The Compact Muon Solenoid Experiment

Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

01 May 2020 (v3, 06 August 2021)

The Phase-2 Upgrade of the Hardware Trigger of CMS at the LHC

Manfred Jeitler for the CMS Collaboration

Abstract

The CMS experiment uses a two-level triggering system consisting of the Level-1, instrumented by custom-design hardware boards, and the High Level Trigger, a streamlined version of the offline reconstruction software running on a computer farm. The upgrade of the collider to the 'High-Luminosity LHC' that will deliver a luminosity of $5 - 7 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to 140–200 pile-up events, also necessitates a substantial upgrade of the trigger system to make optimal use of the data. An important difference from the present system will be that after the upgrade, information from the silicon strip tracker will be available already for the Level-1 Trigger. This will allow CMS to use so-called 'particle flow' objects, i.e. signals seen not only in one subdetector but put together from all available subdetectors, resulting in much sharper cuts on trigger objects. Also, trigger rates will rise both at Level-1 (from 100 kHz to 750 kHz) and at the High-Level Trigger. At the same time, more sophisticated algorithms will be available at Level-1. To make this possible, the latency - the processing time available for arriving at the Level-1 trigger decision - will increase significantly. Machine-learning techniques such as Boosted Decision Trees have already started to be implemented in the trigger electronics and will occupy a more important place in the future. The use of High-Level Synthesis (HLS) tools will allow physicists to formulate trigger requirements in a language closer to that of data analysis. To avoid missing unexpected signatures from New Physics, studies are underway to employ anomaly detection using autoencoders.

Presented at *INSTR20 Instrumentation for Colliding Beam Physics*,
Published in *JINST (Journal of Instrumentation)*, Volume 15, September 2020, doi =
10.1088/1748-0221/15/09/c09009, url = <https://doi.org/10.1088/1748-0221/15/09/c09009>

The Phase-2 Upgrade of the Hardware Trigger of CMS at the LHC

M. Jeitler^a on behalf of the CMS collaboration

*^aInstitute of High Energy Physics of the Austrian Academy of Sciences,
Vienna, Austria*

E-mail: manfred.jeitler@cern.ch

ABSTRACT: The CMS experiment uses a two-level triggering system consisting of the Level-1, instrumented by custom-design hardware boards, and the High Level Trigger, a streamlined version of the offline reconstruction software running on a computer farm. The upgrade of the collider to the “High-Luminosity LHC”, which will deliver a luminosity of $5-7.5 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$, corresponding to 140–200 pileup events, also necessitates a complete upgrade of the trigger system to make optimal use of the data. An important difference from the present system will be the fact that after the upgrade, information from the silicon strip tracker will be available already for the Level-1 Trigger. This will allow CMS to use so-called “particle flow” objects, i.e. signals seen not only in one subdetector but put together from all available subdetectors, resulting in much sharper efficiency turn-on curves for trigger objects. Also, trigger rates will rise both at Level-1 (from 100 kHz to 750 kHz) and at the High-Level Trigger (from 1 kHz to 7.5 kHz). At the same time, more sophisticated algorithms will be available at Level-1. To make this possible, the latency - the processing time available for arriving at the Level-1 trigger decision - will increase significantly. Machine-learning techniques such as Boosted Decision Trees have already started to be implemented in the trigger electronics and will occupy a more important place in the future. The use of High-Level Synthesis (HLS) tools will allow physicists to formulate trigger requirements in a language closer to that of data analysis. To avoid missing unexpected signatures from New Physics, studies are underway to employ anomaly detection using autoencoders.

KEYWORDS: Trigger, LHC, HL-LHC, upgrade

Contents

1	The physics program and data-taking challenges of the “High-Luminosity LHC” (HL-LHC)	1
2	The existing CMS Level-1 trigger system and upgrade plans	3
3	The upgrade strategy	3
3.1	Silicon strip tracker	4
3.2	Calorimeters	6
3.3	Muon system	6
3.4	Long-lived particles and displaced vertices	7
3.5	Anomaly trigger	9
3.6	Scouting	10
4	Conclusions	10

1 The physics program and data-taking challenges of the “High-Luminosity LHC” (HL-LHC)

The physics program of the “High-Luminosity LHC” centers around two main topics:

