

Article

L1 Triggering on High-Granularity Information at the HL-LHC

Louis Portalès [†] 

Laboratoire Leprince-Ringuet, CNRS/IN2P3, Ecole Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France; louis.portales@cern.ch

[†] On behalf of the CMS Collaboration.

Abstract: The CMS collaboration is building a high-granularity calorimeter (HGCAL) for the endcap regions as part of its planned upgrade for the High-Luminosity LHC. The calorimetric data will form part of the Level-1 trigger (hardware) of the CMS experiment, reducing the event rate from the nominal 40 MHz to 750 kHz with a decision time (latency) of 12.5 microseconds. In addition to basic tracking information, which will also be available in the Level-1 trigger system, the use of particle-flow techniques will be facilitated as part of the trigger system. Around 1-million “trigger channels” are read at 40 MHz from the HGCAL, presenting a significant challenge in terms of data manipulation and processing for the trigger system: the trigger data volumes will be an order of magnitude above those currently handled at CMS. In addition, the high luminosity will result in an average of 140 (or more) interactions per bunch crossing that produce a huge background rate in the forward region and these will need to be efficiently rejected by the trigger algorithms. Furthermore, the reconstruction of particle clusters used for particle flow in high hit-rate events presents a complex computational problem associated with the trigger. We present the status of the trigger architecture and design, as well as the algorithmic concepts needed in order to tackle these major issues.

Keywords: CMS; phase-2; upgrade; trigger; HGCAL



Citation: Portalès, L. L1 Triggering on High-Granularity Information at the HL-LHC. *Instruments* **2022**, *6*, 71. <https://doi.org/10.3390/instruments6040071>

Academic Editors: Fabrizio Salvatore, Alessandro Cerri, Antonella De Santo and Iacopo Vivarelli

Received: 21 September 2022

Accepted: 24 October 2022

Published: 31 October 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. The High-Granularity Calorimeters of CMS at the High-Luminosity LHC

The LHC is currently starting its third collision campaign, Run 3, providing the CMS experiment with [1] proton–proton (pp) collisions at a centre-of-mass energy of $\sqrt{s} = 13.6$ TeV. After the completion of Run 3, in 2025, a three-year shutdown of the accelerator is foreseen, to proceed with its High-Luminosity upgrade (HL-LHC), that should provide, by 2029, pp collisions with an instantaneous luminosity of $L = (5 - 7.5) \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, about 4-times larger than the luminosity delivered in previous runs. This upgrade creates significant technical challenges concerning future experiments, that will need to be upgraded (phase-2 upgrade) in order to handle the extreme radiation levels, detector occupancies and collision rates associated with the increased luminosity. In particular, the end-cap calorimeters currently installed in the CMS detector will be replaced. The choice was made to install new highly granular sampling calorimeters, HGCAL [2], used both as an electromagnetic (CE-E), and hadronic (CE-H) calorimeter. The HGCAL layout is illustrated in Figure 1. The CE-E is composed of hexagonal silicon (Si) modules in high-radiation regions, close to the beam pipes and interaction point. Two types of Si modules are employed. In the more forward region, thin (120 μm) high-density modules composed of 432 sensor channels are used, while thicker (200–300 μm) low-density modules with 192 sensors are used further away from the beam. The CE-H uses both Si modules in the high-radiation region, and scintillating tiles with SiPM readout in the lower radiation regions. The HGCAL is arranged in 47 layers, and provides a total of about 6-million individual channels, that will allow precise measurement of particle positions, energy, and timing. Another important aspect of the phase-2 CMS upgrade concerns its Level-1 (L1) hardware trigger electronics, that will be designed to, at least, maintain the trigger efficiencies and thresholds achieved during Run 2 and Run 3. The HGCAL will contribute to the upgraded L1 system through the generation of trigger

primitives (TP): trigger clusters providing information on the shower profile of physics objects, and coarser trigger towers primarily used to reconstruct hadronic objects. In the following, the architecture of the HGCAL TP generation (TPG) system is presented, and its impact is illustrated through the description of some of the L1 trigger algorithms benefiting from it.




Figure 1. Schematic view of one HGCAL endcap slice (right), a representative hadronic calorimeter layer (top left), and electromagnetic calorimeter layer (bottom left).

2. The HGCAL Backend Electronics and Trigger Primitive Generation

The complete HGCAL electronics chain used for the TP generation (TPG) is schematically shown in Figure 2. The calorimeter information is first processed by on-detector electronics (FE). Each low-density (high-density) module is equipped with either three (six) readout chips, HGCROC, that measure individual cell charges at 40 MHz, and group four (nine) of these cells into trigger cells (TCs), such that each module outputs a total of 48 TCs. These TCs are then processed by concentrator chips, the ECON-T, that apply a first pre-selection of the TCs, using one of the algorithms currently considered: the Threshold algorithm, used by default, that removes any TC below a given energy threshold; the Best-Choice algorithm that selects a fixed number of the most energetic TCs, or the Super-Trigger-Cell algorithm that groups the TCs in larger areas. It also compresses the TC data before forwarding it to the backend electronics through low-power GigaBit Transceiver (lpGBT) links at a rate of about 10 Gbps.

