Triggering Tb/s of data: CMS perspective

Santiago Folgueras on behalf of the CMS Collaboration

Funded by the European Union
European Research Council
Established by the European Commission
INTREPID
Universidad de Oviedo

Introduction: how a trigger system works?

- Data explosion and AI applications: our world needs **higher throughput** and **real-time computing** capabilities
- LHC provides ideal benchmark to explore real-time data processing technologies
- Only a handful of the collisions contain **interesting physics**
- **Trigger system** decides, in **real time** if a collision is saved or lost forever

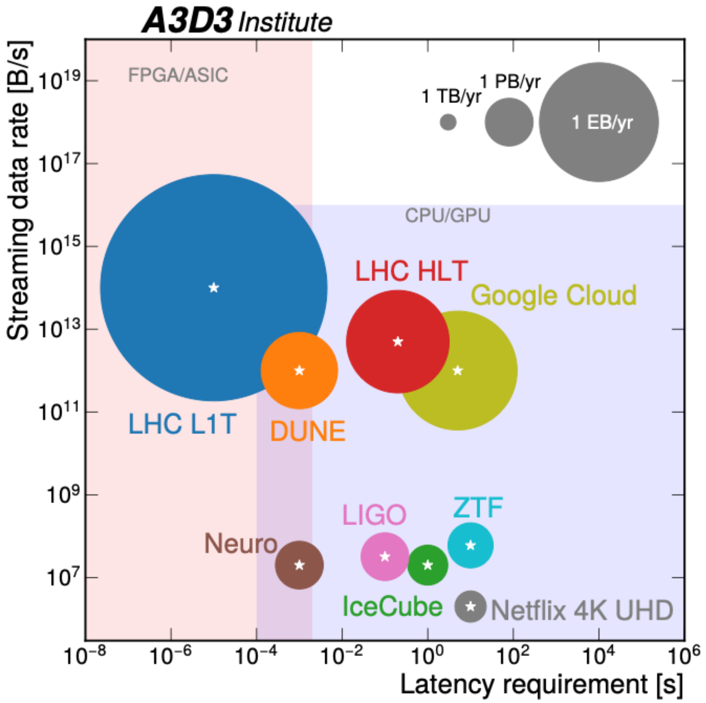

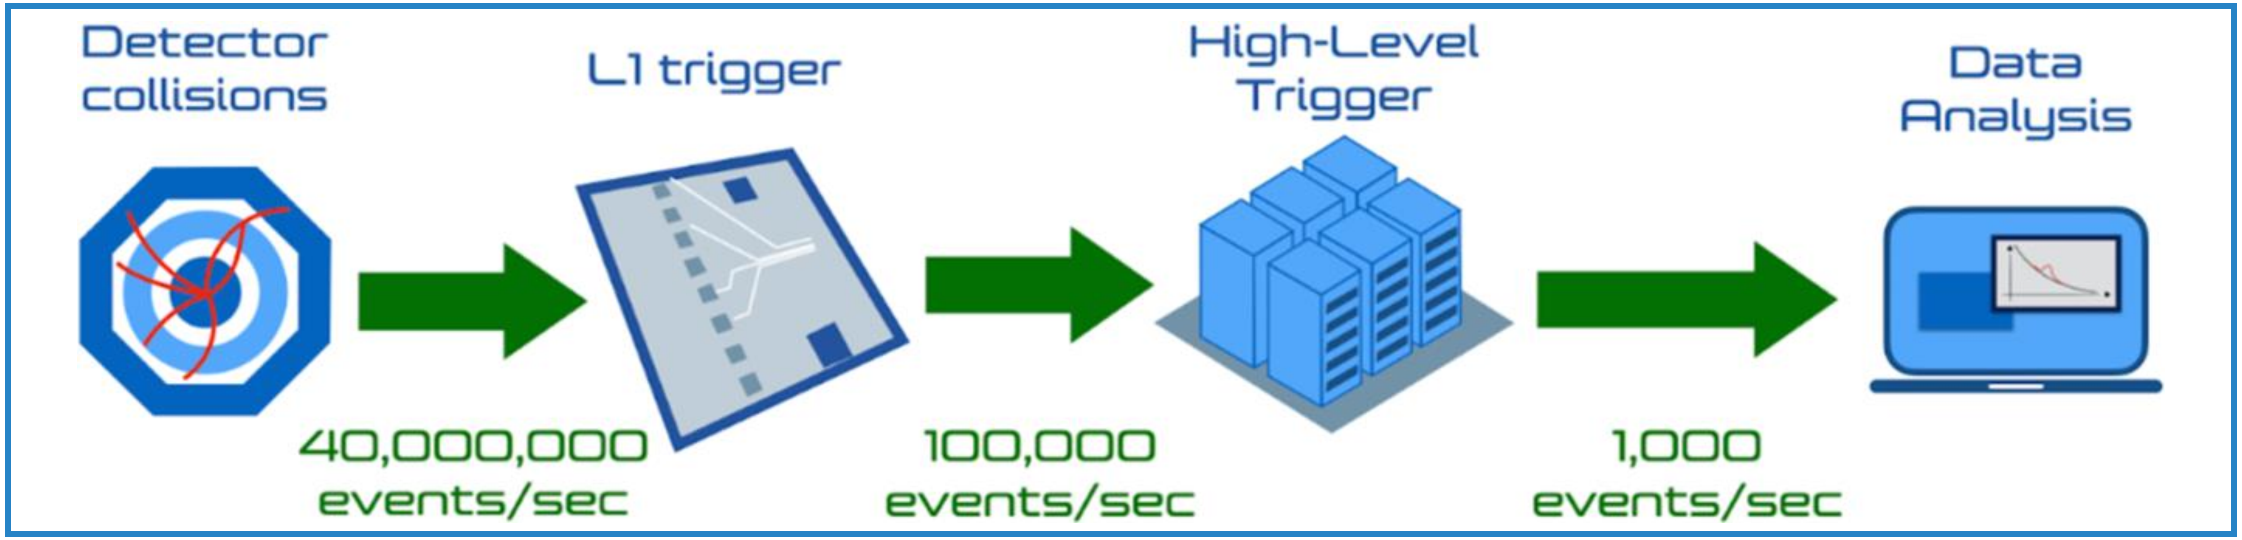

Towards the HL-LHC

- **Preparing for the big upgrade** of the LHC detectors, starting 2030.
- HL-LHC upgrade offers an **unprecedented opportunity** to explore uncharted lands and achieve scientific progress.
  - 10 times more data to what we will have by the end of Run 3 will facilitate a rich physics program.
- **Extend reach of new physics searches**: unexplored signatures (LLPs, HSCPs…) or regions of the phase-space will be within reach.
- **Improve current understanding of the SM and Higgs** sector by improving existing precision measurements and accessing rare decays ($H \to \mu\mu$) or production modes (HH) previously unseen at the LHC.
- However, this physics program will have to overcome **significant challenges** to succeed.

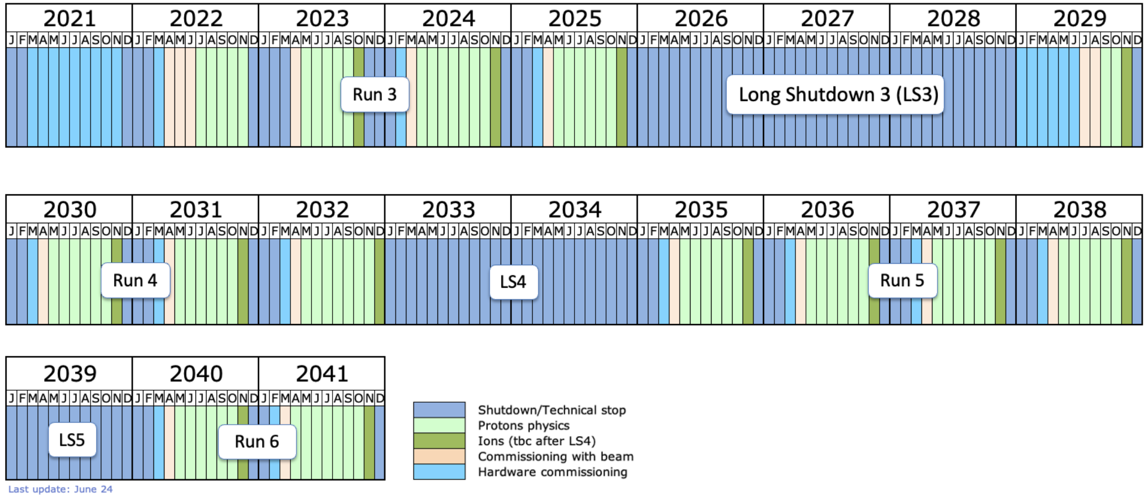

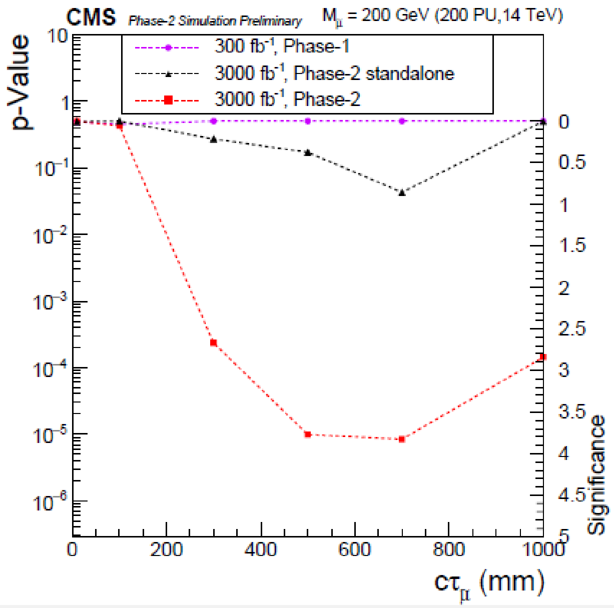

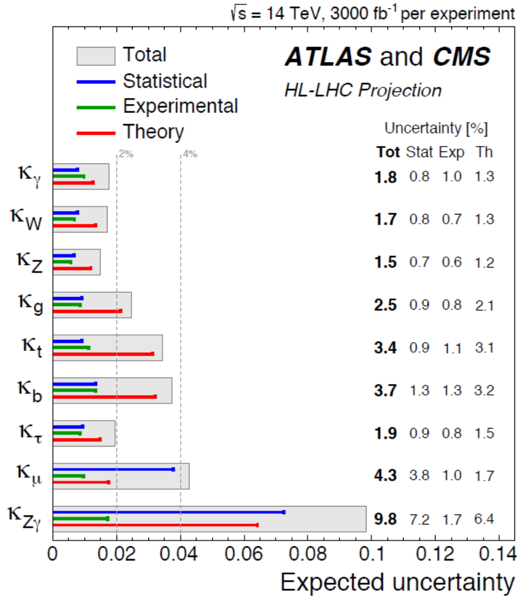

Universidad de Oviedo

The Phase-2 Trigger Upgrade: Strategy

- **Benefit from the upgrade of the CMS detector:** high granularity information and tracking information
- The system allows a **throughput** of > **+64 Tb/s** using top-of-the-line FPGAs and ultra-fast optical links (25 Gbps).
  - Adapt and evolve as needs of experiment change.
  - Increased bandwidth to **750 kHz** at increased latency of **< 12.5 µs**
- Incorporate **sophisticated algorithms and advanced techniques** to extend CMS physics acceptance

t = 0

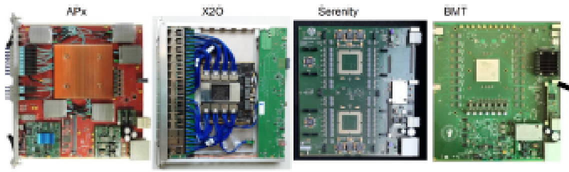

1 small box = 1 FPGA board with AMD VU13P FPGA

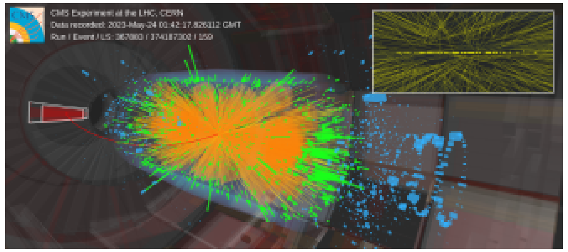

t < 12.5 µs

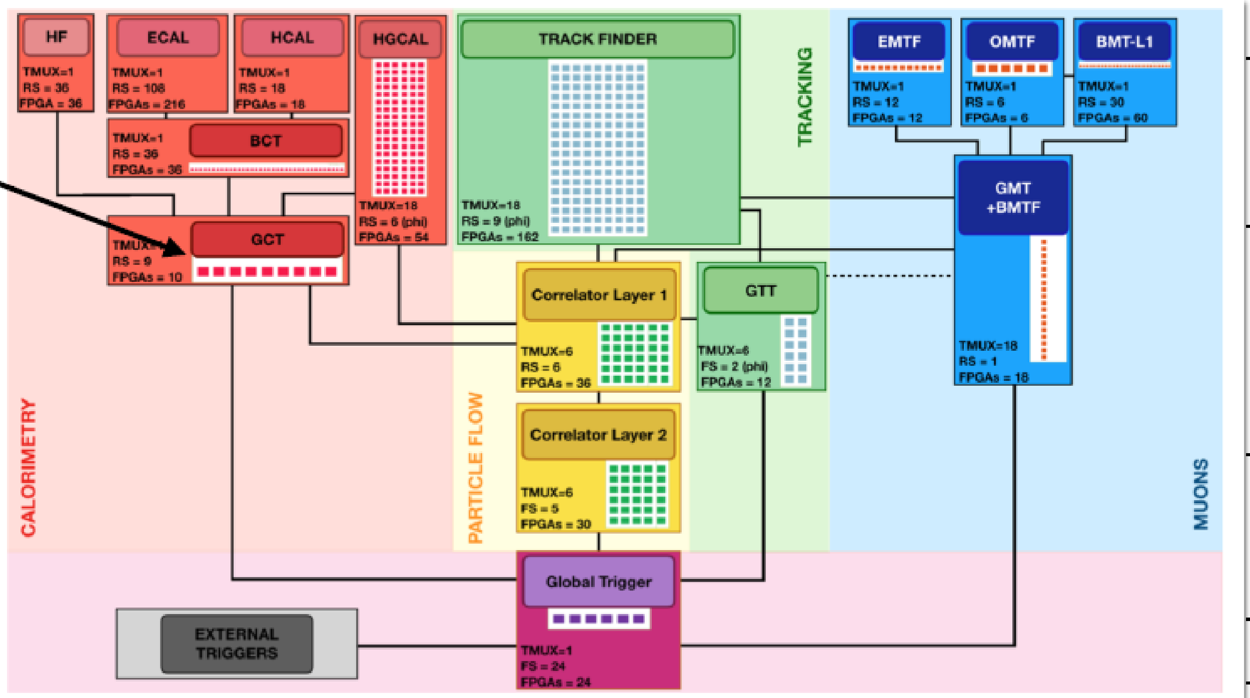

0 µs	Detector hits
5 µs	Clusters & Tracks
6 µs	Particles
7 µs	Event Categorisation
8 µs	1 bit: keep / discard

Universidad de Oviedo

Hardware prototypes

- **Design philosophy:**
  - Custom ATCA-boards. Generic Processing Engines → I/O, FPGA → sophisticated algo, arch flexibility
- **Design evolution:** increased I/O and computing power
  - FPGA : larger A2577 pin package, Xilinx Virtex Ultrascale VU13P
  - Optics : New denser version of on-board fly over Samtec Firefly & QSFP
  - Processors on board running commercial linux for flexible configuration and monitoring

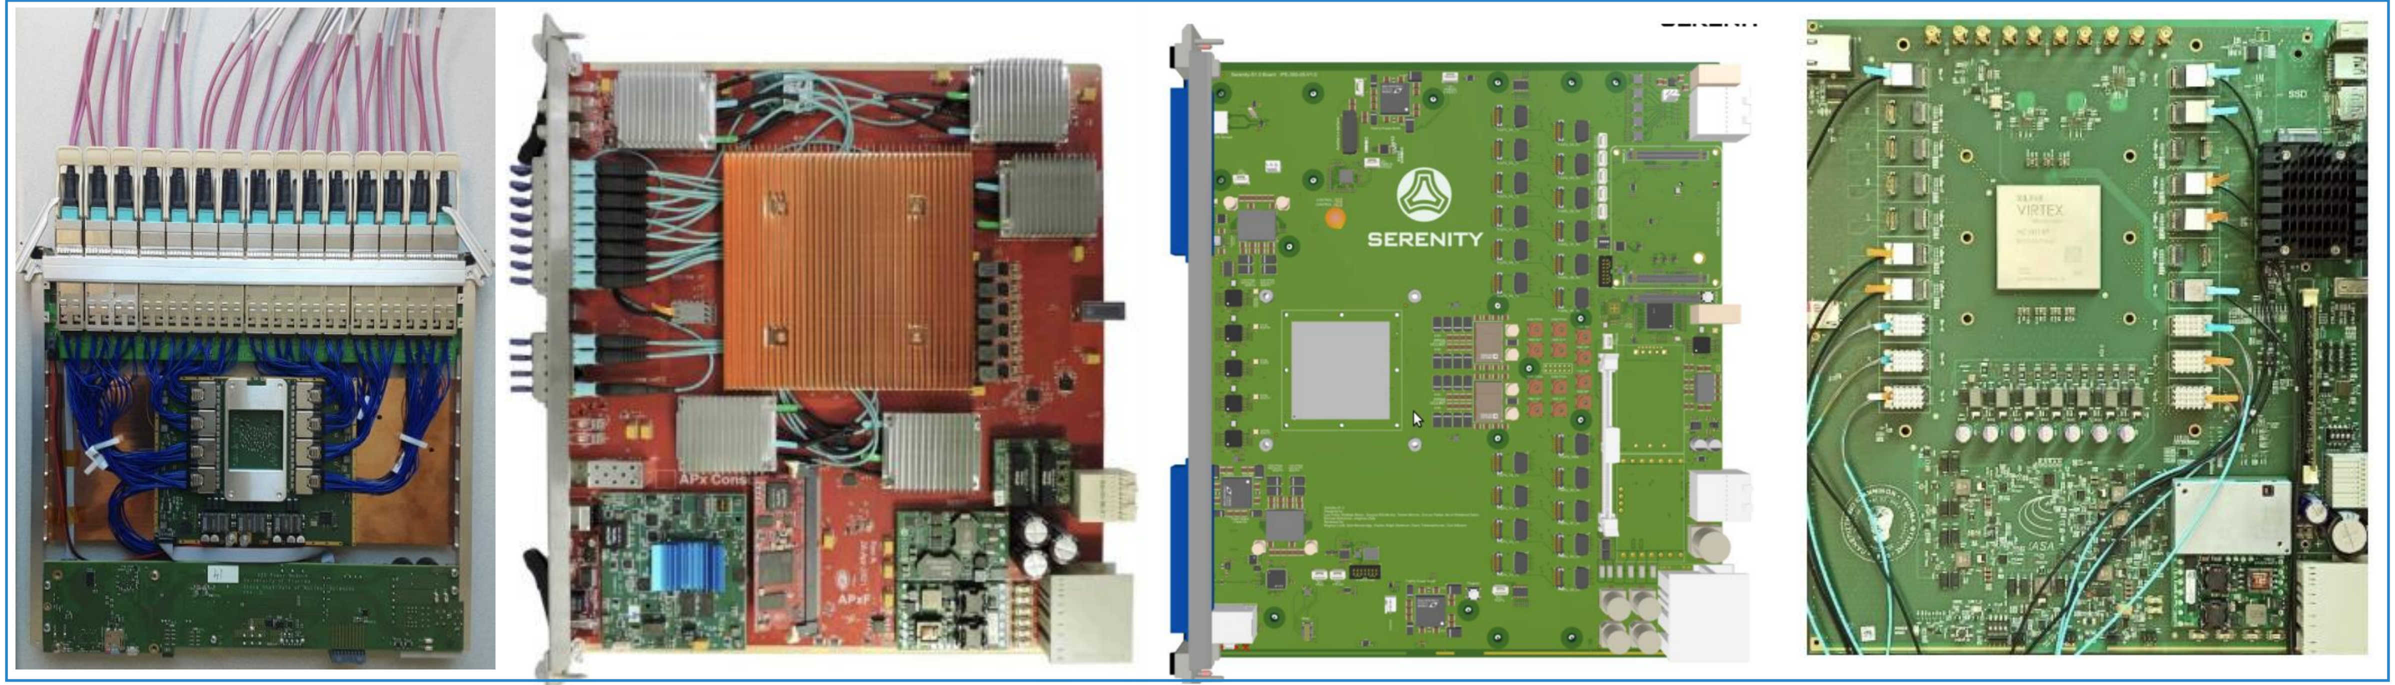

Hardware optics and thermal tests

- **Optical requirements:**
  - Support sufficient signal integrity in both the electrical and optical domains by demonstrating a bit error rate (BER) much better that 10^{-12}
  - Optics should provide sufficient optical margin with a receiver sensitivity better than -6 dBm to ensure operability at end of life (as laser degrades)
  - Tested Samtec Firefly x12 and QSFP (single and double density)
- **Thermal performance**
  - Integrity of the optics and FPGAs

