Available on CMS information server

CMS CR -2021/228

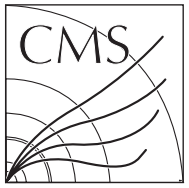

The Compact Muon Solenoid Experiment

Conference Report

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

21 October 2021 (v4, 07 November 2021)

Firmware Architecture of the back-end DAQ system for the CMS High Granularity Endcap Calorimeter detector.

Stavros Mallios for the CMS Collaboration

Abstract

During the High-Luminosity phase of the Large Hadron Collider, the Endcap Calorimeter detectors of the Compact Muon Solenoid experiment will be replaced by the High Granularity Calorimeter. For reading out the new calorimeter, Field Programmable Gate Array firmware was developed targeting the off-detector hardware. The firmware is responsible not only for the readout of the detector but also for its slow control and timing. To facilitate system maintenance, the firmware is optimized to handle all the different Front-End electronics configurations and data rates using a single - highly configurable - design. This manuscript presents the firmware architecture and the implementation.

Presented at *TWEPP 2021 Topical Workshop on Electronics for Particle Physics 24/10/2021*

Prepared for submission to JINST
Topical Workshop on Electronics for Particle Physics - TWEPP2021
20 - 24 September 2021
CERN

Firmware Architecture of the back end DAQ system for the CMS High Granularity Endcap Calorimeter detector.

S. Mallios,[a,1] P. Dauncey[b], A. David[a] and P. Vichoudis[a] on behalf of the CMS collaboration.

[a] *CERN, Geneva, Switzerland*
[b] *Imperial College London, London, England*

E-mail: stavros.mallios@cern.ch

ABSTRACT: During the High-Luminosity phase of the Large Hadron Collider, the Endcap Calorimeter detectors of the Compact Muon Solenoid experiment will be replaced by the High-Granularity Calorimeter. For reading out the new calorimeter, Field Programmable Gate Array firmware was developed targeting the off-detector hardware. The firmware is responsible not only for the readout of the detector but also for its slow control and timing. To facilitate system maintenance, the firmware is optimized to handle all the different Front-End electronics configurations and data rates using a single - highly configurable - design. This manuscript presents the firmware architecture and the implementation.

KEYWORDS: CMS; HGCAL; High Luminosity LHC; High granularity calorimeter; LHC Phase 2; Data acquisition system; Back-end electronics;

[1] Corresponding author.

Contents

1. Introduction

The upgraded High Luminosity LHC, after the third Long Shutdown (LS3), will provide an instantaneous luminosity of 7.5×10^{34} $cm^{-2}s^{-1}$ (levelled), at the price of a dramatic increase of the number of pileup interactions. The upgraded CMS detector [**1**] will be read-out at an unprecedented event rate of 750 kHz and a data rate of up to 50 Tb/s.

To cope with the new challenging conditions, the CMS collaboration will replace the existing electromagnetic and hadronic endcap calorimeters with the new High Granularity Calorimeter (HGCAL). The new detector will feature unprecedented transverse and longitudinal segmentation for both electromagnetic (CE-E) and hadronic (CE-H) compartments and will have the ability to withstand the high radiation levels. The CE-E and a large fraction of CE-H will use silicon hexagonal shaped detector modules as active detector material, while the lower-radiation part of the CE-H will be instrumented with scintillator tiles with on-tile SiPM readout. [**2**]

The firmware architecture presented here is part of the Back-End (BE) Data Acquisition system (DAQ) designed for HGCAL. Each BE DAQ board will receive data from the Front-End (FE) electronics through optical links running at 10.24 Gb/s and process, buffer, merge and forward them to the next stage, the DAQ and Timing Hub (DTH) board [**3**]. It will also configure the FE electronics and distribute the high precision clock and the trigger signals through optical links running at 2.56 Gb/s.

2. The HGCAL Control and Data Acquisition architecture

The main purpose of the DAQ system is to provide the data pathway and time decoupling between the synchronous detector readout and data reduction (Level 1 Trigger - L1T), the asynchronous selection of interesting events (High Level Trigger - HLT), their local storage at the experiment site, and the transfer to Tier-0 for offline permanent storage and analysis.[**4**]

2.1 The front-end electronics

The FE electronics (see Fig. 1), located on or close to the detector, consist of a mixture of analog and digital components in charge of collecting, processing and digitizing signals from the sensors themselves, and ancillary logic to control these processes and to monitor the detector. The custom rad-hard HGCAL Readout Chips (HGCROC) [**5**] were designed to collect and digitize the detector sensors. Data from multiple HGCROCs will be concentrated using dedicated endcap concentrator (ECON) ASICs (Application specific integrated circuit). The ECON chip comes in two flavours: one for the DAQ data (ECON-D) and the other for the trigger data (ECON-T). The concentrated data will be sent to the BE via the radiation hard, low power lpGBT (low-power Gigabit Transceiver) chipset [**6**] through fibres (9k total) running at 10.24 Gb/s. The lpGBT chipset will also receive and distribute clock, configuration, and fast commands data from the back end through a downlink running at 2.56 Gb/s.

