

Overview of the HL-LHC Upgrade for the CMS Level-1 Trigger

Simone Bologna^{*†}

University of Bristol, UK

E-mail: simone.bologna@cern.ch

In view of the High-Luminosity LHC, the Compact Muon Solenoid (CMS) experiment is planning to entirely replace its trigger and data acquisition system. Novel design choices are being explored through ATCA prototyping platforms and newly available interconnect technologies providing links up to 28 Gb/s. Higher-level trigger object reconstruction is performed through large scale FPGAs (such as Xilinx UltraScale) handling over 50 Tb/s of fine granularity detector data with an event rate of 750 kHz.

Topical Workshop on Electronics for Particle Physics TWEPP2019

2-6 September 2019

Santiago de Compostela - Spain

^{*}Speaker.

[†]On behalf of the CMS Collaboration

1. Introduction

In 2026 LHC will enter its High-Luminosity phase (HL-LHC) [1], delivering an instantaneous luminosity of $7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ along with 200 average simultaneous interactions per bunch crossing (pileup). The increased luminosity will enable the Compact Muon Solenoid experiment (CMS) to perform precision standard model measurements and search for rare new physics phenomena. In order to meet these goals, CMS will undergo a large program of upgrades. A new level-1 trigger (L1T) able to run sophisticated algorithms similar to the one used in the high-level trigger will be deployed [2]. This new system will enable CMS to maintain, and in some areas improve, the same trigger thresholds and physics acceptance as in the previous runs.

The HL-LHC CMS detector upgrade, called Phase-2 CMS, features a new tracking system and endcap calorimeter. Pixel sensors and doublets of silicon strips [3] are employed in the upgraded tracker. The High-Granularity Calorimeter (HGCAL) is used as endcap calorimeter in Phase-2 CMS endcap calorimeter. HGCAL is a sampling calorimeter made of silicon sensors and plastic scintillators [4]. HGCAL has longitudinal and transverse granularity in order to enable for accurate 3D positioning and track matching. A more detailed description of the Phase-2 CMS detector upgrades can be found in [5].

2. Overview of the Phase-1 and Phase-2 level-1 trigger architectures

An overview of the CMS Phase-1 L1T system used during LHC Run-2 is shown in Figure 1a. The system is made of around 100 FPGAs used to analyse and distribute around 5 Tb/s of data. The

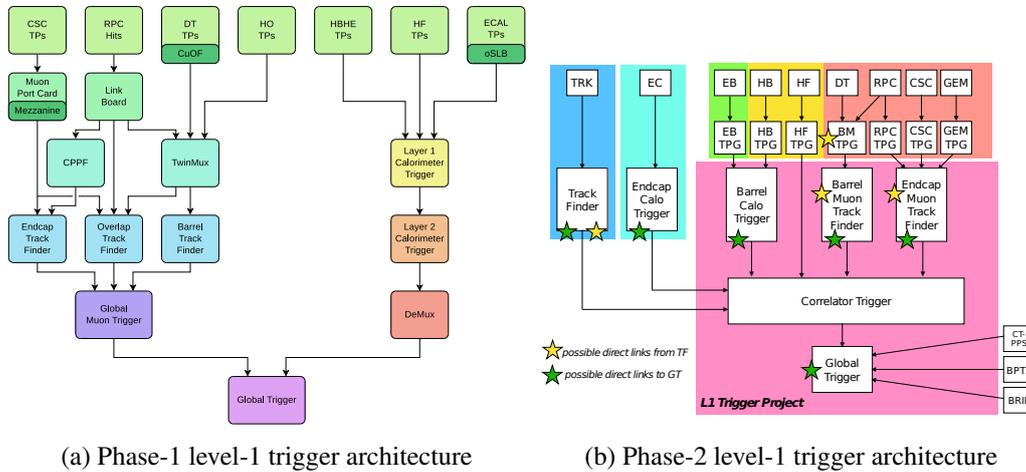


Figure 1: Overview of the Phase-1 and Phase-2 level-1 trigger architectures. In the former, inputs from the calorimeter and muon chamber are sent to the respective trigger subsystems, which are responsible for computing trigger objects. In the latter, three new systems are introduced: the endcap calorimeter trigger, responsible for processing data from HGCAL; the track finder, which reconstructs tracks in the tracker; the correlator trigger, which runs particle flow algorithms.