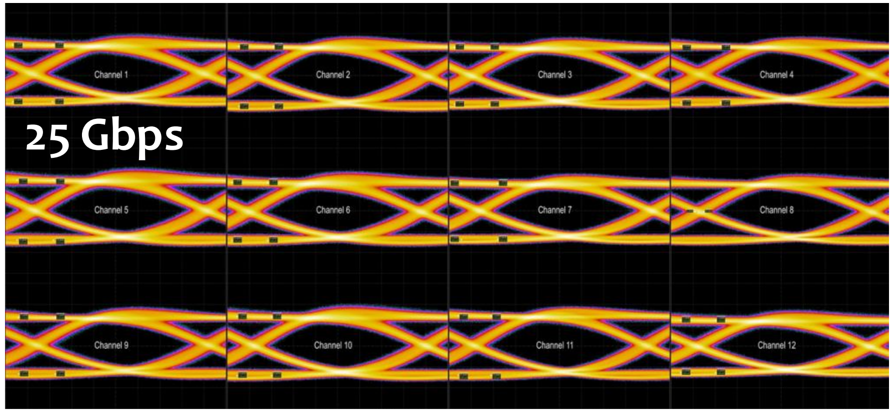

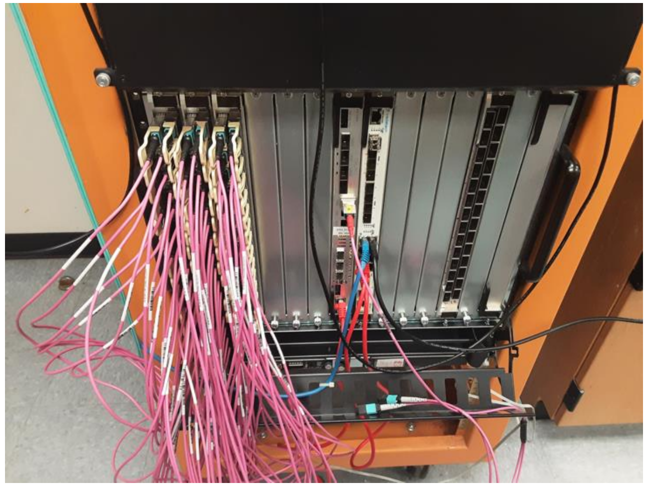

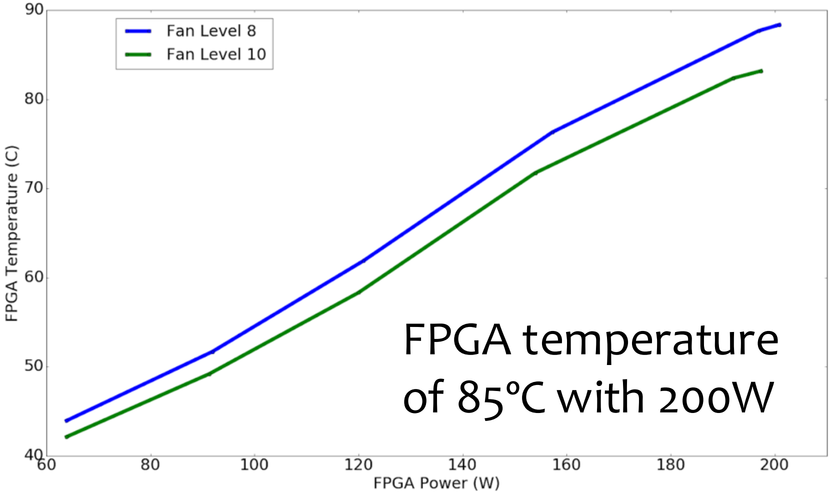

Algorithms for the Level-1 trigger

- Extensive use of tracking to reach near offline performance (sharper efficiency turn-on curves) + reconstruction of Primary Vertex.
- Exploit complementarity of different object flavor:
  - 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

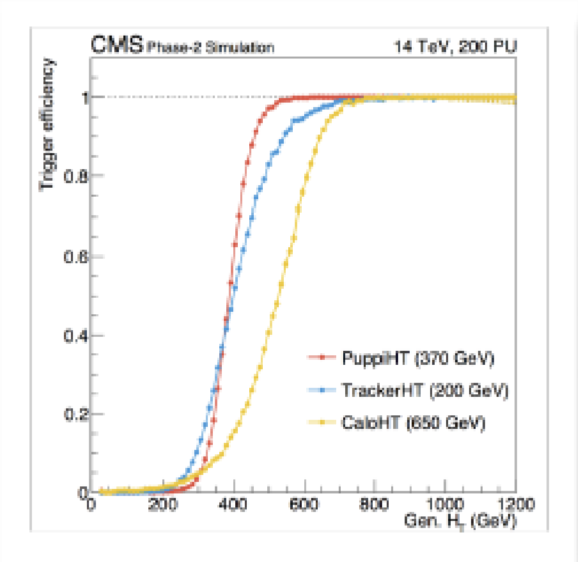

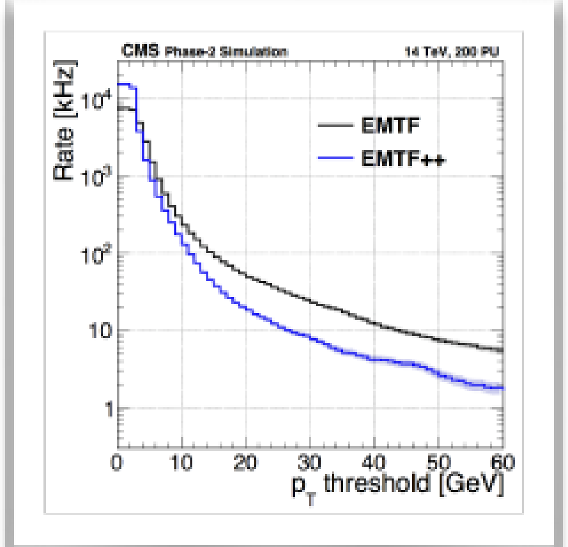

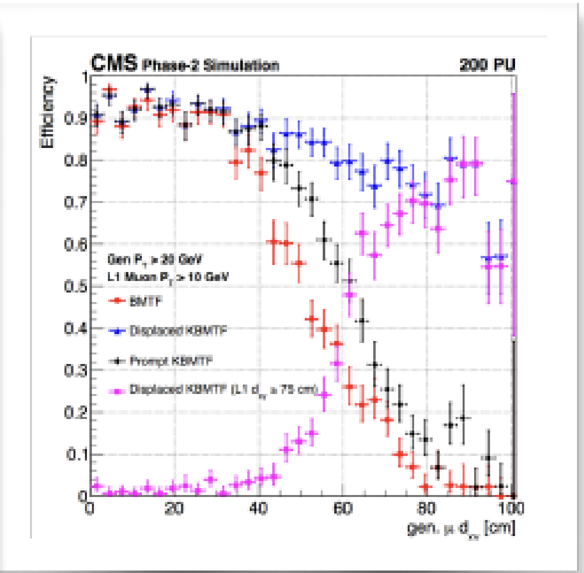

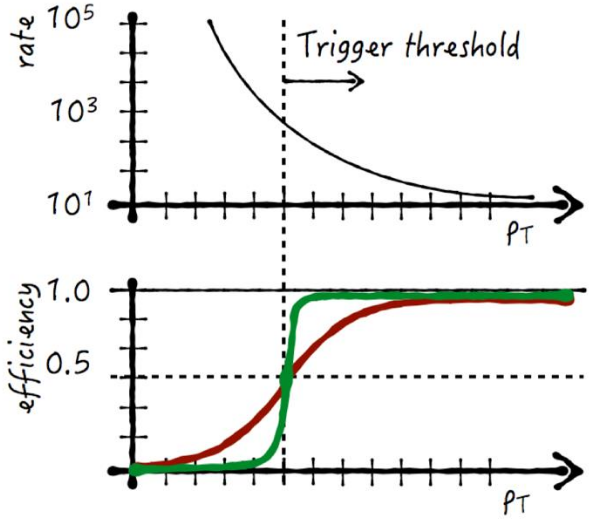

Universidad de Oviedo

Extensive use of ML algorithms

Each small box is one Xilinx Ultrascale+ FPGA

These have NNs and/or BDTs inside

hls4ml Conifer

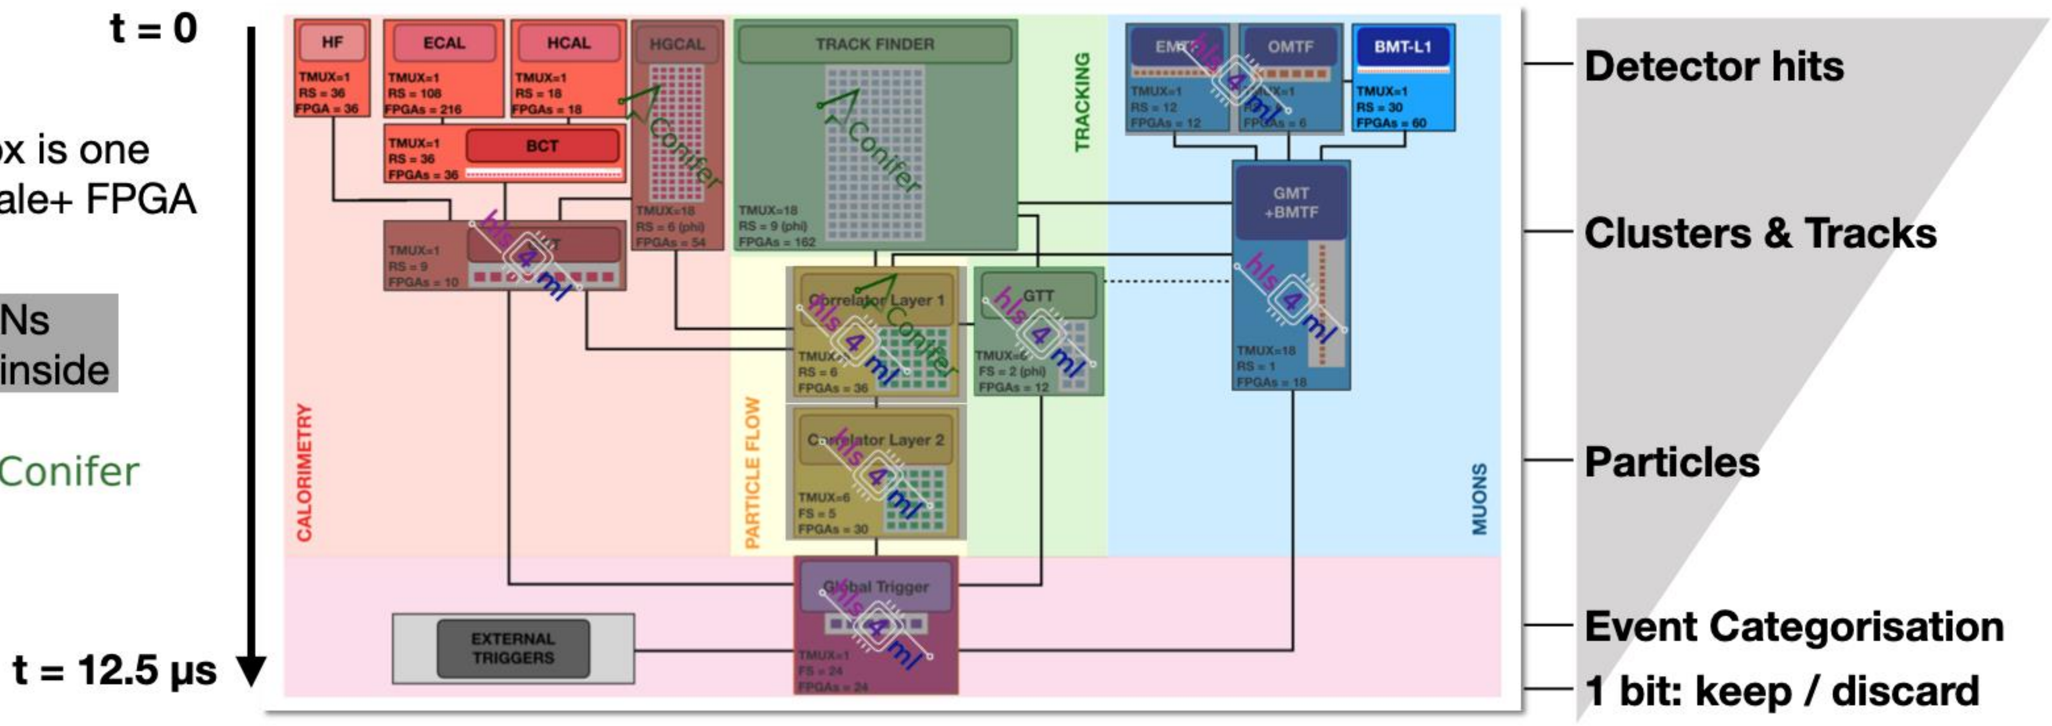

Algorithm into firmware: latency and resource utilization

- **Firmware design and integration:**
  - Algorithm developed mostly in C → High Level Synthesis (HLS). Using Vivado HLS, Vitis HLS
  - Many tools available for Machine Learning inference: hls4ml, Conifer for BDT evaluation
  - New fixed-point arithmetic in C++ [taken from Xilinx libraries] → emulator firmware
  - Continuous integration of the firmware in repository
  - Verify timing, resources utilization & latency: all using less than 50% resources, whole system evaluated to 8.6 μs