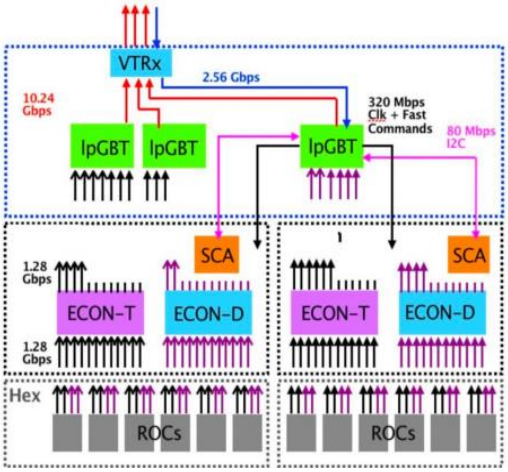

Figure 1. The FE electronics diagram.

2.2 The back-end electronics

The BE DAQ boards (based on the hardware platform developed by the Serenity collaboration [**7**]) will be located in the service cavern (an environment with minimal magnetic field and negligible radiation). Each board receives detector data from the ECON-Ds via 108 lpGBT links running at 10.24 Gb/s and concentrates and forwards them to the Central DAQ via 12 SLINKs (SlinkRocket is a custom protocol used to transmit data form the BE DAQ card to the DTH card [**8**]) running at 25 Gb/s. It also provides the required buffering, to avoid data loss due to rate fluctuations, and a configurable lpGBT-to-SLINK mapping to mitigate the variations

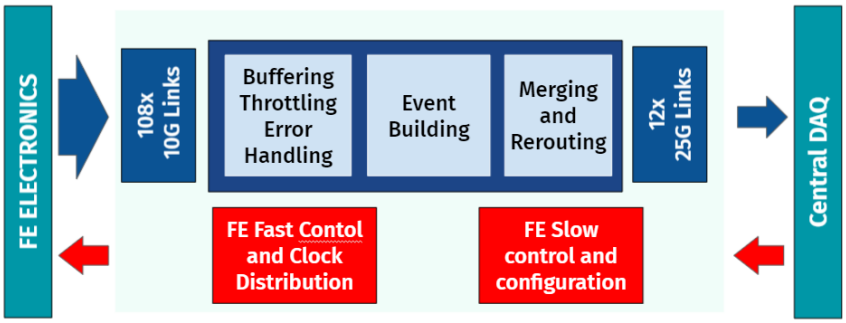

Figure 2. The HGCAL BE DAQ board architecture. The DAQ receives, buffers and processes data from up to 108 lpGBTs through the uplink path (blue). In addition, it configures and monitors the FE electronics (Slow control) and distributes the clock and fast control signals from the central timing, control and distribution system (Fast Control), through the downlink path (red).

in rate/channel. The BE DAQ is also responsible for configuring the FE electronics and distributing the high precision clock and the trigger signals to the on-detector electronics (see Fig. 2).

The central DAQ hardware (DTH board) distributes clock signals, synchronization commands, Level-1 triggers, and slow control to BE DAQ. It also receives event fragments from the BE DAQ, and subsequently concentrates, buffers and transmits them to the data-to-surface network.

3. The HGCAL BE DAQ system

3.1 The FE-to-BE data path (uplink)

Each BE DAQ FPGA will be capable to receive and process data from up to 108 lpGBTs. The FE-to-BE mapping varies significantly, resulting in different number of ECON-Ds per lpGBT link, due to the non-uniform occupancy in the endcaps. Some ECON-Ds will require two lpGBT links while in other cases up to six ECON-Ds can send data through one lpGBT link. A single robust firmware design was developed to handle all cases to avoid the maintenance complications of multiple firmware designs. A total of 54 configurable processing block instances are receiving and buffering the data from the 108 lpGBT links, with each unit assigned to two lpGBT channels and capable of handling data from 1 up to 12 ECON-Ds. A single FPGA can process data from up to 648 ECON-Ds. The ECON-D packets received by each processing unit are then readout to event buffers. The event frames from multiple event buffers are then merged to larger packets and send to the DTH through 12 SLINKs running at 25Gb/s.