system reduces the detector readout rate from 40 MHz to a maximum of 100 kHz with a latency of $3.8 \mu\text{s}$. Trigger Primitives (TPs) from the calorimeter and muon systems are sent to their respective

trigger subsystems to compute trigger objects; the global trigger sends the event to the high-level trigger if at least one trigger condition is fulfilled.

For the first time, the CMS Phase-2 L1T introduces inputs from the tracker and 3D energy clusters from HGAL. An overview of the system is shown in Fig. 1b: the Phase-2 L1T has a maximum accept rate of 750 kHz, a latency of 12.5 μ s, and an input bandwidth of around 50 Tb/s, dominated by 40 Tb/s from barrel ECAL crystal TPs. As in the Phase-1 system, inputs from the subdetectors are sent to their respective trigger subsystems, which preprocess the data for the correlator trigger and compute standalone trigger objects, i.e. trigger object that are reconstructed using TPs from a single subdetector. The correlator trigger (CT) receives data from all the trigger subdetectors and runs sophisticated particle-flow algorithms [6]. For robustness, trigger subsystems directly send standalone trigger objects to the global trigger. Time multiplexing [7] is widely used in the Phase-2 trigger: processors analyse every N-th event, with N being the time multiplexing period; to enable the system to constantly process data, consecutive events are sent to consecutive boards in a round-robin fashion.

3. Hardware and firmware R&D

By limiting the flavour of electronics boards used and using the same communication links and protocol, the CMS Phase-2 L1T gains in simplicity of commissioning, monitoring, and control of the system. Generic and highly-configurable boards based on the ATCA standard are being designed and studied. By mounting large FPGAs and many 25 Gb/s optical links using Samtec Firefly Modules [8], the target boards can be used in a wide spectrum of applications. The two main R&D hardware platforms that are being pursued to equip the system are Serenity and Advanced Processor.

Serenity (Fig. 2a) is a modular flexible board mounting FPGAs on daughter cards via an interposer module. It features around 100 optical links able to transmit data up to 28 Gb/s. Interposers leave great space for FPGA choice, allowing for FPGAs from different vendors to be used on the same hardware platform. Slow control is performed via a Com-Express Type 10 [9] mezzanine. Communication with the board is done over PCIe and AXI bus using IPBus [10], a customisation of the IP protocol that enables users to send commands for control and monitoring. Prototype boards have been produced and used as a testing platform for algorithms. Eye scans have shown reliable performance up to 28 Gb/s and thermal tests have been run on both simulation and hardware, showing acceptable results.

The **Advanced Processor (APx)** boards (Fig. 2b) mount high-end Xilinx FPGAs directly on the board. Similarly to Serenity, around 100 optical links running up to 28 Gb/s are available on the device. An Embedded Linux Mezzanine, based on a Xilinx ZYNQ system-on-chip, is used for board control. Eye scans up to 28 Gb/s have been run successfully and thermal tests have been run on the device by stimulating it with high load currents and clock speeds.