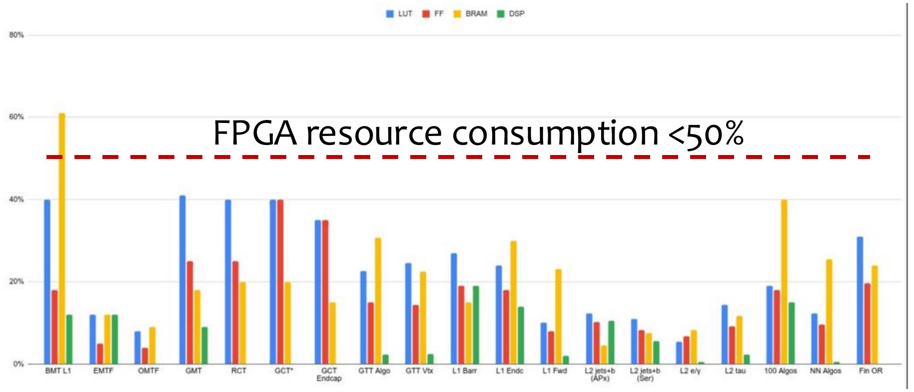

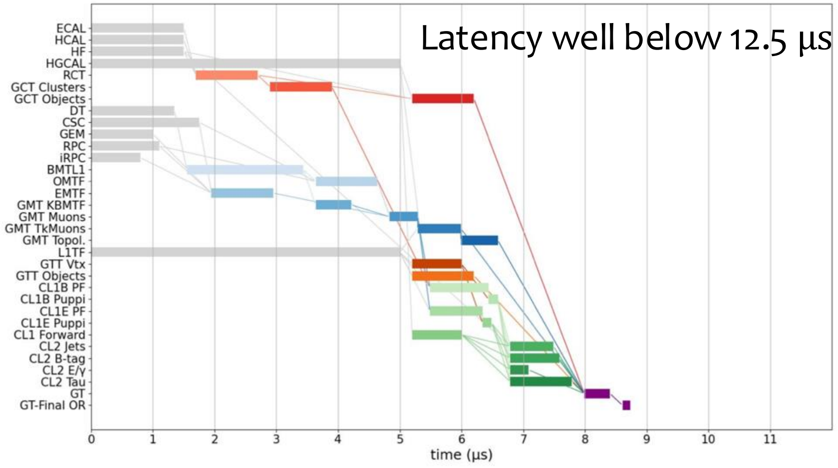

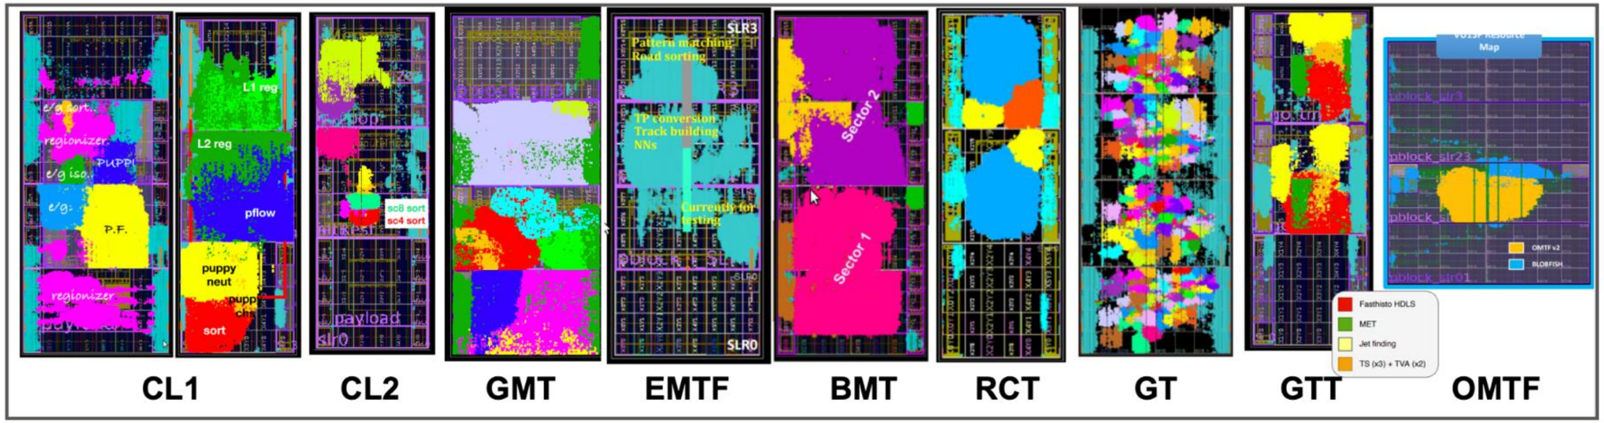

Testing new ideas during Run 3

- With almost one and half year to go, Run-3 has already surpassed Run-2 luminosity
  - **Almost 170 pb^{-1} recorded**
- Successful feedback loop into the current system: the Run-3 system now features new algorithms, optimisation techniques, hardware, inspired from the phase-2 upgrade project:
  - LLPs triggers: displaced muons, muon showers, delayed jets…
  - 40 MHz scouting (real-time data analysis)
  - Inclusion of the first anomaly detection trigger on live data: AXOL1TL and CICADA
- System exceeding original design. **Having a flexible design is an advantage!**

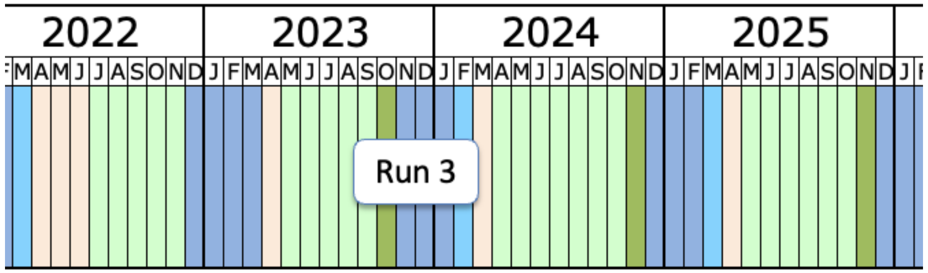

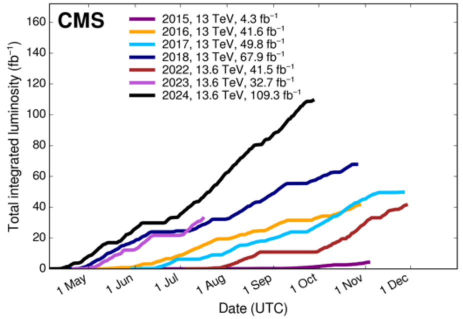

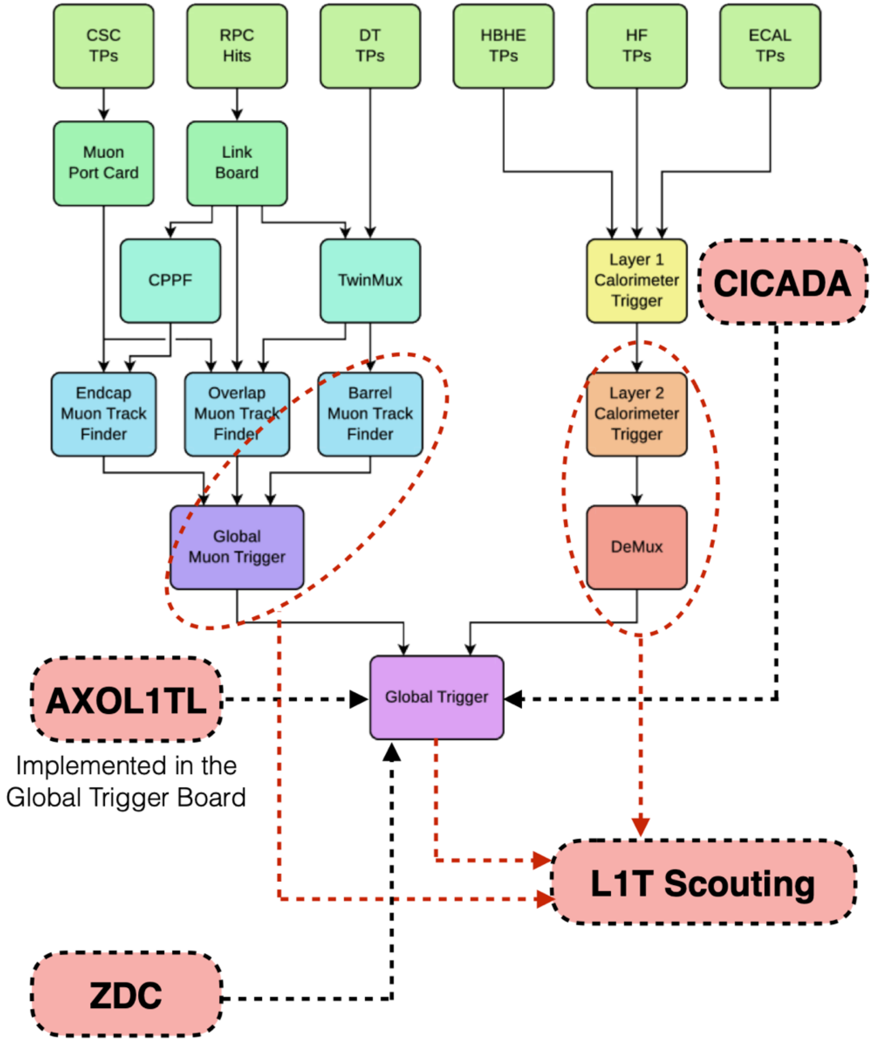

Displaced/delayed jets

- **ECAL** measures arrival time of objects with precision of ~200 ps (for energy deposits >50 GeV). Tau seeding at L1 and trackless jets at HLT
- Use **HCAL** time information at the L1 trigger level to identify delayed jets (>6ns). Prompt veto applied
- High multiplicity at the **muon system** for long-lifetimes

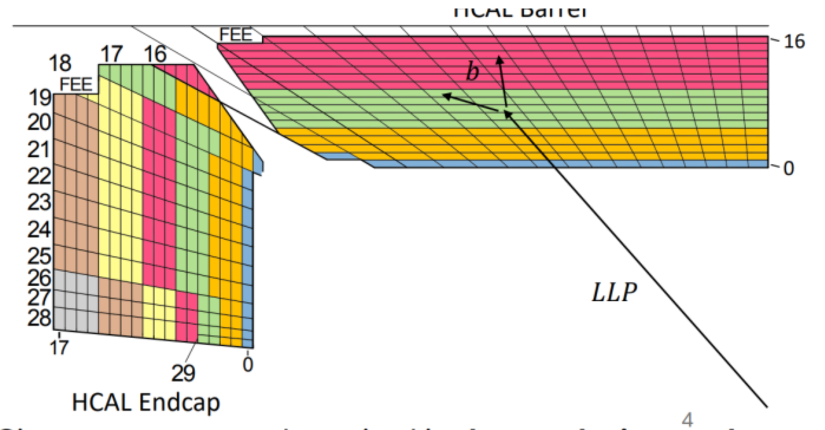

Signatures: energy deposited in **deep calorimeter layers**

Triggered on October 8th 2022

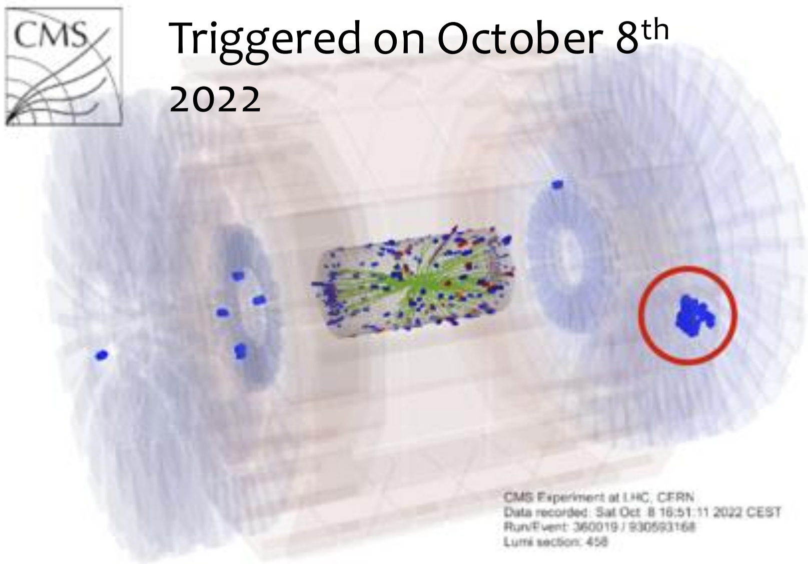

Muon (*more later*)

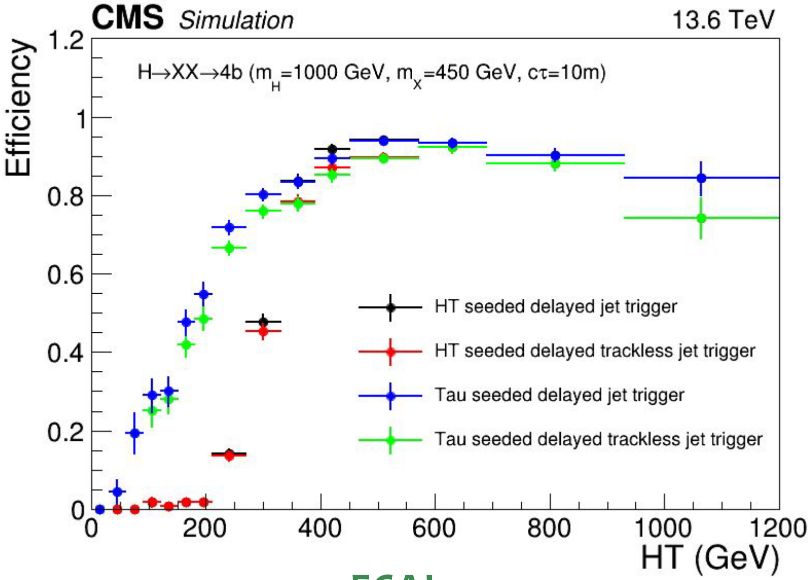

ECAL

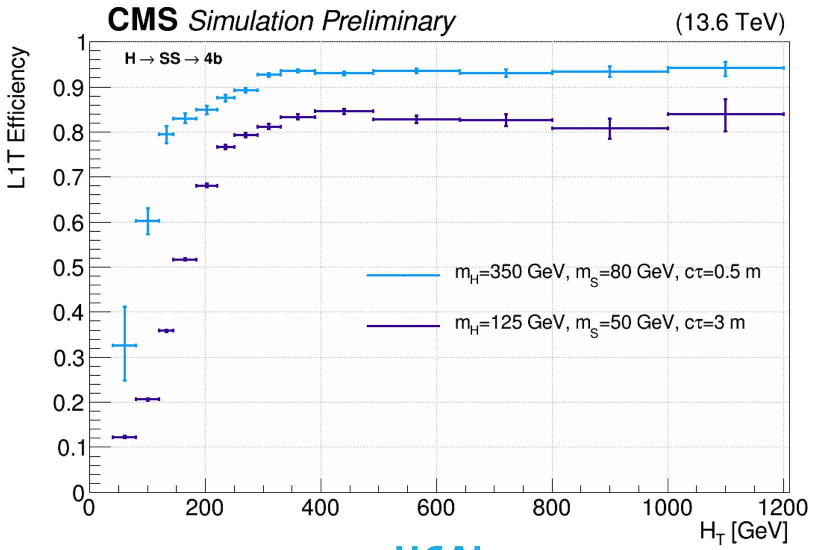

HCAL

ML at L1: Anomaly detection

- Where's the new physics? To find anything, you need a trigger
  - If we knew what we were looking for, we'd build a trigger for it!
- Cast a wide, model-independent net
  - Learn what an average event looks like, pick things that are rare
  - Autoencoder, trained on random beam events
  - Reconstruction error is a metric for anomalous-ness
- AXOL1TL & CICADA
  - Low-level variables (L1T or Calorimeter objects)
  - Outputs an anomaly metric to keep the event or not

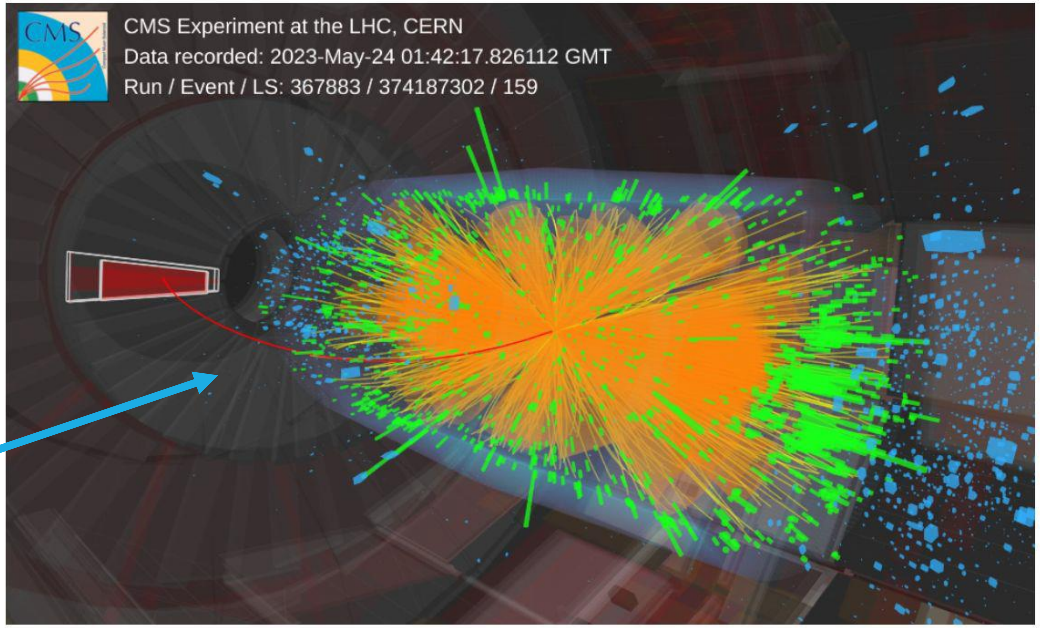

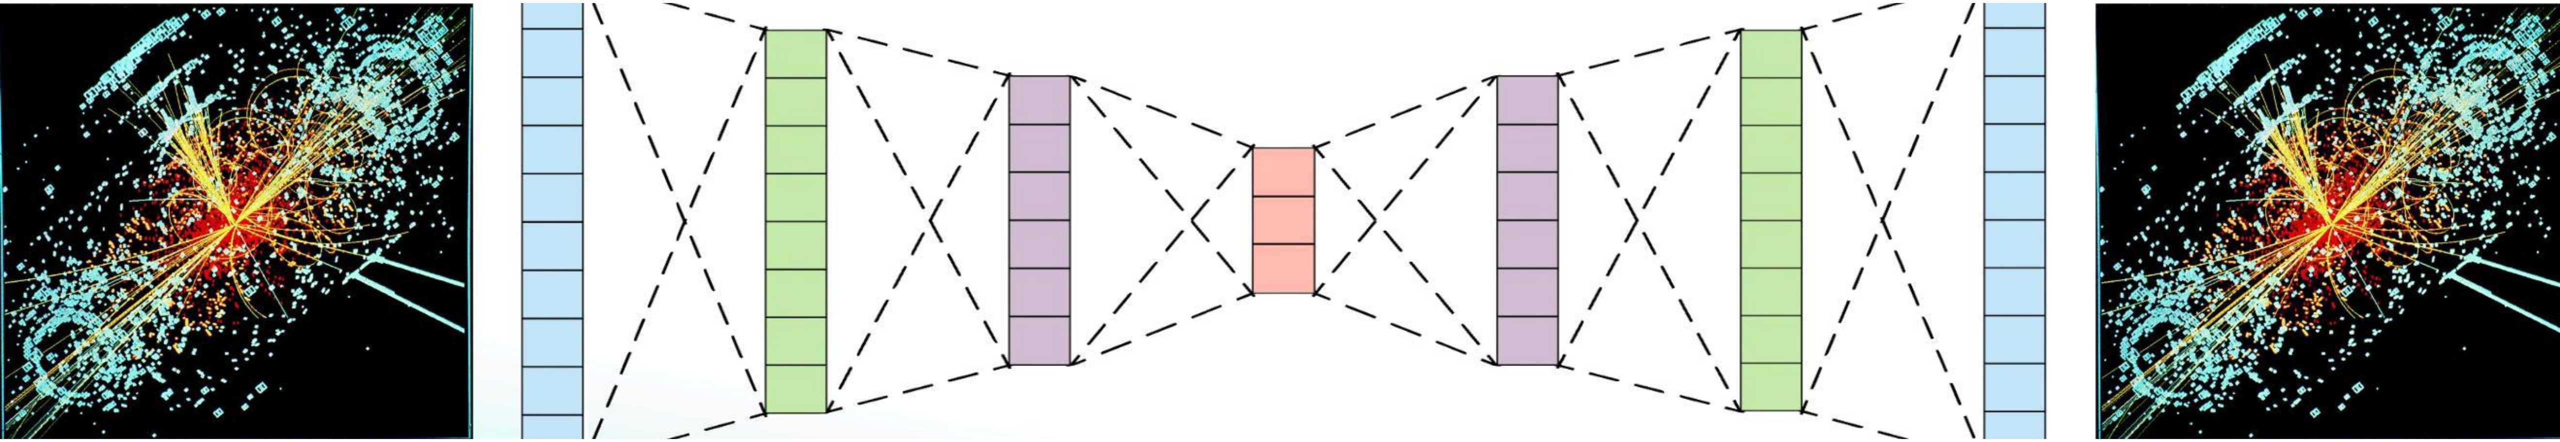

https://arxiv.org/abs/2403.16134
https://cds.cern.ch/record/2904692/

Triggerless analysis (aka scouting)

