Available on CMS information server

CMS CR -2022/102

The Compact Muon Solenoid Experiment

Conference Report

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

12 July 2022 (v2, 13 July 2022)

A 40 MHz Level-1 trigger scouting system for the CMS Phase-2 upgrade

Dinyar Rabady for the CMS Collaboration

Abstract

The CMS Phase-2 upgrade for the HL-LHC aims at preserving and expanding the current physics capability of the experiment under extreme pileup conditions. A new tracking system incorporates a track finder processor, providing tracks to the Level-1 (L1) trigger. A new high-granularity calorimeter provides fine-grained energy deposition information in the endcap region. New front-end and back-end electronics feed the L1 trigger with high-resolution information from the barrel calorimeter and the muon systems. The upgraded L1 will be based primarily on the Xilinx Ultrascale Plus series of FPGAs, capable of sophisticated feature searches with resolution often similar to the offline reconstruction. The L1 Data Scouting system (L1DS) will capture L1 intermediate data produced by the trigger processors at the beam-crossing rate of 40 MHz, and carry out online analyses based on these limited-resolution data. The L1DS will provide fast and virtually unlimited statistics for detector diagnostics, alternative luminosity measurements, and, in some cases, calibrations. It also has the potential to enable the study of otherwise inaccessible signatures, either too common to fit in the L1 trigger accept budget or with requirements that are orthogonal to “mainstream” physics. The requirements and architecture of the L1DS system are presented, as well as some of the potential physics opportunities under study. The first results from the assembly and commissioning of a demonstrator currently being installed for LHC Run-3 are also presented. The demonstrator collects data from the Global Muon Trigger, the Layer-2 Calorimeter Trigger, the Barrel Muon Track Finder, and the Global Trigger systems of the current CMS L1. This demonstrator, as a data acquisition (DAQ) system operating at the LHC bunch-crossing rate, faces many of the challenges of the Phase-2 system, albeit with scaled-down connectivity, reduced data throughput and physics capabilities, providing a testing ground for new techniques of online data reduction and processing.

Presented at *15th Pisa Meeting 2022 15th Pisa Meeting on Advanced Detectors*

A 40 MHz Level-1 trigger scouting system for the CMS Phase-2 upgrade

Rocco Ardino[a,b], Christian Deldicque[a], Marc Dobson[a], Sabrina Giorgetti[a,b], Gaia Grosso[a,b], Thomas James[a,c], Emilio Meschi[a], Dinyar Rabady[a,*], Attila Racz[a], Hannes Sakulin[a], Petr Zejdl[a], for the CMS collaboration

[a] *CERN, Esplanade des Particules 1, Meyrin, 1211, Switzerland*
[b] *Universita di Padova, Via VIII Febbraio, 2, Padova, 35122, Italy*
[c] *Imperial College, Exhibition Rd, South Kensington, London, SW7 2BX, United Kingdom*

Abstract

The CMS Phase-2 upgrade for the HL-LHC aims at preserving and expanding the current physics capability of the experiment under extreme pileup conditions. A new tracking system incorporates a track finder processor, providing tracks to the Level-1 (L1) trigger. A new high-granularity calorimeter provides fine-grained energy deposition information in the endcap region. New front-end and back-end electronics feed the L1 trigger with high-resolution information from the barrel calorimeter and the muon systems. The upgraded L1 will be based primarily on the Xilinx Ultrascale Plus series of FPGAs, capable of sophisticated feature searches with resolution often similar to the offline reconstruction. The L1 Data Scouting system (L1DS) will capture L1 intermediate data produced by the trigger processors at the beam-crossing rate of 40 MHz, and carry out online analyses based on these limited-resolution data. The L1DS will provide fast and virtually unlimited statistics for detector diagnostics, alternative luminosity measurements, and, in some cases, calibrations. It also has the potential to enable the study of otherwise inaccessible signatures, either too common to fit in the L1 trigger accept budget or with requirements that are orthogonal to “mainstream” physics. The requirements and architecture of the L1DS system are presented, as well as some of the potential physics opportunities under study. The first results from the assembly and commissioning of a demonstrator currently being installed for LHC Run-3 are also presented. The demonstrator collects data from the Global Muon Trigger, the Layer-2 Calorimeter Trigger, the Barrel Muon Track Finder, and the Global Trigger systems of the current CMS L1. This demonstrator, as a data acquisition (DAQ) system operating at the LHC bunch-crossing rate, faces many of the challenges of the Phase-2 system, albeit with scaled-down connectivity, reduced data throughput and physics capabilities, providing a testing ground for new techniques of online data reduction and processing.

