

## The CMS Trigger system for the HL-LHC

---

**Alexandre Zabi**<sup>†,\*</sup>

*Laboratoire Leprince-Ringuet CNRS/IN2P3 - Ecole Polytechnique,  
Institut Polytechnique de Paris, Plaiseau, France*

*E-mail:* [Alexandre.Zabi@cern.ch](mailto:Alexandre.Zabi@cern.ch)

The High-Luminosity LHC will open an unprecedented window on the weak-scale nature of the Universe, providing data to perform high-precision measurements of the Standard Model, as well as searches for new physics beyond the Standard Model. Such precision measurements and searches require information-rich datasets with a statistical power that matches the high luminosity provided by the upgrade of the LHC. Efficiently collecting these datasets will be a challenging task, given the harsh environment of up to 200 proton-proton interactions per LHC bunch crossing. For this purpose, CMS is designing an efficient two-level trigger system: the Level 1 Trigger (L1T), implemented in advanced hardware, and the High Level Trigger (HLT), a streamlined version of the CMS reconstruction software running on a computer farm. The L1T will include tracking information and high-granularity calorimeter information for the first time. The current conceptual system design is expected to take full advantage of FPGA and link technologies over the coming years, providing a high-performance, low-latency computing platform for large throughput and sophisticated data correlation across diverse sources. The envisaged L1T system will closely replicate the full offline object reconstruction, to perform more sophisticated and optimized selection. The higher luminosity, event complexity and input rate of 750 kHz present an unprecedented challenge to the HLT, which aims to achieve a similar efficiency and rejection factor as today, despite the higher pileup, and a purer preselection. The introduction of a heterogenous platform combining CPUs and GPUs is described. The expected performance of the upgraded trigger system is presented.

*40th International Conference on High Energy physics - ICHEP2020  
July 28 - August 6, 2020  
Prague, Czech Republic (virtual meeting)*

---

\*Speaker

<sup>†</sup>On behalf of the CMS Collaboration

## 1. Introduction

The High-Luminosity LHC (HL-LHC) [1] presents the opportunity for a very rich and ambitious physics program, exploiting an integrated luminosity of 3000–4000 fb<sup>-1</sup>. The LHC will undergo major upgrades of its components leading to an increase of the instantaneous luminosity to  $5 \times 10^{34}$  cm<sup>-2</sup>s<sup>-1</sup>, five times the accelerator's original design value. In its "ultimate" configuration, the HL-LHC will reach a peak instantaneous luminosity of  $7.5 \times 10^{34}$  cm<sup>-2</sup>s<sup>-1</sup>, increasing the average number of proton-proton collisions per bunch crossing (pileup) to around 200. The ultimate performance of the HL-LHC would enable the collection of 400 to 450 fb<sup>-1</sup> of integrated luminosity per year, potentially providing a total of 4000 fb<sup>-1</sup> to each of the CMS and ATLAS experiments. The CMS detector requires a trigger and data acquisition system with exceptional performance to collect the required information-rich datasets with these challenging running conditions. Along with the sub-detector upgrades [2], a complete replacement of the trigger, including the Level-1 (L1) and the high level trigger (HLT), and data acquisition (DAQ) system, with increased throughput, is planned. The Phase-2<sup>1</sup> upgrade of the trigger and DAQ system will keep a two-level strategy, while increasing the L1 maximum rate to 750 kHz to maintain the acceptance for physics. The total latency will be increased from 3.8 μs to 12.5 μs to allow, for the first time, the tracker and high-granularity calorimeter information to be included. Trigger data analysis will be performed through sophisticated algorithms, including widespread use of Machine-Learning in large FPGAs [3]. Similarly to Phase-1, the Phase-2 HLT will have access to the full granularity of the detector with the target for average timing per event of 500 ms (measured on a 2018 HLT node). The selection algorithms will perform a rate reduction leading to an output bandwidth of 7.5 kHz. The most promising avenue of development is the use of heterogenous hardware; for example, porting part of the reconstruction to run on GPUs [4].

