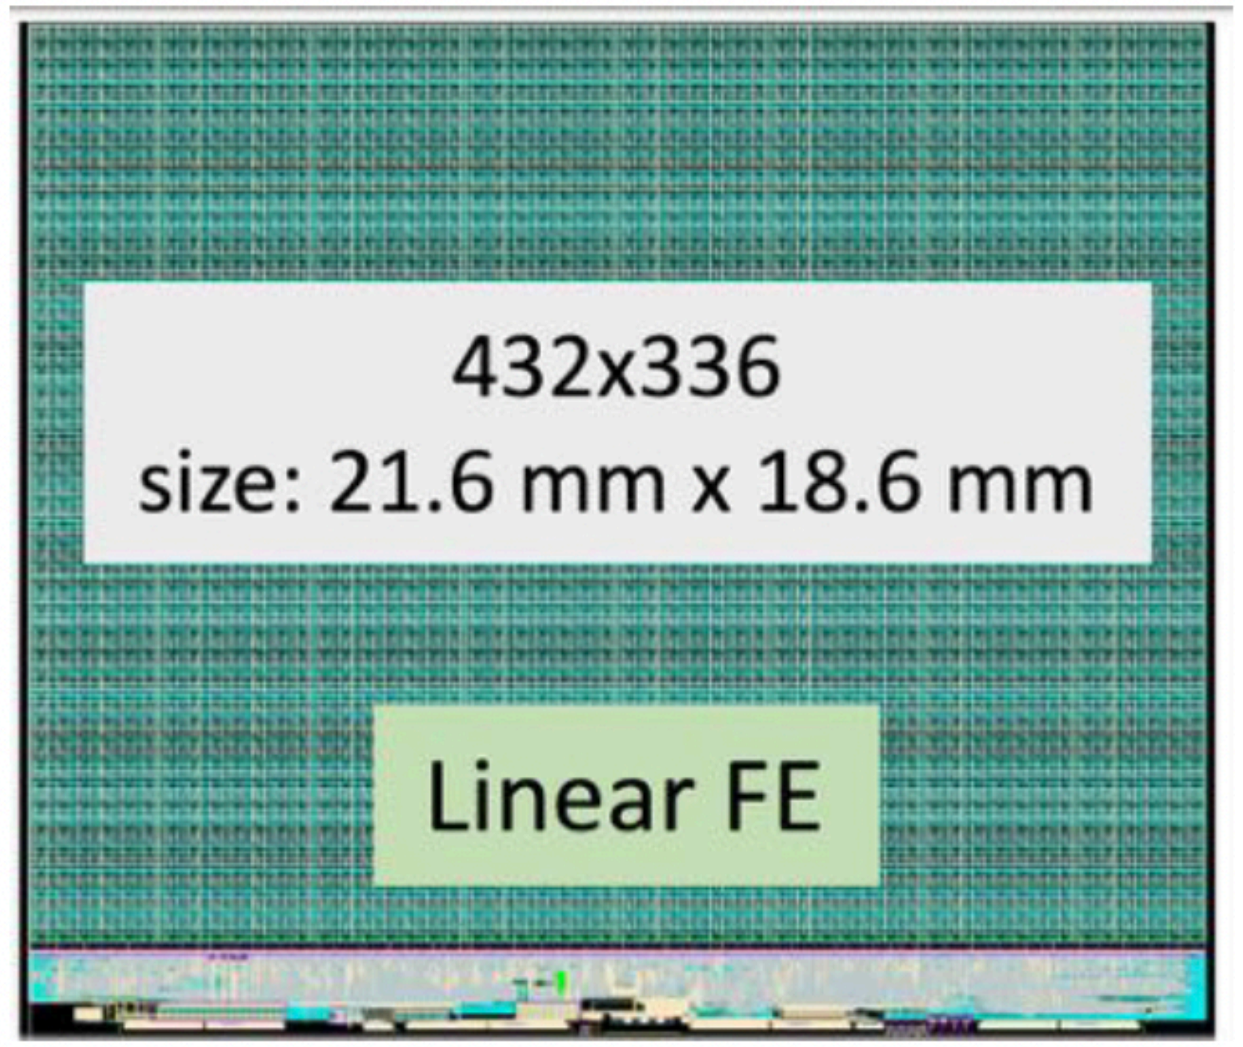

Source: CMS

HL-LHC Outer Tracker

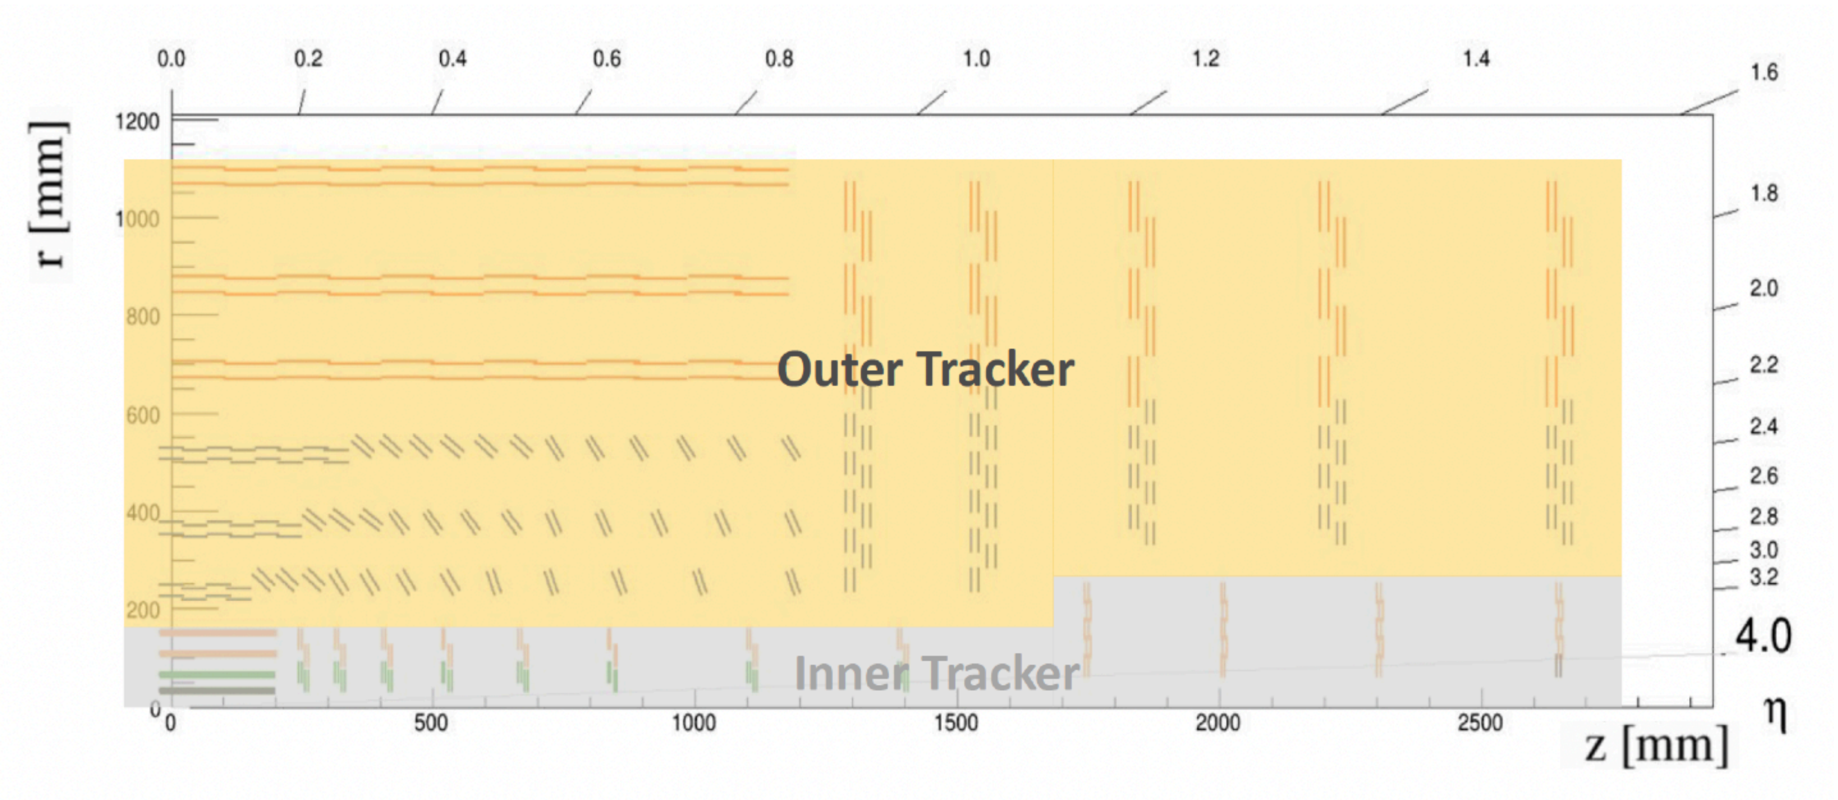

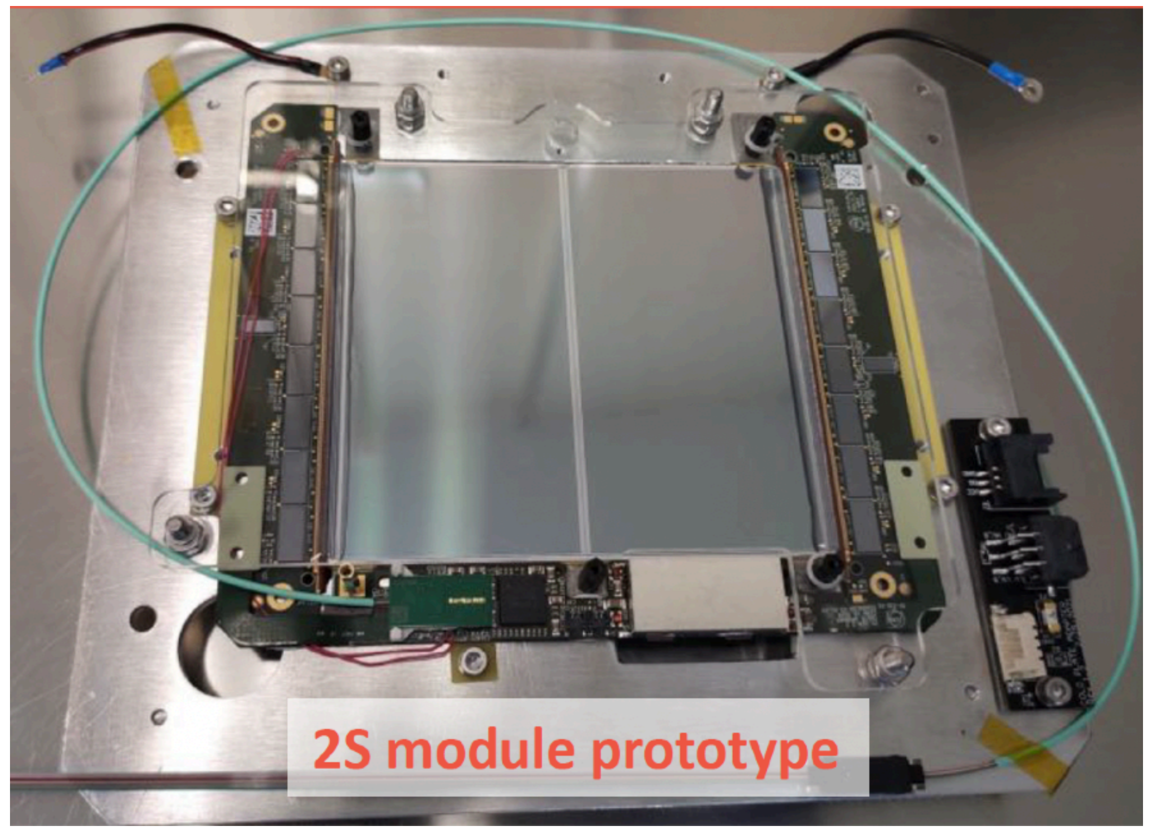

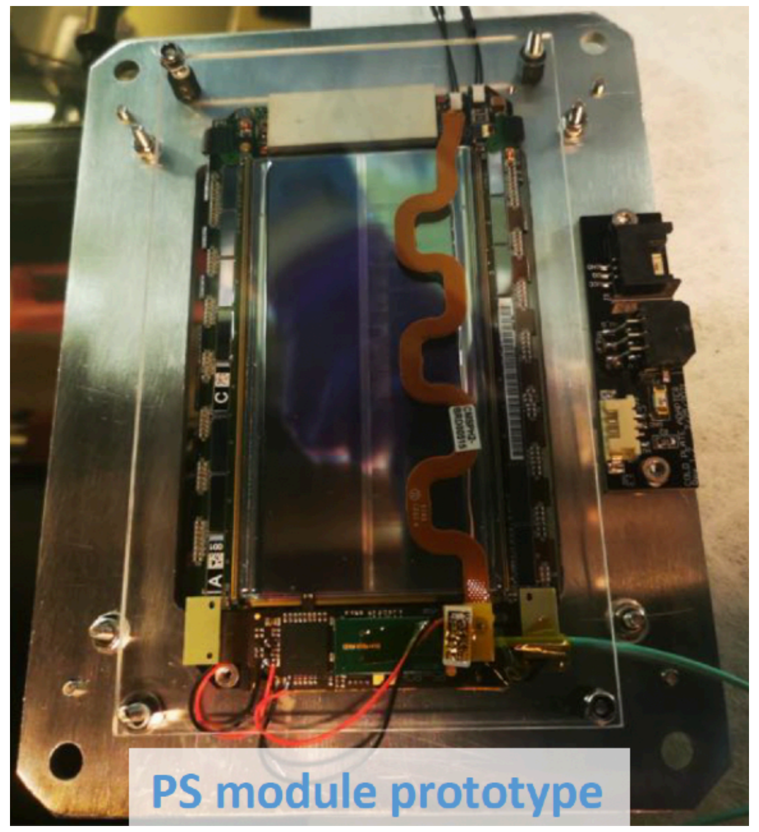

Based on silicon strip and strip+pixel sensors

- 2S modules of 10x10 cm^2
- PS modules 5x10 cm^2
- Both types implement a novel idea to reduce the transmitted trigger data: **sensor doublet**
  - Filters out energy deposits not compatible with particle trajectories with pt<2 GeV
  - Factor 10 data reduction