Keywords: Data acquisition, Trigger, Online data processing

1. Introduction

The High-Luminosity LHC (HL-LHC) [1] is expected to increase the delivered instantaneous luminosity to $5 \cdot 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, five times the accelerator’s original design value. In its “ultimate” configuration, the HL-LHC will reach a peak instantaneous luminosity of $7.5 \cdot 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, increasing the average number of proton-proton collisions per bunch crossing (pileup) to around 200. To cope with these extreme conditions the Compact Muon Solenoid (CMS) experiment [2] will undergo a significant upgrade: a new tracking system will be equipped with a hardware track finder which will provide tracks to the Level-1 trigger; in the endcaps a high-granularity calorimeter will provide fine-grained energy deposits; front-end electronics will be upgraded for the barrel calorimeter and muon system, in order to transmit higher-resolution data to the Level-1 trigger. Given the total event size of 7.5 MB, reading out the entire detector at the 40 MHz bunch crossing rate will not be possible due to limitations in readout, storage, and analysis capabilities. For this reason, CMS will continue to use two trigger levels: a Level-1 trigger based on field-programmable gate-arrays (FPGAs), selecting events at 750 kHz, and a high-level trigger running on a farm of compute nodes performing the second level of selection in software. The upgraded Level-1 trigger [3], illustrated in the left-hand part of figure 1, will be able to execute sophisticated algorithms that were thus far not feasible to perform in the hardware trigger, such as vertex finding, particle flow with pile-up per particle identification (PUPPI), and a Kalman Filter track reconstruction, allowing the Level-1 trigger to approach offline resolutions. The CMS data taking operation is based around selecting promising events using the two-stage trigger system and reading out the entire detector at full granularity if such events are found. While this full detector granularity is required for many analyses, the study of some other physics processes can potentially benefit from an analysis of the full available dataset at the LHC bunch crossing frequency, albeit with the resolution of the Level-1 trigger.

The proposed architecture and hardware for a scouting system, which will receive data from the Level-1 trigger over spare output links and perform a quasi-online analysis on them in a heterogeneous computing farm, will be presented. Such a

*Corresponding author
Email address: dinyar@cern.ch (Dinyar Rabady)

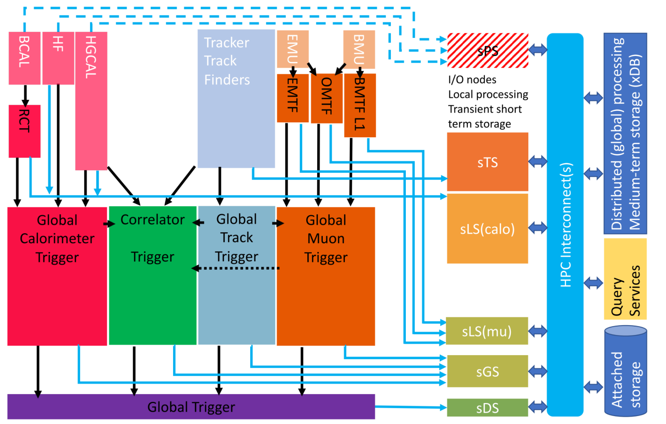

Figure 1: Architecture of the Phase-2 Level-1 trigger and scouting system. Scouting I/O nodes will receive data from spare L1 trigger outputs and propagate them to a distributed computing farm via a high performance computing interconnect. Analysis results will then be sent further for long-term storage. The system is stageable by design, already being able to provide utility in the smallest version where it receives data only from the global trigger stages (sDS and sGS).

