



The Compact Muon Solenoid Experiment
Conference Report

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



12 October 2021 (v3, 02 March 2022)

TEPX as a high-precision luminosity detector for CMS at the HL-LHC

Mykyta Haranko, Georg Auzinger for the CMS Collaboration

Abstract

The CMS BRIL project upgrades its instrumentation for the Phase-2 detector to provide high-precision luminosity and beam-induced background measurements. A part of the CMS Inner Tracker - the Tracker Endcap Pixel Detector (TEPX) - will allocate a fraction of the read-out bandwidth for luminometry. The implications of the proposed approach are highlighted. A dedicated luminosity trigger and clock distribution system is introduced and a test implementation on a demonstrator system is described. A demonstrator of the real-time on-FPGA pixel cluster counting algorithm is also described.

Presented at *TWEPP 2021 Topical Workshop on Electronics for Particle Physics*

PAPER

TEPX as a high-precision luminosity detector for CMS at the HL-LHC

To cite this article: M. Haranko *et al* 2022 *JINST* 17 C03001

View the [article online](#) for updates and enhancements.



The Electrochemical Society
Advancing solid state & electrochemical science & technology

242nd ECS Meeting

Oct 9 – 13, 2022 • Atlanta, GA, US

Abstract submission deadline: **April 8, 2022**

Connect. Engage. Champion. Empower. Accelerate.

MOVE SCIENCE FORWARD



Submit your abstract



TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2021
20–24 SEPTEMBER, 2021
ONLINE

TEPX as a high-precision luminosity detector for CMS at the HL-LHC

M. Haranko* and G. Auzinger on behalf of the CMS collaboration

European Organization for Nuclear Research (CERN), Geneva, Switzerland

E-mail: mykyta.haranko@cern.ch

ABSTRACT: The CMS BRIL project upgrades its instrumentation for the Phase-2 detector to provide high-precision luminosity and beam-induced background measurements. A part of the CMS Inner Tracker — the Tracker Endcap Pixel Detector (TEPX) — will allocate a fraction of the read-out bandwidth for luminometry. The implications of the proposed approach are highlighted. A dedicated luminosity trigger and clock distribution system is introduced and a test implementation on a demonstrator system is described. A demonstrator of the real-time on-FPGA pixel cluster counting algorithm is also described.

KEYWORDS: Data processing methods; Detector control systems (detector and experiment monitoring and slow-control systems, architecture, hardware, algorithms, databases); Pattern recognition, cluster finding, calibration and fitting methods; Data acquisition concepts

*Corresponding author.

Contents

1	Introduction	1
2	Back-end electronics	1
3	Dedicated trigger and timing distribution infrastructure	2
4	Luminosity processing	4
4.1	Pixel cluster counting algorithm	4
4.2	Algorithm performance	5
	Summary	6

1 Introduction

The CMS detector [1] will undergo a so-called Phase-2 upgrade [2] to meet the challenges imposed by the HL-LHC [3] operational conditions. Many of the instruments used by the CMS Beam Radiation Instrumentation and Luminosity (BRIL) project during runs 2 and 3 will have to be replaced [4] for HL-LHC. One of the important subsystems in the scope of the BRIL Phase-2 luminosity upgrade is the Tracker Endcap Pixel detector (TEPX) [5], which will be located in the very forward extremity of the Inner Tracker (IT) volume. It will be composed of four large double disks per CMS end, each made up of five rings with a varying number of pixel modules. The occupancy in TEPX will be relatively low with respect to other parts of the IT (less than 10^{-3} [5]), allowing the usage of the headroom in readout bandwidth for the luminosity measurement by sending 75 kHz of dedicated luminosity triggers in addition to the nominally 750 kHz of Level-1 (L1) triggers for physics. The data corresponding to those luminosity triggers can then be separated from the normal data acquisition stream for online processing in a luminosity-dedicated system. In addition, the innermost ring of the last disk in $|z|$ (D4R1) is beyond $|\eta| = 4$ and, because of the very few tracking points at that pseudorapidity, it is outside the nominal acceptance for efficient tracking. This will provide a possibility to exclusively use this part of the detector for luminosity and beam-induced background measurements utilizing the full trigger bandwidth.