- Storing and analysing events at L1 or HLT (x100 smaller event size)
- Crucial for very low-mass bump-hunt searches, compressed spectra or b-physics

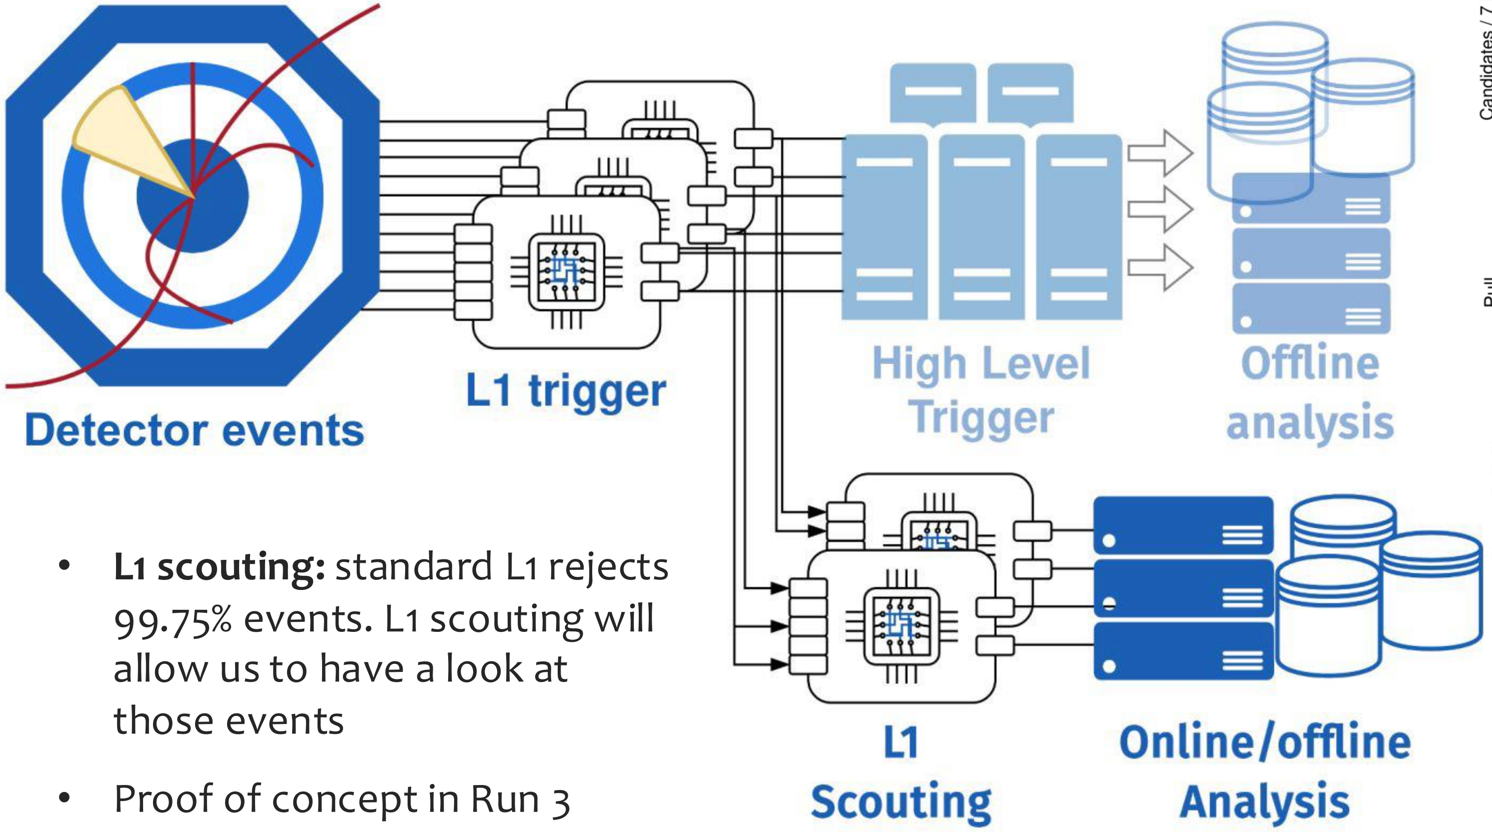

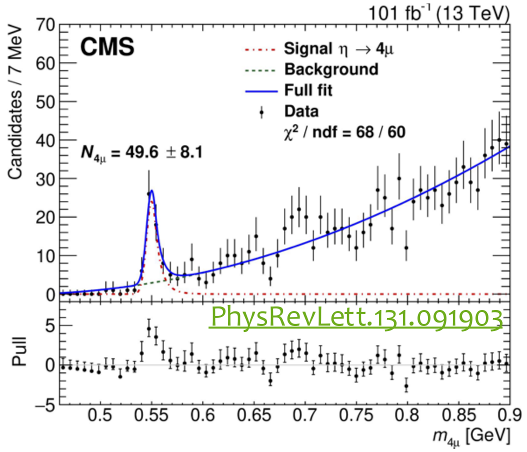

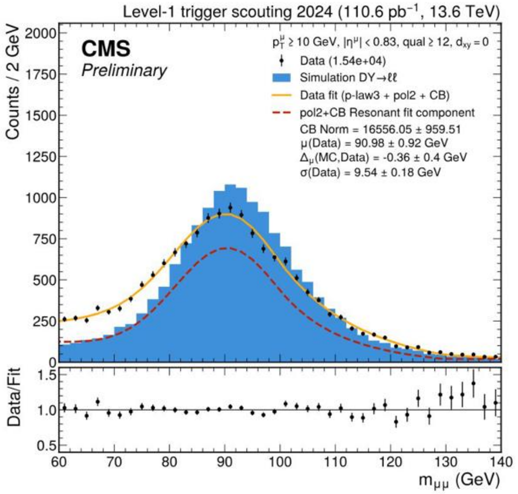

- **L1 scouting:** standard L1 rejects 99.75% events. L1 scouting will allow us to have a look at those events
- Proof of concept in Run 3

Level-1 Data Scouting rack

Impact of trigger design beyond HEP

- **Impact on society**
  - Massive surge of data and AI applications. The need of processing large amounts of data is an ever-increasing challenge.
  - HEP experiments provide the perfect test bed for advanced AI algorithms developments, real-time data processing and low-power solutions
- **Developing ideas for CMS trigger and beyond:** NextGen and INTREPID projects
  - Enhance the triggers and the data collection and processing, and thus the scientific potential, of ATLAS and CMS in the HL-LHC phase **beyond the currently projected scope**.

INTREPID

NexTGen Next Generation Triggers

- **Driving a lot of attention**
  - from national and international funding agencies and industrial partnerships (CMS is working with Amazon, Google, Micron...)
  - Emerging applications outside HEP: data reduction onboard satellites, quantum control systems, brain implants…
  - Custom silicon for Machine Learning is big industry trend - acceleration of specific workloads

Conclusions

- The CMS trigger system for HL-LHC will process data at ~64 Tb/s using top-of-the-line FPGAs and high-speed links
- Level-1 Hardware trigger with enhanced capabilities complying with physics requirements using sophisticated ML-based algorithms
- Modular and flexible design to adapt for future ideas using custom ATCA boards
- Hardware demonstration ongoing and some tests in Run-3 data taking
- Future designs are showing exciting prospects, even beyond HEP

Triggered on October 8th 2022

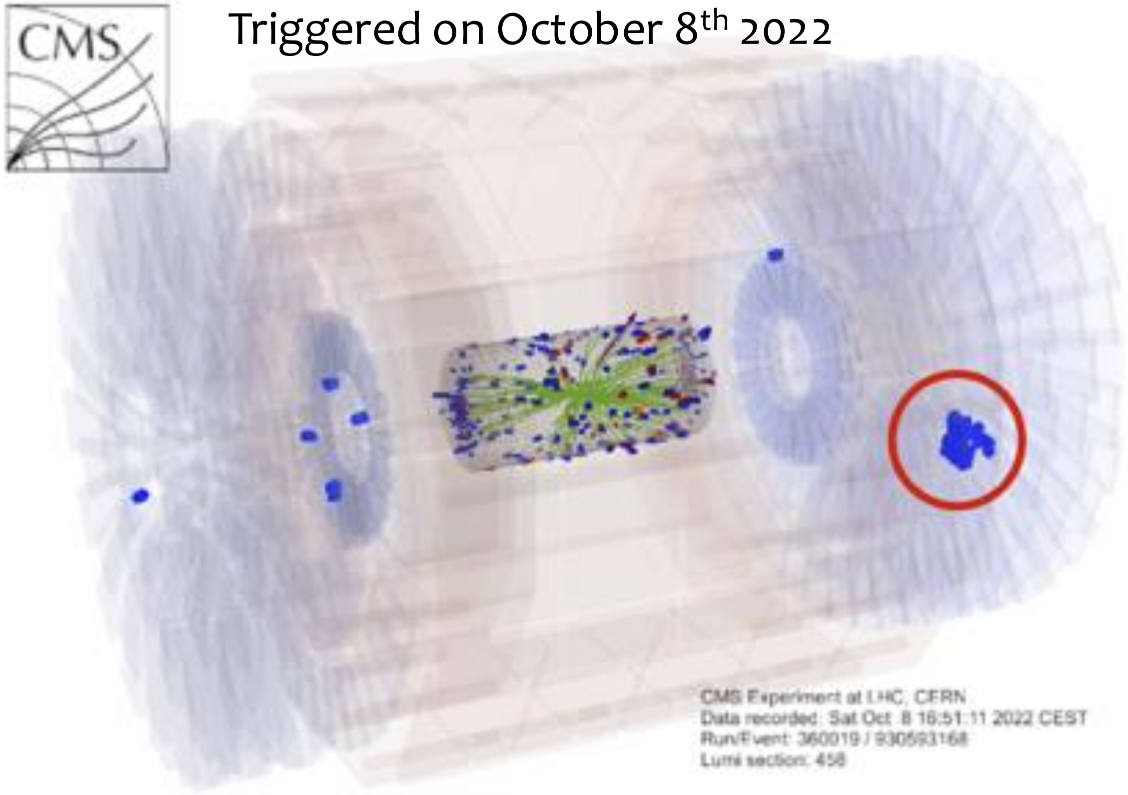

Funded by the European Union
European Research Council
Established by the European Commission
INTREPID

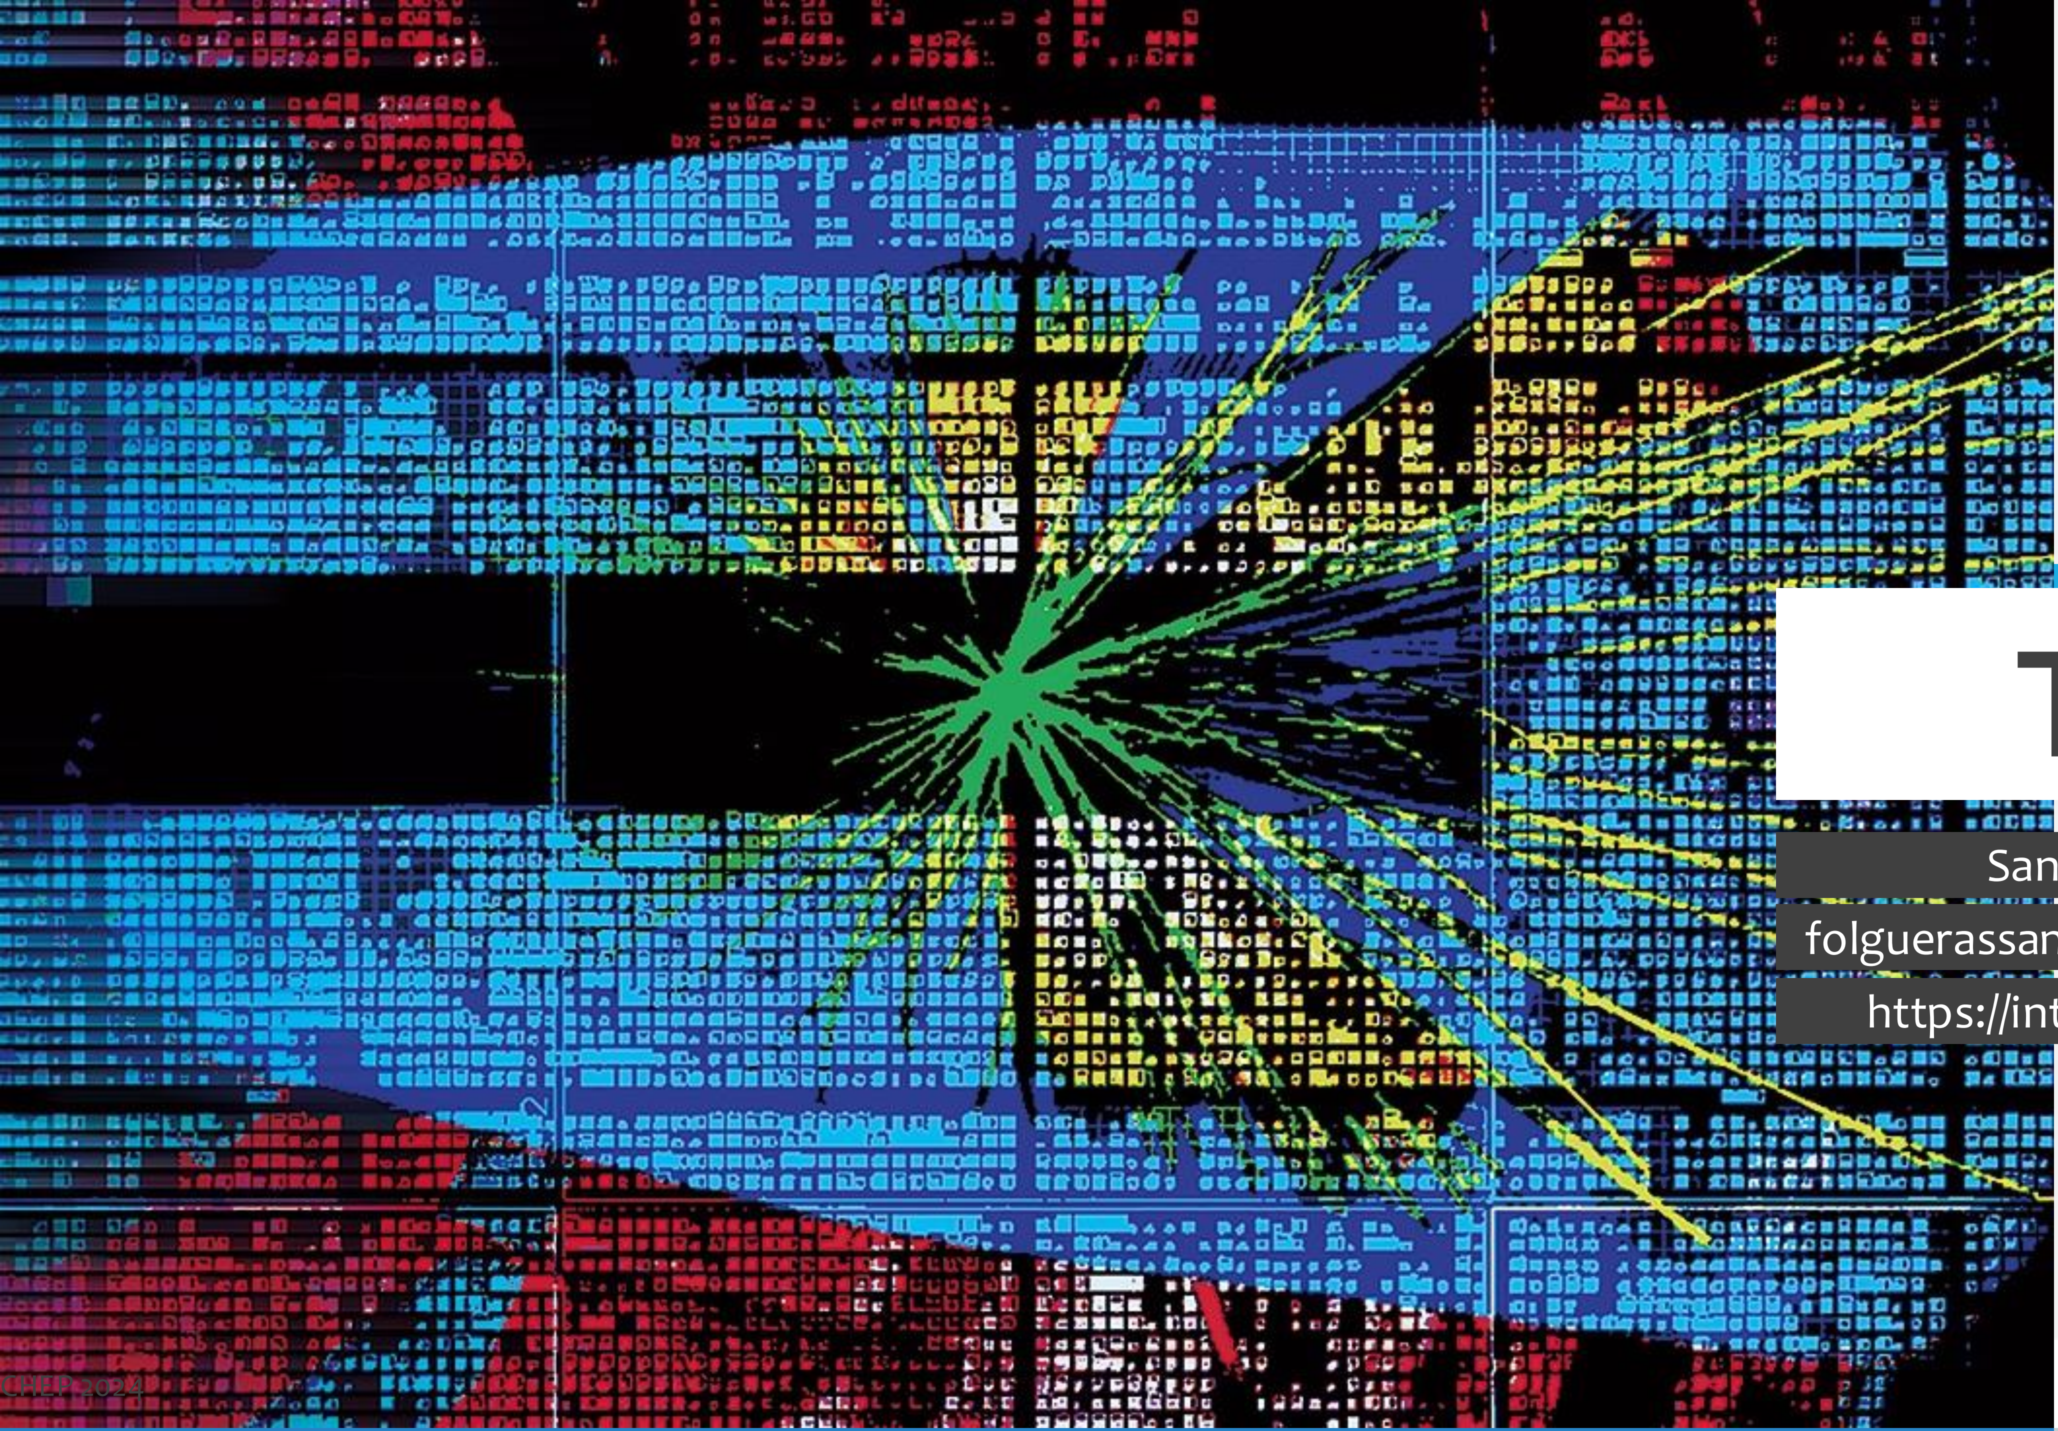

Thanks!