scouting system will be operated largely independently from the standard trigger and data acquisition chain. In some cases the data gathered purely by the scouting system may be sufficient to obtain new physics results, while in other cases it may lead to hints that could be used to design a dedicated trigger algorithm to further investigate. For example, in dark sector searches, where the models predict a wide range of signatures with low signal rates, requiring superior trigger efficiency, the scouting, applying no other thresholds than the implicit ones, can be used for early identification of promising potential signals. Once a potential signal is found by the scouting system, it may be possible to develop a dedicated Level-1 trigger algorithm for the corresponding signature. In contrast, searches for the rare $W \rightarrow \pi\pi\pi$ decay could be performed exclusively within the scouting system, profiting from the capability to record all bunch crossings.

Additionally the virtually unlimited dataset afforded by the scouting system can be used to significantly improve the diagnostic capability of the Level-1 trigger, as well as for instant online luminosity measurements based on certain physics processes.

2. Architecture

The proposed scouting system will make use of spare optical outputs of Level-1 trigger boards, receiving trigger data using the same 25 Gb/s serial interconnect technology and common link protocol as utilised within the trigger itself. These outputs will be captured by dedicated FPGA boards that work as the interface between the synchronous trigger domain and the asynchronous domain of the scouting data taking, but will also perform pre-processing such as zero-suppression or recalibration of the incoming data. The first stage in the scouting data processing will be performed in so-called I/O nodes directly connected to the data taking boards, making use of distributed algorithms to extract features while data are still buffered in short-term memory. These nodes may also be equipped with GPUs or other accelerators. Once features deemed interesting are detected they, or even the associated "full" events, will be streamed over a high performance computing interconnect to a dedicated processing farm. An interesting possibility at this stage is to stream "mini events" without any preselection on to the processing farm, to make use of the high statistics in detector diagnostics or luminosity monitoring.

The L1 scouting computing farm will make use of distributed stream processing for feature reconstruction and extractions and finally put its output in a database for medium-term storage that will allow analysis by query, while only the analysis results will be sent to permanent storage.

While the scouting system can in principle take data from the full Phase-2 Level-1 trigger chain, see figure 1, it is designed to be deployed in stages. The baseline proposal is to receive data from the global trigger stages, this means the input (sGS) and output (sDS) of the Global Trigger. In further steps muon tracks and calorimeter objects (sLS) as well as tracker tracks (sTS) could be received as well. The ultimate stage would comprise the reception of the calorimeter trigger primitives themselves (sPS).

3. The scouting board for the CMS Phase-2 upgrade

The L1 scouting project has identified the DAQ800 board as a suitable hardware platform for the Phase-2 scouting system. This board [4] is being developed as a dedicated readout card for the CMS Phase-2 upgrade and boasts two powerful Xilinx VU35P FPGAs, each chip connected to 6x4 FireFly connectors that are used to provide 24x 25 Gb/s input bandwidth, and to 5 QSFP connectors that provide 5x 100 Gb/s output bandwidth. In its configuration as a readout board, the DAQ800 receives data via a custom synchronous link protocol (SlinkRocket), aggregates the data and transmits them via TCP/IP to the receiver units of the central CMS DAQ system.

For the purposes of data scouting the sender module will be reused almost without modification, while the receiving side will be replaced by a module capable of receiving data over the asynchronous serial link protocol used by the L1 trigger. As the receiving bandwidth is slightly larger than the sending one, a mild zero suppression scheme will need to be implemented before the sender module.

4. Demonstration during LHC Run-3

For Run-3 of the Large Hadron Collider (LHC) a demonstrator system will allow the evaluation of ideas and observe the system behaviour with real data. This demonstrator system receives data from the Phase-1 [5] upgraded Global Trigger (uGT), upgraded Global Muon Trigger (uGMT), Calorimeter Trigger, and the Barrel Muon Track Finder (BMTF). As shown in figure 2 the data from these systems is received in FPGA-based processing boards and subsequently transmitted to computing nodes (DSBU) that can perform "event" building and further processing.