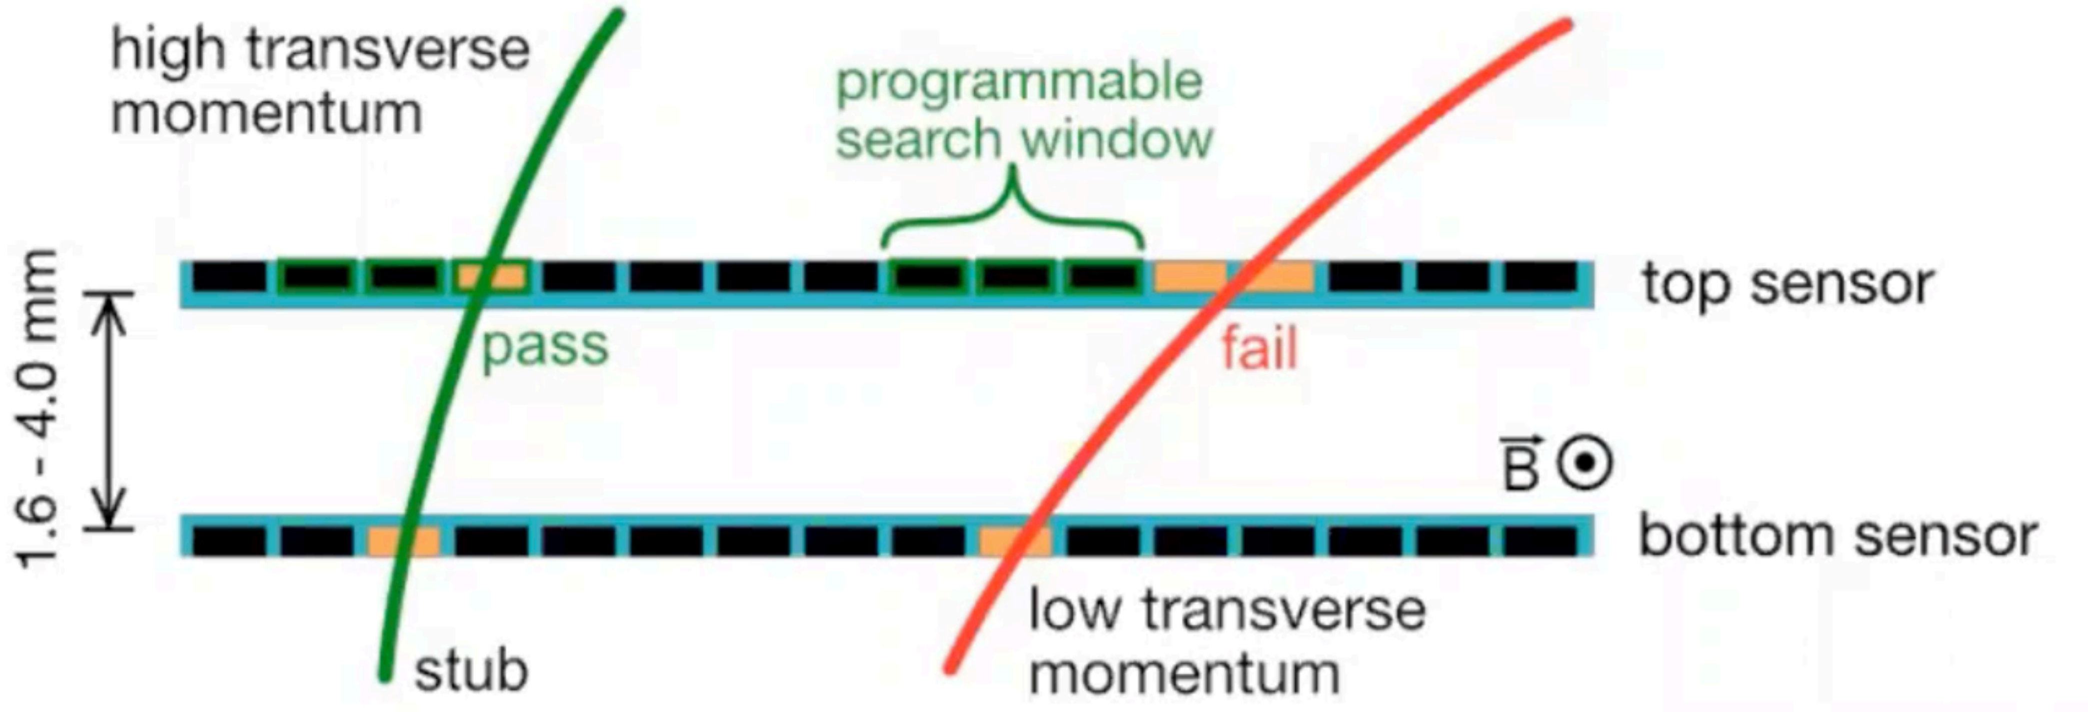

Source: CMS

HL-LHC Outer Tracker - Read-out

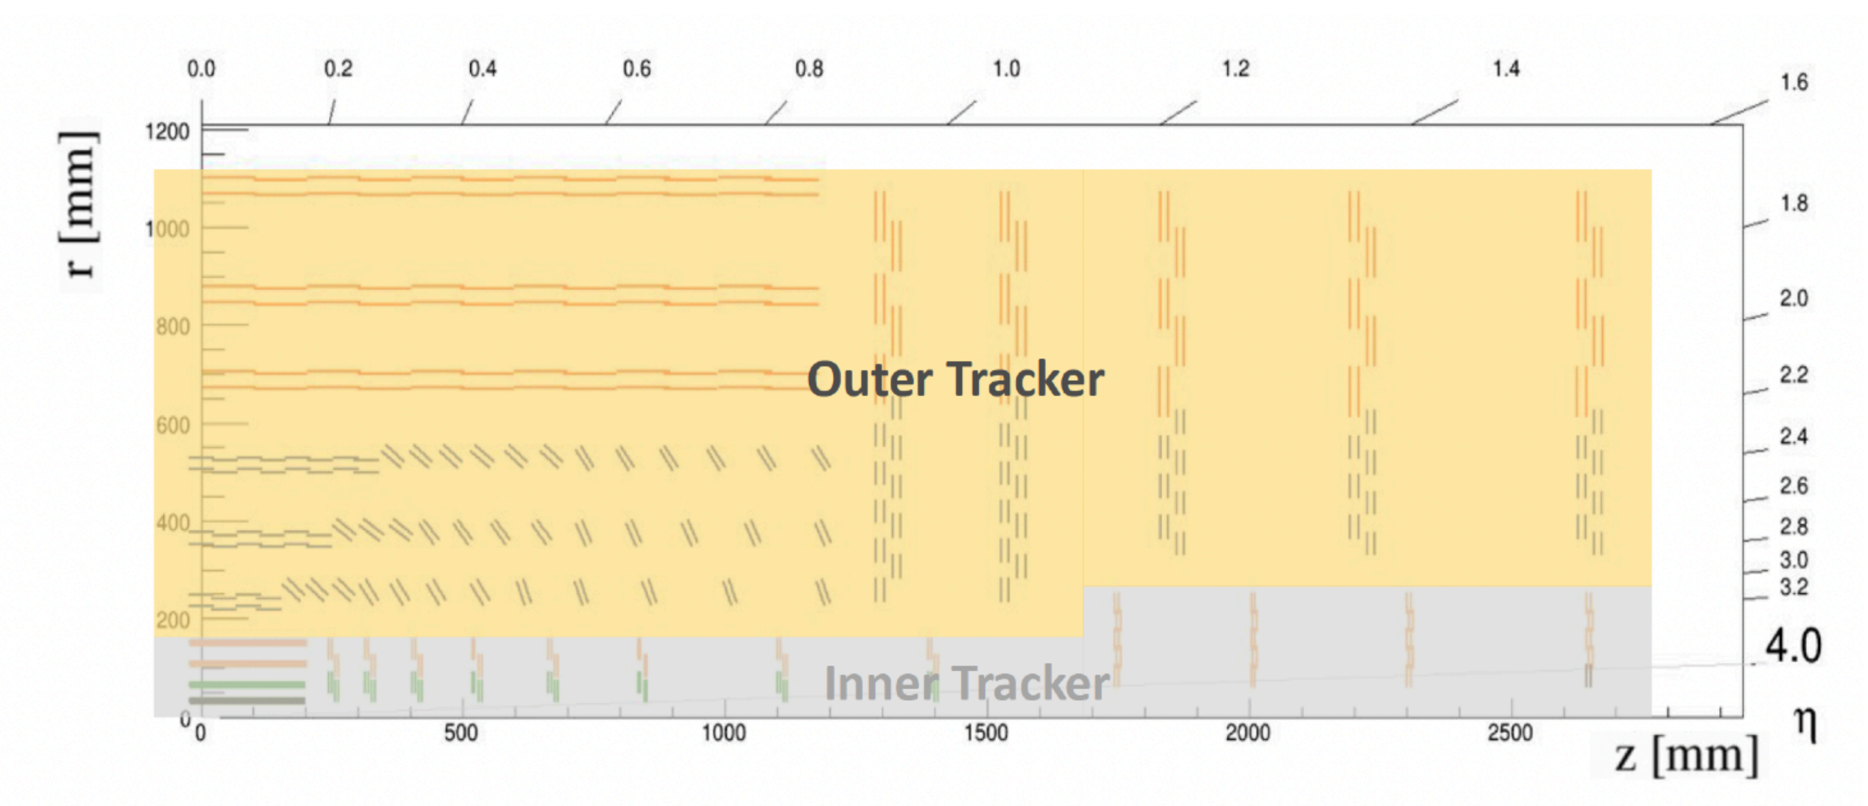

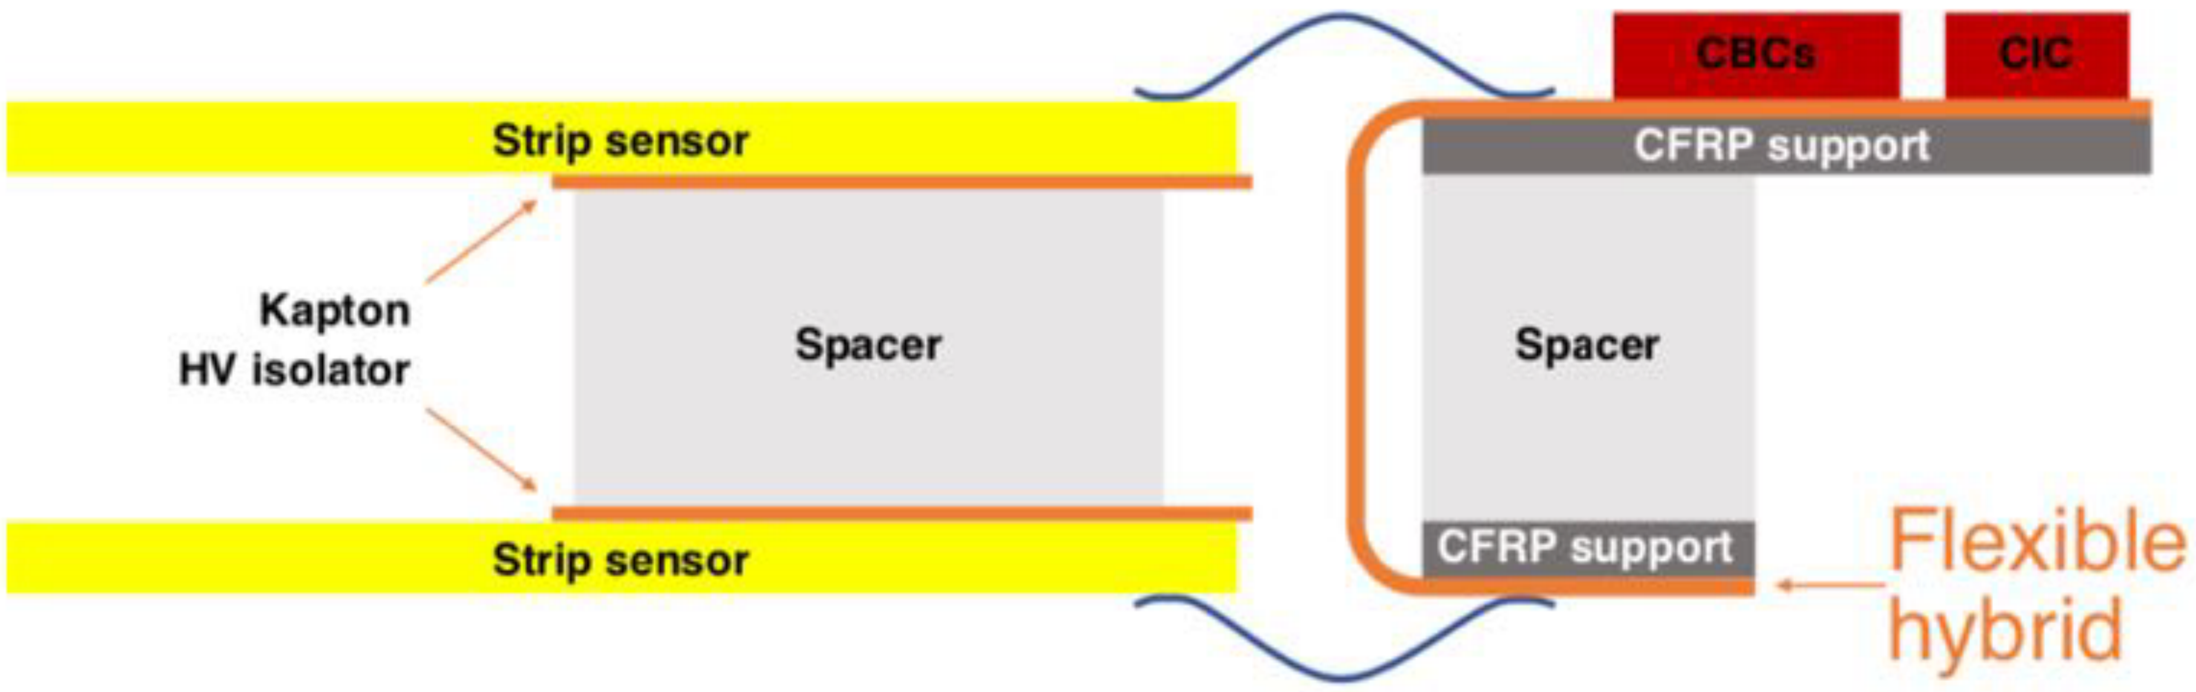

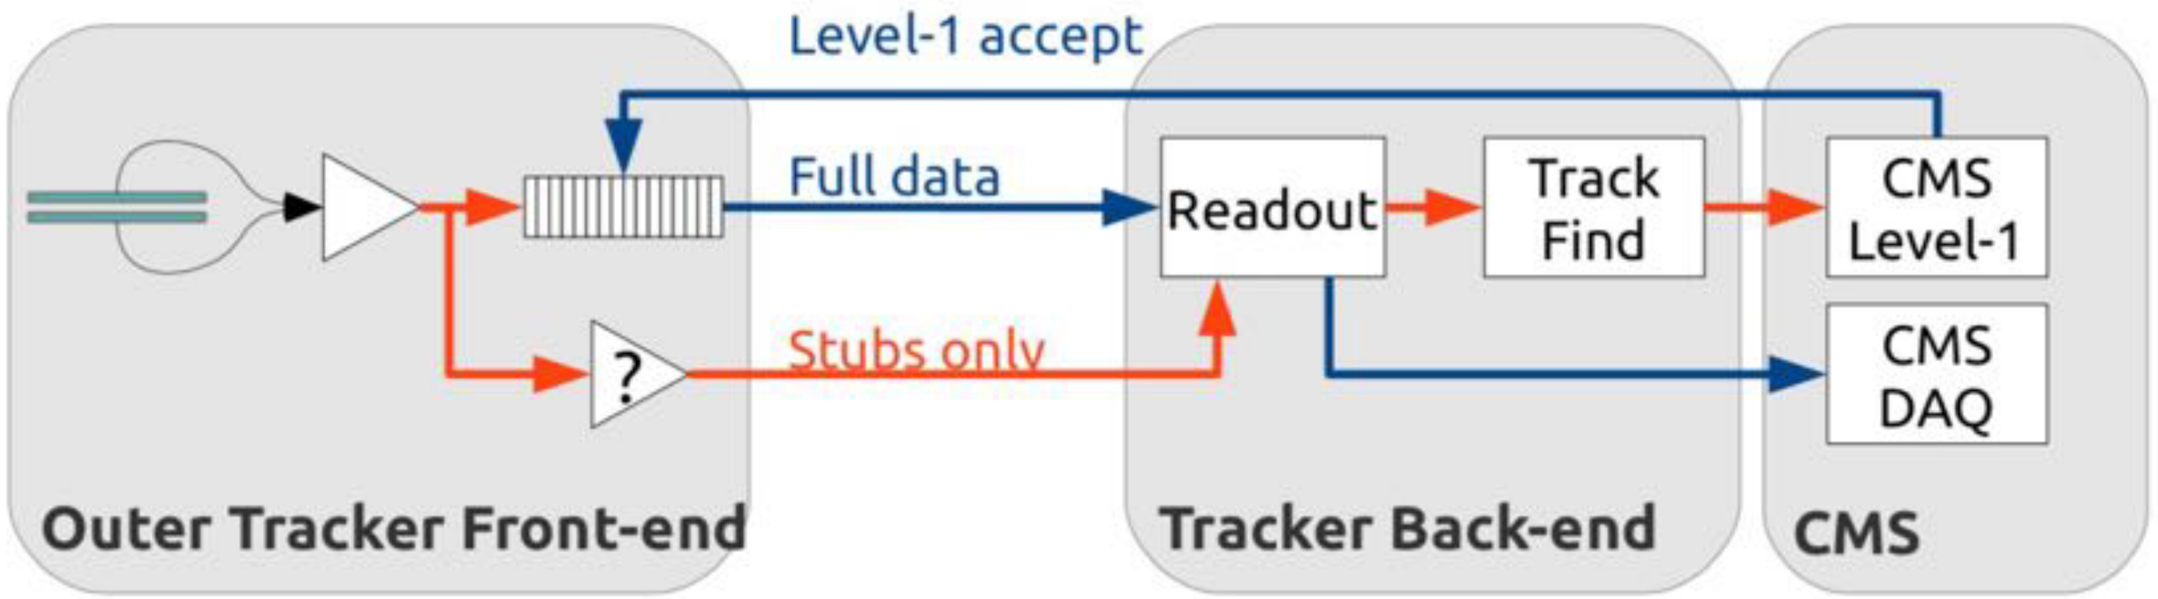

- PS and 2S modules have custom read-out electronics
- Both modules have a common custom chip to gather the information from the read-out chip and receive the L1 Trigger information
- Data is transmitted via optical fibre using a CERN developed driver VTRx+ using a custom protocol LpGBT