Santiago Folgueras
folguerassantigo@uniovi.es
https://intrepid.uniovi.es/

Funded by the European Union
European Research Council
Established by the European Commission
INTREPID

Introduction: the CMS Trigger System

Data is selected for offline analysis 2-tiered trigger system

Level 1 Trigger (L1T)

- Hardware system run on FPGAs
- Designed to reduced rate from 40 MHz to 110 kHz
- Fixed latency of 4 µs

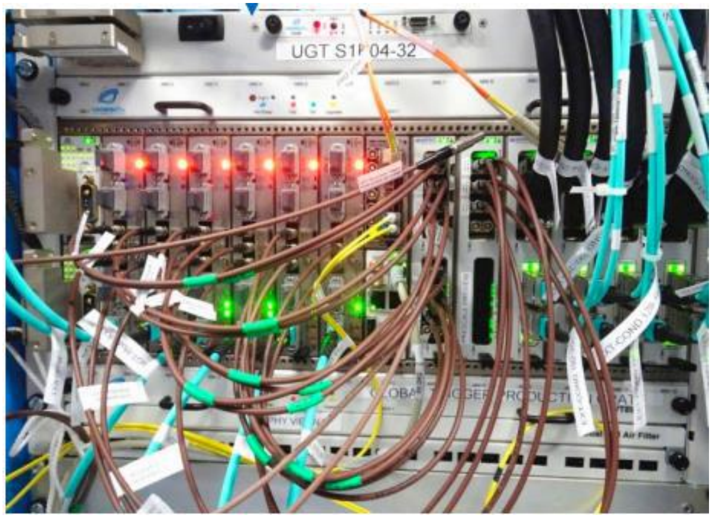

High Level Trigger (HLT)

- Software system run on CPU/GPU farm
- Designed to further reduce rate to 1-5 kHz
- Latency: 200-300 ms

Global event reconstruction (Particle-Flow) at Level-1

- Availability of tracks & high-granularity calorimetry
  - Implement global event reco @L1 and pileup mitigation
- **Challenge: can we run full PF+PUPPI at L1? YES!**
- Demonstrated a working PF+PUPPI algorithm:
  - Hugely reduces the event complexity and allows for a lot of flexibility in downstream design
  - L1 Algorithms looks like offline reconstruction
  - PF+PUPPI developed with Vivado HLS (*written by physicists + engineers*)

- What are the advantages of particle flow?

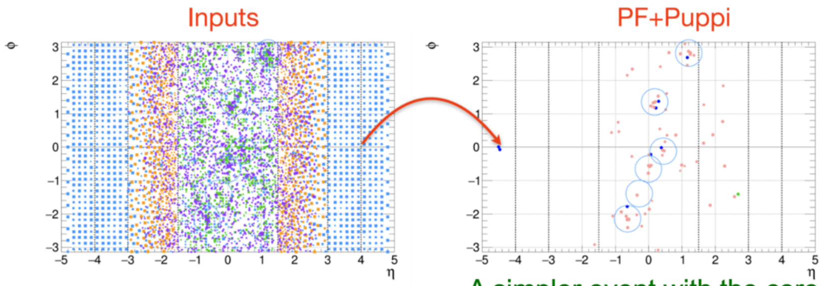

A simpler event with the core physics preserved

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

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

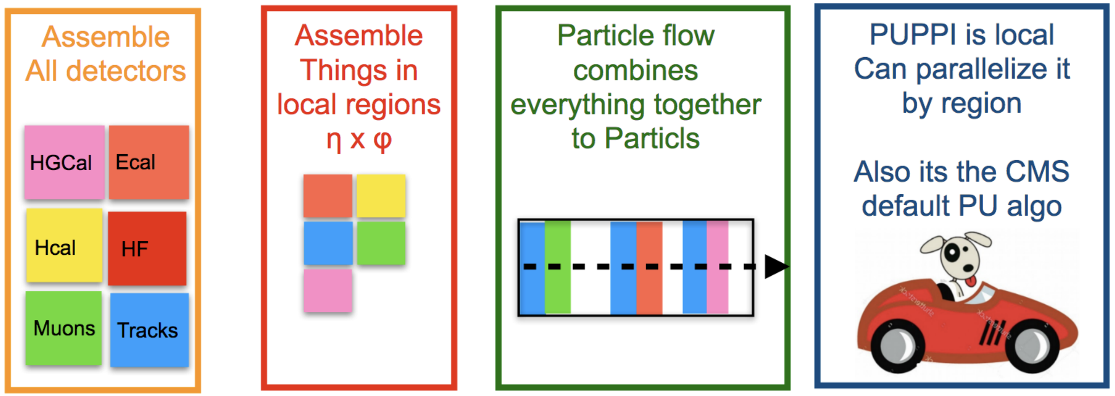

HW demostrator

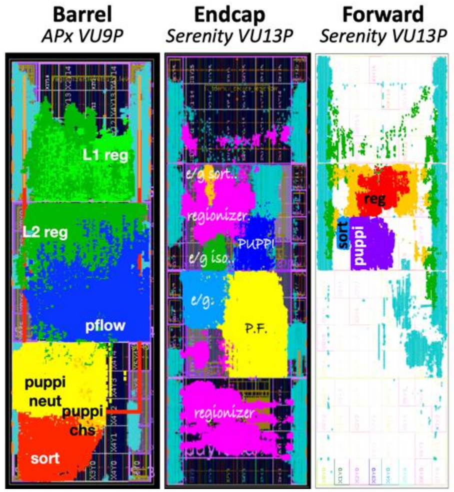

Recent development highlights (with ML)

- **NN Vertex Finding:**
  - Combination of dense BDTs and CNN to perform Vertex Finding and Track-to-Vertex association
  - Firmware quantised and pruned to fit within FPGA
  - ***Improved performance wrt to baseline (reduction in the tails of the residual by 50%)***

- **b-tagging:**
  - Training NN to ID jets from b-quarks
  - Runs on PUPPI particles within each jet and discriminate between b-quark jets and those from light quarks and gluons

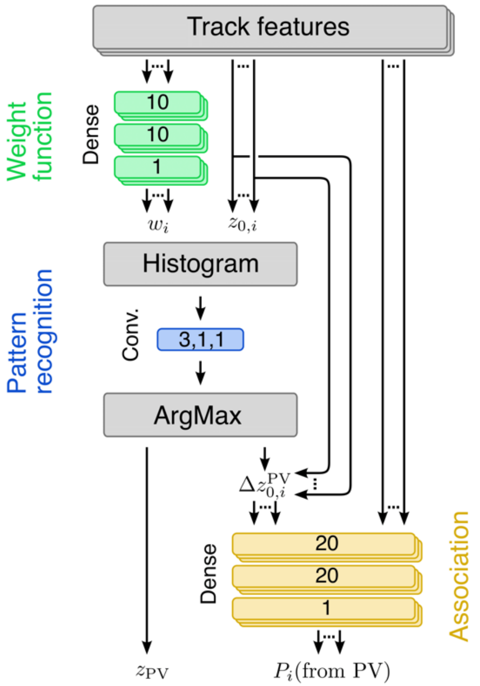

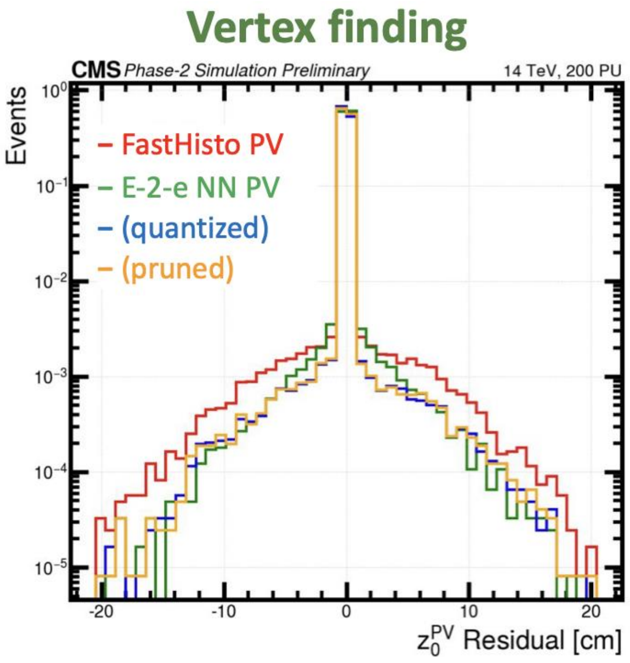

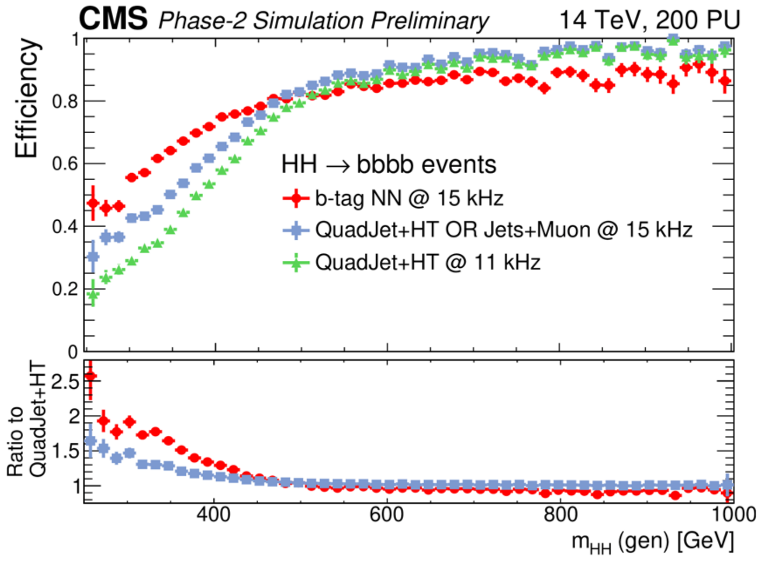

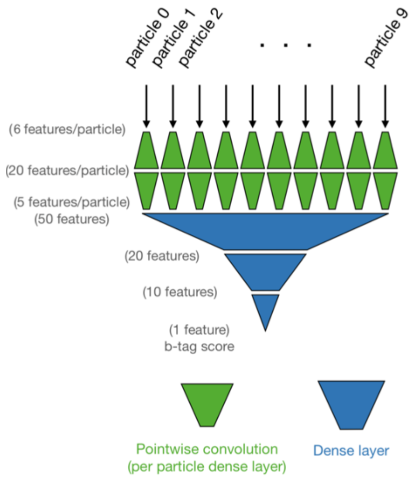

Recent development highlights (with ML)

- **Electron-ID**
  - New Composite-ID, combines information about tracks and clusters in the HGCAL into a single model for matching and identification
  - A single BDT model: controlling the identification of track and calorimeter deposit and the tightness of the matching. ▸ 10% more efficiency for the same rate

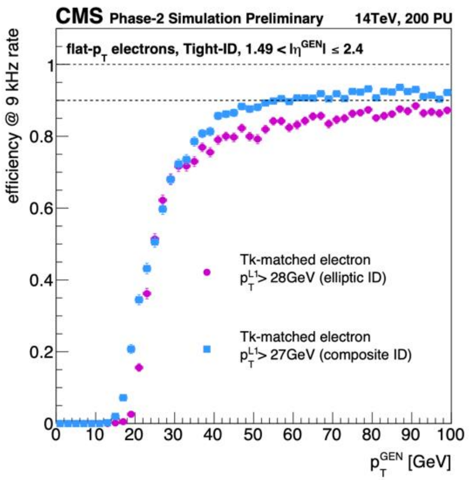

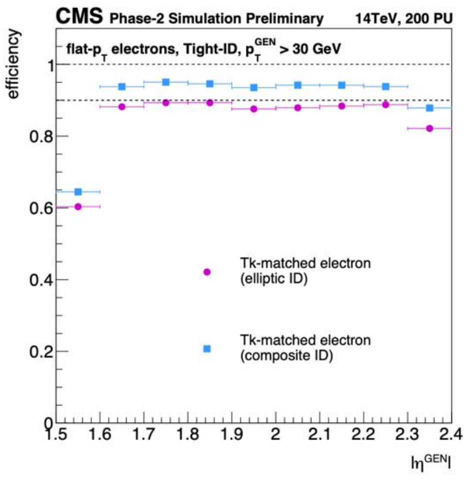

	LUT	FF	BRAM	DSP
e/γ IP	3.1%	0.4%	0.0%	1.6%
Total	24.4%	17.6%	29.5%	14.3%

- **Tau reconstruction: Tauminator**
  - Training dedicated CNN to reconstruct and identify Tau-induced signal in calorimeters (5x9)
  - Elegant way to deal with different geometries in Barrel (Crystals) and EndCap (HGCAL 3D clusters).

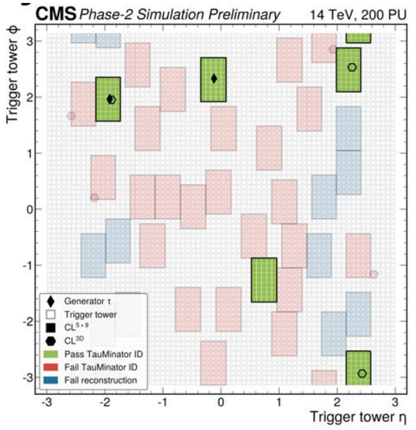

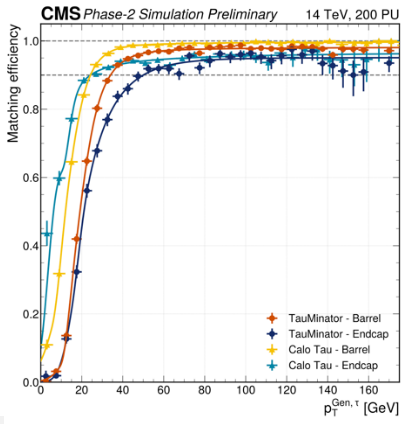

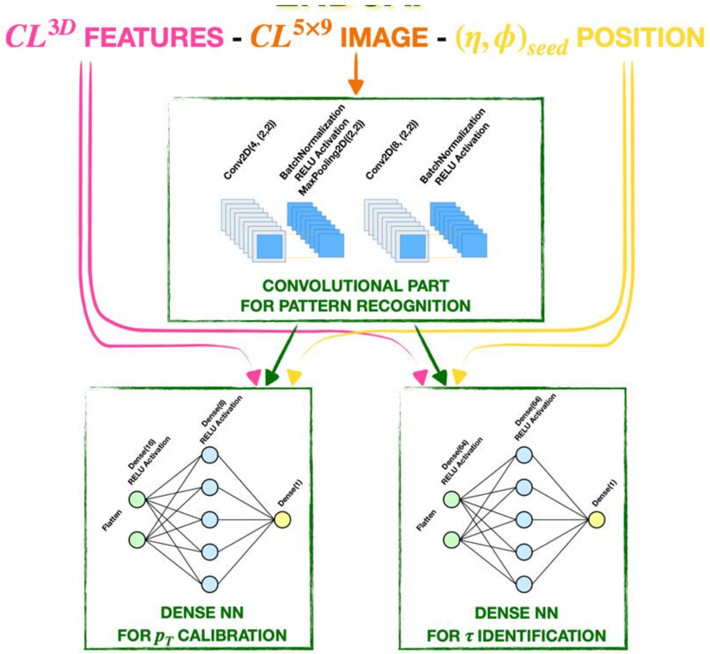

Recent development highlights (with ML)

- **SeededConeJets:**
  - Jet finding based on PF candidates
  - Iterative approach computing distance between each particle and jet radius (SC4 or 8), compute jet axis and energy.
  - Jet matching anti-kt jets