VHDL and High-Level Synthesis (HLS) have been used for the development of FPGA firmware. HLS enables users to develop firmware in high-level languages and it has being used for rapid prototyping of firmware in order to deal with the increased amount of data and complexity of the algorithms. Thanks to the higher abstraction level and simplicity, HLS enables users to write adaptable firmware without any HDL background, as a consequence a larger number of people can

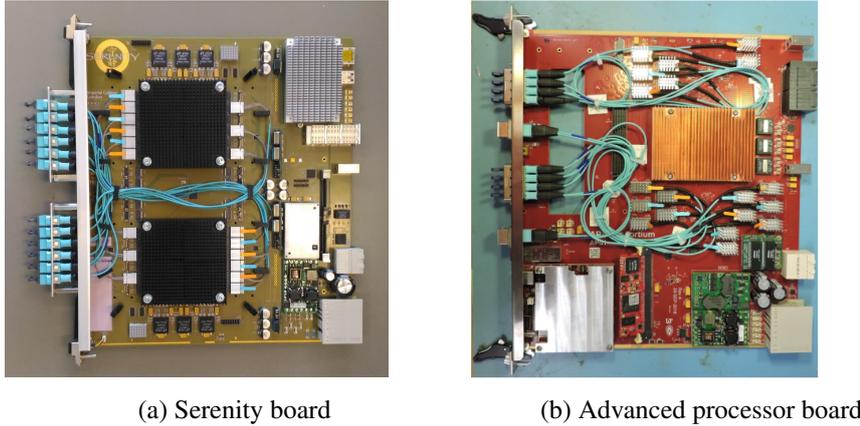


Figure 2: Pictures of the Serenity and Advanced Processor boards.

develop trigger firmware. The infrastructure firmware is separated from the algorithms and has been centrally developed in VHDL by experts, providing a uniform interface to online software. The separation has enabled users to develop their own firmware and test it on systems such as Amazon AWS FPGAs before actually integrating it with the trigger infrastructure.

4. Trigger subsystems and performance

The L1T **track finder** receives track stubs made of hit doublets from the silicon strips. The stub angle is used to estimate the transverse momentum (p_T) of the track and reduce the data rates: only 3% of particles have transverse momentum greater than 2 GeV, therefore introducing a p_T threshold greatly reduces the required bandwidth. At pileup 200, around 15000 stubs are sent to the trigger, corresponding to around 200 100-bit tracks on average. The track finder employs 162 boards and 18 crates, each crate processing a single event, with a time multiplexing period of 18 bunch crossings. The output bandwidth of the track finder is 3.6 Tb/s. The track finding algorithms under study employ a two-step filtering strategy: first, a fast stub filtering algorithm is run; second, a more accurate track finding technique is used to eliminate fake tracks and compute the final track parameters. Three FPGA-based algorithms are under consideration: a combination of a Hough transform and a Kalman filter; a seeding step finding tracklets from adjacent stub pairs and a χ^2 fit; a combination of tracklets and Kalman filter from the previous two algorithms. These algorithms have been implemented in both hardware and software and have been shown 95% efficiency in track finding over the $|\eta| < 2.4$ range.

The Phase-1 calorimeter trigger receives trigger towers from ECAL, corresponding to energy sums over groups of 5×5 PbWO₄ crystals. The upgraded **barrel calorimeter trigger** (BCT) has access to improved granularity from ECAL by receiving crystal-level information, instead of sums over groups of 5×5 PbWO₄ crystals as in the current system. Additionally the BCT receives HCAL trigger primitives carrying depth info from 7 segments. BCT computes clusters and finds standalone trigger objects. Two layers form this trigger subsystem: the first one, made of 36 boards, finds proto-clusters in calorimeters regions; the second one, made of 3 boards, stitches the proto-clusters in region boundaries, computes the final clusters, and finds standalone trigger objects. The total

output bandwidth from the BCT is 2.3 Tb/s. Half of the electromagnetic and the full hadronic section of HGICAL will be used in the **HGICAL trigger** to compute 3D clusters in a time-multiplexed system with a period of 18 bunch crossings. Clusters passing a certain p_T threshold will be sent to the upstream systems: a p_T cut of 1 GeV reduces the number of cluster to a maximum of 400. The HGICAL trigger is expected to output 1.6 Tb/s. The improved granularity enables the calorimeter trigger to have an identification efficiency for standalone electrons and photons up to 99% while halving the rate, compared to the Phase-1 system. Furthermore, it has been shown that it is possible to reduce the rate by a factor of 10 by matching electrons to tracks, at the cost of decreasing the efficiency by 10%.