The corresponding data rate of each processing unit will differ from FPGA to FPGA which makes the mapping of the 54 Processing units to the 12 SLINKs complicated. A configurable interconnect logic offers a flexible mapping of the 54 Event buffers to 12 fixed rate (25 Gb/s) output links, to balance out the very inhomogeneous input rate (see Fig. 3). This interconnect block allows to configure each SLINK to forward data from 1 up to 12 event buffers.

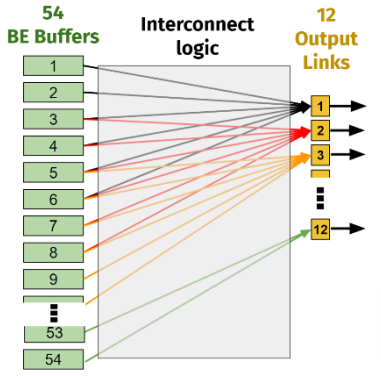

Figure 3. The partially programable input-to-output routing

3.2 The Control links (downlink)

The BE DAQ board, in addition to moving event data from the FE to the central DAQ is also responsible for the configuration and monitoring of the FE electronics (Slow control) and the distribution of the clock and fast control signals from the central timing, control and distribution system (TCDS), through the downlink path.

3.2.1 Slow Control (SC)

The BE DAQ board is responsible for sending the configuration and monitoring data to all the various FE ASICs. There are two different “flavours” of FE electronic paths: the IC/EC (Internal/Extrenal Control) path carrying configuration data to the FE electronics in the silicon part of the detector, and the HDLC (High-level Data Link Control) path driving the dedicated SCA (Slow Control Adapter) ASICs in the scintilator part. In both cases multiple configuration commands are stored in dedicated SC buffers before sending to the FE. A transactor handles the communication between the SC buffers and the SC formater. The FE responses are also stored in dedicated buffers accessed by the software (see Fig. 4).

Due to the sheer number of SC channels handled by each FPGA and the limited resources, it was impossible to instantiate one SC block per channel. To solve this problem, a

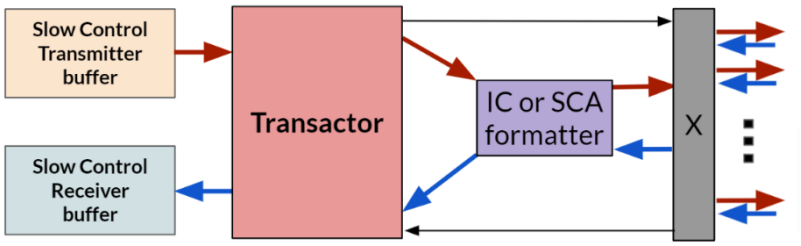

Figure 4. Simplified diagram of the Slow Control distribution block. There are 8 block instantiations driving 16 IC channels each and 8 block instantiations driving 40 SCA channels each.

cross point switch was added to allow a single SC block to drive multiple SC channels offering broadcasting capabilities and a tradeoff between FPGA resource usage and configuration time. There are currently 8 IC blocks, driving 16 channels each, and 8 SCA blocks, driving 40 channels each, but in the final system this will be optimized to offer the best FPGA resource usage to configuration time ratio.

3.2.2 Fast Control (FC) and auto-recovery

The BE DAQ board receives the clock and fast control commands from the DTH boards using through the ATCA backplane. A dedicated FC block receives the FC data, encodes them to 8-bit system specific fast commands and transmits them to all FC channels.

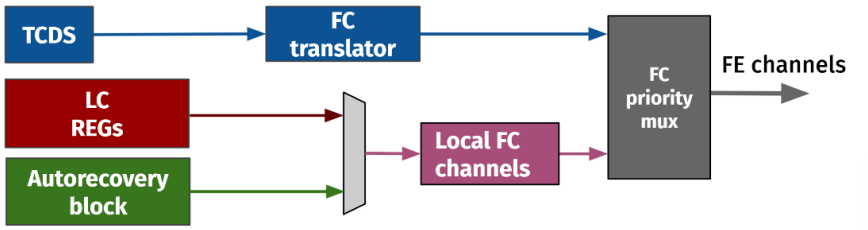

Figure 5. A simplified diagram of the fast control blocks. The FC blocks receives the fast commands from the central DAQ and the local fast commands from the dedicated registers or from the auto-recovery blocks and forwards them to the FE through the downlink.