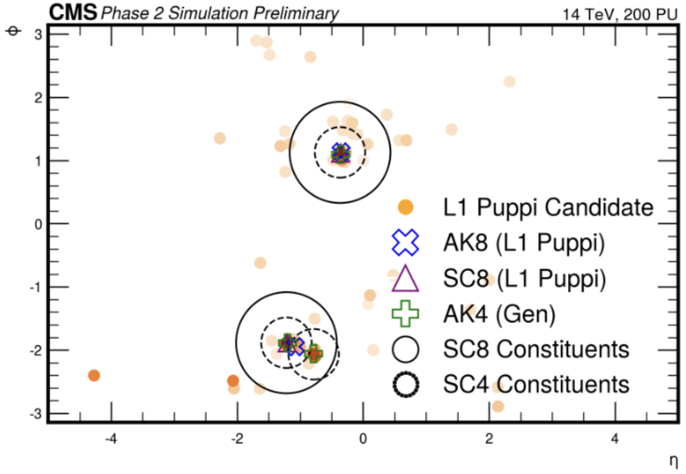

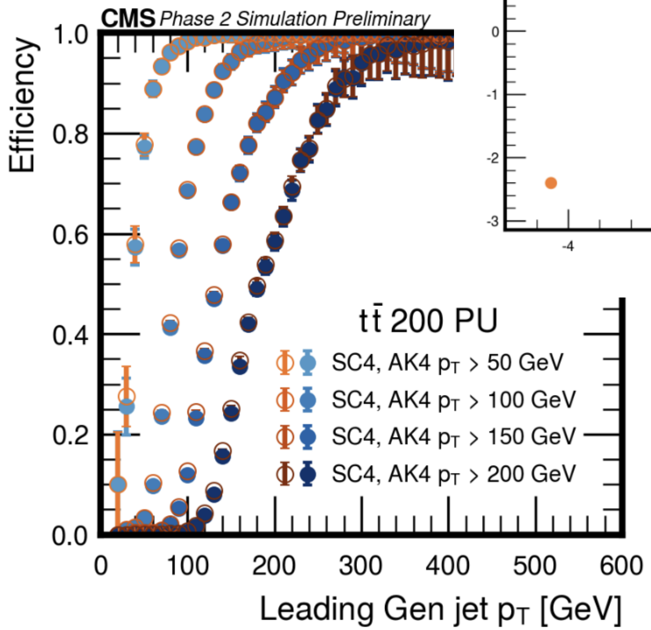

- **Continual learning**:
  - Elegant way to deal with changing detector conditions (ageing, noise, LHC interfill, etc.)
  - Train a model with a continuous stream of data. Learns from a sequence of partial experiences rather than all the data at once.
  - Update model to changing conditions without large MC production.
  - Method tested on Vtx reconstruction

Non-CL Top-Up Model

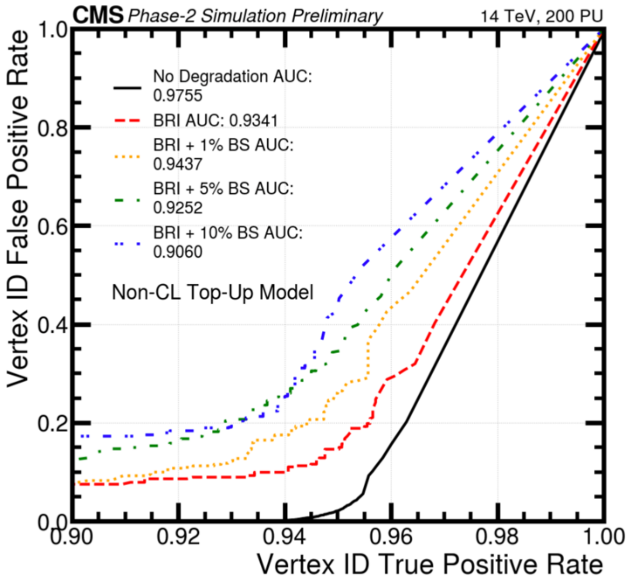

CL Top-Up Model

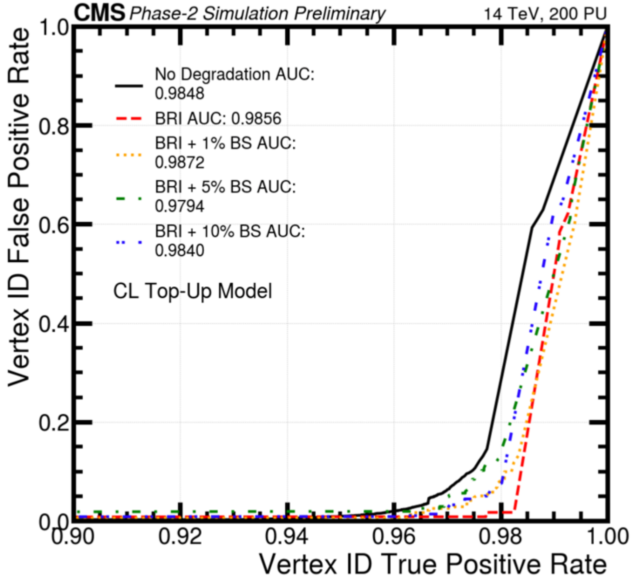

HL-LHC: challenges

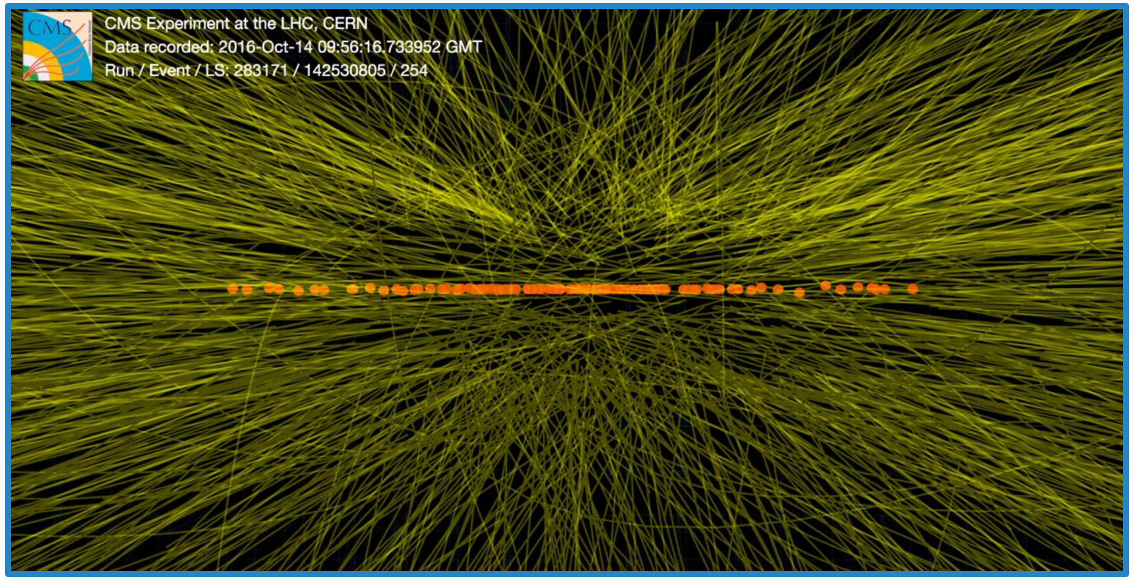

- **Expected pileup** (PU): ~140 (nominal HL-LHC lumi)
- Motivates/requires:
  - Improved granularity wherever possible
  - Novel approaches to in-time Pile Up mitigation: Precision Timing detectors (30ps)
  - A complete renovation of the Trigger and DAQ systems for better selectiveness, despite the high PU.

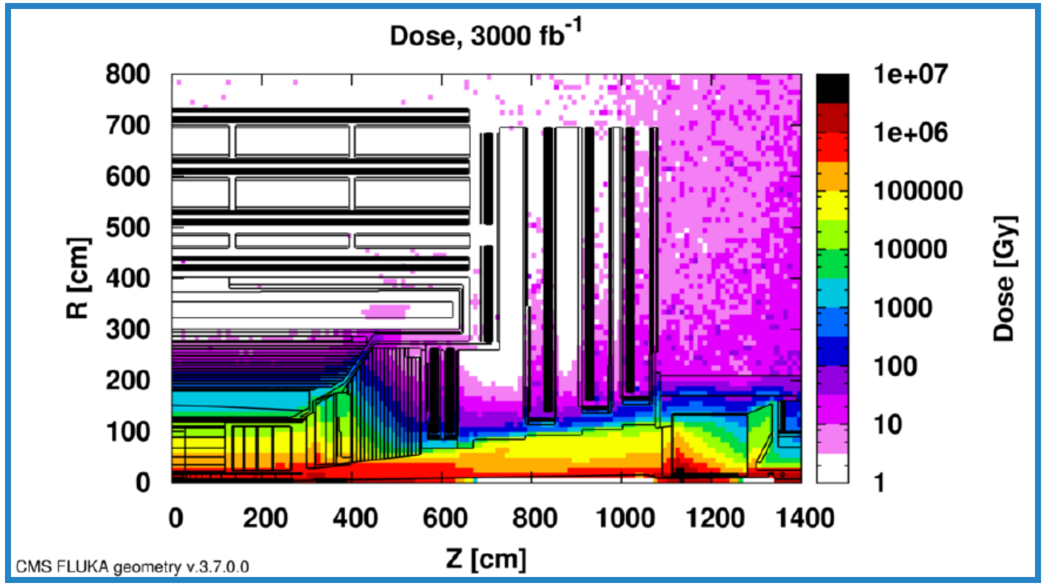

- **Radiation damage / accumulated dose** in detectors and on-board electronics may result in a progressive degradation of the performance.
- Maintain detector performance in harsh conditions:
  - The complete replacement of the Tracker and Endcap Calorimeter systems.
  - Major electronics overhaul and consolidation of the Barrel Calorimeters and Muon systems

From ML to FPGA

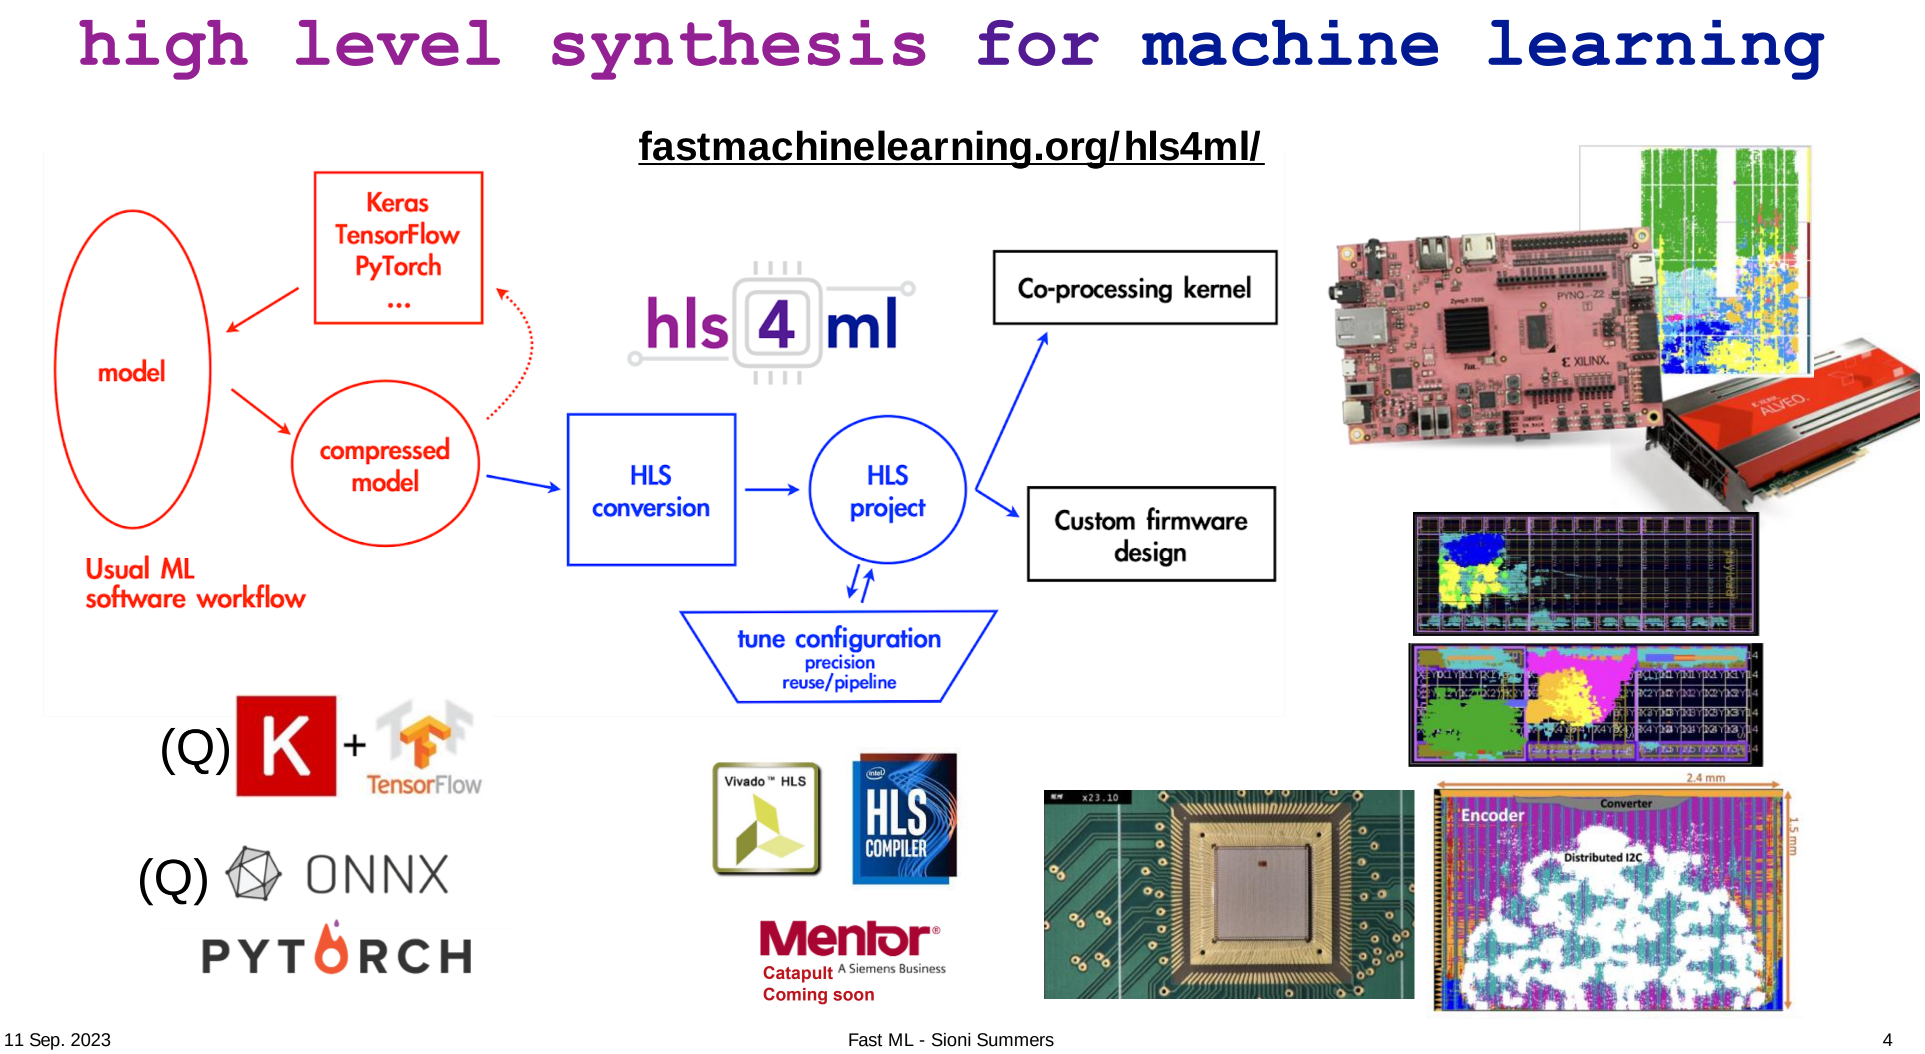

The Phase-2 Trigger Upgrade: Physics case

Improve precision of SM tests (*i.e. Higgs couplings,* m_W)

Target unobserved SM processes (*i.e.* $H \to HH;\ H \to cc$)

Search for deviations at high momenta (*i.e. Effective Field Theories*)

Probe new phase space (*i.e. Long-lived particles*)

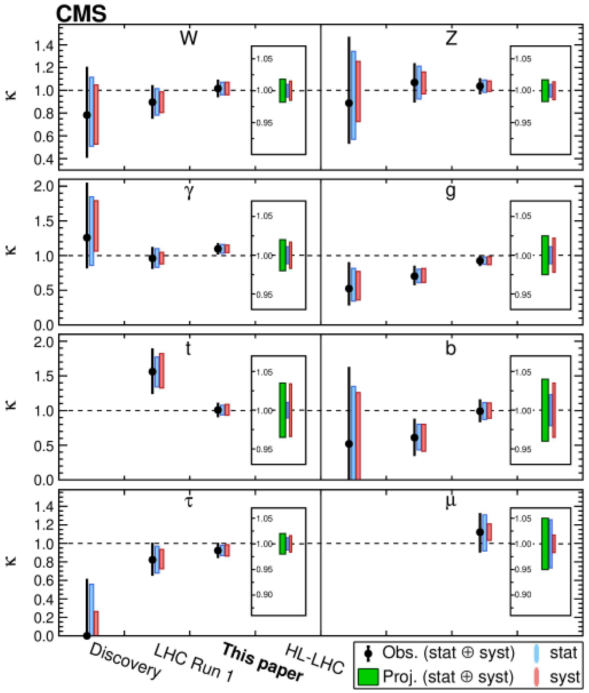

Nature 607, 60–68 (2022)