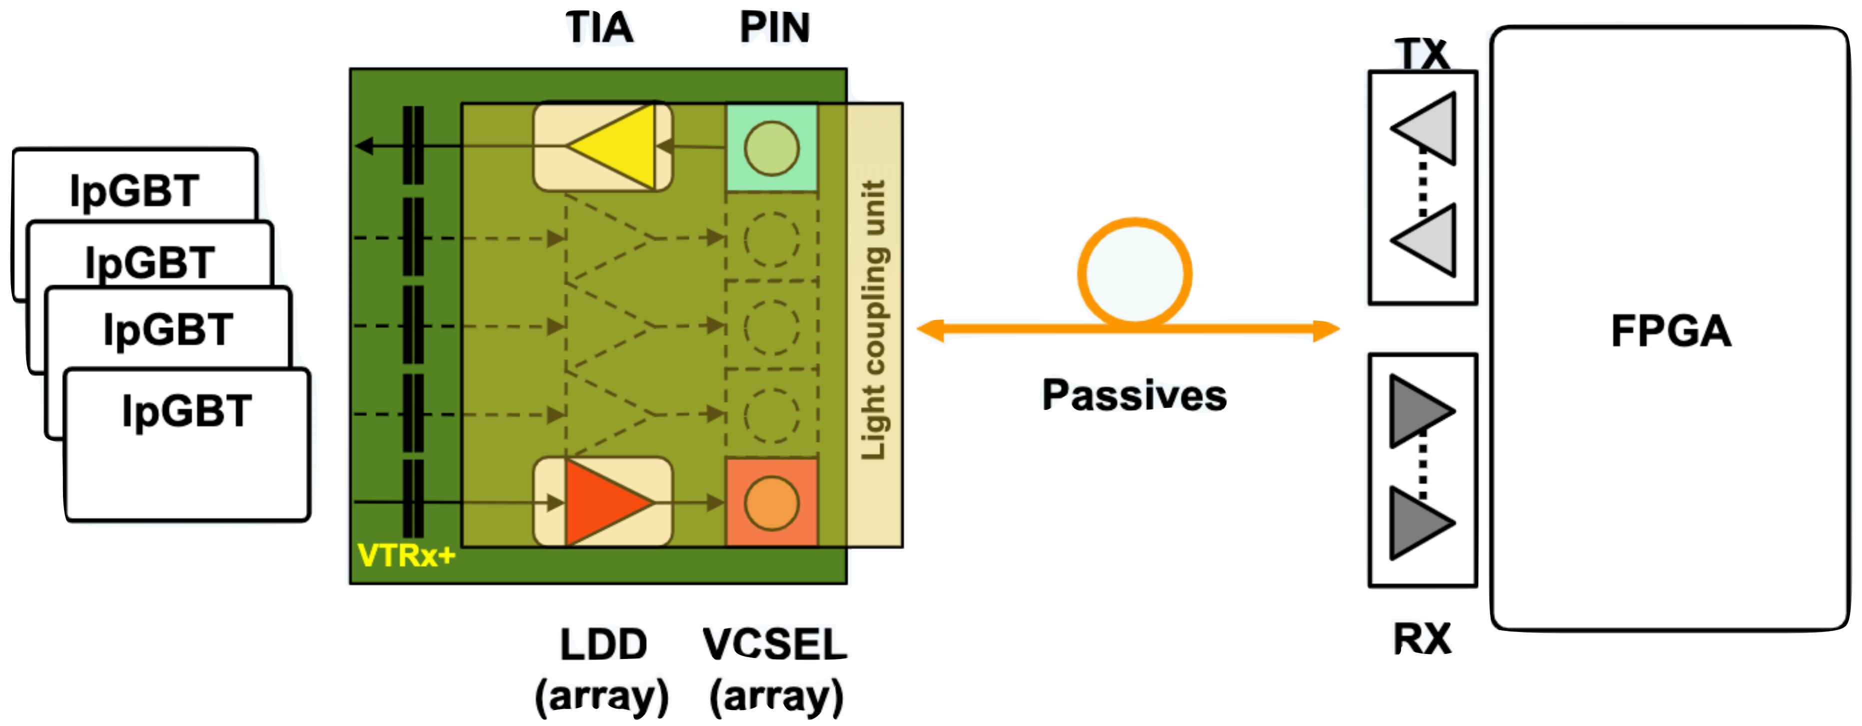

Source: CMS

HL-LHC Outer Tracker - Trigger path

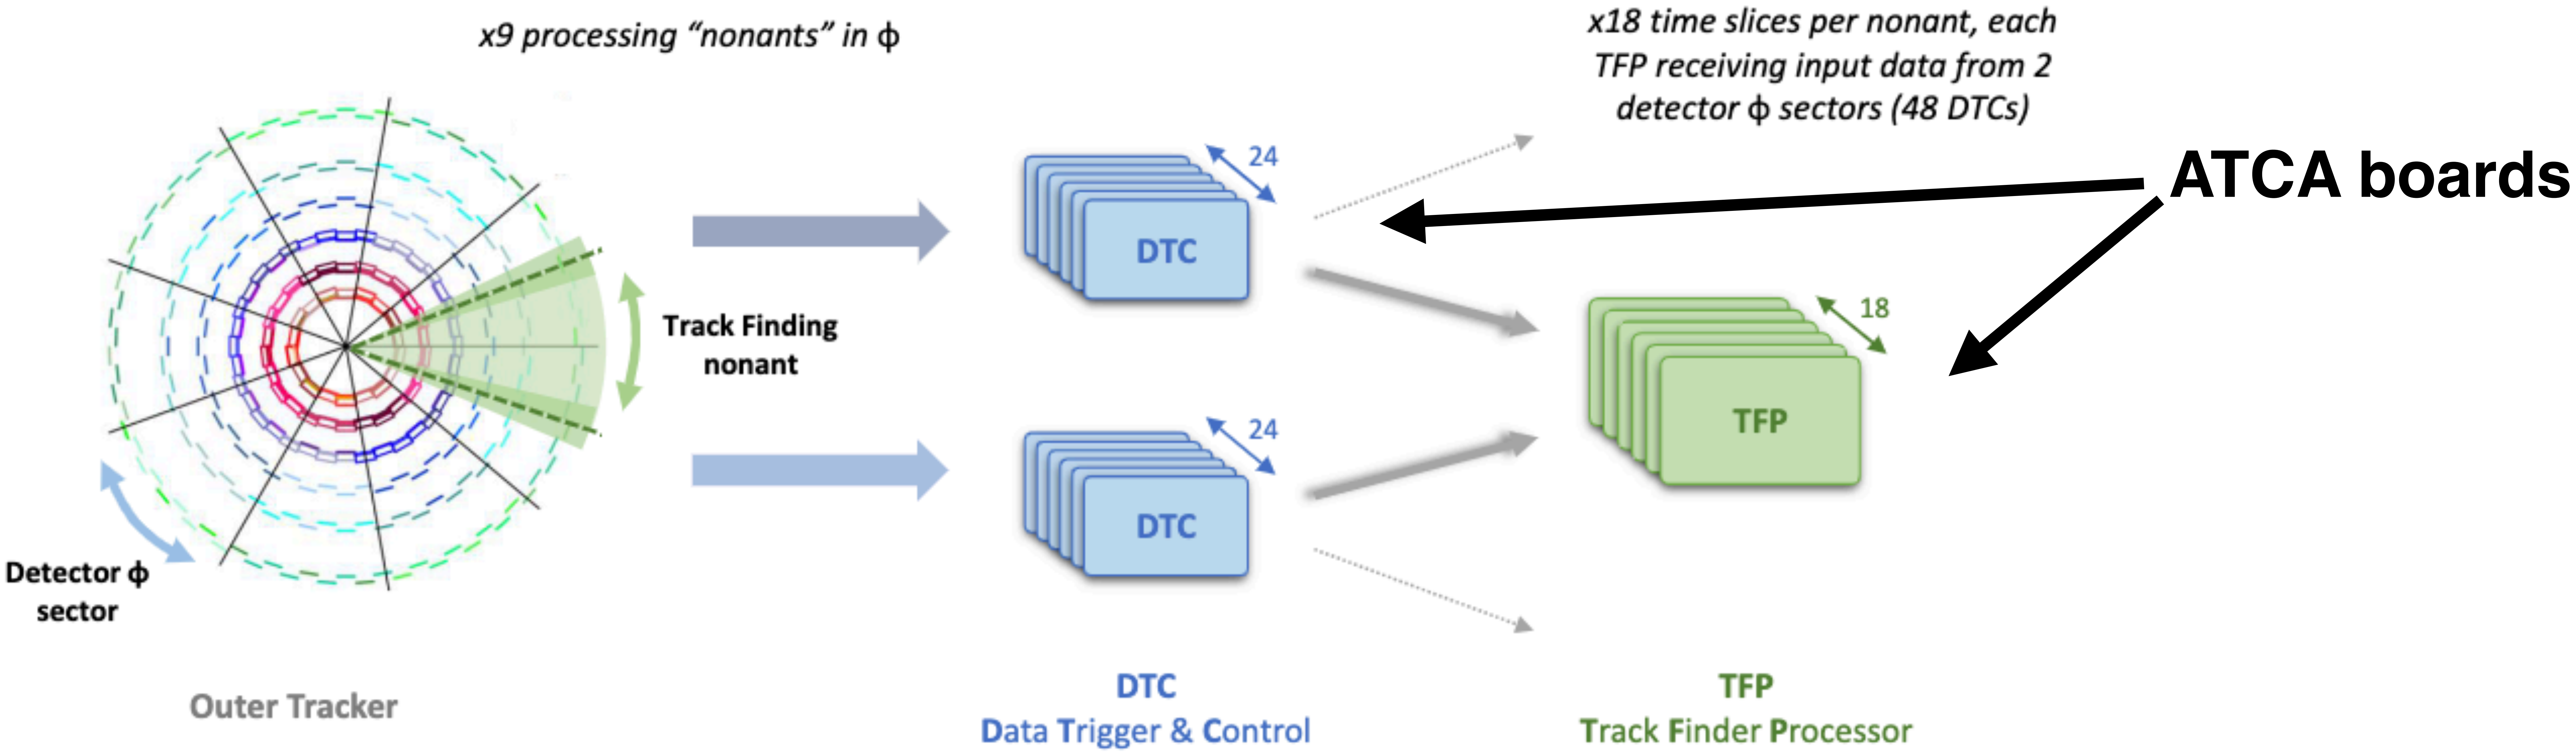

- Trigger data flows from the front end electronics to the **track finder board** (TFP) to generate trigger primitives (**tracks**)
- Track reconstruction at trigger Level 1 in less than **5 μs**
- Data from outer tracker in |η|<2.4 are sent to a data carrier board (DTC), which sends the data to the track finder board (TFP) with a **time multiplexing of x18**
- Each TFP board receives an event every 450 ns

Source: CMS

ATCA boards

- Based on Advanced Telecommunication Computing Standard: ATCA
- Each board is 280 mm deep and 322 mm high
- All the **back-end systems** and **L1 Trigger** will be implemented in ATCA boards
- Boards are installed in a 14-slot ATCA crate
- Tracker back-end will use custom **ATCA boards**:
  - **Serenity**: Outer Tracker DTC
  - **Apollo**: Inner Tracker DTC and TFP
  - **DTH**: Data concentrator, and timing/trigger deliverer

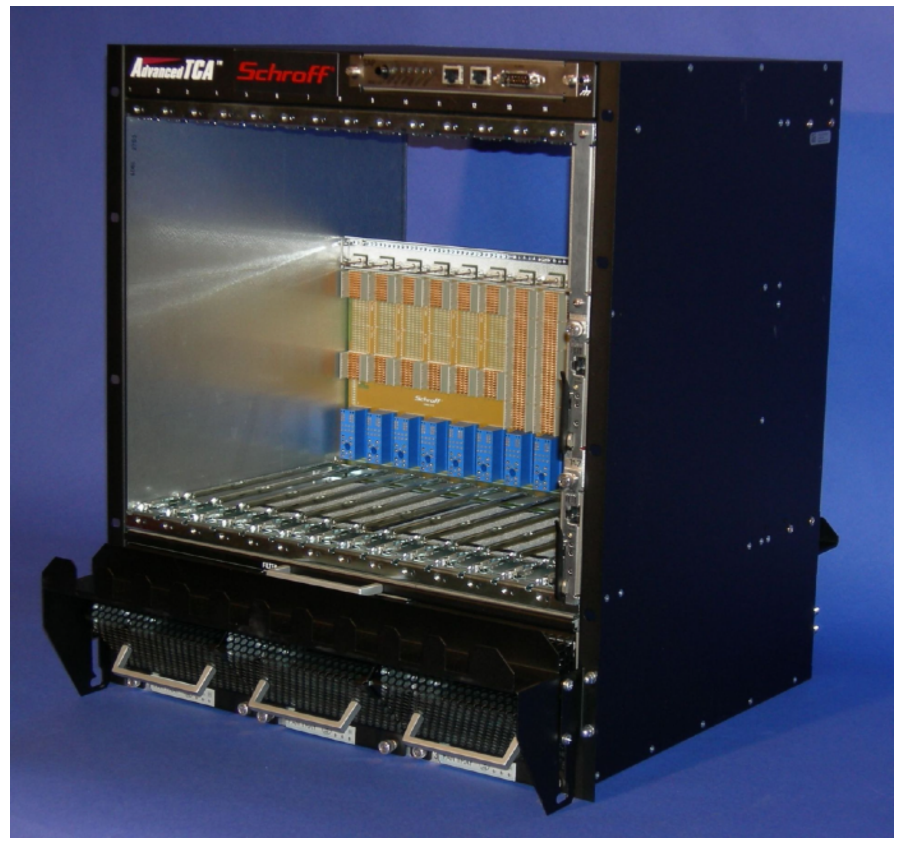

Serenity

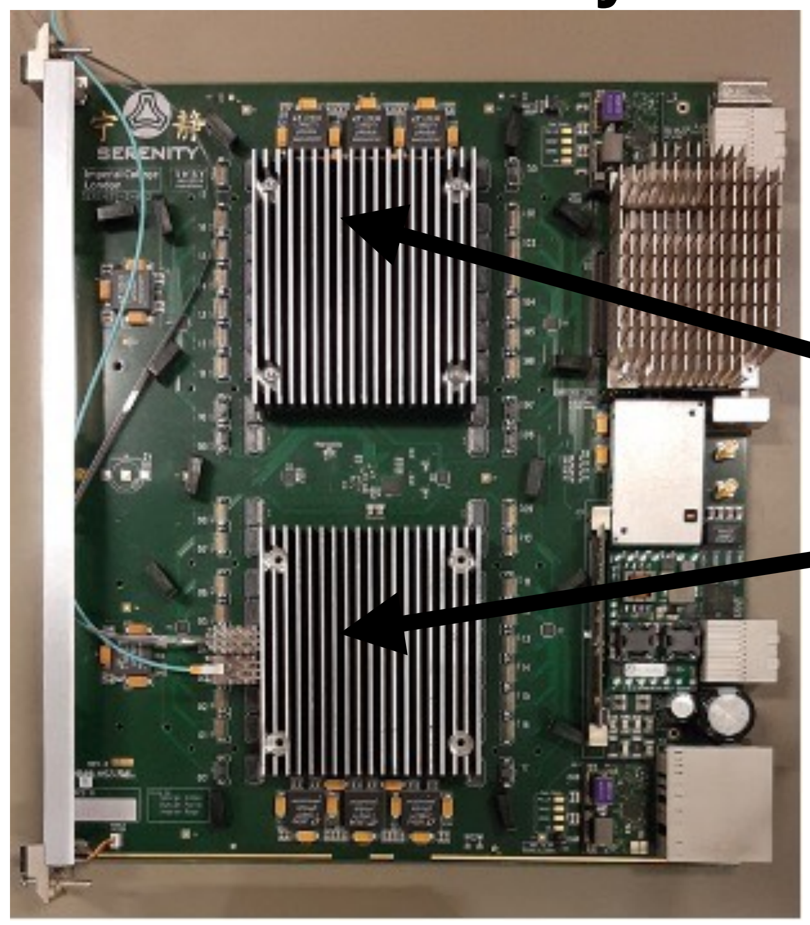

FPGAs (Under the heat sinks)