The off-detector backend electronics for TPG comprises two main blocks, referred to as Stage 1 and Stage 2, both implemented on dedicated Serenity ATCA boards equipped with Xilinx VU13P FPGAs, and supporting up to 120 input and output links. A total of 42 boards (14 per 120° sector) per endcap are considered for Stage 1. These will be tasked with unpacking raw data received from the ECON-T, building partial Trigger Towers (TT), sorting and truncating lists of TCs, and time-multiplexing the data before being sent to Stage 2. The Stage 2 logic is implemented on 18 boards for each of the 120° sectors, and is used to build 3D clusters and full trigger towers from the objects provided by Stage 1 and to properly format these objects before forwarding them to the L1 trigger. More details on the Stage 1 and Stage 2 algorithm blocks are given in Sections 2.1 and 2.2, respectively.




Figure 2. Schematic representation of the trigger primitive generation front-end and back-end hardware chain.

2.1. Stage 1 Algorithms

Figures 2 and 3 references the different algorithm blocks implemented in each Stage 1 FPGA. These blocks can be grouped into four primary components:

- **Data aligner and unpacker:** uniformizes the variable-size data sent from the front-end electronics and groups it per bunch crossing (BX);
- **Sum tower energies:** builds CE-E and CE-H partial TT (pTT) from module energies;
- **TC sorting and truncation:** groups TCs into 3D regions; 42×2 bins in the $(r/z, \phi)$ plane per 120° sector, and returns a sorted list the most energetic TCs in each bin;
- **Data formatter and unpacker:** Reformats the TCs and pTT data and time-multiplexes it before forwarding it to Stage 2.




Figure 3. Schematic view of the current HGCAL TPG Stage 1 algorithmic blocks.

Particular care has been taken in the development of the sorting and truncation algorithm, whose design is optimized for its implementation in FPGA. The choice is made to use Kenneth Batcher’s odd–even mergesort algorithm [3]. The algorithm’s logic is depicted in Figure 4. The input data are handled in a highly parallelizable manner through a network of small sorter and merger blocks. The sorting network structure depends only on the number of inputs, which will be fixed for each bin, allowing for hard-coded implementation in firmware. The number of output TC per bin is also a fixed parameter, allowing for the truncation to be performed on-the-fly in a resource-efficient manner, by removing all comparators in the network pointing to truncated output nodes only.

The Stage 1 algorithm implementation has gone through several optimizations. In particular, the mapping between the front-end electronics and the Stage 1 FPGA is defined in a way that balances the load per FPGA, each processing an equivalent number of TC per bunch crossing. As the available bandwidth between Stage 1 and Stage 2 limits the number of TC per output link to 420 with the current VU13P FPGA, the number of output TC in each bin is also optimized. A default truncation profile is defined as a function of r/z , in a so-called pileup-driven way, which aims at truncating as little information as possible, achieving more than 99% TC selection efficiency. Alternatively, two signal-driven approaches are considered: a flat truncation profile, conserving an equal number of TC in each bin, and corresponding to an optimal distribution of the available hardware resources, and an area-weighted truncation, where more TC are removed in the smaller low- r/z bins. The impact of the truncation profile on e/γ clustering performance is illustrated in Figure 5. It was found that the profile choice does not impact the cluster reconstruction efficiency. However, the stricter signal-driven profiles were found to yield better energy resolution, due to the removal of more information from the pileup.




Figure 4. Example of an odd–even mergesort network taking as input the energies of 16 TC and returning the 5 most energetic. The blue vertical links represent comparators in use, where the higher-energy TC is moved up. The dashed red vertical lines represent comparators that are removed to reduce the output size.




Figure 5. Impact of the truncation threshold definition on e/γ clustering performance. (a) Reconstruction efficiency as a function of $|\eta|$. The $|\eta|$ range is reduced with respect to the HGCal acceptance to remove possible border effects. (b) Momentum resolution as a function of the reconstructed cluster’s p_T .

2.2. Stage 2 Algorithms

The data processed in the Stage 1 blocks are sent to the Stage 2 FPGA, whose algorithmic blocks are illustrated in Figure 6. Similarly to Stage 1, four distinct blocks are considered for Stage 2, including a data unpacker and packer, responsible for formatting the information received from Stage 1 and sending to the L1 trigger, respectively. The unpacked data follow one of two paths depending on their nature. The partial towers follow the Tower energy path, in which they are summed into trigger tower maps of 20×24 bins in the (η, ϕ) plane, while the sorted and truncated TC collections follow the clustering path.