CMS-EXO-19-019

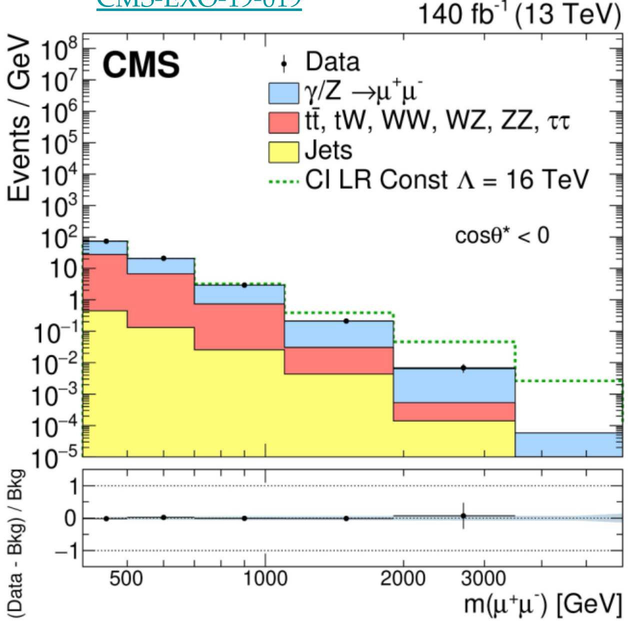

$$\mathcal{L}_{eff} = \mathcal{L}_{SM}^{(4)} + \sum \frac{C_x}{\Lambda^2} O_{6,x} + h.c.$$

e.g. for cτ = 5 cm, <βγ> ~ 30

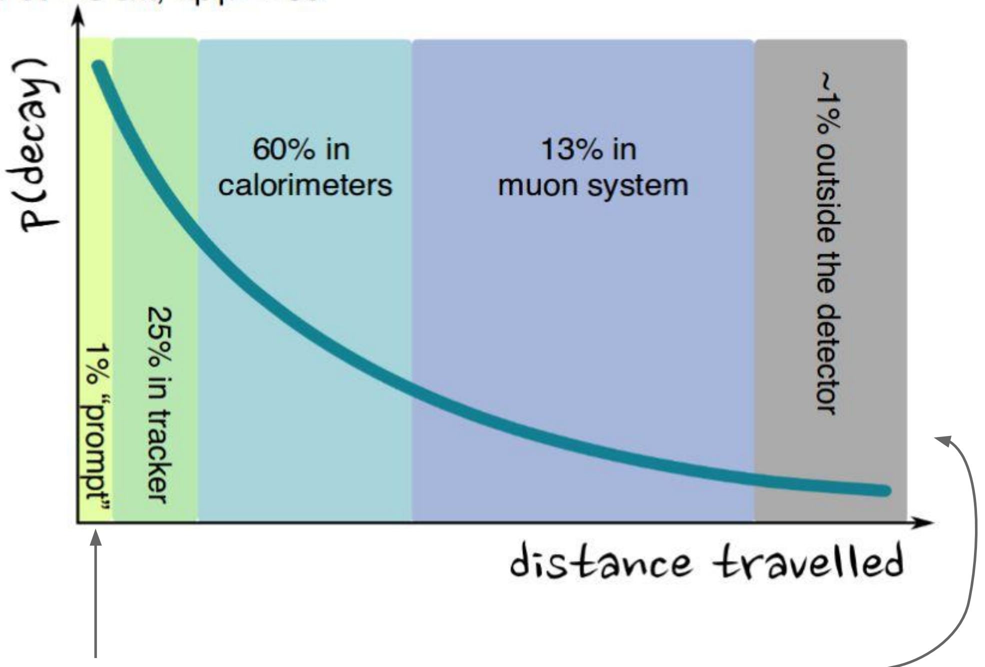

Bulk of searches

Run 3 at a glimpse

- With almost one and half year to go, Run-3 has already surpassed Run-2 luminosity
  - Almost 170 pb^{-1} recorded
- **New strategies** have been deployed both at L1T and HLT
- Excellent opportunity to extend physics reach and try new ideas to guide our path in the future
  - New capabilities to trigger on long-lived particles
  - Anomaly detection
  - Triggerless readout (scouting)
  - Increased GPUs usage
  - Extensive use of ML techniques

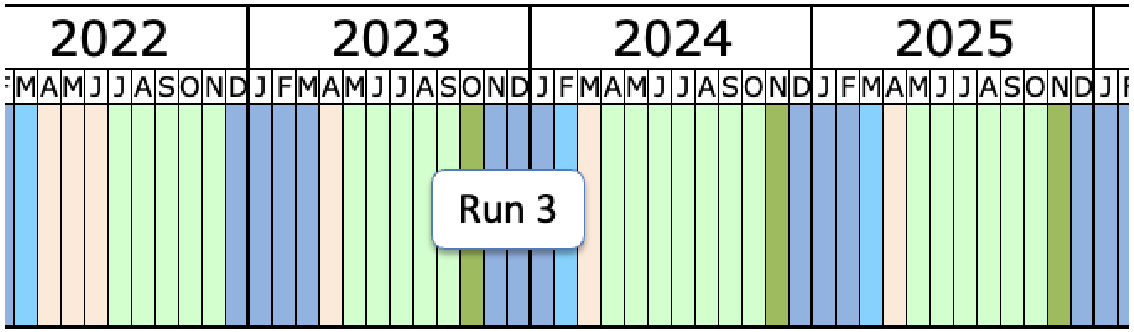

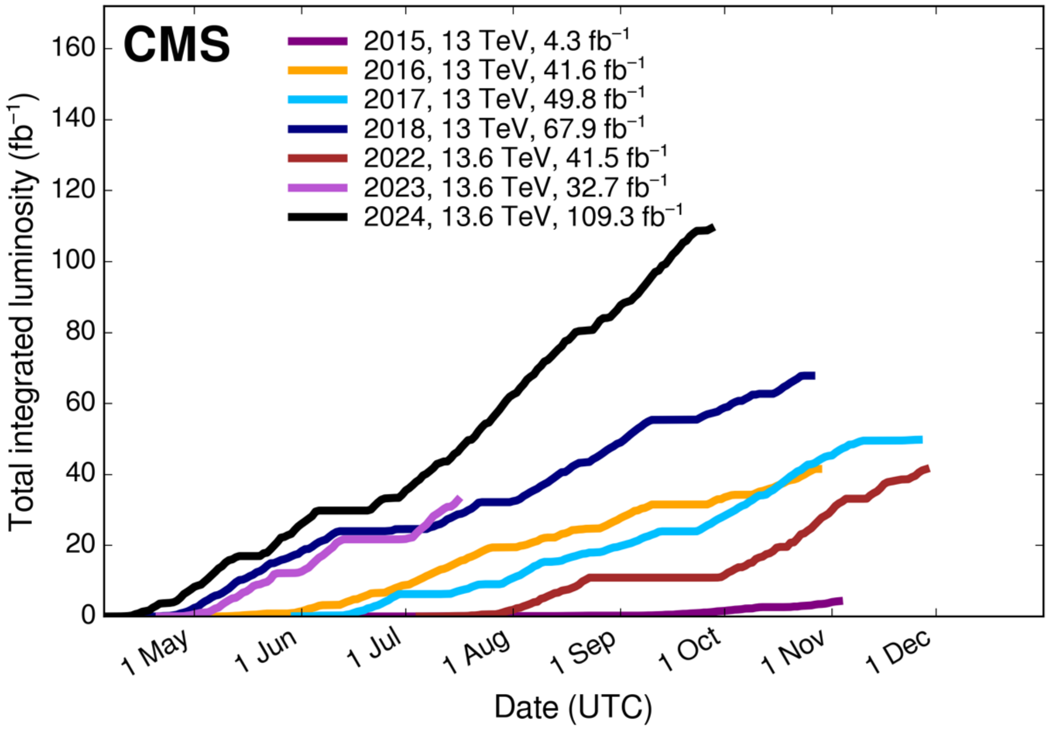

ML at HLT

- Tau @HLT
  - Reconstruction: Hadron plus strip
  - DeepTau identification: CNN+DNN based tagger

CMS DP-2024/042

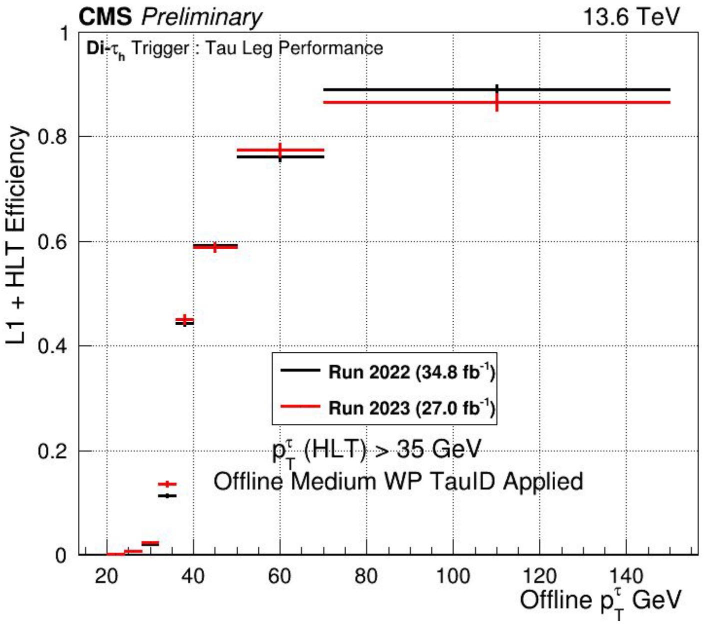

- ParticleNet b-jet tagger @HLT. GNN-based
  - Jets treated as a permutation-invariant point cloud
  - Performance gain, especially for HH processes

CMS DP-2023/021

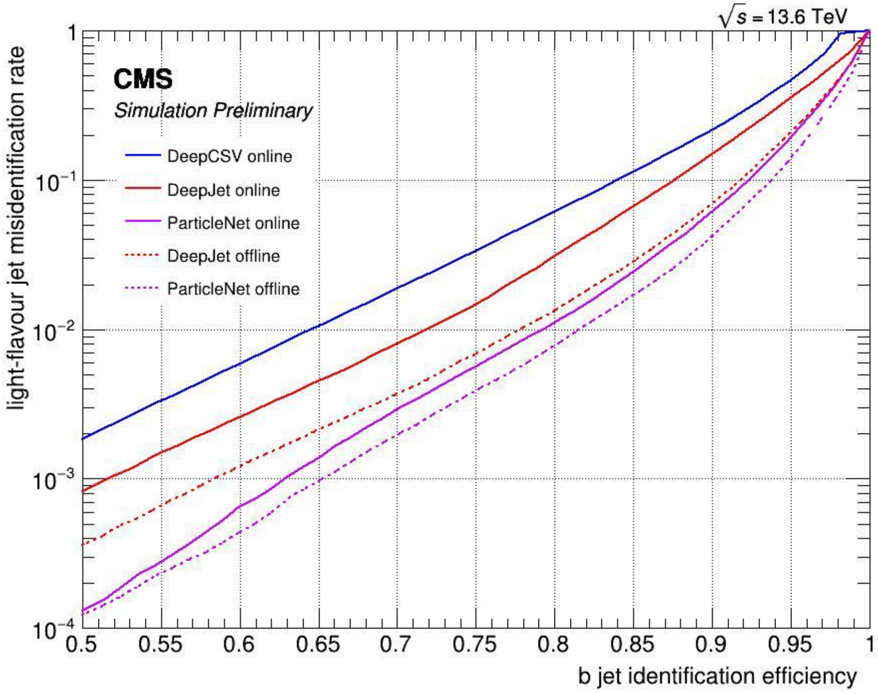

Long-lived particle triggers

- Many models predict the existence of **long-lived particles** (LLPs)
  - Many Exotic scenarios not envisioned when the trigger system was being designed!
- LLPs transit layers at later times, timing information
- LLPs decay far from the interaction point and show displaced signatures
  - Dedicated trigger paths exploiting unique features
  - Displaced jets in the tracker, calorimeters, or muon systems
- Strategies adopted mainly at HLT for Run 3
  - Some ideas already at L1
- Run 3 is the perfect benchmark for “crazy” ideas for HL-LHC

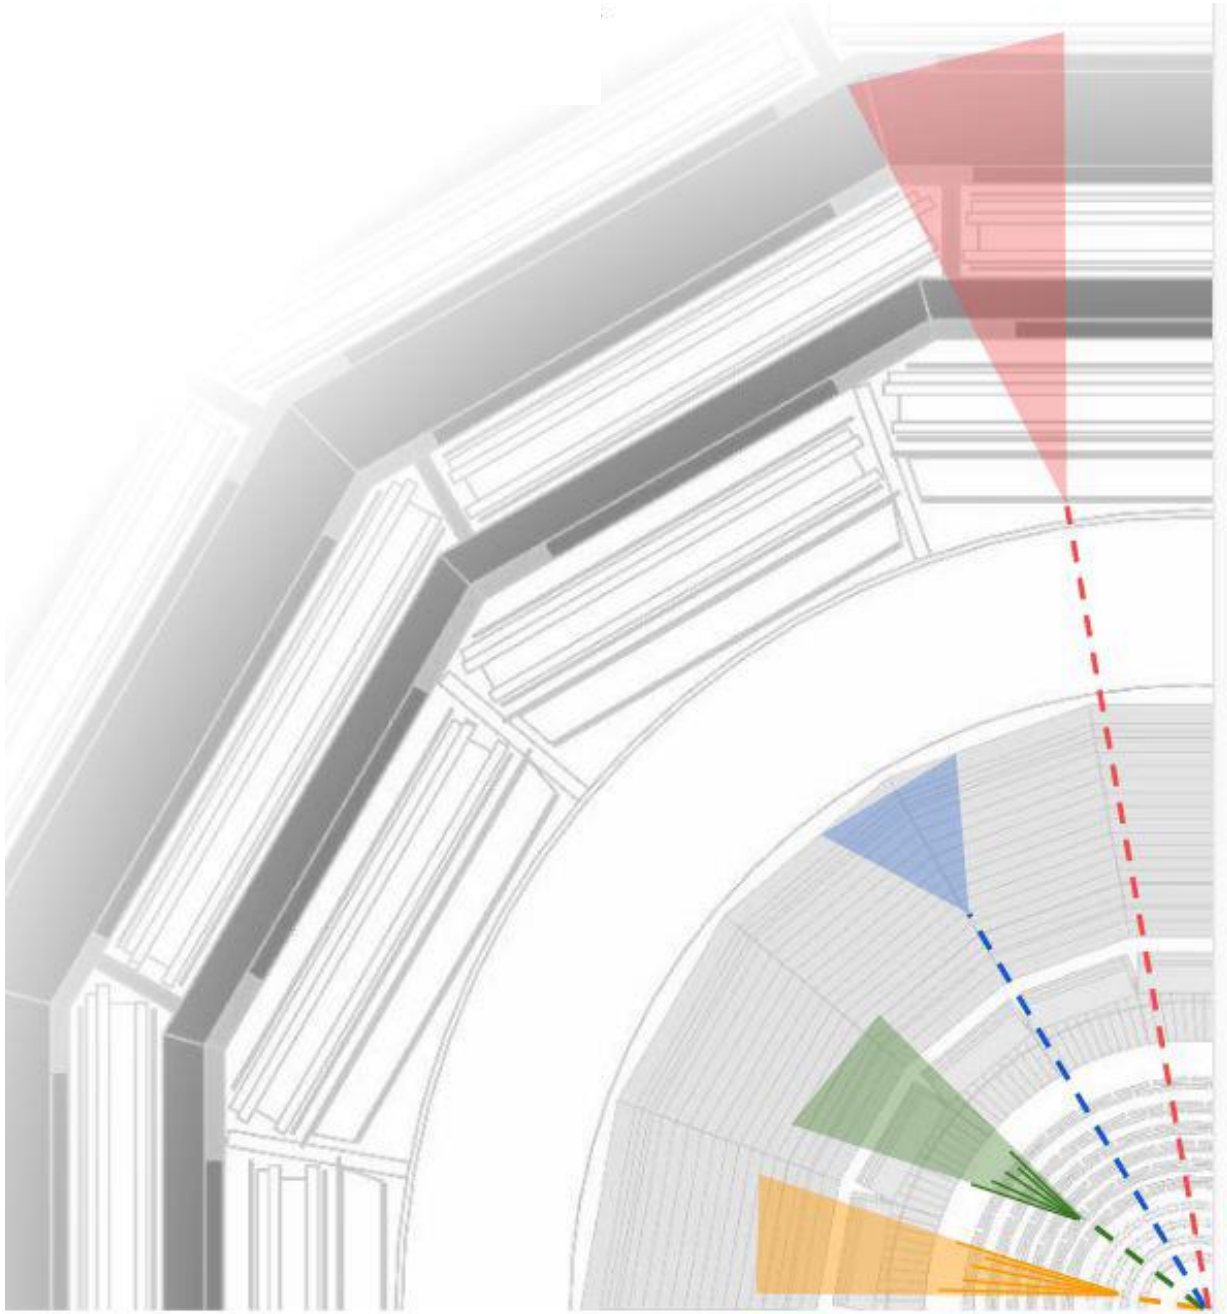

First Run 3 search: displaced dimuons at 13.6 TeV

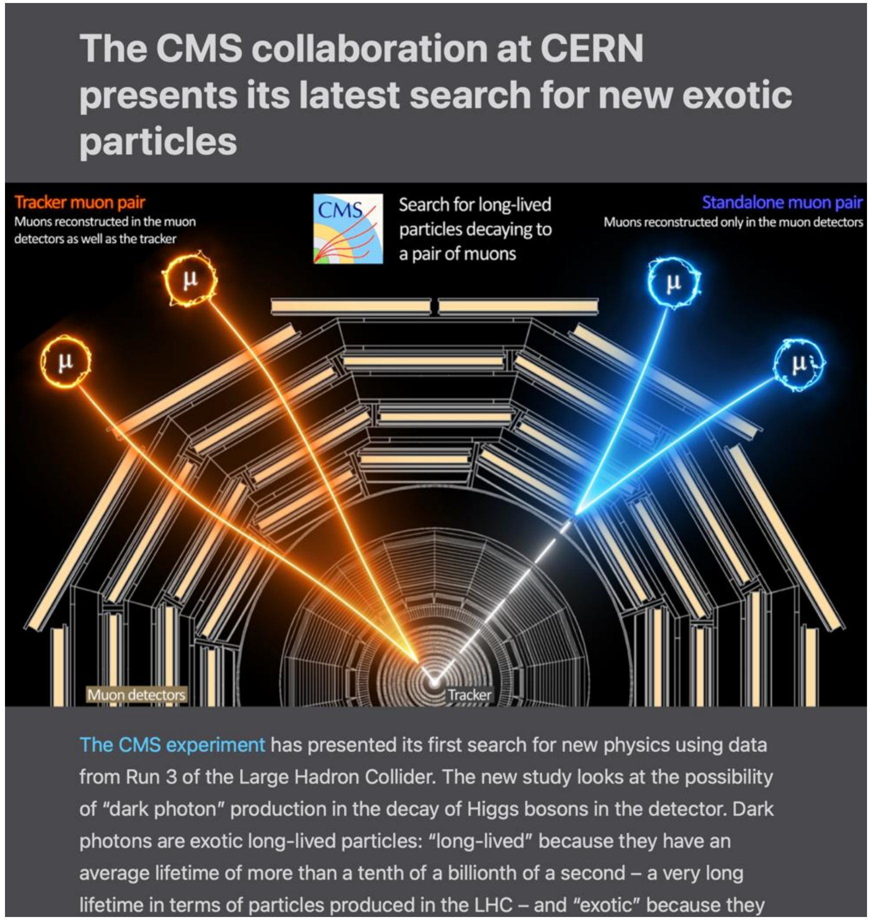
The CMS collaboration at CERN presents its latest search for new exotic particles

The CMS experiment has presented its first search for new physics using data from Run 3 of the Large Hadron Collider. The new study looks at the possibility of "dark photon" production in the decay of Higgs bosons in the detector. Dark photons are exotic long-lived particles: "long-lived" because they have an average lifetime of more than a tenth of a billionth of a second – a very long lifetime in terms of particles produced in the LHC – and "exotic" because they

https://cms.cern/news/long-lived-particles-light-lhc-run-3-data

https://home.cern/news/news/physics/cms-collaboration-cern-presents-its-latest-search-new-exotic-particles