2 Back-end electronics

The back-end system of the IT will consist of a set of so-called IT Data, Trigger and Control (IT-DTC) boards connected to the front-end electronics using optical links. The hardware platform for the IT-DTC is the so-called Apollo board [6], which follows Advanced Telecommunications Computing Architecture (ATCA) and among other elements hosts a System-on-Chip (SoC) and two processing FPGAs. The Xilinx Virtex Ultrascale+ VU13P FPGA is currently the prime candidate

for the processing FPGAs. The back-end card will control and monitor the front-end modules, and acquire the triggered event data. Through the Data Acquisition (DAQ) and Trigger and Timing Control and Distribution System (TCDS) Hub (DTH) placed in each crate, physics data from the IT-DTCs will be forwarded to the central CMS DAQ system.

According to its name, the DTH is also responsible for the interface with the CMS TCDS2 [7]. The luminosity triggers will be generated by the so-called BRIL Trigger Board (BTB) and distributed to all CMS subsystems via the aforementioned TCDS2 system. The front end itself is totally agnostic to the type of trigger. The decision to forward an event fragment to either the central DAQ or the luminosity processing system will be taken at the IT-DTC level according to the trigger type sent by the TCDS2 system. Figure 1 illustrates the data flow within the TEPX system based on the kind of trigger.




Figure 1. Block diagram of the data flow based on trigger types in the TEPX system. The luminosity trigger and data paths are shown in red [4].

3 Dedicated trigger and timing distribution infrastructure

In order to monitor the beam-induced background under various conditions, the TEPX D4R1 system is planned to run even when the LHC beams are not qualified and/or the CMS detector is not acquiring data. The bunch clock frequency is not fixed during some of these periods and the CMS systems are switched to a local clock source, thus making it impossible to use the bunch clock provided by the TCDS2 system for the measurements. As illustrated in figure 2, the BTB will be placed in a dedicated crate together with a DTH card in order to be synchronized with the rest of the CMS systems. The BTB will recover the bunch clock and orbit signal from the LHC machine interface [7], synchronize the slow TCDS2 data stream (LHC fill, CMS run, luminosity section and luminosity word identification numbers) to the LHC clock and orbit signal, and send a TCDS2-like control stream to the DTH in the D4R1 crate, thus, in this case, acting as local instance of the top node in the TCDS2 system (also known as the TCDS2 Captain). In addition to the above, the BTB will also receive technical triggers from the Beam Pick-up Timing system for Experiments (BPTX) [4] and, while also using them for luminosity trigger generation, encode them onto the optical link and send to the CMS Global Trigger (GT). Similar to Apollo in terms of its structure — the Serenity ATCA board [8] will serve as the hardware platform for the BTB.

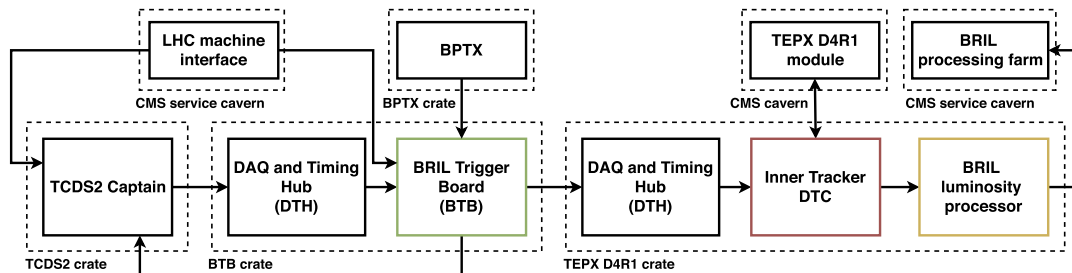


Figure 2. Schematic diagram of the clock and trigger distribution to the TEPX D4R1 system.