## 2. The CMS detector for Phase-2

In order to fully exploit the HL-LHC running period, major consolidations and upgrades of the CMS detector are planned [2]. Given the high particle multiplicity expected, the performance required on event object reconstruction to achieve the extraction of physics signatures relies on the implementation of higher granularity detectors along with robust readout electronics. The CMS collaboration plans to replace both the Strip and Pixel tracking detectors, with an Inner Tracker featuring small-size pixel sensors and an Outer Tracker equipped with strip and macro pixel sensors, extending their coverage to  $|\eta| = 3.8$ . The Outer Tracker will implement stacked strip modules, reducing the hit multiplicity and allowing track candidates for the trigger (L1 tracks) to be reconstructed up to  $|\eta| = 2.4$ . The readout electronics for the barrel calorimeters will be replaced to achieve finer granularity and provide timing information. The endcap calorimeters will be replaced by the high-granularity calorimeter (HGCAL), implementing over 6 million readout channels. This sampling calorimeter will provide shower separation and identification adapted to harsher conditions in the forward region of the detector. The muon detection system redundancy achieved through the combination of drift tubes (DTs), resistive plate chambers (RPCs), and cathode strip chambers (CSCs) will remain with consolidated electronics. Additional improved RPC (iRPC)

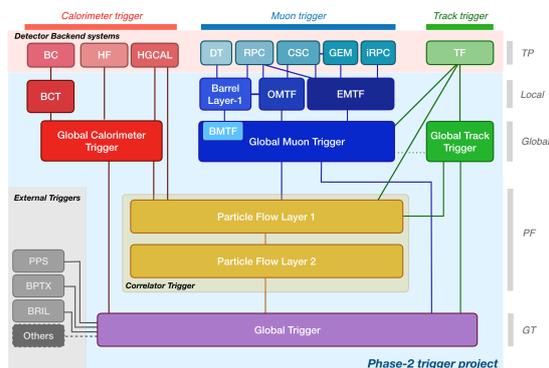
<sup>1</sup>Phase-2 refers here to the HL-LHC running period starting in 2025, while Phase-1 refers to the ongoing run.

chambers and gas electron multiplier (GEM) chambers will be installed to extend the coverage up to  $|\eta| = 2.4$  and  $2.8$ , respectively. A minimum ionizing particle timing detector placed in front of the barrel and endcap calorimeters will provide precise timing measurement of charged tracks.

### 3. The Level-1 trigger Phase-2 upgrade

The Phase-2 upgrade of the L1 trigger system is designed not only to maintain the efficiency of the signal selection to the level of the Phase-1 performance but also to significantly enhance, or enable, the selection of any possible new physics manifestations that could lead to unconventional signatures [3]. High-precision measurements of physics processes will benefit from the extension of the available phase space such as enhanced trigger coverage in the forward region of the detector or the ability to exploit fully hadronic final states. Moreover, a longer latency will enable higher-level object reconstruction and identification, as well as the evaluation of complex global event quantities and correlation variables to optimize physics selectivity. The implementation of sophisticated algorithms using particle-flow (PF) reconstruction techniques or Machine-Learning based approaches can now be contemplated. In addition, the design includes a dedicated scouting system streaming data from key parts of the trigger at 40 MHz, via FPGAs into HPC resources. The scouting system provides unprecedented flexibility for parasitic debugging and commissioning of new ideas and is also being investigated for physics channels which are impossible through traditional triggering techniques.

The Phase-2 Level-1 Trigger system will need to handle an increased data volume from higher granularity detectors and a higher particle multiplicity. The total amount of input data which is required to be processed is over 60 Tb/s in comparison with the 2 Tb/s for the Phase-1 system. The upgrade project includes a program of R&D to produce the required prototype electronics based on modern technology [3]. Generic high input/output processing boards based on the Advanced Telecommunications Standard (ATCA) have been designed and equipped with Xilinx Virtex Ultrascale+ (VU9P) FPGAs (providing 8 times more computing resources than the Virtex 7 family used in Phase-1). The boards feature high-speed serial optical links running at up to 28 Gb/s (compared to the 10 Gb/s of the Phase-1 system) to transport the large data volumes.