Apollo

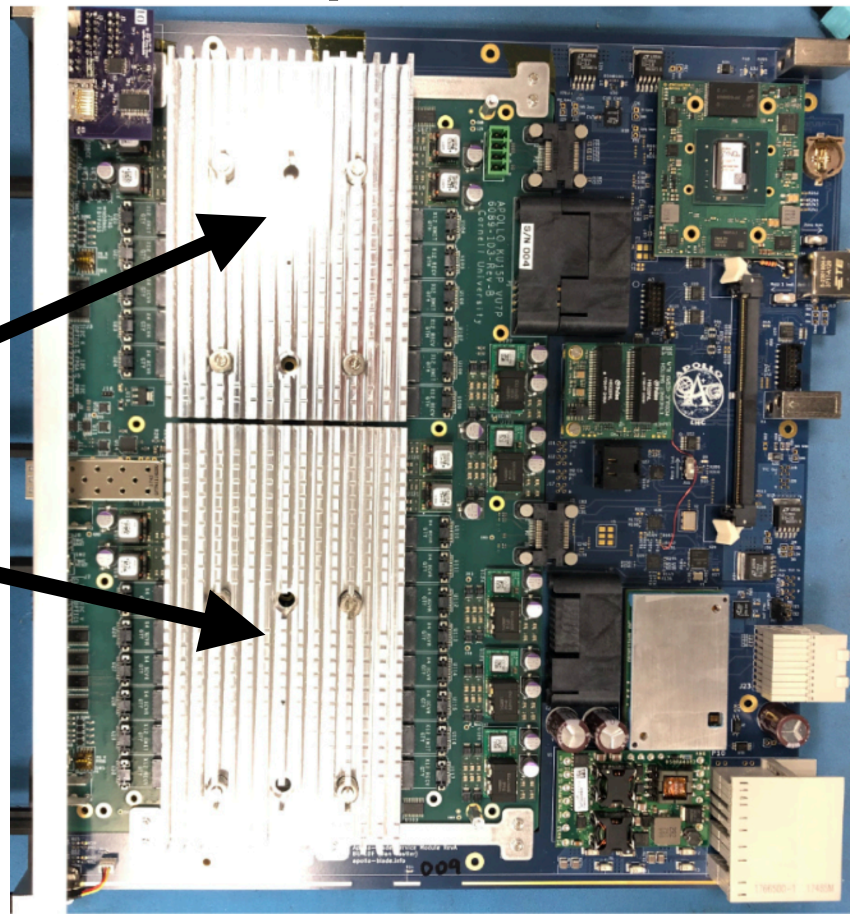

DTH

Prototype ATCA boards

Source: CMS

HL-LHC CMS Tracker data acquisition

Track Trigger Algorithm

Each event could have **10k stubs** and **O(100) reconstructed tracks** → harsh **combinatoric problem**

Algorithm under development: Tracklet + Kalman filter

1. Pattern based on “**tracklet**” seeding
2. **Kalman filter** for identify best stub candidates and track parameters
3. Boost decision tree to evaluate **track quality**

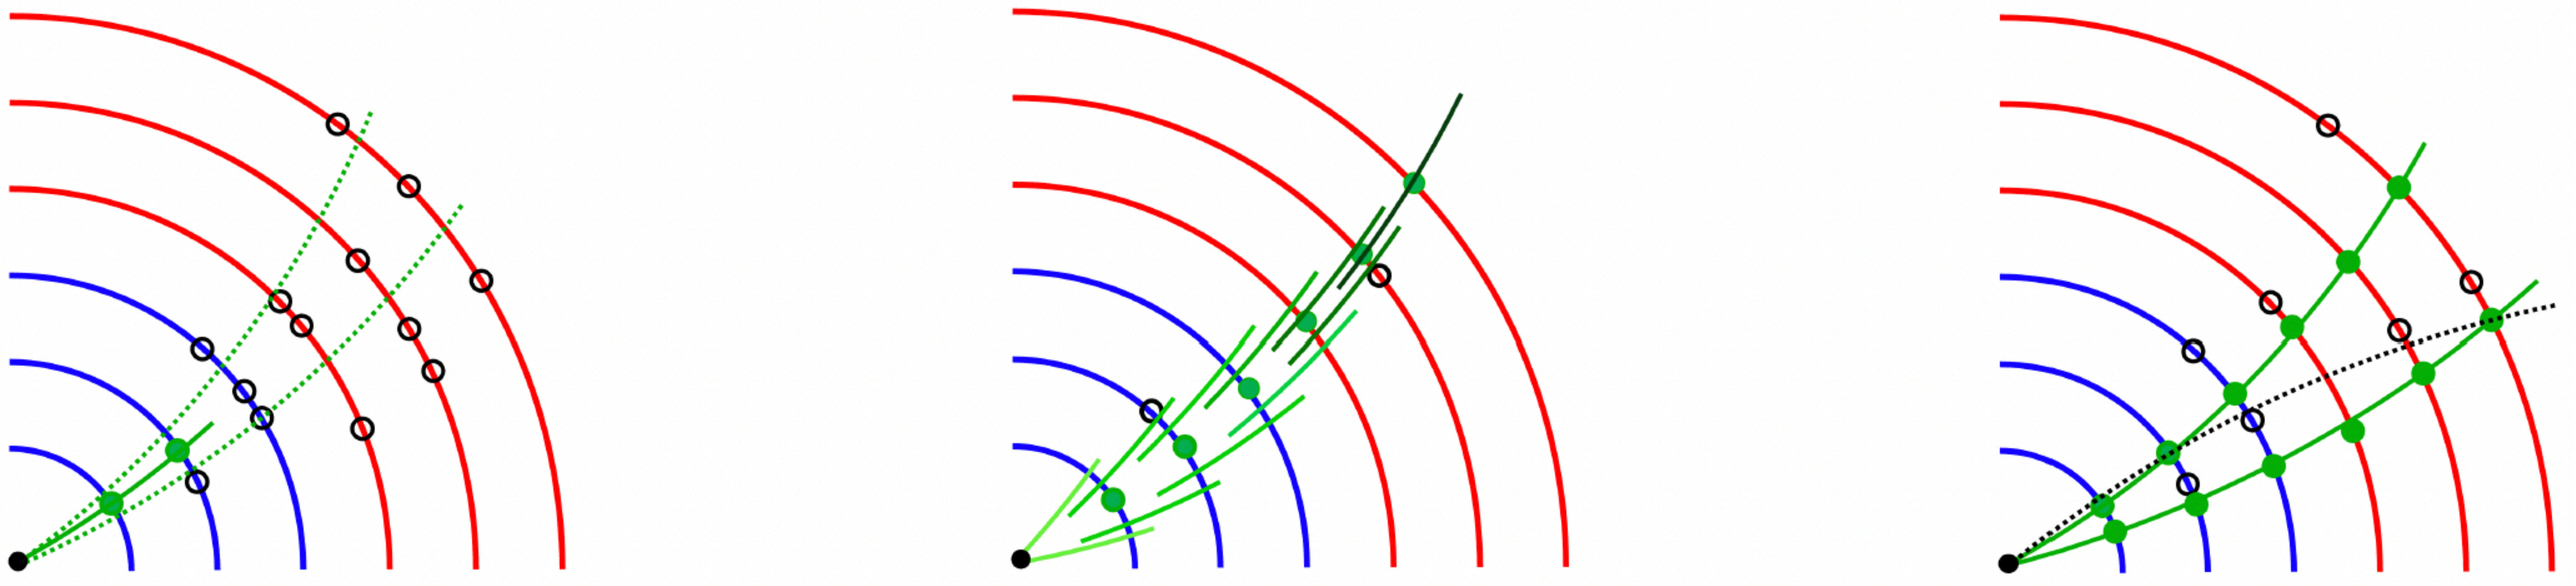

1. 2. 3.

Track Trigger Algorithm Firmware

- Implemented in **HLS and VHDL** in the FPGA of the TFP board
- Pipeline made of processing modules (red) and memory modules (blue)
- Horizontal scalability (**parallelisation**)
- Kalman filter and Track quality written in VHDL
- Targeting 240 MHz
- Full algorithm needs two VU13P FPGA
- Implementation of some stages with ML models under test