Prototype firmware for the BTB has been implemented and tested using the prototypes of the DTH and Serenity cards. For the tests, a CMS-like clock was supplied from the DTH, whereas one of the onboard PLLs was configured to internally generate a free-running bunch clock with ramping frequency, thus emulating the LHC clock (extracted from the real accelerator operation log). The ramping tests were carried with ramps obtained from the recorded LHC ramp by multiplying the frequency steps by 1, 10, 100, and 1000 in order to exercise the system performance under more extreme conditions than required for operation. The measured clock frequencies as a function of time are shown in figure 3 (left), the LHC clock frequency steps were multiplied by 10 in this case. In the firmware, the so-called TCDS2 synchronization block was extracting the required data from the original TCDS2 stream and synchronizing them to the LHC clock domain. The synchronized TCDS2 stream was sent through an optical loop-back link to the firmware module emulating the DTH placed in the D4R1 crate. Both the original TCDS2 stream and the one received from the loop-back fiber were compared in the checker block, which contained two monitoring instances that independently analyzed both TCDS2 streams. The analyzers stored all the received information in FIFOs and when the full luminosity word was ready for readout in both FIFOs, a dedicated state machine compared their contents and counted the number of mismatches. After the ramping test, a set of errors was injected on purpose into the synchronized TCDS2 stream to verify the performance of the checker block. The measured mismatch count between the two data streams is shown in figure 3 (right). The checker block performs as expected and is able to detect any corruptions of the synchronized TCDS2 stream with respect to the original one. No mismatch was detected during the clock frequency change.

The study described above has shown the feasibility of implementing the clocking infrastructure proposed for D4R1. In the nearest future, the prototype firmware is planned to be extended by the luminosity trigger generation algorithm and by the logic performing the sampling and re-transmission of the technical triggers from the BPTX system.




Figure 3. Measured CMS (green) and LHC (red) clock frequency (left) and mismatch count as measured by the checker block (right) as a function of time during a simulated LHC ramp. The arrows indicate times when errors are injected into the simulated TCDS2 stream [4].

4 Luminosity processing

On the processing side, the luminosity event fragments from the IT-DTC will be received by dedicated luminosity processing blades also based on the Apollo board, but utilizing independent firmware to perform real-time pixel cluster counting on its processing FPGAs. The IT-DTCs and luminosity blades will connect to each other through optical links on a one-to-one basis. After unpacking the event data on the luminosity blades, the per-chip events will be distributed to instances of the pixel cluster counting algorithm. Cluster counts for each event will be then grouped and summed per TEPX quarter ring and the total counts for each quarter ring will be transferred to the histogramming instances. The latter will transfer the accumulated histograms to the on-board SoC and further to the BRIL computing farm at the end of each integration period for luminosity calculation.

4.1 Pixel cluster counting algorithm

The BRIL-operated D4R1 modules will be receiving triggers at a maximum rate of 1 MHz at an average number of interactions per bunch crossing (pileup) of 200, which defines the minimal required event processing rate. D4R1 will contain only 20 modules (80 chips) per z end that will be read out by a single IT-DTC. The IT-DTCs corresponding to the rest of TEPX will send event data from up to 704 chips per IT-DTC at a significantly lower trigger rate. Therefore, the FPGA pixel clustering module has to be fast enough to handle the D4R1 data and light in terms of FPGA resources utilized to handle the data from the rest of TEPX.

A test algorithm has been developed to prove the feasibility of the aforementioned task. Two hits were defined to form a cluster if they are horizontally, vertically, or diagonally adjacent. Thus, it is enough for them to share a single corner to form a cluster. The algorithm does not use the hit charge information from the front end and does not store cluster geometry information, such as cluster position or size, thus performing only cluster counting. The algorithm is split into six sequential processing stages operating at 320 MHz clock rate and performs tasks such as the chip data decoding, clustering of 4×4 pixel fragments, checking cluster isolation, merging

of vertically adjacent clusters and merging of horizontally adjacent clusters. The latter three stages keep internal cluster counts for filtering nonadjacent clusters. The last processor in this chain performs accumulation of the internal counts. The load between processor input buffers is dynamically balanced by the back-pressure mechanism to compensate for events containing larger clusters.