Furthermore, the FC block allows for local generation of fast commands. There are dedicated local-FC registers driving each FC channel accesible by software. The autorecovery block can also access the FC channels and send autorecovery fast commands to specific FE chips, a feature particularly important for the quick recovery of individual FE ASICs temporarily out of sync (see Fig. 5).

4. FPGA selection and implementation

During the firmware development, extensive studies were carried out — driven by the evolving system requirements and electronics cost — to identify the best and most cost-effective solution for the BE DAQ system. To do so, special care was taken while designing the firmware allowing device retargeting by simply changing a few constants and constraints. Additionally, a sophisticated testbench was designed to allow feeding realistic input data and

comparing the firmware behavior in respect with that of the system emulator (running in software), thus facilitating the validation of the firmware functionality. After trying various

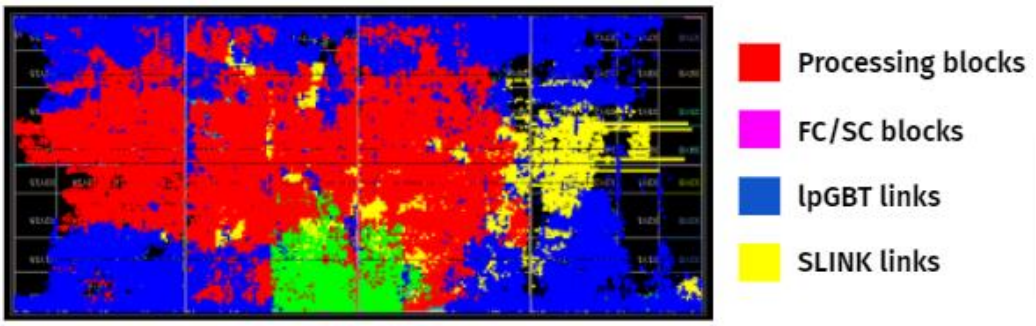

Figure 6. The VU13p device view with the resources utilised by the different BE DAQ firmware blocks highlighted in different colors.

Serenity board configurations (dual KU15P, dual VU7P), the studies have shown that the optimum solution is Serenity boards equipped with a single Xilinx VU13P FPGA, each serving 108 input links and 12 output links. In Fig. 6 you can see the VU13P device with the resources utilized by the different BE DAQ firmware blocks, highlighted in different colors.

5. Conclusion and outlook

This firmware work is driving a full BE DAQ architecture review, improving upon the TDR expectations and towards pre-production in 2022. This was achieved thanks to the implementation & simulation strategies selected. Special care has been taken during the firmware development to facilitate targeting different FPGA devices. After careful studies, the best option has been identified to be a single-FPGA board equipped with a VU13P Xilinx Ultrascale+ FPGA. The behaviour of the firmware must now be validated in hardware, once the VU13P-based prototype will become available. Meanwhile, hardware tests will take place using previous generation Serenity hardware.

References

[1] CMS Collaboration, *The CMS experiment at the CERN LHC*, JINST 3 (2008) S08004, doi:10.1088/1748-0221/3/08/S08004

[2] CMS Collaboration, *The phase 2 Upgrade Endcap Calorimeter*, CMS Technical Design Report, CERN-LHCC-2017-023, CMS-TDR-019, CERN, Geneva, Nov 2017.

[3] J.-m. O. Andre, U. Behrens, A. Bocci, J. Branson, S. Cittolin, D. Da Silva Gomes et al., *Design and development of the DAQ and Timing Hub for CMS Phase-2*, Tech. Rep. CMS-CR-2018-307, CERN, Geneva, Oct 2018

[4] C. Collaboration, *The Phase-2 Upgrade of the CMS DAQ* Interim Technical Design Report, Tech. Rep. CERN-LHCC-2017-014. CMS-TDR-018, CERN, Geneva, Sep 2017.

[5] G. Bombardi et al., *"HGCROC-Si and HGCROC-SiPM: the front-end readout ASICs for the CMS HGCAL"* 2020 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC).

[6] LpGBT design team. LpGBT Manual. 2019. url: https://lpgbt.web.cern.ch/lpgbt/index.html.

[7] CMS Collaboration, *Serenity - An ATCA prototyping platform for CMS Phase-2*, CMS-CR-2018-327, Oct 2018.

[8] DAQ-DTH design team. SlinkRocket Manual. 2021. url: https://cms-daq-dth-p1.web.cern.ch/slinkrocket