The **muon trigger** features improved spatial and timing resolution and receives hits from both tracker and muon systems. More powerful algorithms, such as the Kalman filter in the barrel region and improved pattern matching in the endcap are used to improve the muon trigger resolution. The upgraded muon trigger is also capable of triggering on displaced particles.

The **correlator trigger** (CT) receives data from all the trigger subsystem mention above; its input bandwidth is 8.5 Tb/s. The CT runs sophisticated particle flow and pileup subtraction techniques [11] that are currently used only offline. A time-multiplexed, two-layered architecture is under consideration for this system: the first layer build particle flow candidates by matching calorimeter and tracking information; the second one reconstructs trigger objects from particle flow inputs. Vertexing performance of the system has been studied, showing feasibility in hardware, with at least 85% vertex reconstruction efficiency at pileup 200. Particle flow techniques have shown a tenfold reduction in the rate of missing transverse momentum triggers and a great improvement in efficiency, demonstrating the potential of this algorithm.

5. Conclusions

By receiving the full granularity data from most of the CMS Phase-2 detector, employing powerful hardware able to process great amounts of data, running more sophisticated algorithms on upgraded hardware able to efficiently reject pileup, the Phase-2 CMS L1T can run a trigger menu made of triggers with thresholds similar to Phase-1 while keeping the total accept rate under the 750 kHz limit. If the CMS L1T were not to be upgraded, the same menu would require an accept rate of 4 MHz. The upgrade enables CMS to retain its sensitivity to important standard model processes such as the Higgs physics and to extend its physics reach through the usage of triggers on more exotic objects.

References

- [1] G Apollinari et al. *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. URL: <http://cds.cern.ch/record/2116337>.
- [2] CMS Collaboration. *The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report*. Tech. rep. CERN-LHCC-2017-013. CMS-TDR-017. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283192>.

- [3] CMS Collaboration. *The Phase-2 Upgrade of the CMS Tracker*. Tech. rep. CERN-LHCC-2017-009. CMS-TDR-014. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2272264>.
- [4] CMS Collaboration. *The Phase-2 Upgrade of the CMS Endcap Calorimeter*. Tech. rep. CERN-LHCC-2017-023. CMS-TDR-019. Technical Design Report of the endcap calorimeter for the Phase-2 upgrade of the CMS experiment, in view of the HL-LHC run. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2293646>.
- [5] J Butler et al. *Technical Proposal for the Phase-II Upgrade of the CMS Detector*. Tech. rep. CERN-LHCC-2015-010. LHCC-P-008. Geneva: CERN, 2015. URL: <https://cds.cern.ch/record/2020886>.
- [6] CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector. Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.CMS-PRF-14-001. CMS-PRF-14-001-004. 10 (2017), P10003. 82 p. DOI: 10.1088/1748-0221/12/10/P10003. URL: <https://cds.cern.ch/record/2270046>.
- [7] G. Hall. “A time-multiplexed track-trigger for the CMS HL-LHC upgrade”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 824 (2016). Frontier Detectors for Frontier Physics: Proceedings of the 13th Pisa Meeting on Advanced Detectors, pp. 292–295. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2015.09.075>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900215011304>.
- [8] *Samtec Firefly web page*. URL: <https://www.samtec.com/optics/optical-cable/mid-board/firefly>.
- [9] *Com Express Type 10 website*. URL: <https://www.congatec.com/en/products/com-express-type-10.html>.
- [10] Thomas Stephen Williams. *IPbus A flexible Ethernet-based control system for xTCA hardware*. Tech. rep. CMS-CR-2014-334. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/2020872>.
- [11] Daniele Bertolini et al. “Pileup Per Particle Identification”. In: *JHEP* 10.arXiv:1407.6013 (2014), 59. 23 p. DOI: 10.1007/JHEP10(2014)059. URL: <https://cds.cern.ch/record/1745357>.