Figure 6. Schematic view of the current HGCal TPG Stage 2 algorithmic blocks.

The clustering path is composed of several distinct steps. The first one, referred to as Histogramming, consists of sorting the TC into 42×108 fine bins in the $(r/z, \phi)$ plane, in three partially overlapping 180° sectors in each endcap. The resulting histogram is then smeared both in the r/z and ϕ directions, reducing the impact of statistical fluctuations. Local energy maxima are then identified, and those above a predefined energy threshold are used to define cluster seeds. The clusters are built by aggregating all the TC within a layer-dependent distance ΔR^2 of each seed, defined in the $(x/z, y/z)$ plane. Finally, the cluster properties required by the L1 trigger algorithms, regarding both their kinematics and quality, are evaluated and optimally encoded, considering the bandwidth limitation before being forwarded to the L1 trigger.

The TC clustering efficiency depends on several parameters of the algorithms, for instance, the choice of energy threshold for cluster seed identification. The impact of this seeding threshold on cluster reconstruction is illustrated in Figure 7a, while Figure 7b shows its impact on trigger rate. The default threshold definition considered includes an area-based reweighting of the seed energies paired with a low-energy threshold of $8.5 \text{ mip} \times \sin\theta_{\text{TC}}$ (referred to herein as mipT), where 1 mip represents the energy deposit expected from a minimum ionizing particle. This threshold is defined so that the energy required to define a seed is larger in regions where the bins cover a larger portion of the detector. This allows the obtainment of a constant pileup density as a function of $|\eta|$, but yields significant cluster reconstruction efficiency loss at low $|\eta|$.

By instead using a tighter, but constant, energy threshold, the efficiency loss is reduced, reaching close to perfect efficiency for values below 20 mipT, with no significant increases in the trigger rate.




Figure 7. Impact of the seeding energy threshold definition on the e/γ cluster performance. The black lines show the performance achieved with an $|\eta|$ -dependent area-weighted threshold, while the dashed-colored lines illustrate the performance obtained with constant thresholds. (a) Cluster reconstruction efficiency as a function of $|\eta|$. The $|\eta|$ range is reduced with respect to the HGAL acceptance to remove possible border effects (b) trigger rate as a function of the offline threshold on the reconstructed object's p_T .

3. HGAL at Level-1 Trigger

The HGAL TPG comprises only a part of the larger-scale upgrade of the L1 trigger system for HL-LHC [4], depicted in Figure 8, that will benefit from upgraded hardware. In turn, far more room will be available compared with Run 2 and 3 to implement complex L1 algorithms, using topological information from the events and advanced machine-learning algorithms, adapted for direct implementation into the L1 trigger firmware, but also from improved granularity for the barrel calorimeters primitives, a muon trigger whose range will be extended up to $|\eta| = 2.8$, and for the first time access to tracking information, up to

$|\eta| = 2.4$. The trigger clusters and towers that will be provided will therefore be used in several ways, either to build standalone objects: electrons, photons, hadronically decaying τ leptons, calorimeter jets and energy sums (H_T), up to $|\eta| = 3$ in the Global Calorimeter Trigger, or in combination with information from the muon and track trigger blocks as inputs to more advanced object reconstruction algorithms, such as Particle-Flow [5] or PUPPI [6] that have thus far been limited to offline reconstruction, in the Correlator Trigger. Illustrative examples are given in the following sections.




Figure 8. Phase-2 L1 trigger blocks. The blocks receiving HGCAL trigger primitives as input are highlighted.

3.1. Particle-Flow and PUPPI

Particle-Flow is the reconstruction algorithm used by CMS both offline and at HLT since Run 1, which combines information from all subdetectors to reconstruct physics objects. Its offline implementation is not applicable at L1 due to the extreme 40 MHz rate of collision events; however, a re-designed and simplified version will be implemented.

Similarly, PUPPI (PileUp Per Particle Identification), is the main pileup mitigation algorithm used offline in CMS. It employs vertexing information to remove contributions from charged pileup objects, and event-wide information to weight-down contributions from neutral pileup objects. This event-wide information is difficult to use at L1, again necessitating an adapted version of the algorithm that relies only on local information.