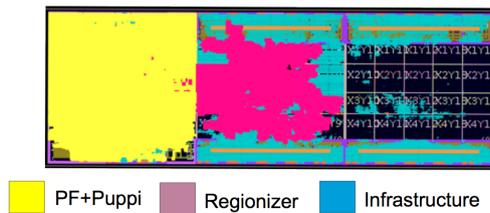
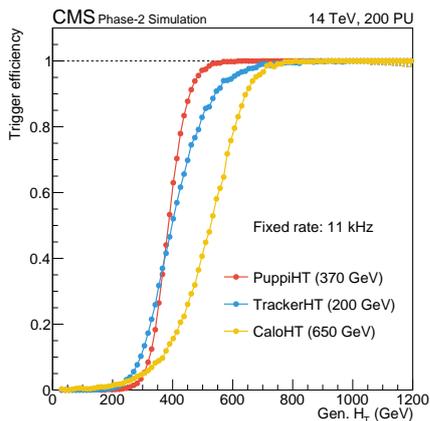


**Figure 1:** Functional diagram of the CMS L1 Phase-2 upgraded trigger design. The calorimeter trigger is composed of the barrel calorimeter trigger (BCT) and the global calorimeter trigger (GCT), receiving inputs from the barrel (BC), endcap (HGCAL) and forward (HF) calorimeters. The muon trigger is composed of a barrel layer-1 and muon track finders processors: BMTF, OMTF and EMTF, for each pseudorapidity regions: barrel, overlap and endcap respectively, and receiving inputs from drift tubes (DT), resistive plate chambers (RPC), cathode strip chambers (CSC), and gas electron multipliers (GEM). The global muon trigger (GMT) matches muons with tracks from the track finder (TF). The event vertex is reconstructed in the global track trigger (GTT) and the correlator trigger (CT) implements the particle-flow reconstruction. The global trigger (GT) issues the final L1 trigger decision.

The conceptual design of the Phase-2 Level-1 Trigger system is the result of several considerations: the design has to efficiently distribute and process the input trigger primitives, provision appropriate resources and interconnections and retain enough headroom for future flexibility and

robustness to evolve with running conditions and physics needs. The high-level functional diagram of the system is shown in Fig. 1. The system features four distinct and independent trigger processing paths with a calorimeter trigger, a muon trigger, a track trigger and a particle-flow trigger. This division reflects the need to generate complementary types of trigger objects to achieve the best physics selectivity. The key design feature is the implementation of a correlator trigger combining all detector information and running sophisticated algorithms. The final trigger decision is performed at the global trigger level. This architecture meets additional constraints, such as the allowed maximum FPGA occupancy remaining below 50% (to ensure future flexibility in the design of algorithms) and the total latency remaining under  $9.5 \mu\text{s}$  (to retain 20% contingency).

The trigger algorithms are designed with the extensive use of tracking information to reach near offline performance. The availability of fully reconstructed tracks translates into sharper turn-on efficiency curves. Algorithm implementation in firmware relies heavily on High-Level-Synthesis, allowing for faster turn-around and the development of new approaches, such as those based on Machine-Learning techniques. With the availability of tracking and high granularity detector data, global event reconstruction algorithms such as particle-flow can be implemented. Particle-flow reconstruction [5] has been successfully used by CMS in offline data analyses and the HLT. Additionally, the reconstructed event primary vertex from tracks is used by the PUPPI [6] algorithm to filter particles based on a measure of their probability of coming from pileup. The combination of PF and PUPPI leads to a large reduction of the event complexity, while preserving the core physics information. This translates into a smaller bandwidth and reduced FPGA resource utilization. Trigger objects are formed from PF and PUPPI candidates, such as the  $H_T$  trigger algorithm, for which efficiency curves are shown in Fig. 2 (left). This ambitious prototype algorithm was implemented in firmware and demonstrated in hardware using Vivado-HLS targeting a Xilinx-VU9P FPGA. The algorithm uses less than 50% of the FPGA resources (see Fig. 2 (right)) and a latency of  $0.7 \mu\text{s}$  meeting the requirements of the project.