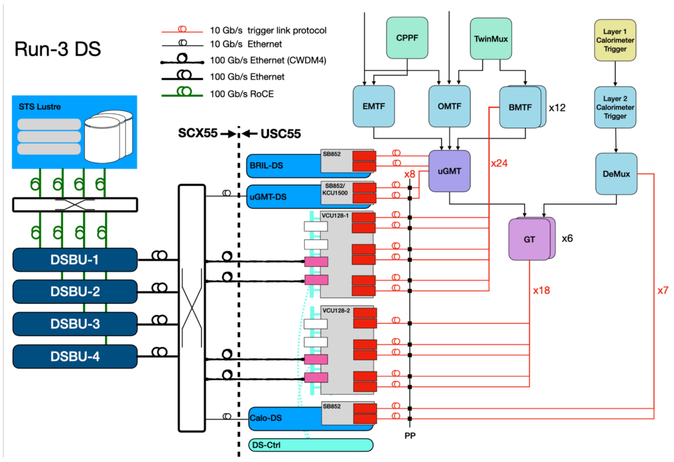

Figure 2: Architecture of the Run-3 scouting system demonstrator. Level-1 trigger data is received from the upgraded Global Trigger (uGT), upgraded Global Muon Trigger (uGMT), the calorimeter trigger (DeMux), and the Barrel Muon Track Finder (BMTF) using a mixture of I/O nodes. These data are then propagated to compute nodes (DSBUs) via Ethernet where they can be further processed and subsequently sent to long-term storage.

The Run-3 demonstrator is a heterogeneous system, composed of three classes of receiver boards: The Xilinx KCU1500, a development kit hosting the KU115 FPGA that can receive eight 10 Gb/s optical links and interfaces to a host computer via PCIe. It then utilises DMA to transfer the received data to the host machine, from where it is transferred on to the associated DSBU machine. The KCU1500 has already been used in a smaller scale demonstrator system at the end of Run-2, receiving inputs from the uGMT [6].

The second board is the Micron SB852, a PCIe card equipped with a Xilinx VU9P FPGA supporting the Micron Deep Learning Accelerator (MDLA). Similar to the KCU1500 it provides eight 10 Gbps optical input links and transfers data to the host PC via DMA.

The third board used in the Run-3 scouting demonstrator is the Xilinx VCU128, a development kit equipped with a VU37P FPGA and providing 24 input links at 25 Gb/s and 4 output channels at 100 Gbps, essentially supplying the features of one half of a DAQ800 board. Note, however, that the input links will be operated at 10 Gbps as this is the link speed used by the Phase-1 Level-1 trigger.

4.1. Results

Data from both the uGMT and the calorimeter trigger path has been collected using two KCU1500 boards since early 2021. This timespan covered periods of cosmic data taking with and without the CMS magnet switched on, as well as the LHC beam test in October 2021.

4.1.1. Independent verification of the muon impact parameter

Cosmic muons recorded without activated magnetic field afford an opportunity to verify the new impact parameter (d_{xy}) assignment by the BMTF using geometrical arguments. As the cosmic muon passes through the CMS detector, it creates two tracks from the point of view of the trigger: One track in the top half of the detector, and a second track in the bottom half of the detector roughly 25-50 ns later, and therefore recorded in the subsequent, or next but one, bunch crossing, see figure 3.

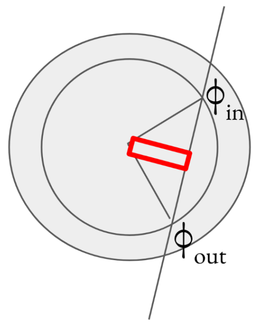

Figure 3: Illustration of the passage of a cosmic muon through the CMS detector. Highlighted in red is the impact parameter d_{xy}.