- Precision measurements of standard model parameters, including but not restricted to the parameters of the Higgs boson
- Searches for physics beyond the Standard Model of particle physics (“BSM physics” or “New physics”)

Both of these targets require a much higher integrated luminosity than what could be achieved with the present systems. For CMS to arrive at a total integrated luminosity of up to 4000 fb^{-1} over the lifetime of the HL-LHC, present plans aim at increasing the instantaneous luminosity from the original LHC design value of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ to up to $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. At a bunch-crossing rate of 40 MHz, this means that about 200 proton-proton interactions will happen at each bunch crossing (in other words, the “pileup” will be 200; see Fig. 1). While for some types of Level-1 triggers the number of triggered events is proportional to the pileup, for other types the rate increase at higher pileup may be strongly non-linear (see Fig. 2), which poses a serious challenge to the detector readout and data acquisition systems.

If the experiments were to keep their present triggering systems when the LHC collider reaches this performance, they would have to significantly increase their triggering thresholds. This would result in a large loss in sensitivity in the low-mass range.

Figure 1: An event from a test run showing 86 reconstructed vertices (shown as orange dots). Pileup at the HL-LHC may be still higher by more than a factor of 2.

Figure 2: While for certain single- and double-object Level-1 trigger types rates increase roughly proportional to pileup (left), for some other triggers (in particular, sums of energies or missing transverse momentum, or multi-object triggers) rates may rise in a strongly non-linear way (right). The plots show Run-2 conditions, when the instantaneous luminosity was about $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

Moreover, experiments will have to look for novel event signatures. No convincing indications of New Physics have been found with the present systems and this suggests that relying exclusively on conventional signatures may not lead to discoveries in the future either. However, physics processes not observed so far might require looking for them in a different way. One important example of this is the fact that so far most searches assumed that new particles created at the proton interaction vertex would decay almost immediately into standard-model particles or invisible particles (such as the “Lightest Supersymmetric Particle” or LSP). However, there are now numerous theories according to which metastable particles might be created that could give rise to displaced decay vertices (dubbed “long-lived particles”; cf. [1]). Trigger and analysis strategies will have to allow for the fact that the interaction point and the decay vertex need not necessarily coincide as otherwise such kinds of events could be lost.

2 The existing CMS Level-1 trigger system and upgrade plans

The CMS experiment uses a two-level triggering approach, with a hardware-based Level-1 trigger (L1) that uses only data from the muon and calorimeter systems, and a High-Level trigger (HLT) implemented in a computer farm [2]. Since the startup of the LHC, the CMS Level-1 trigger rate was limited to 100 kHz while the High-Level trigger further reduced the data amount to a final data taking rate of about 1 kHz. Certain improvements in the Level-1 trigger were implemented after Run-1 (2011-2012) of the LHC (the so-called “Phase-1 trigger upgrade” of CMS [3] carried out during and after “Long Shutdown 1” of the LHC) but the mentioned trigger rate limits were not changed.

After Run-3 of the LHC (currently scheduled for the period 2021-2024) major upgrades in the accelerator and the detectors will take place during “Long Shutdown 3” [4]. At this time, also the Level-1 trigger systems will be completely replaced in the framework of the “Phase-2 trigger upgrade”. Advances in detector and trigger electronics and in computing systems will make it possible to increase the L1 rate from 100 kHz to 750 kHz while the HLT output rate is planned to grow by the same factor from 1 to 7.5 kHz. Retrieving the full detector information for a much higher percentage of bunch collisions will make it possible to conduct targeted searches for very specific event signatures.

Apart from the output rates, an important parameter of the L1 trigger is the latency, the time available for taking the trigger decision and during which the full data must be stored in the detector electronics. Larger latency entails costs for the on-detector storage in terms of financing, power and cooling but at the same time makes it possible to carry out more complicated calculations and retrieve more detector data as input. For the Phase-2 upgrade, the latency will be increased from 3.8 to 12.5 μ s. This will be particularly important in order for the L1 trigger to include the information from the strip tracker and high-granularity calorimeters within the framework of a particle-flow approach as described below.