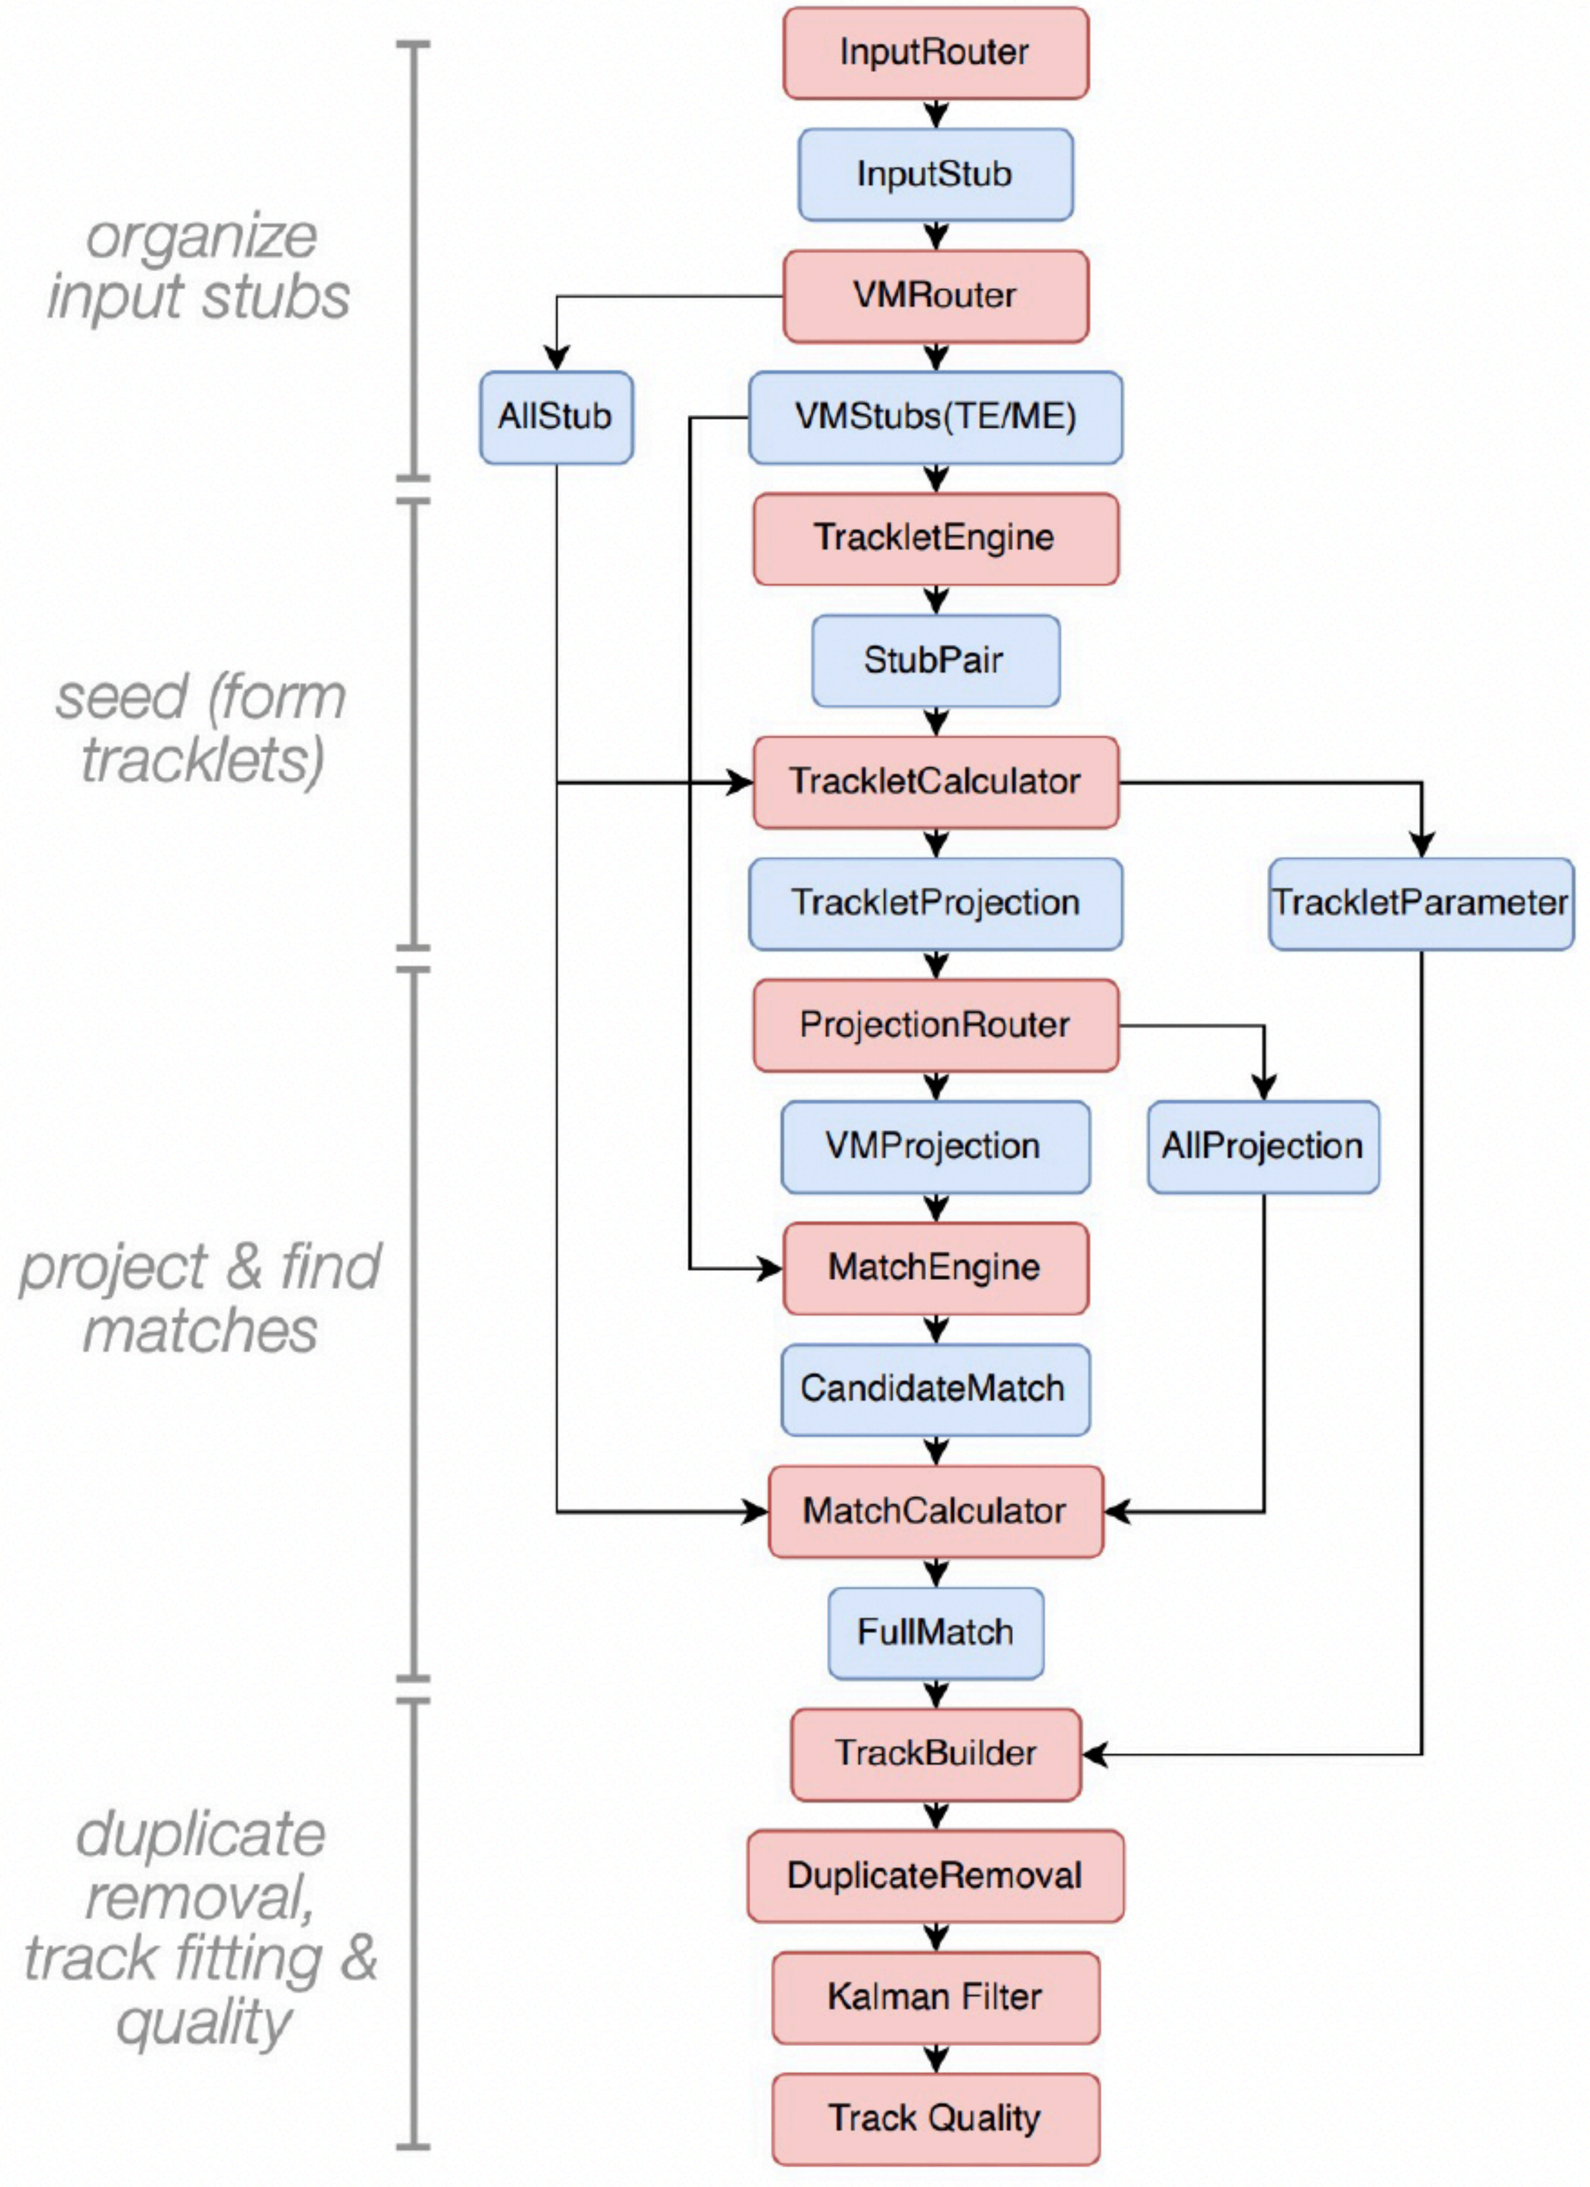

Track Fitter Performance

High efficiency across η

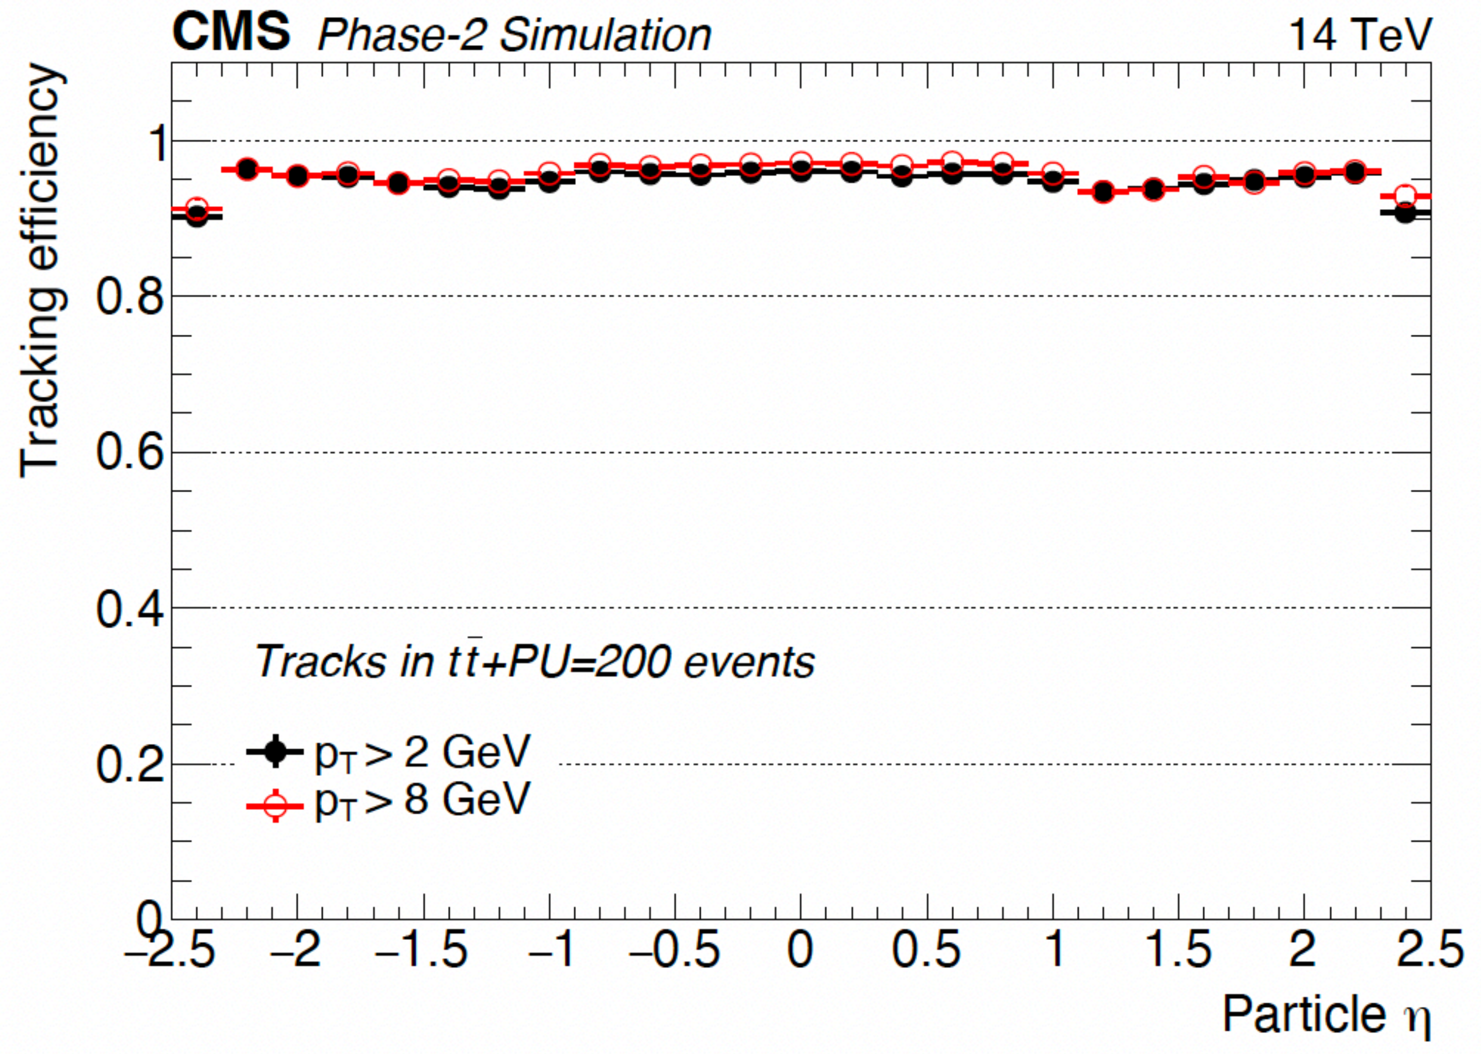

Muon L1 reconstruction

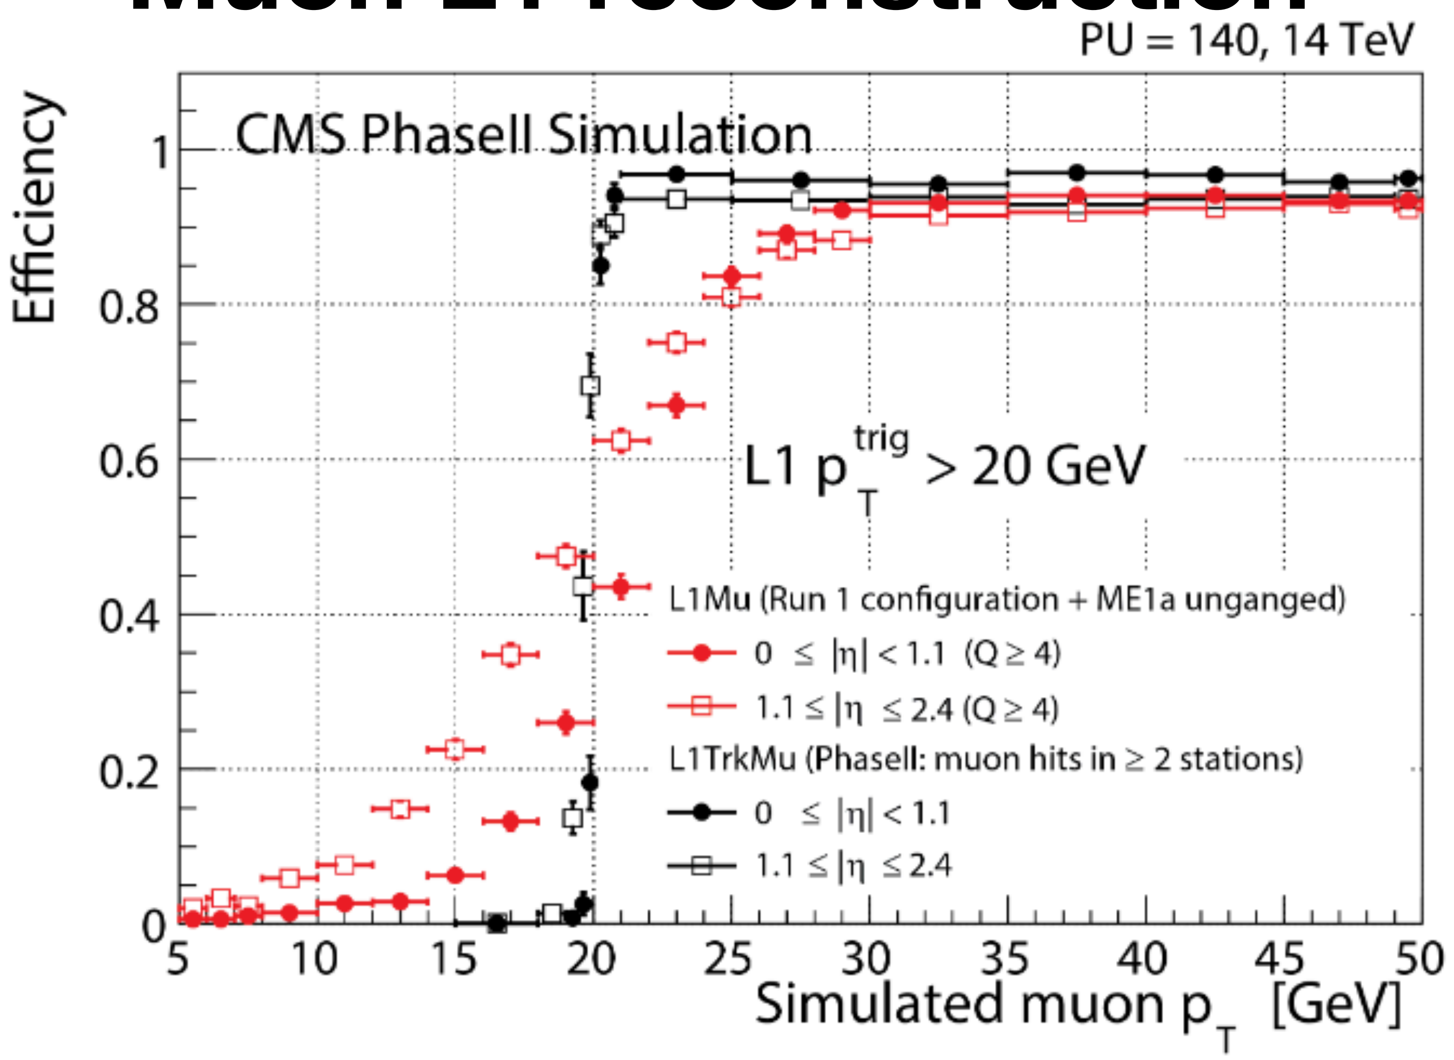

- The reconstruction of tracks in L1 Trigger at HL-LHC allows to keep good detector performance in the presence of high pile-up
  - Primary vertices, jets reconstruction at L1
  - Better resolution, e.g. for muons, implies lower trigger t