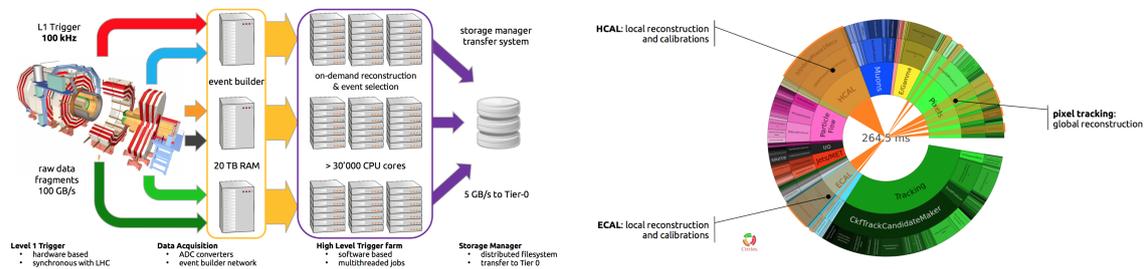


**Figure 2:** Left: The signal efficiency is shown for  $\bar{t}\bar{t}$  events selected by  $H_T$  triggers utilizing different sets of inputs at a fixed trigger rate of 10.5 kHz. The performance of the  $H_T$  trigger quantity computed with PUPPI jets, tracker-only jets and calorimeter jets are compared. Top: The FPGA floorplan for the prototype particle-flow algorithm, showing how the elements of the algorithm have been placed in the FPGA fabric (Xilinx Virtex (VU9P)).

#### 4. The Phase-2 Upgrade of the High Level Trigger

The HLT system for Phase-2 [4] is aiming to achieve a rejection factor of 100:1 while facing harsher LHC running conditions and increased L1T output rate (750 kHz instead of 100 kHz

during Phase-1). The event reconstruction will be performed on data delivered by a more advanced detector with enhanced granularity and increased acceptance as described in Section 2. The average processing time per event should be kept under 500 ms <sup>2</sup> while this HLT processing time is expected to increase with instantaneous luminosity and pileup. Current studies are ongoing using extrapolation from 2017 data and are indicating that the processing time does not scale linearly at higher pileup. The current HLT system implements a CPU farm composed of more than 30,000 cores (see Fig. 3). The software event reconstruction and selection algorithms run on commercial servers. The software includes about 4000 CMS Software (CMSSW) modules and exploits multi-threading (processing multiple events concurrently). The average HLT processing time per event is 264 ms, calculated on 2018 CPU and considering an average pileup of 50. The challenges imposed by the HL-LHC running conditions require an upscaling of the HLT system by a factor 18 in terms of computing power and a factor 25 more data throughput (65 GB/s instead of 2.5 GB/s at the start of Run-2) [4].



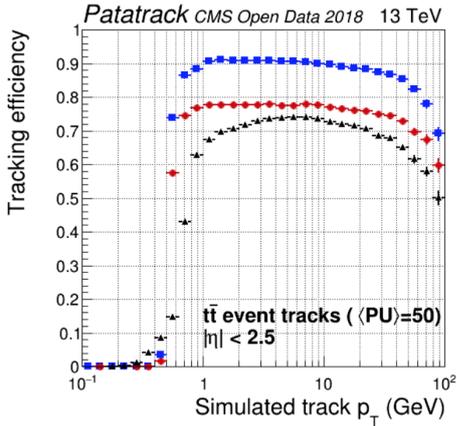
**Figure 3:** Left: Current CMS trigger and data acquisition system. Right: Online reconstruction processes running at the HLT. Highlighted are the processes considered to be offloaded to GPUs.