3 The upgrade strategy

The most fundamental change in the triggering approach will be the inclusion of the silicon strip tracker into the Level-1 trigger, which will require the replacement of the tracker detector. Using

the tracker information together with the data from the calorimeters and the muon system will make it possible to select events with much higher efficiency by using a “particle-flow” approach already at the first trigger level.

Apart from including the strip tracker into the trigger, the readout granularity of trigger objects will be improved for calorimeter objects. The muon system will be extended by additional detectors, which will also be included in the trigger. The architecture of the upgrade trigger is shown in Fig. 3 and more detailed information can be found in the Technical Design Report of the Phase-2 Upgrade of the CMS Level-1 Trigger [5].

Figure 3: Architecture of the CMS L1 trigger after the Phase-2 upgrade (from [5]). Data from the calorimeter, the muon and the silicon strip tracker subsystems will be first reconstructed inside the corresponding detector subsystem and then combined in the particle flow electronics, the so-called “Correlator Trigger”, into trigger objects corresponding to physical particles. Logic combinations between different trigger objects are then made in the Global Trigger. In order to optimize efficiency for events where the data from one detector subsystem are sufficient to reconstruct a particle, direct connections between the detector subsystems and the Global Trigger that bypass the Correlator Trigger are also foreseen. This will also make it easier to cope with potential initial problems in carrying out the full particle flow reconstruction in the L1 trigger. (Figure from [5]).

3.1 Silicon strip tracker

Due to the large data volume of the silicon strip tracker, it would not be possible to read out all the information at a rate of 40 MHz and send it to the trigger electronics. Therefore, the on-detector electronics will perform a preselection. The on-detector electronics generates fragments of tracks (so-called “stubs”) computed from the data from closely-spaced silicon sensor modules that make up a tracker layer. The bending between the two sides of a module is used to discard track fragments below a certain minimum transverse momentum and forward only those above this threshold to the

track finder electronics (see Fig. 4). Fig. 5 shows the simulated distribution of the numbers of tracks when applying a threshold of $p_T > 2$ GeV.

Track candidates for the trigger (L1 tracks) will be reconstructed up to $|\eta|=2.4$ in pseudorapidity. In the track finder, trigger primitives will be reconstructed by a hybrid algorithm performing a Kalman filter fit on associated stubs that together form a larger track fragment (track “seed” or “tracklet”).

Figure 4: In order to reduce the data sent to the trigger electronics to a manageable volume, only track fragments (“stubs”) with a curvature corresponding to a minimum transverse momentum of 2 GeV are read out from the detector and forwarded to the trigger. (Figure from [4]).

Figure 5: Most tracks seen by the silicon tracker have very low transverse momentum (p_T). The plot shows the maximum number of prompt tracks with $p_T > 2$ GeV found in any single of the nine ϕ sectors of the L1 track finder for simulated $t\bar{t}$ events overlaid with an average pileup of 200. (Figure from [5]).

The subsequent combination with information from the full CMS detector in the particle flow layer will make it possible to improve the efficiency, lower trigger rates and fight the adverse effects of pileup [7]. Fig. 6 illustrates this by showing the L1 trigger turn-on curve for $t\bar{t}$ events (trigger efficiency vs. simulated true total transverse momentum H_T).

Figure 6: Monte-Carlo simulation results for the performance for triggering $t\bar{t}$ events (L1 trigger turn-on curve) on total transverse momentum (H_T) with the calorimeter only, just using the tracker information, or using the combined information from all subdetectors entering the particle-flow calculation (“PuppiHT”). (Figure from [5]).

3.2 Calorimeters

The endcap calorimeters of the CMS detectors suffer from radiation damage and will be completely replaced by a radiation-hard silicon and scintillator based “High Granularity Calorimeter” (HGCAL) with over 6 million readout channels in Long Shutdown 3, one quarter of which is shown schematically in Fig. 7. This will also open up new possibilities for triggering. Differently from the present system, this sampling calorimeter features strong longitudinal segmentation with 14 layers in the electromagnetic section and 22 layers in the hadronic section entering the L1 trigger with individual signals, which will be helpful in disentangling pileup events. The HGCAL will provide shower separation and identification. The high 3D-granularity will make it possible to use a particle-flow reconstruction approach.