Give access to hadronic channels:
e.g. $\phi \to KK$, thus to $B_s \to \phi\phi$

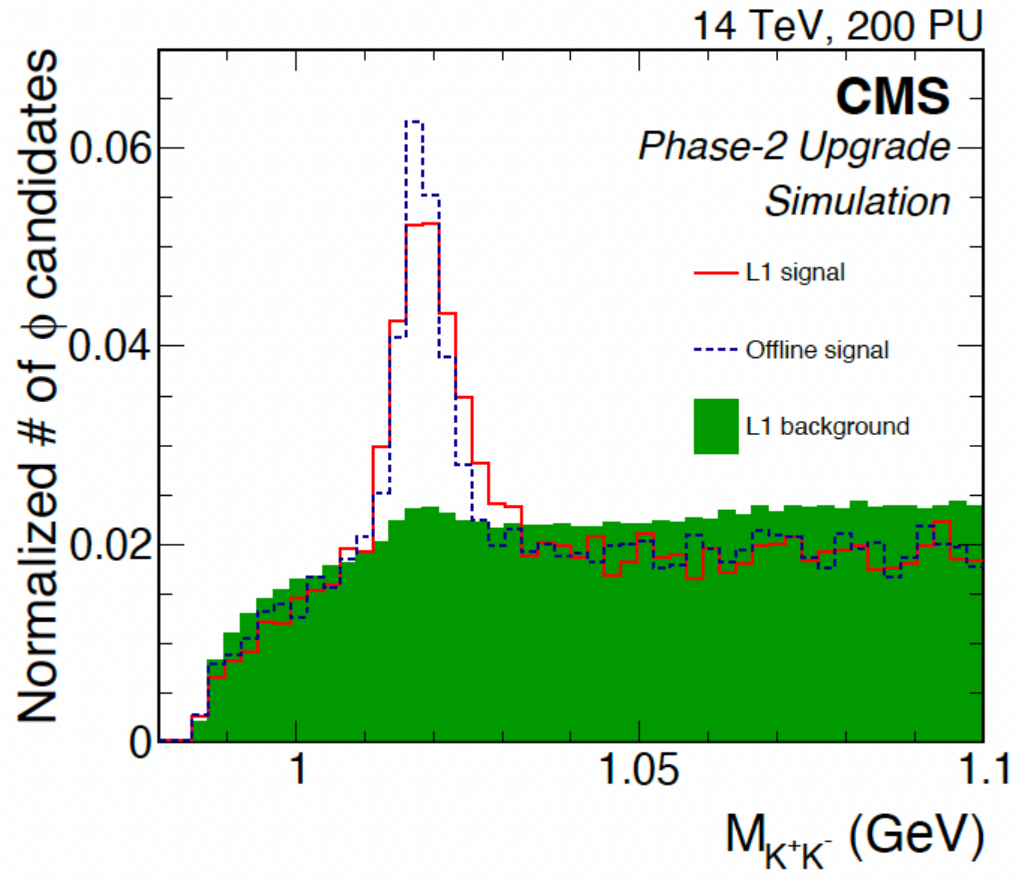

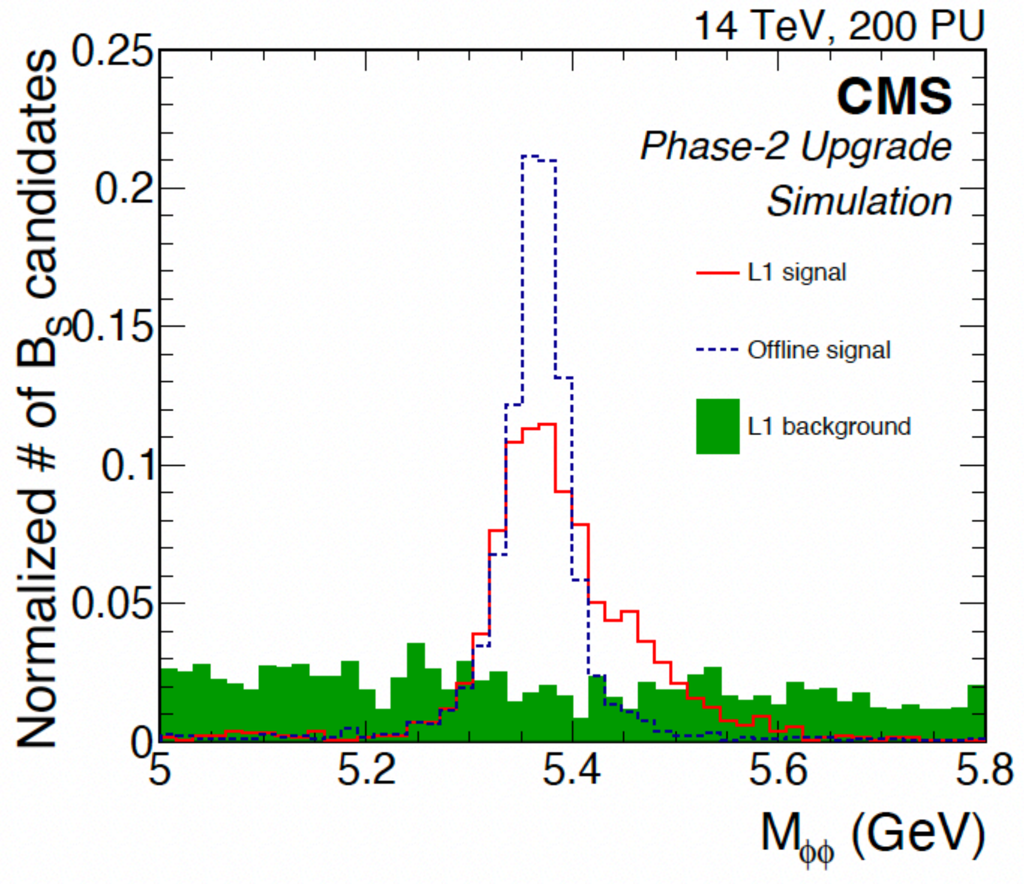

CMS-TDR-014

Summary

- The data acquisition system is fundamental in each HEP experiment
- LHC provides high proton bunch collision rate which generates **high data volumes**
- A **two-level triggering system** is used to find the signature of interesting physics events
- In 2029 **LHC** will start a new phase with **higher luminosity** (higher number of collisions per bunch crossing)
- CMS will update all the sub-detectors and the data acquisition system to cope with the higher pileup
- Details of the future **Tracker** acquisition system have been shown

Additional slides

Source: CMS

Example: Jet Reconstruction

- Clustering of particles into jets and tag the quark origin, e.g. c, s, b, etc
- Implemented using HLS4ML toolkit
- Fast and performant jet reconstruction for FPGA
- Latency: 750 ns for 12 jets, performance similar to standard algorithm