The main strategy used at HLT is to deploy the same framework and the same algorithms as the ones used in the offline data analyses. This allows new triggers (or physics seeds) to be rapidly developed and maintained while ensuring reproducibility and quick assessment of trigger efficiencies. The same strategy is employed for the Phase-2 upgrade of the HLT system, while the infrastructure is considering innovative solutions to optimise processing time. As other experiments at the LHC, CMS has used CPUs to process events at the HLT <sup>3</sup> during Run-1 and Run-2. The ongoing R&D aims at using coprocessors as offload engines for specific online reconstruction algorithms. A heterogeneous architecture would combine both CPUs and GPUs, where 80% of the online processing could be offloaded to GPUs. A first demonstration of this original architecture has been conducted and is referred to as "patatrack" [8]. This demonstrator performs pixel-based tracks and vertices reconstruction on a GPU platform. The workflow consists of copying the raw data from CPU to GPU, where multiple kernels run to perform the various steps of the tracking reconstruction algorithms before copying the tracks and vertices back to the host CPU. More details on the technical implementation of the patatrack demonstrator can be found in Ref. [8]. Other target HLT processes, such as ECAL and HCAL local reconstruction and calibration could also run on GPUs (see Fig. 3). A first estimation indicates that a single GPU can perform 60% more processing than a given HLT node equipped with 32 cores while using 25% of the same power. Based on these

<sup>2</sup>Estimated from CPUs available in 2018. It is expected that at the beginning of Run-4, CPUs should be around 2 to 3 times faster for the same price, and a further improvement will come from the use of GPUs.

<sup>3</sup>ALICE started using GPUs in 2010, LHCb is planning to include GPUs during Run-3

results, the current CMS HLT could offload 20% of the online reconstruction to GPUs. Plans to deploy this heterogenous architecture during Run-3 are in motion, allowing CMS to gain running experience. The superior computing power deployed by the GPUs compared to CPUs is used to perform efficiently sophisticated algorithms such as tracking reconstruction, resulting in improved performance for physics. Because the processing time is reduced by 60%, complex combinations of quadruplet and triplet pixel hits, including intermediate steps to remove fakes before final fitting would lead to a significant improvement in reconstruction efficiency as displayed in Fig. 4.



**Figure 4:** Pixel tracking reconstruction efficiency as a function of the simulated track  $p_T$  with the patatrack demonstrator (see text). The performance of tracks reconstructed with a combination of triplets and quadruplets (blue) and with quadruplets only (red) are compared with the standard online reconstruction in 2018 (black). These results were obtained with  $t\bar{t}$  simulated events with an average pileup of 50.

## 5. Conclusion

The CMS experiment is proposing solid solutions to the trigger and data acquisition challenge imposed by the extreme HL-LHC running conditions. The Phase-2 Level-1 trigger upgrade project is constructing a flexible and modular architecture with enhanced capabilities complying with the physics requirements. Sophisticated algorithms have been prototyped in FPGAs and exploit target hardware with demonstrated functionalities. The HLT proposes an innovative heterogenous architecture combining CPUs and GPUs along with a new programming model. Both projects intend to perform extensive demonstration of part of their system already during Run-3 of the LHC. Combination of Level-1 and HLT systems will provide CMS with event selection capabilities beyond the standard realm of triggering techniques.

## References

- [1] *High-Luminosity Large Hadron Collider (HL- LHC)* G. Apollinari et al., CERN Yellow Rep. Monogr. 4 (2017) 1.
- [2] *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, CMS, CERN-LHCC-2015-010, CMS-TDR-15-02.
- [3] *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CMS, CERN-LHCC-2020-004, CMS-TDR-021.
- [4] *The Phase-2 Upgrade of the CMS DAQ Interim Technical Design Report*, CMS, CERN-LHCC-2017-014, CMS-TDR-018.
- [5] *Particle-flow reconstruction and global event description with the CMS detector*, CMS, JINST 12 P10003 (2017) .
- [6] *Pileup Per Particle Identification*, D. Bertolini, P. Harris, M. Low, and N. Tran, JHEP 10 (2014) 059.
- [7] *Patatrack Heterogeneous Computing 2018 Demonstrator: Pixel Tracks*, CMS, CERN-CMS-DP-2018-059
- [8] *Heterogeneous reconstruction of tracks and primary vertices with the CMS pixel tracker*, A. Bocci, M. Kortelainen, V. Innocente, F. Pantaleo and M. Rovere. arXiv:2008.13461, 2020.