The electromagnetic (ECAL) and hadronic (HCAL) calorimeters in the barrel region will not be replaced but their read-out electronics will be upgraded to supply the trigger with fine-grain information from the individual crystals of the ECAL (differently from the present system, where blocks of 5×5 crystals are combined into one “trigger tower”) and with timing information. This will allow the calorimeter trigger to achieve better efficiency at about half the trigger rate.

3.3 Muon system

The muon system will be improved by installing additional detectors in the forward region: so-called “improved RPC” (iRPC) chambers (Resistive Plate Chambers) and gas electron multiplier (GEM) chambers will extend the coverage up to $|\eta| = 2.4$ and 2.8, respectively. Both the on-detector and off-detector electronics of the muon system will be upgraded. New track finding algorithms will be deployed already in Run 3 (see below).

Figure 7: Longitudinal cross section of the upper half of one endcap calorimeter (High-Granularity Calorimeter) for the CMS upgrade. It consists of an electromagnetic (CE-E) compartment followed by a hadronic (CE-H) compartment. The CE-H is instrumented partly with silicon sensors and partly with scintillator tiles.

3.4 Long-lived particles and displaced vertices

Since the beginning of CMS, searches for new particles mostly assumed that they would either decay promptly (so quickly and so close to the interaction point that their lifetime or displacement could not be resolved) or pass undetected through CMS and be accessible only by way of the indirect signature of missing transverse momentum. No New Physics has been found in this way and there is an increasing number of theoretical arguments and calculations indicating that new particles might be relatively long-lived so that displaced secondary decay vertices might be resolved. While a number of analyses are targeting such scenarios, the L1 trigger would have suppressed the major part of such types of events so far.

Fig. 8a illustrates how standard-model particles could appear at displaced vertices, originating from the decay of invisible or visible, relatively long-lived BSM (Beyond the Standard Model) particles but also shows the case of visible, long-lived (metastable), slowly moving new particles (HSCP, “Heavy Stable Charged Particle”) or of disappearing tracks due to relatively long-lived, charged BSM particles. Fig. 8b shows how an almost straight track with high transverse momentum coming from a displaced vertex would appear as a low-momentum track when reconstructed under the assumption that it originates at the beam interaction point.

Track-finding algorithms must therefore be modified to allow for tracks not originating at the interaction vertex. At present, the firmware for the muon track finder in the barrel is already being modified to calculate the transverse moment p_T by using a Kalman filter approach both with and without the assumption that a track originates at the interaction vertex. This new algorithm will already be employed in LHC Run 3 (scheduled to start in 2021) and will allow CMS to carry out first targeted searches for long-lived particles to achieve physics results as well as to collect valuable

Figure 8: Displaced vertices and long-lived particles in the CMS detector.

experience for the upgrade.

Charged particles moving much slower than the speed of light (HSCPs) would escape the CMS legacy trigger even when originating at the interaction point due to the wrong timing in the muon system. Fig. 9 shows the improvement from new algorithms expected for the trigger efficiency for such particles.

Figure 9: For the existing trigger, the efficiency drops strongly for particles moving slower than at two thirds of the speed of light (red curve). An additional HSCP trigger algorithm for muons can strongly enhance the triggering efficiency at lower velocities. (Figure from [5]).

3.5 Anomaly trigger

A justified argument often brought forward is that we do not know what “New Physics” looks like and might therefore be missing signals when developing trigger algorithms based on specific hypothetical scenarios. A way to counter this problem is to look not for events with a more or less specific topology but rather search for events that appear unlike everything we know. This can be achieved by using “anomaly detection” based on machine-learning techniques such as autoencoders. The underlying idea is here to reduce the data volume (“encode” the input) in a “latent space” and then attempt to reconstruct the original data by “decoding” this compressed information (see Fig. 10). The encoding and decoding algorithms can be developed by using machine learning and a distance function that shows how well the reconstructed events match the original data. In CMS, a similar approach has already been used not in the trigger but for the assessment of data quality [8]. In high-energy physics this approach is already being applied for off-line analysis but CMS is now investigating how to implement the neural networks trained by machine learning into the FPGAs (Field Programmable Gate Arrays) of the online trigger electronics, where resource and latency constraints are much more stringent than in the offline world.

Figure 10: In order to detect anomalies, data are compressed and encoded in a reduced “latent space” and subsequently again decoded and compared to the original input data. “Typical” data similar to the ones in the training sample will show little differences while data of a strongly different kind, to which the neural network was not exposed during training, will feature more significant differences.