Using the fact that the magnetic field is off and therefore the track is expected to traverse the detector in a straight line, the relation $d_{xy} \propto 520 \cdot \cos((\phi_{\text{in}} - \phi_{\text{out}})/2)$ correlates the impact parameter and azimuthal coordinates of the incoming and outgoing cosmic muon leg, where the factor 520 is the distance of the CMS muon system from the center of the detector in centimetres. The good agreement between the BMTF-assigned impact parameter and the predicted value based on the azimuthal coordinates is shown in figure 4, the divergence of the model from the data is caused by the fact that the impact parameter is limited to two bits in hardware and therefore saturates at value 3.

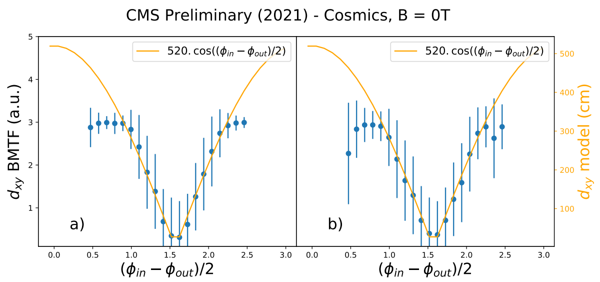

Figure 4: Correlation between the difference in azimuthal coordinates of the incoming and outgoing leg of a cosmic muon and the impact parameter (d_{xy}) assigned by the BMTF to the incoming (a) and outgoing (b) leg when the CMS magnet is switched off. The value plotted for d_{xy} is the average of the values of all entries for the given bin.

The scouting system allows such studies to be easily performed due to its ability to record all bunch crossings. In contrast, while many CMS sub-detectors currently read out a window of bunch crossings around the triggered event, the collection of such events relies on the trigger firing and not being subject to so-called trigger rules suppressing the readout. Furthermore it is unlikely it will be possible to routinely read out bunch crossing windows after the Phase-2 upgrade due to the large volumes of data involved.

4.1.2. *First beam data in preparation for Run-3*

In October 2021 the LHC performed beam tests for which both the muon detectors and the calorimeters could be switched on. Using the endcap muons received from the uGMT the beam halo could be observed. The beam halo consists of muons created by the interaction of the LHC beam with elements of the accelerator or beam gas and moves in parallel to the beamline. For figure 5, pairs of muons were selected in opposite endcaps, with opposite charge, and with $\Delta BX = 2$, corresponding to 50 ns.

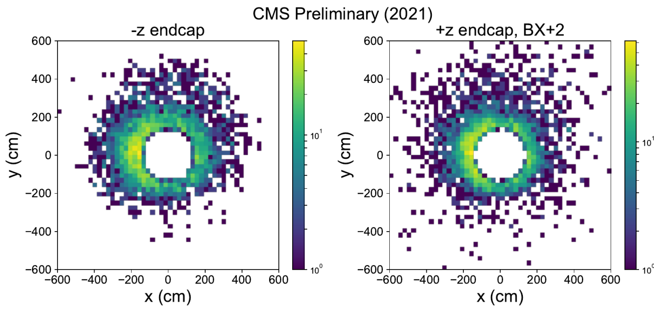

Figure 5: Muon occupancy in the negative (-z) and positive (+z) muon endcaps. Halo muons from the beam travelling from the negative to the positive side were selected by requiring a muon track in the negative side paired with a second track with opposite charge 50 ns later in the positive side.

Only bunch crossings with non-colliding bunches, sufficiently "distant" from collisions were selected. In this fill, beam 2 (travelling from the negative z side to the positive z side in CMS coordinates) exhibited a large halo. Similar to the impact parameter analysis discussed previously, these data would be challenging to assemble using the standard CMS readout system.

While the calorimeter scouting system was still in the process of being commissioned at the time of the beam test, first data could already be taken. Figure 6 shows the occupancy of Level-1 trigger reconstructed tau candidates, corresponding to what was seen in the dedicated Level-1 trigger monitoring systems.