4.2 Algorithm performance

In order to qualify the performance of the algorithm, 500 events simulated with an average pileup of 200 for the whole TEPX were split into 2804763 single chip events with at least one hit, and were injected into the clustering algorithm sequentially with a mean trigger rate of 1 MHz. The fraction of events where the calculated cluster count did not match the one obtained from the reconstruction algorithm used in the CMS offline reconstruction has been categorized, identifying that in 2.56% of events the charge thresholds have been applied in the offline algorithm resulting in cluster splitting. In 0.01% of events the maximal cluster size was exceeded for the offline algorithm, also resulting in cluster splitting. These two types of events were not considered as failures of the FPGA algorithm. Finally, only in less than 0.01% of events the cluster counts were mismatched due to unidentified reasons, with a typical mismatch by one cluster.

The effect of various pileup conditions on the mean cluster count was studied by injecting the simulated event data for average pileup values in the range from 1 to 200 with a mean trigger rate of 1 MHz, as shown in figure 4 (left) for D4R1. As expected, the value grows linearly with pileup. In order to prove that the algorithm is capable of handling the required event rates, the mean event processing time was measured as a function of pileup and was later converted into the maximum trigger rate, as shown in figure 4 (right). The maximum trigger rate at an average pileup of 200 is 1.33 ± 0.44 MHz. This satisfies the requirements and also provides sufficient margins. A relatively wide standard deviation of the event processing time, and consequently the processing rate, is caused by the more time-intensive reconstruction of larger clusters.




Figure 4. Mean cluster count per event (left) and calculated maximum trigger rate (right) at different pileup conditions for disk 4 ring 1 (D4R1) of TEPX. The error bars on the right plot represent the standard deviation of the distribution, while the blue dashed line indicates a trigger rate of 1 MHz [4].

The number of algorithm instances sufficient to process data from a single IT-DTC (up to 704 chips) was placed on the VU13P FPGA. The resulting implementation did not contain any failing timing path and utilized only 45% of logic resources and 27% of memory resources. The present implementation of the algorithm can be easily placed on the target FPGA, and sufficient space is available for further improvement of the algorithm.

Summary

A prototype version of the FPGA pixel cluster counting algorithm has been implemented and tested. The developed algorithm is fully able to perform online cluster counting at the required event rate of 1 MHz and provides cluster counts consistent with the ones obtained from the offline reconstruction algorithm, with a low fraction of mismatches. Sufficient FPGA resource margin is available for further improvement of the algorithm. In particular, charge information and geometrical properties of clusters will be considered in the final implementation. In order to prove the possibility of running TEPX D4R1 during times when the LHC bunch clock frequency is not fixed, a prototype firmware for the BRIL trigger board has been developed proving the feasibility of the proposed clock distribution infrastructure.

References

- [1] CMS collaboration, *The CMS experiment at the CERN LHC*, 2008 *JINST* **3** S08004, <https://cds.cern.ch/record/1129810>.
- [2] CMS collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, Technical Proposal CERN-LHCC-2015-010, LHCC-P-008, CMS-TDR-15-02, <http://cds.cern.ch/record/2020886> (2015).
- [3] G. Apollinari et al., *High Luminosity Large Hadron Collider HL-LHC*, <https://cds.cern.ch/record/2120673> (2015) <https://doi.org/10.5170/CERN-2015-005.1>.
- [4] CMS collaboration, *The Phase-2 Upgrade of the CMS Beam Radiation Instrumentation and Luminosity Detectors*, Tech. Rep., CERN, Geneva, <https://cds.cern.ch/record/2759074> (2021).
- [5] CMS collaboration, *The Phase-2 Upgrade of the CMS Tracker*, Technical Design Report CERN-LHCC-2017-009, CMS-TDR-014, <https://cds.cern.ch/record/2272264> (2017).
- [6] E.S. Hazen et al., *The APOLLO ATCA Platform*, *PoS TWEPP2019* (2020) 120.
- [7] CMS collaboration, *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*, Tech. Rep., CERN, Geneva, <https://cds.cern.ch/record/2759072> (2021).
- [8] A. Rose et al., *Serenity: an ATCA prototyping platform for CMS Phase-2*, *PoS TWEPP2018* (2019) 115.