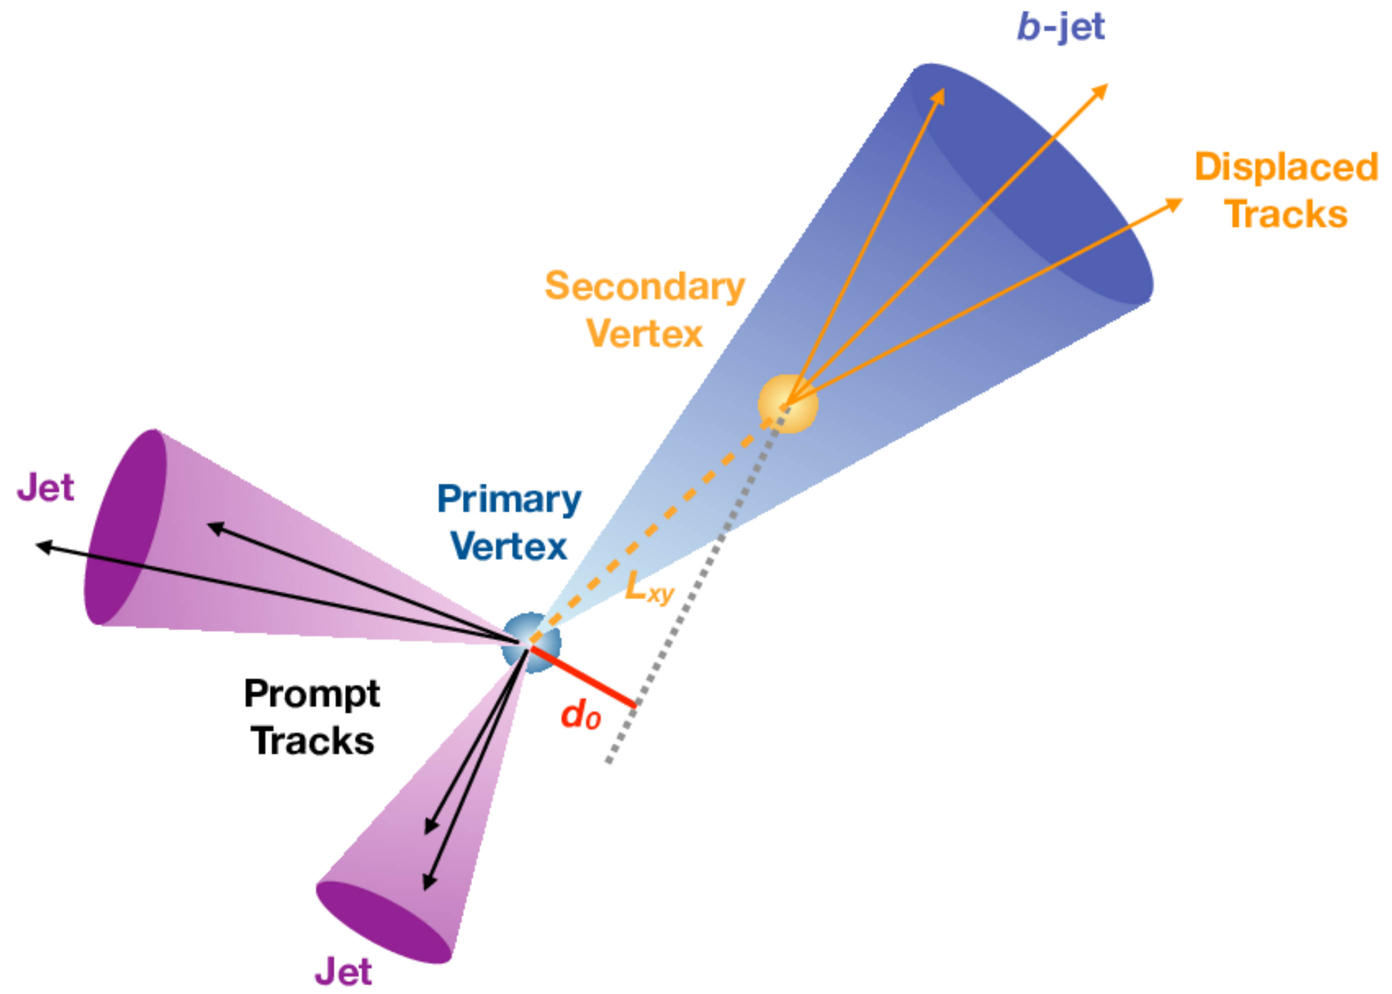

Source: CERN

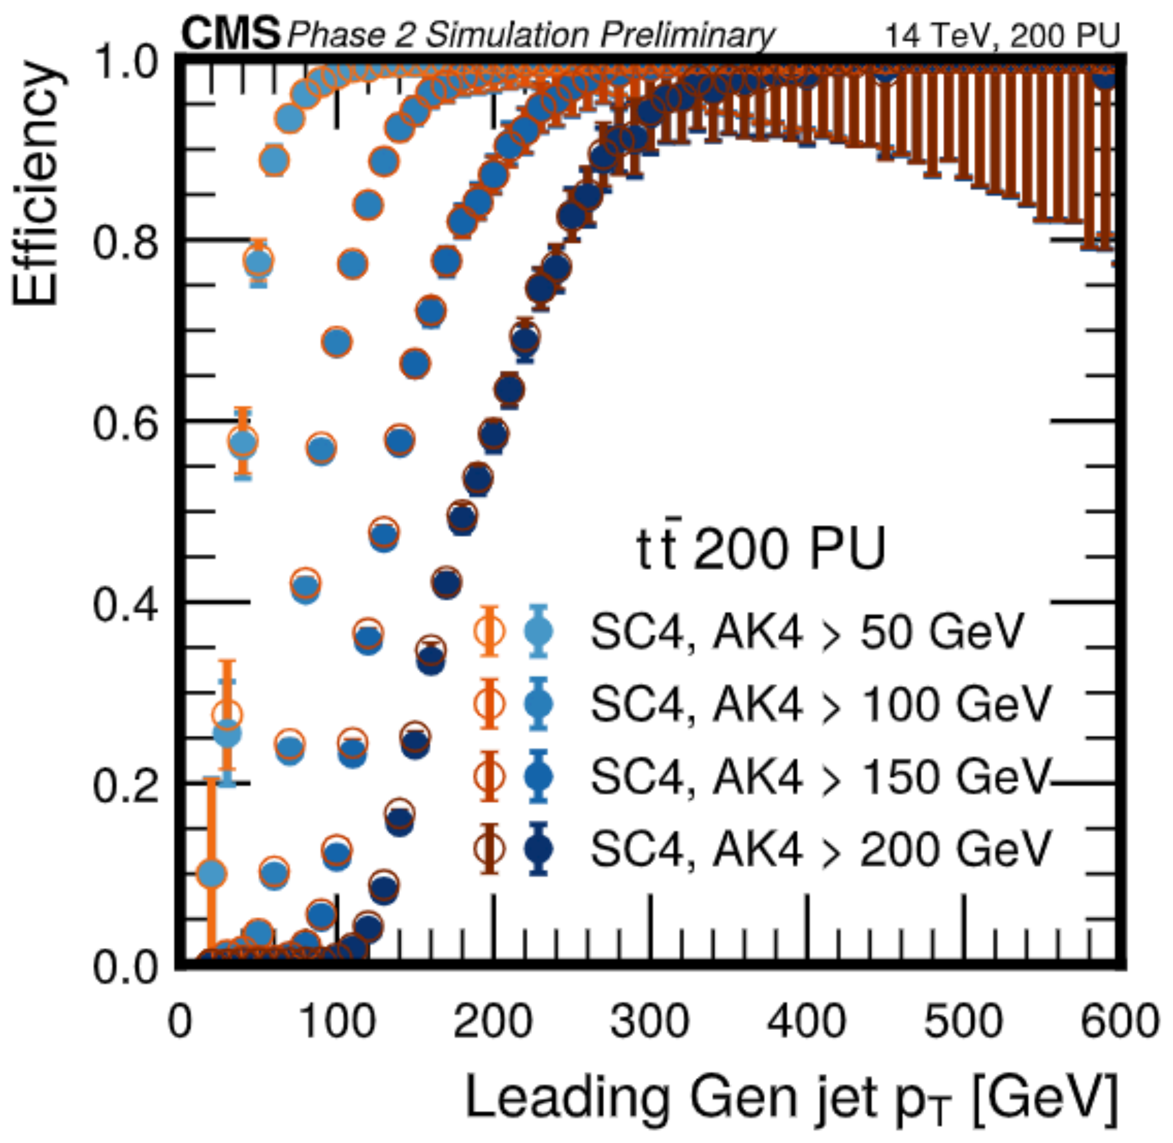

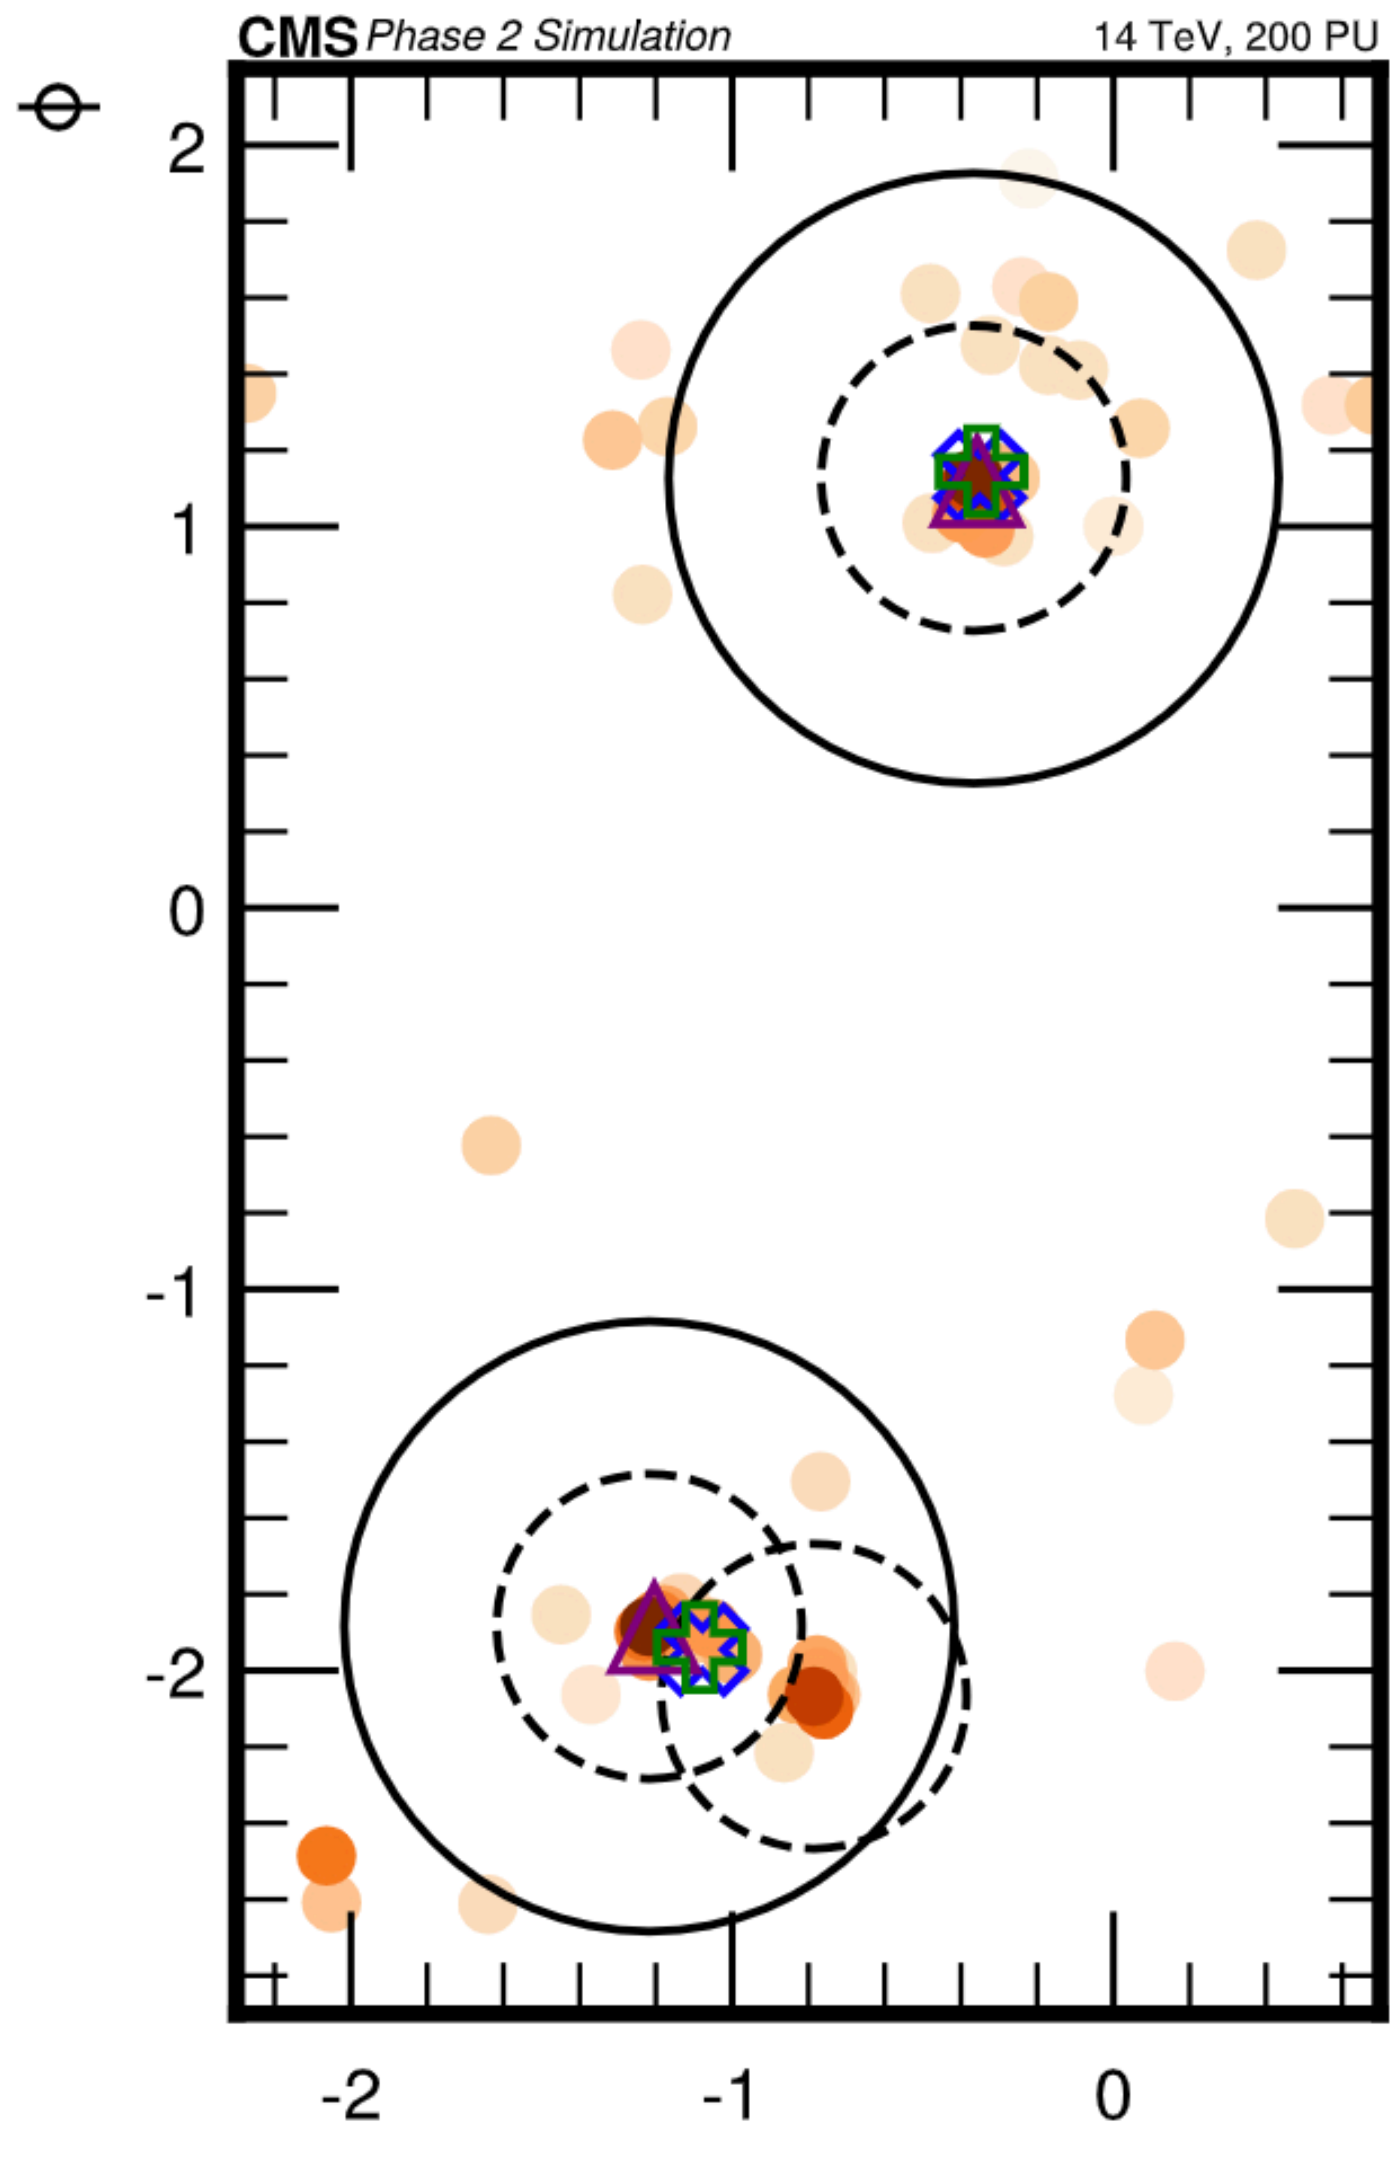

arXiv:2310.08062