Such a system might allow the trigger system to discard well-known, frequent standard model physics events while keeping events that look very different. The training of the autoencoder can be achieved in a purely data-driven way by using the copious (unlabeled) experimental data taken already because any kind of so far unobserved, New Physics events necessarily must be rare and will therefore have a negligible effect on the training. This is being done now to start using the system in the upcoming LHC Run 3. For HL-LHC, experimental data with the expected amount of pileup are not yet available and the training is based on Monte Carlo simulations.

So far, an autoencoder has been trained with experimental minimum-bias data [5] and its efficiency in detecting various kinds of Higgs decays has been successfully demonstrated (see

Fig. 11).

Figure 11: Trigger efficiency and rate of several types of Higgs boson decay signals from an autoencoder trigger trained with minimum-bias events collected by CMS (from [5]).

3.6 Scouting

A complementary approach to recording data within a certain bandwidth budget is to reduce not the number of events but their size. In CMS, this technique is referred to as “scouting” (in ATLAS, a similar approach is called “trigger-level analysis”). Scouting has already been performed by CMS for certain events triggered by the L1 trigger for which the full event data could not be recorded because of the limitations of data acquisition and storage.

In 2018, towards the end of LHC Run 2, first tests were made with “40-MHz scouting” using muon data, where only the data available to the L1 trigger are used and a reduced record of events not even triggered by the L1 trigger can be stored and analyzed. For Run 3, this system will be gradually extended. A typical use case for scouting could be the generic search for two-jet resonances, where a very large data volume is needed to suppress background.

4 Conclusions

In order to significantly improve the precision of Standard-Model physics measurements and to find evidence of New Physics, which has turned out to be more elusive than originally hoped for, the luminosity of the LHC collider will be significantly increased after Run 3. This creates new challenges for the LHC detectors and their trigger systems. In CMS, large parts of the detector and the trigger electronics will be replaced or extended during Long Shutdown 3. Not

only will the hardware be upgraded but also algorithms and their implementation in firmware and software will be using a completely new approach. Novel approaches are, among others, triggers allowing for displaced vertices, anomaly triggers and 40-MHz scouting. For this, changes and further developments are already planned for the upcoming Run 3 starting in 2021. Some of these might already yield physics results and at the same time provide valuable technical information to be incorporated into the detector and system upgrade for the transition to HL-LHC during Long Shutdown 3 currently scheduled for the years 2025-2027.

References

- [1] J. Alimena et al., *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, arXiv:1903.04497 [hep-ex], 2019; <https://arxiv.org/abs/1903.04497>
- [2] CMS collaboration, *The TriDAS Project Technical Design Report, Volume 1: The Trigger Systems*, CERN/LHCC 2000 - 38, CMS TDR 6.1; December 15, 2000; <http://cmsdoc.cern.ch/cms/TDR/TRIGGER-public/trigger.html>
- [3] CMS collaboration, *CMS Technical Design Report for the Level-1 Trigger Upgrade*, CERN-LHCC-2013-011, CMS-TDR-12, 2013; <https://cds.cern.ch/record/1556311>
- [4] CMS collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, <https://cds.cern.ch/record/2020886/files/LHCC-P-008>
- [5] CMS collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger: Technical Design Report*, CERN-LHCC-2020-001, CMS-TDR-20-001, 2020; to be published.
- [6] CMS collaboration, *The Phase-2 Upgrade of the CMS Trigger: Interim Report to the LHCC*, CERN-LHCC-2017-013, CMS-TDR-17-004, 2017.
- [7] D. Bertolini, P. Harris, M. Low, and N. Tran, *Pileup Per Particle Identification*, JHEP 10 (2014) 059, doi:10.1007/JHEP10(2014)059, arXiv:1407.6013.
- [8] Pol, Adrian Alan et al., *Anomaly detection using Deep Autoencoders for the assessment of the quality of the data acquired by the CMS experiment*, 23rd International Conference on Computing in High Energy and Nuclear Physics, CHEP 2018, Sofia, Bulgaria, 9 - 13 Jul 2018, pp.06008; doi: 10.1051/epjconf/201921406008.