4.2. *Future plans*

In section 4.1 we have shown the feasibility of receiving data from spare Level-1 trigger outputs and using them for later analysis. To demonstrate the architecture envisaged for Phase-2, it is planned to transfer data from the VCU128 board directly to a compute node on the surface using TCP/IP. As a first step we are currently in the process of installing this board in the service cavern of the CMS experiment to join data taking. Initially, data will be transferred via already proven DMA, in order to validate the interface to the Level-1 trigger. Once this is verified, the TCP/IP-based transmission will be demonstrated.

A promising avenue was shown to be the application of machine learning algorithms at the receiver FPGA level in order to e.g., perform fast recalibration of trigger quantities. To demonstrate this the Micron SB852 board will be used, taking advantage of the proprietary MDLA framework.

While a possible scheme for online stream processing has been demonstrated in [7], it is planned to explore the possible use of the CMS offline software framework (CMSSW [8]) to steer the processing of scouting data. To this end, the data from the scouting system will be reformatted to be suitable for ingestion into the framework. The event-based processing model of CMSSW will be reinterpreted to operate on a full orbit per event to better suit the structure of the data taken by the L1 scouting system.

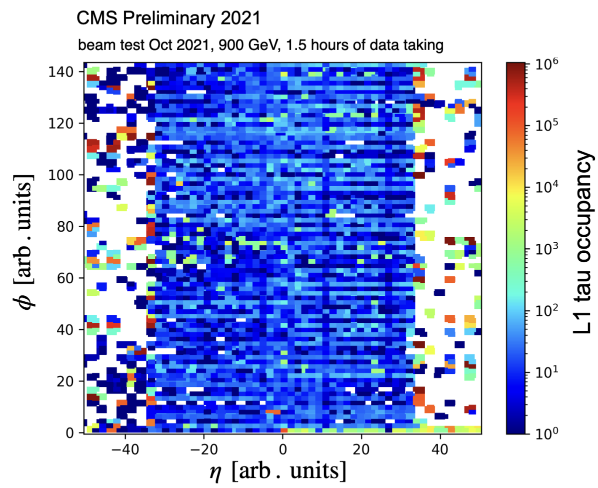

Figure 6: Occupancy of Level-1 trigger reconstructed tau leptons as recorded by the L1 scouting system.

References

[1] O. Aberle, et al., High-Luminosity Large Hadron Collider (HL-LHC): Technical design report, CERN Yellow Reports: Monographs, CERN, Geneva, 2020. doi:10.23731/CYRM-2020-0010.
URL `https://cds.cern.ch/record/2749422`

[2] CMS Collaboration, The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment, JINST 3 (2008) S08004. 361 p, also published by CERN Geneva in 2010. doi:10.1088/1748-0221/3/08/S08004.
URL `https://cds.cern.ch/record/1129810`

[3] CMS Collaboration, The Phase-2 Upgrade of the CMS Level-1 Trigger, Tech. rep., CERN, Geneva (Apr 2020).
URL `https://cds.cern.ch/record/2714892`

[4] CMS Collaboration, The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger, Tech. rep., CERN, Geneva (Mar 2021).
URL `https://cds.cern.ch/record/2759072`

[5] CMS Collaboration, Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV, JINST 15 (2020) P10017. 63 p. arXiv:2006.10165, doi:10.1088/1748-0221/15/10/P10017.
URL `https://cds.cern.ch/record/2721198`

[6] G. Badaro, et al., 40 MHz Level-1 Trigger Scouting for CMS, EPJ Web Conf. 245 (2020) 01032. 7 p. doi:10.1051/epjconf/202024501032.
URL `https://cds.cern.ch/record/2754090`

[7] M. Migliorini, et al., A horizontally scalable online processing system for trigger-less data acquisition, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1036 (2022) 166869. doi:https://doi.org/10.1016/j.nima.2022.166869.

[8] CMS Collaboration, CMS Offline Software.
URL `http://cms-sw.github.io/`