Tracker material budget

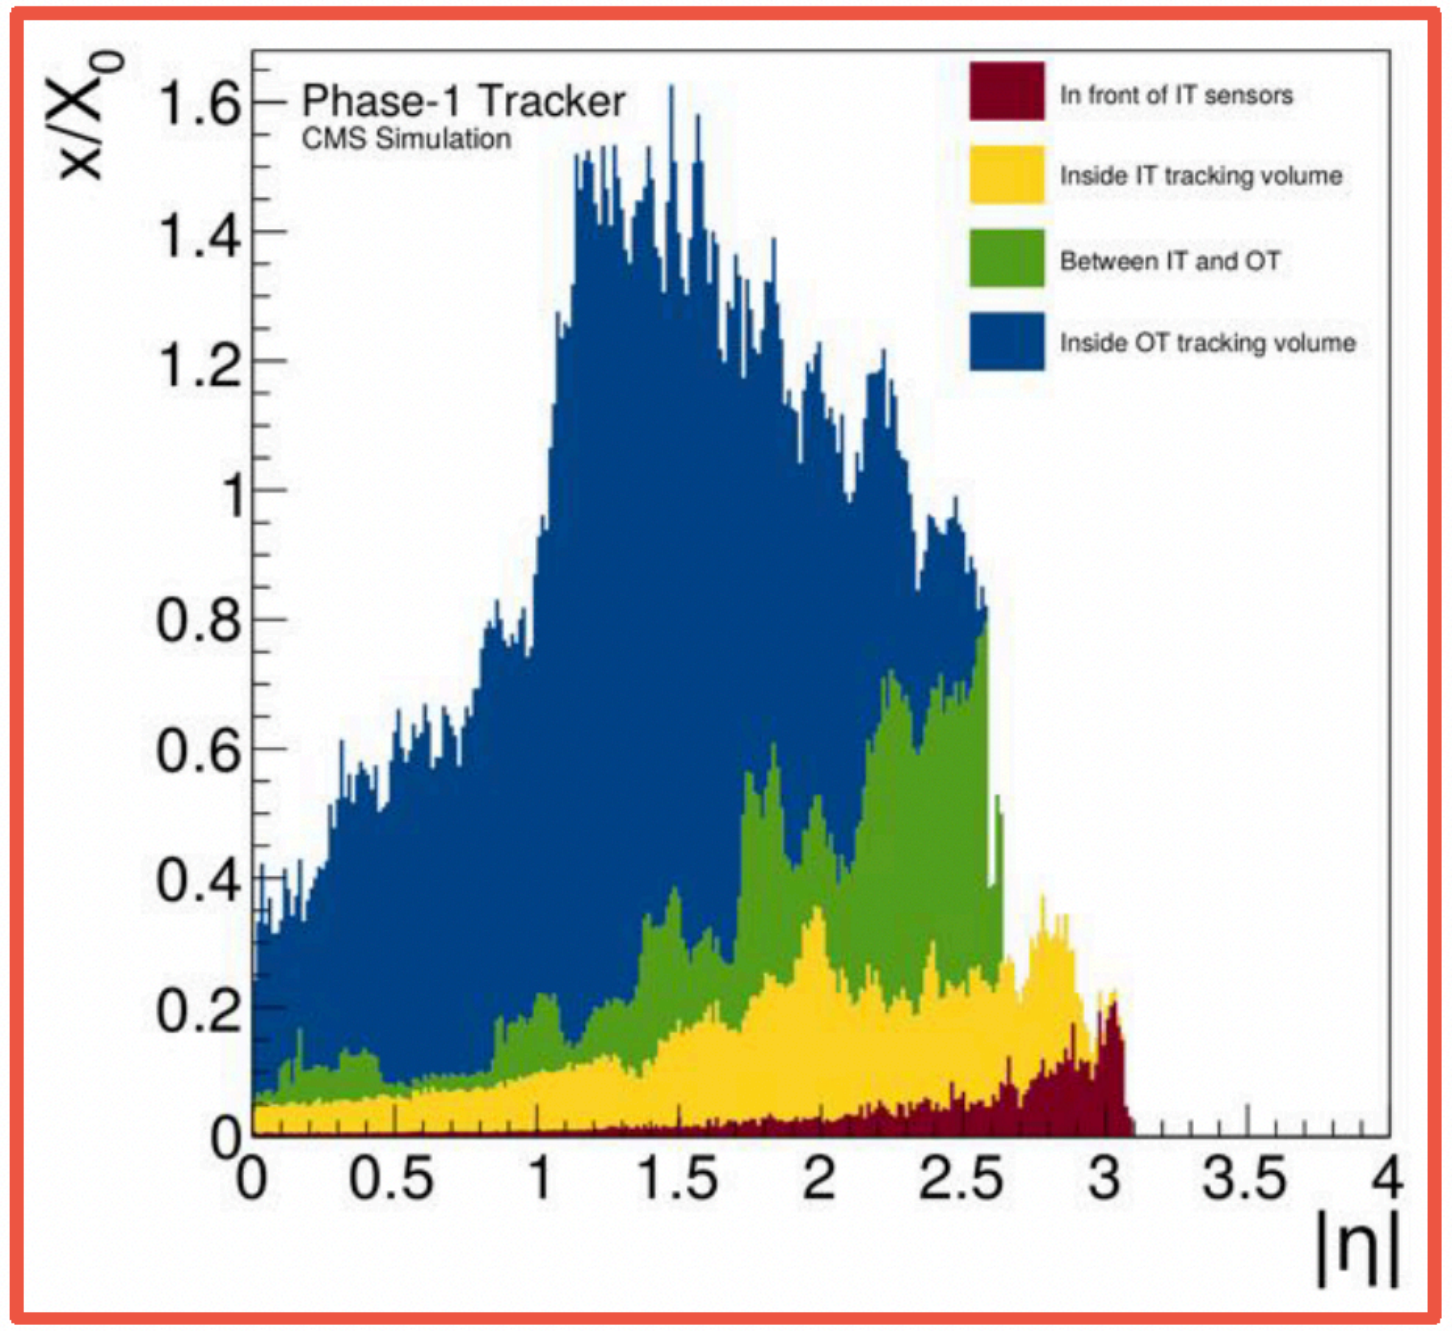

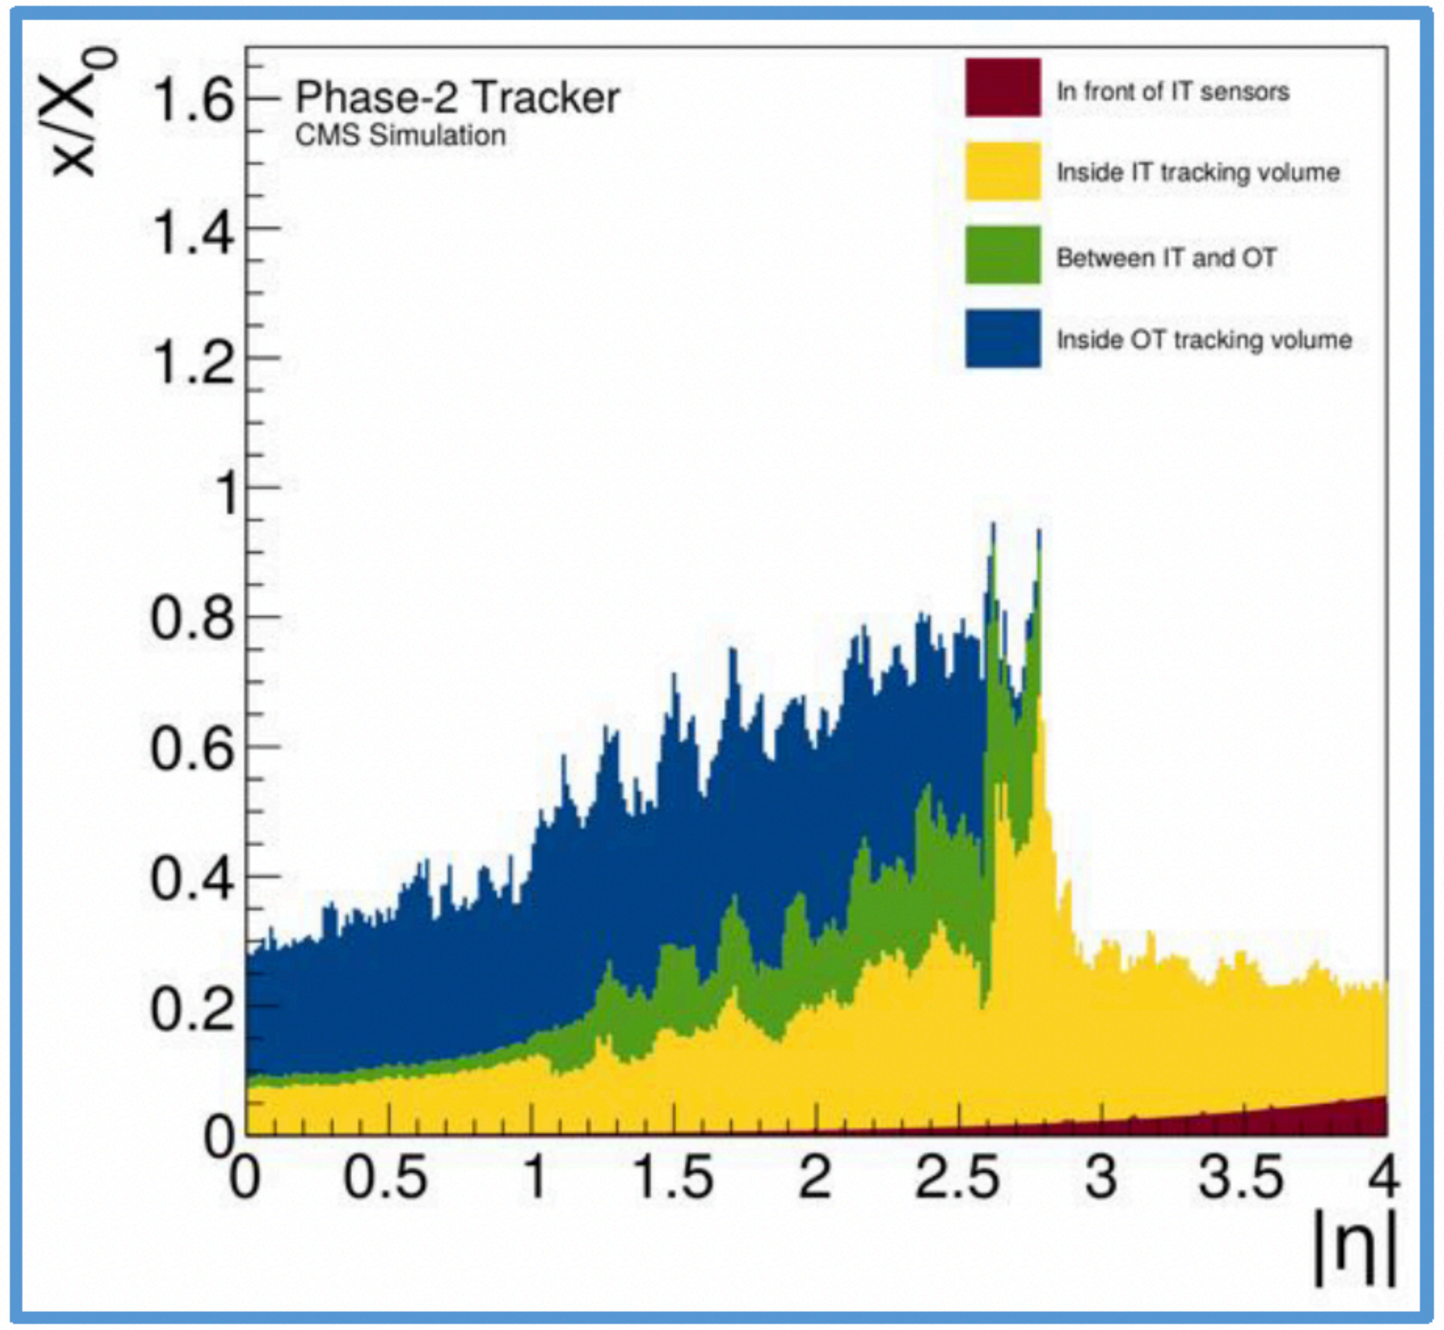

Source: CMS

HL-LHC

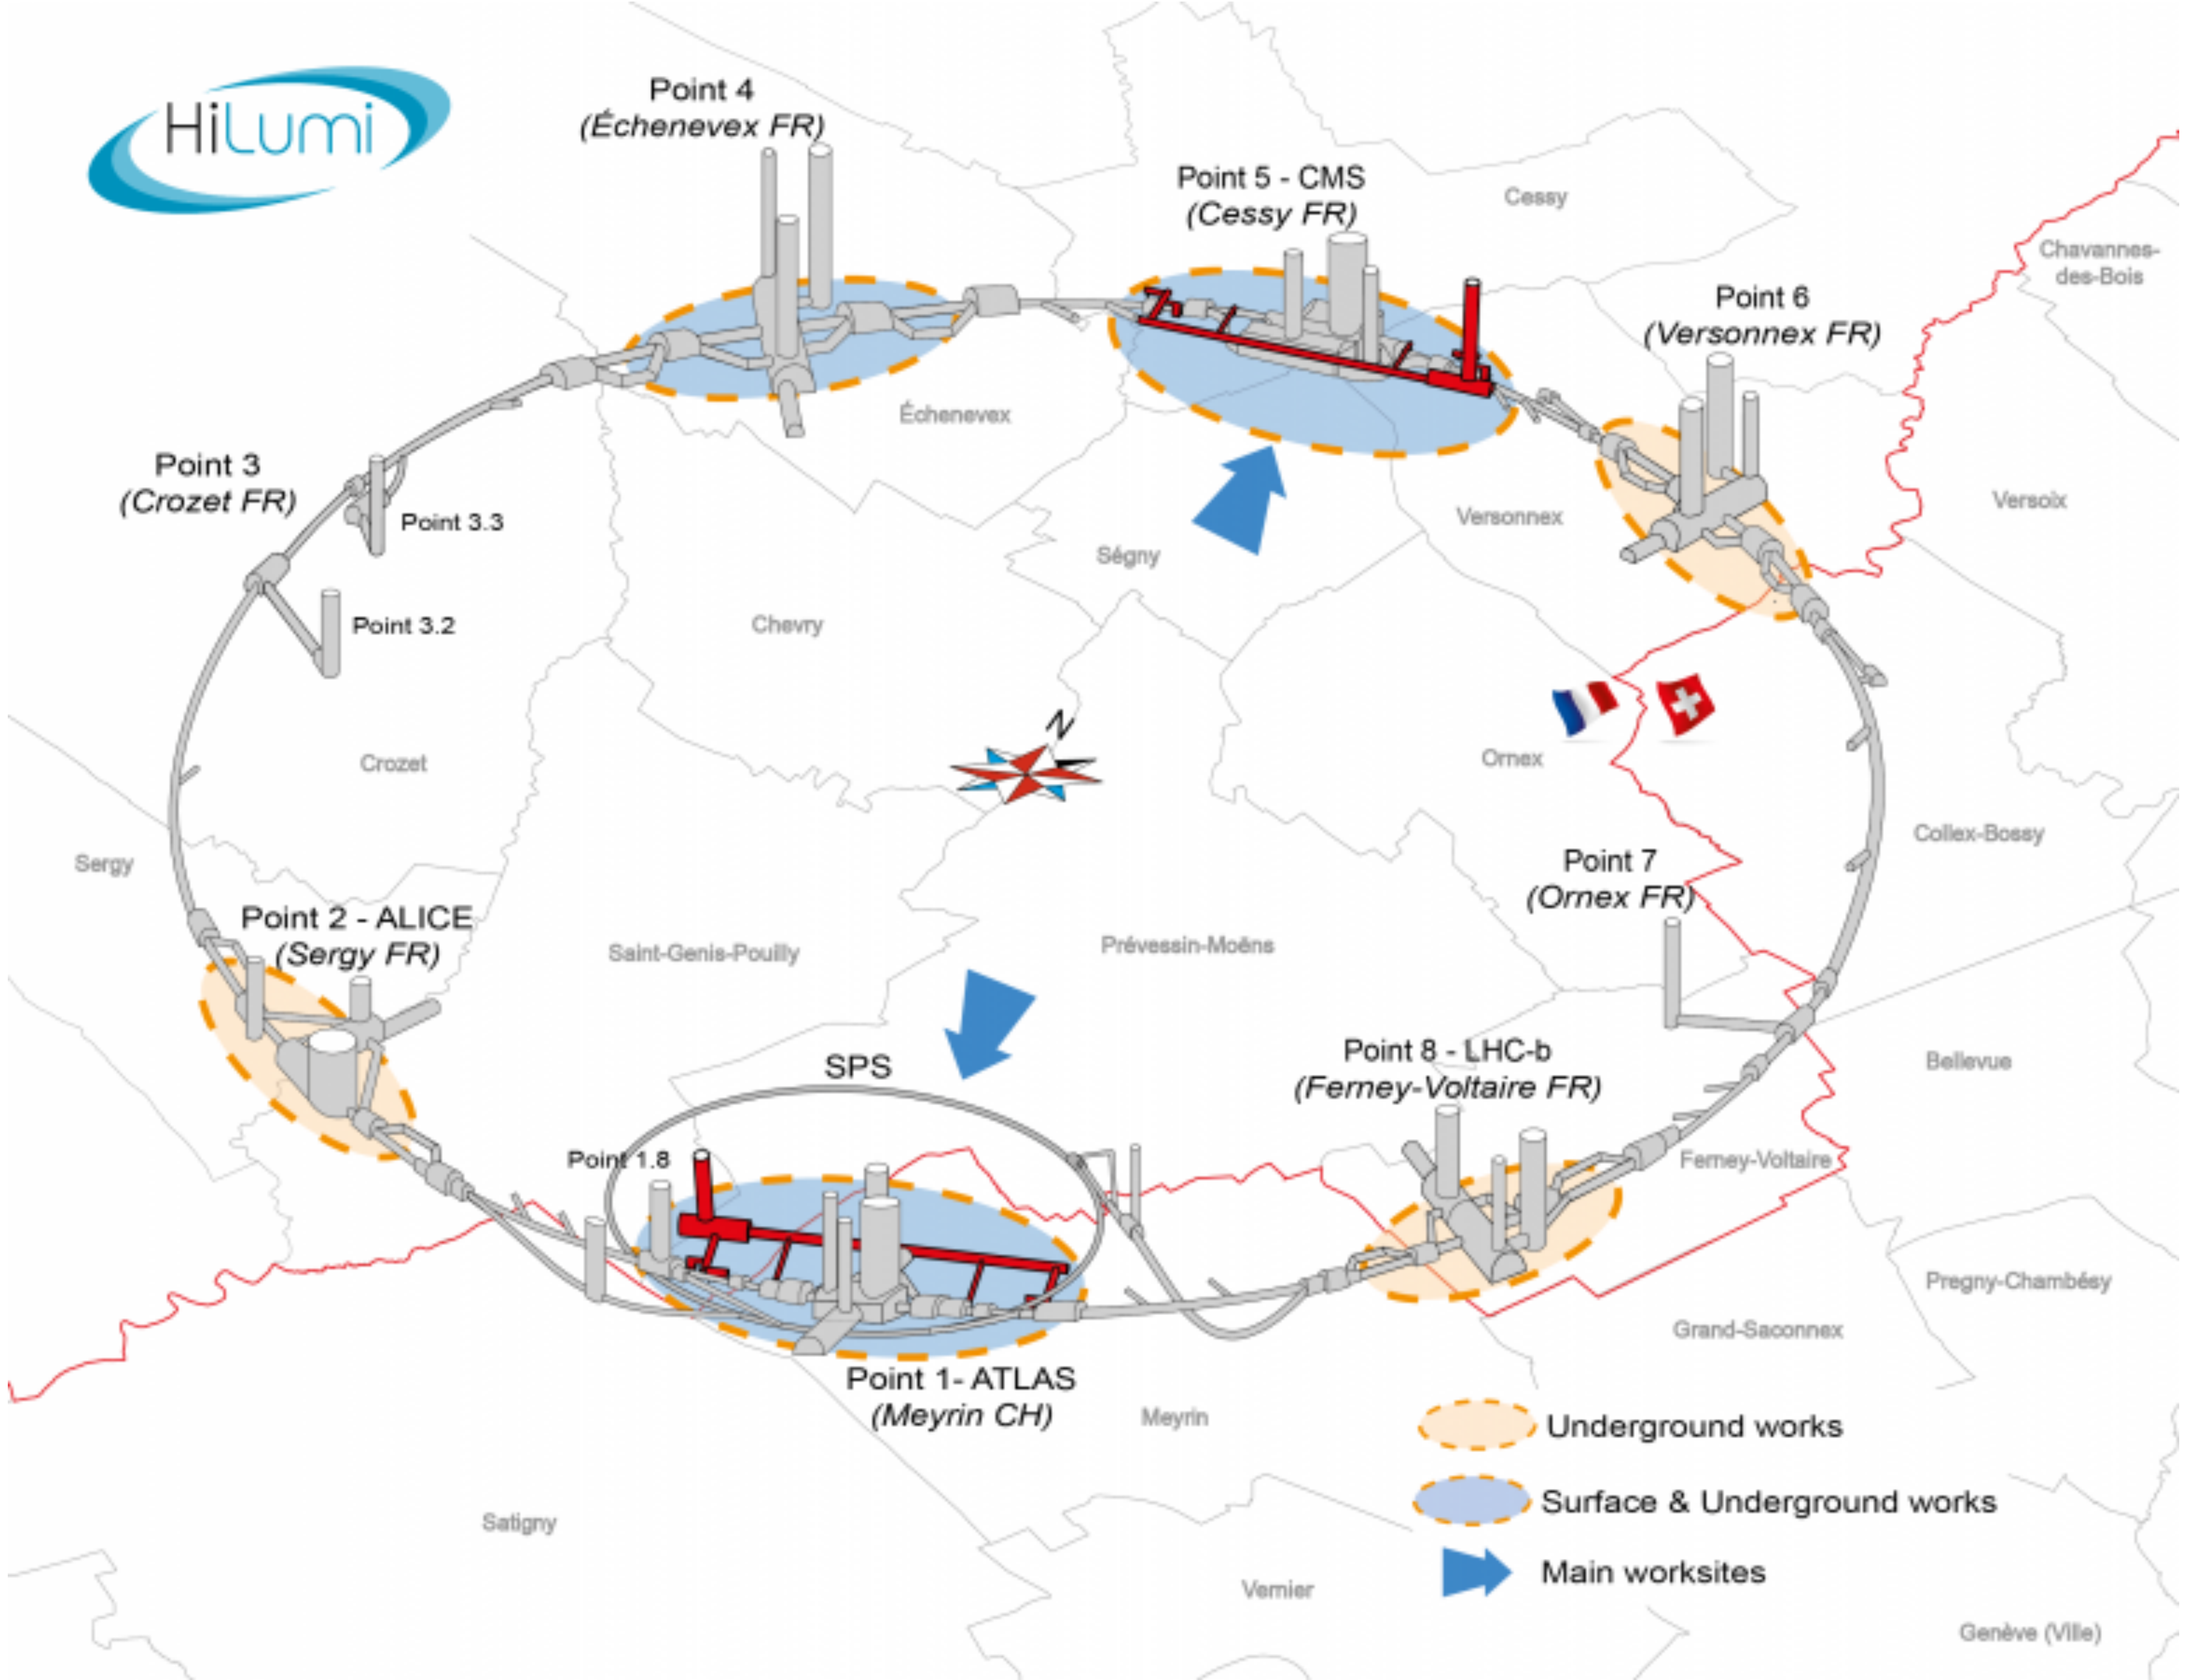

Source: LHC