These algorithms will be used within the track trigger and HGCAL acceptance, up to $|\eta| = 2.4$, and will allow researchers to reconstruct clean, pileup-ridden trigger objects, in many cases with significantly lowered trigger thresholds compared to more standard approaches based on calorimeter information alone, as illustrated in Figure 9a, where one can see that the PUPPI H_T L1 seed allows reduction of the L1 threshold by more than 300 GeV, and provides a sharper turn-on for the same L1 rate of 20 kHz, compared with the similar calorimeter-only seed. It is however interesting to note that some standalone calorimeter objects can be used with competitive threshold, for instance, shown in Figure 9b, the efficiencies for forward calorimeter jets are comparable to that of PUPPI jets in the same $|\eta|$ region, owing to the excellent performance of HGCAL.




Figure 9. Performance comparison between standalone calorimeter trigger hadronic objects and objects built using the PFlow and PUPPI algorithms: (a) Trigger efficiency at fixed L1 rate as a function of generator-level H_T . The corresponding L1 H_T thresholds are referenced in the legend; (b) Trigger efficiency at fixed L1 rate as a function of the generated jet’s p_T . The corresponding L1 p_T thresholds are referenced in the legend.

3.2. Track-Matched e/γ Objects

As shown in Figure 10a, HGCALE will allow the obtainment of comparable trigger efficiency for e/γ objects in the forward region to that expected in the barrel region. In addition, it will be possible to match the calorimeter clusters to tracks, allowing to trade some efficiency off against more manageable rates as shown in Figure 10b.




Figure 10. Performance comparison between standalone and track-matched e/γ trigger objects: (a) Trigger efficiency as a function of the generator-level $|\eta|$; (b) Trigger rate as a function of the threshold on the L1 objects’ p_T .

3.3. Calorimeter τ

The standalone forward calorimeter τ reconstruction presents a good example of the advanced machine-learning-based algorithms that will be implemented at L1. As input, this algorithm receives 3D calorimeter clusters provided by the HGCal TPG blocks, passes them through dedicated Boosted Decision Trees to suppress pileup clusters using cluster shape variables, and then calibrates the clusters, subtracting the remaining contribution from pileup, and applying energy corrections. An additional processing step is considered, that allows the identification of the τ lepton decay modes, using a Random Forest classifier. As illustrated in Figure 11, this algorithm shows promising results, and should efficiently complement the PUPPI Tau algorithm for high p_T objects.




Figure 11. Standalone hadronic τ performance: (a) trigger efficiency as a function of the generator-level τ lepton p_T for 1-prong (red) and 3-prong τ leptons; (b) L1 trigger rates as a function of the τ lepton L1 E_T , for inclusive (black), 1-prong (red), 1-prong + π^0 (blue) and 3-prong (yellow) τ lepton decays.

4. Conclusions

The HGCal will contribute to the CMS L1 trigger upgrade for HL-LHC by providing trigger primitives: calorimeter clusters and towers. Several improvements were recently applied to the trigger primitive generation, specifically regarding the generation of trigger clusters. A new algorithm was introduced for the sorting and truncation of TC sent from the front-end electronics, well suited for the hardware architecture under consideration. Moreover, the TC truncation parameters were tuned to yield improved energy resolution for the trigger objects. The clustering algorithm that follows was also optimized, allowing for a significant improvement in reconstruction efficiency for e/γ clusters.

The HGCal trigger primitives will be widely influential on the L1 trigger, both for electromagnetic and hadronic objects. In conjunction with tracking information, they are anticipated to extend the trigger capability to high pseudorapidities with excellent performance. In addition, they will allow the construction of robust standalone calorimeter objects, that are expected to efficiently complement the PFlow and PUPPI objects at high p_T .

The firmware implementation and hardware tests are still in progress for both the HGCal TPG and L1T electronics, with a large portion of the algorithms already being implemented and tested. Further optimizations and new algorithms benefitting from the improved hardware capabilities are still being considered, and could be implemented before the commencement of HL-LHC data collection.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The author declares no conflict of interest.

References

1. CMS Collaboration. The CMS Experiment at the CERN LHC. *J. Instrum.* **2008**, *3*, S08004.
2. CMS Collaboration. *The Phase-2 Upgrade of the CMS Endcap Calorimeter*; Technical Design Report CMS-TDR-019; CERN: Geneva, Switzerland, 2017.
3. Batcher, K.E. Sorting Networks and their Applications. In Proceedings of the AFIPS Spring Joint Computer Conference, Austin, TX, USA, 30 April–2 May 1968; Volume 32, pp. 307–314.
4. CMS Collaboration. *The Phase-2 Upgrade of the CMS Level-1 Trigger*; Technical Design Report CMS-TDR-021; CERN: Geneva, Switzerland, 2020.
5. CMS Collaboration. Particle-flow reconstruction and global event description with the CMS detector. *J. Instrum.* **2017**, *12*, P10003. [[CrossRef](#)]
6. CMS Collaboration. Pileup mitigation at CMS in 13 TeV data. *J. Instrum.* **2020**, *15*, P09018. [[CrossRef](#)]