With a strong Spanish contribution:

Universidad de Oviedo

IFCA Instituto de Física de Cantabria

Ciemat Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

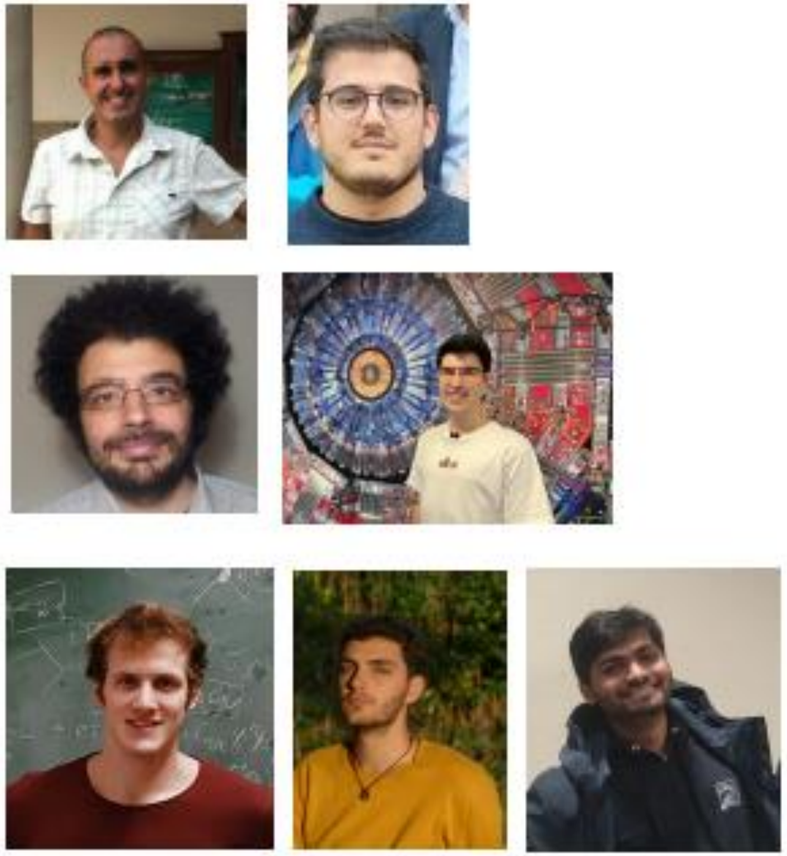

Displaced dimuons at 13.6 TeV. New triggers

- Use the 2022 dataset (**36.7 fb^{-1}**) recorded with **new LLP triggers** with thresholds **down to $p_T(\mu)$ > 10 GeV**
  - Re-optimized L1 triggers, including p_T without beam spot constraint, and new reconstruction algorithms.
  - Use d_{xy} information at trigger level to control the background rate.
- **Factor 2-4 more signal efficiency**
- Despite **2.5 smaller dataset**, comparable (or better) sensitivity w.r.t. 13 TeV result.

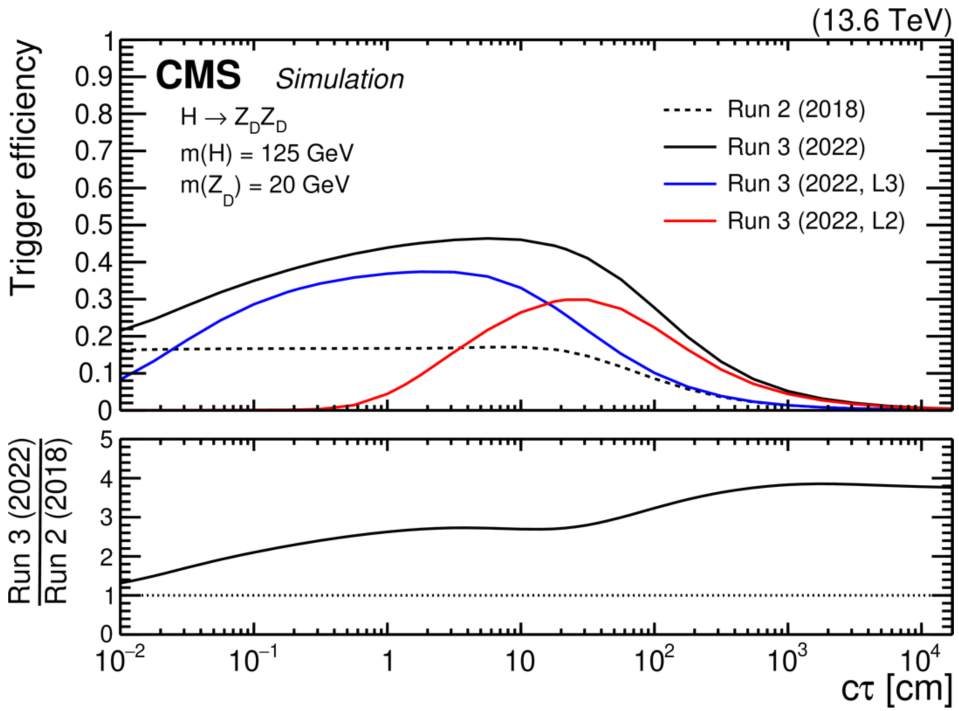

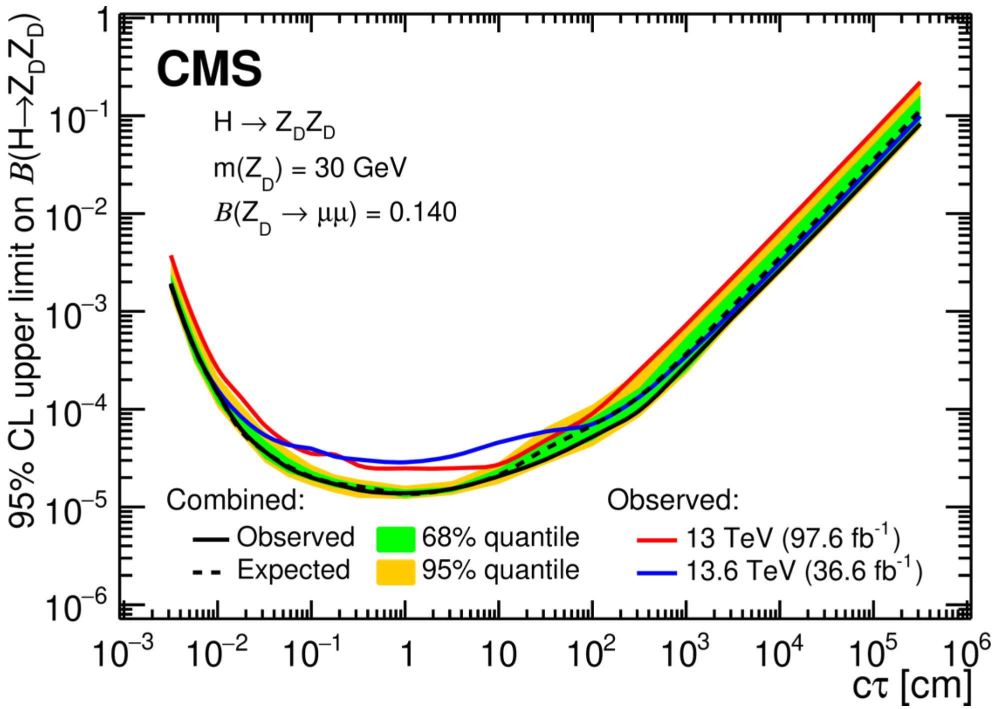

Multithreading and GPUs

- Multithreading (MT) is key to fully exploit HLT farm computational power
  - inter-event, intra-event, in-algorithm parallelism;
  - usage of “data handles” to define the data dependency among modules;
  - lower memory usage

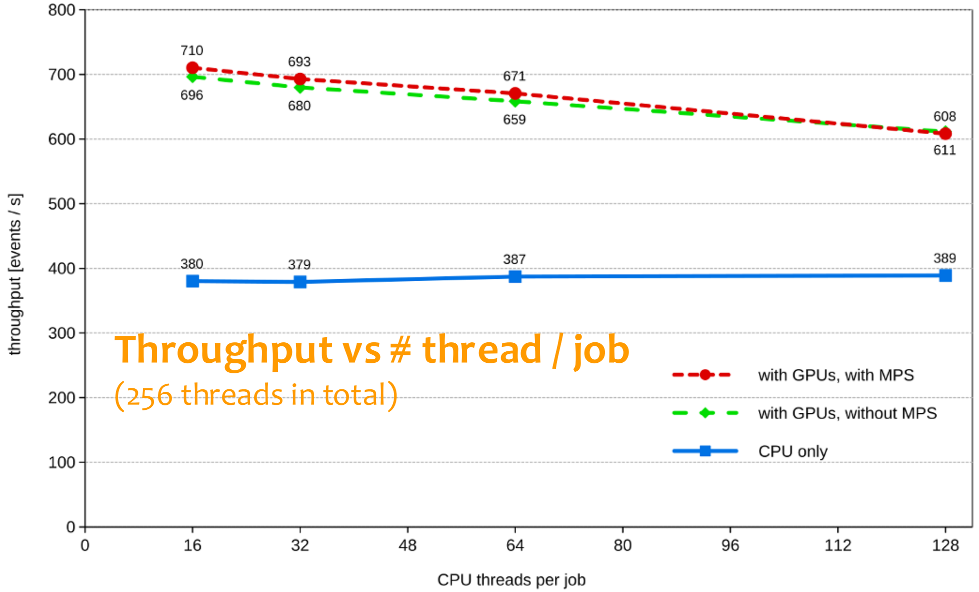

- CMS HLT farm heterogeneous since 2022 (AMD CPU + Nvidia T4):
  - 40% of HLT reconstruction ported to GPU
    - Pixel local reconstruction
    - Pixel tracking and vertexes
    - ECAL local reconstruction
    - HCAL local reconstruction

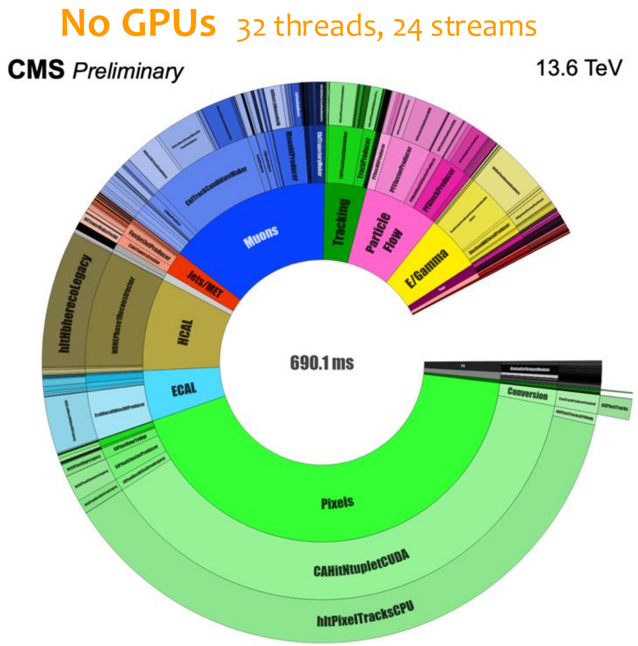

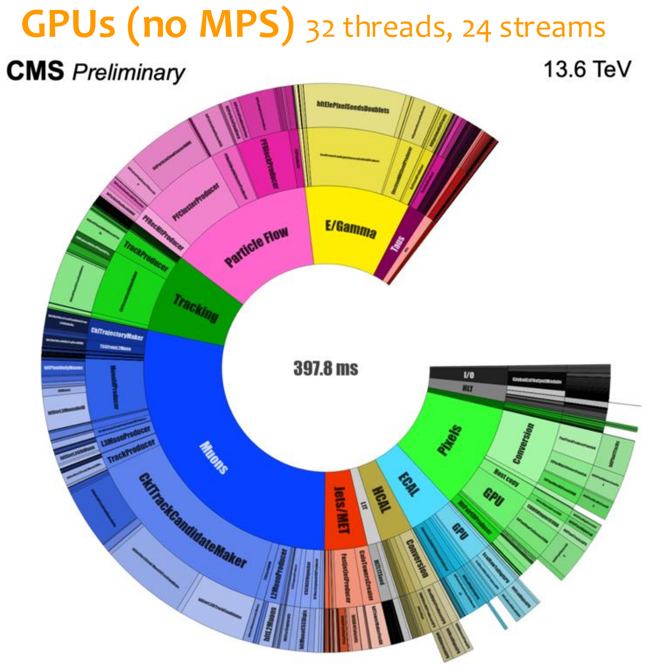

INTREPID project

INnovativeTRiggEr techniques for beyond the standard model PhysIcs Discovery at the LHC

- Improve muon trigger reconstruction with advance techniques based on machine learning: Graph Neural Network
  - Work already started with the overlap muon track-finder, first version of the network, using every detector layer as a nodes and Δϕ and Δη as edge parameters
- Considering AI accelerators (AI Xilinx Versal Chip)
  - Provide the necessary throughput and latency for triggering?

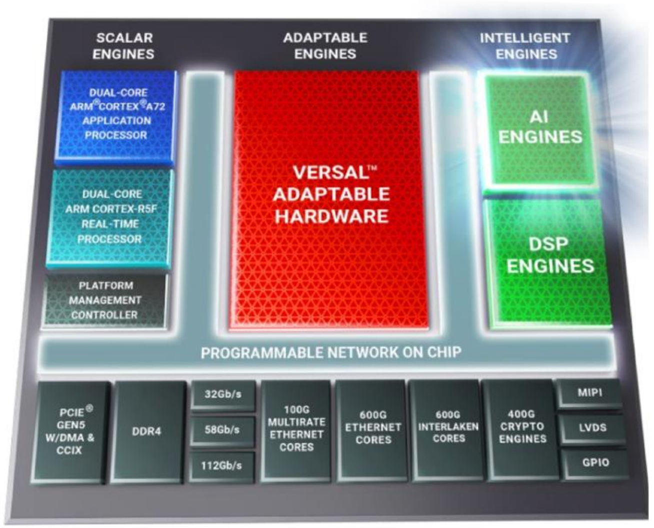

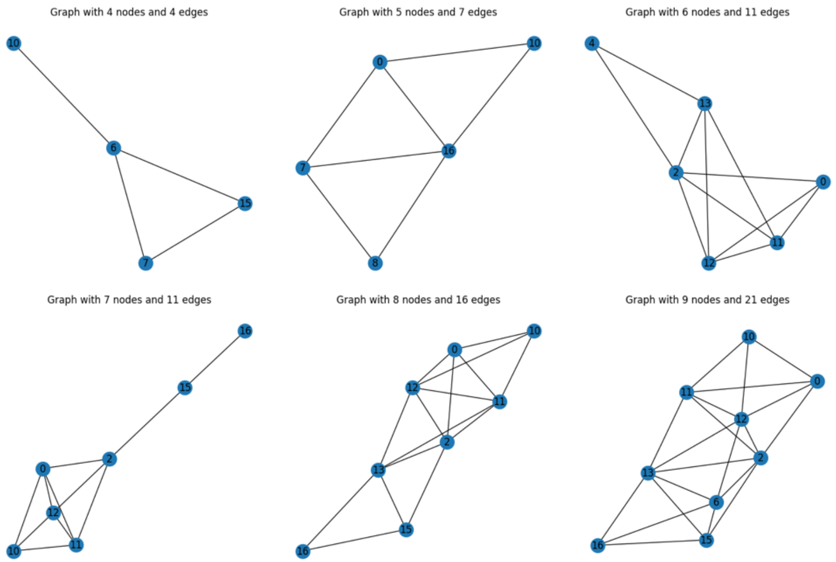

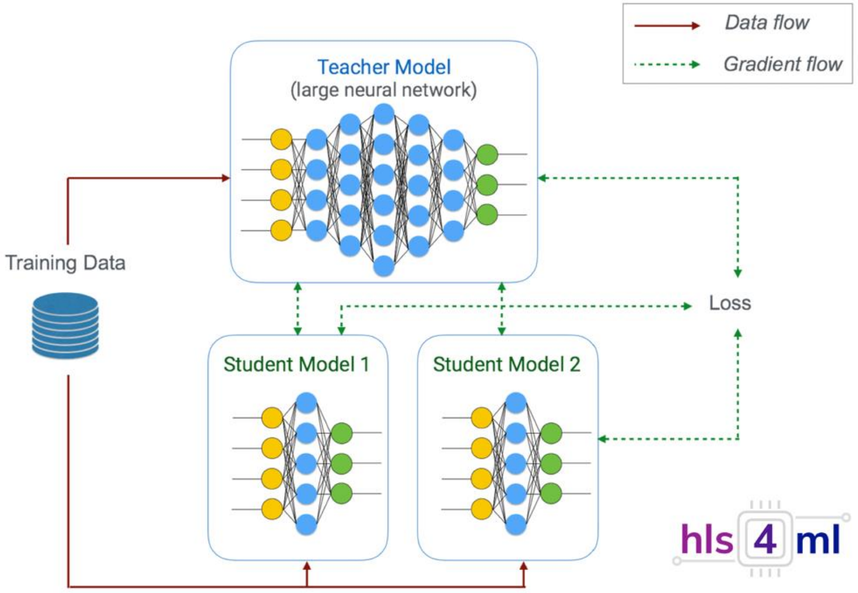

The Next-Generation Trigger Project

Innovative computing technologies for data acquisition and processing for the HL-LHC and beyond

- Enhance the triggers and the data collection and processing, and thus the scientific potential, of ATLAS and CMS in the HL-LHC phase **beyond the currently projected scope**.
    - Accelerate the evaluation and introduction of novel computing, engineering and scientific ideas already with demonstrators for Run3, but with main focus on HL-LHC
    - Provide a major push to the work already ongoing in the experiments, by enabling lines of research **currently not feasible within existing financial, human and technology constraints**
    - Provide **critical insight to develop data flows** for the even more ambitious objectives of a future collider, such as the Future Circular Collider (FCC) currently in its Feasibility Study phase
- CERN involvement to **ensure that other current & future CERN experiments benefit from the results** in terms of computing frameworks and theoretical modelling.
- All project results (IP) will belong to CERN and will be released under a valid open policy and IP generated will be released under appropriate open licenses in compliance with the **CERN Open Science Policy**.

NexTGen Next Generation Triggers

https://nextgentriggers.web.cern.ch/