Published by IOP Publishing for Sissa Medialab

Received: *November 3, 2024*
Revised: *January 1, 2025*
Accepted: *February 21, 2025*
Published: *March 18, 2025*

Topical Workshop on Electronics for Particle Physics
University of Glasgow, Scotland, U.K.
30 September–4 October 2024

Back-end DAQ system prototype testing and integration on a full detector test system for the CMS HGCAL detector

M. Rosado[a,b,c,*] **A. Şarkışla**[a,d] **S. Mallios**[e] **B. Akgün**[d] **A. David**[a] **N. Roma**[b,c] **P. Tomás**[b,c] **and P. Vichoudis**[a] **on behalf of the CMS collaboration**

[a] *Experimental Physics Department, CERN, Esplanade des Particules 1, 1211 Geneva, Switzerland*
[b] *Electrical and Computer Engineering Department, Instituto Superior Técnico, Universidade de Lisboa, Av. Rovisco Pais Nº 1, 1049-001 Lisboa, Portugal*
[c] *INESC-ID, Rua Alves Redol 9, 1000-029 Lisboa, Portugal*
[d] *Physics Department, Boğaziçi University, Bebek, 34342 Beşiktaş, İstanbul, Turkey*
[e] *Physics Department, Imperial College London, Prince Consort Road, London, SW7 2BW, U.K.*

E-mail: martim.rosado@cern.ch

Abstract: The CMS Collaboration will replace its current endcap calorimeters with the new high granularity calorimeter (HGCAL) for operations at the HL-LHC. The HGCAL back-end DAQ (Data Acquisition) system comprises 96 FPGA-based ATCA (Advanced Telecommunications Computing Architecture) boards, each processing data from 108 input optical fibres operating at 10 Gb/s. This paper describes in detail the architecture and prototyping of the elementary unit in the back-end DAQ system of the HGCAL. We then describe and evaluate its integration and performance in a full detector test system. The resulting system provides an average data acquisition throughput at the detector's nominal rate of 750 k events per second.

Keywords: Data acquisition circuits; Digital electronic circuits; VLSI circuits

*Corresponding author.

https://doi.org/10.1088/1748-0221/20/03/C03028

Contents

1 Introduction

The Phase 2 upgrade of the Large Hadron Collider (LHC) motivates the replacement of the Compact Muon Solenoid (CMS) detector [1] endcap calorimeters with the new High-Granularity Calorimeter (HGCAL) [2]. The HGCAL back-end Data Acquisition (DAQ) system [3] is responsible for aggregating front-end data, transmitted via $O(9000)$ optical fibres operating at 10 Gb/s, at a mean collision (event) acquisition rate of 750 kHz. This required bandwidth corresponds to an average data rate of 15 Tb/s and event size of 2.5 MByte. After undergoing an event-building process in the back-end DAQ system, the data are transmitted to the central CMS DAQ system through $O(1100)$ optical fibres working at 24 Gb/s.

Accordingly, the HGCAL back-end DAQ will be implemented in 96 ATCA boards (Serenity [4]) each of which processes on average 160 Gb/s of data at an operating frequency of 320 MHz (210 kb of data per event). Each Serenity board receives 108 input optical fibres and conveys its output data via 12 output optical fibres.

This work describes the architecture of the elementary DAQ unit of the system, the Capture Block, able to perform the required event-building procedure in the HGCAL data while coping with the several configuration scenarios demanded by the heterogeneity of the HGCAL front-end electronics. Furthermore, its testing and validation are also discussed, based on the integration of Capture Blocks in a minimum dimension full detector test system, deployed in beamtest experiments at the European Organization for Nuclear Research (CERN) Super Proton Synchrotron (SPS) North Area, able to acquire data through the complete HGCAL data chain.

2 The HGCAL back-end DAQ system

Each of the 96 Serenity boards in the system hosts a VU13P Xilinx Field Programmable Gate Array (FPGA) with an identical gateware design being responsible for performing three main tasks:

1. Distribute the clock and fast command signals to the front-end electronics, which control the data acquisition process,
2. Distribute slow commands to the front-end electronics, which configure and monitor the Application Specific Integrated Circuits (ASICs) in the front-end, and
3. Process the collision data received from the front-end electronics and transmit the resulting data to the central CMS DAQ system, the main focus of this paper.

2025 JINST 20 C03028

A simplified dataflow of the HGCAL back-end DAQ system is depicted in figure 1. The front-end data are transmitted to the Serenity boards from $O(30\,000)$ data E-link Concentrator ASIC - DAQ (ECON-D) through $O(9000)$ input optical fibres, each working at 10 Gb/s. Each fibre pair conveys data from up to 12 ECON-Ds [5] that are transmitted over the 14 e-links [6] available on 2 low-power Giga-Bit Transceivers (lpGBTs) [7] ASICs. The number of ECON-D serviced by each fibre pair depends on the detector occupancy of the region of the serviced front-end electronics, with each ECON-D using up to 6 e-links to transmit data. Moreover, the e-links of an ECON-D may straddle the 2 lpGBTs of the respective fibre pair. This has the consequence that fibre pairs must be processed together.

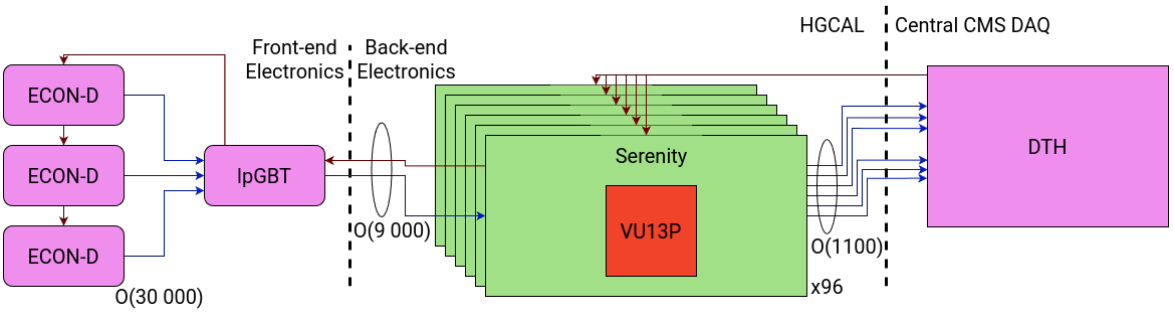

Figure 1. Simplified dataflow of the HGCAL DAQ system. Data signals are highlighted in blue (flowing left to right) and the clock and fast control signals in red (flowing right to left).

To use the same gateware design across the 96 Serenity boards, the processing circuitry handling each pair of fibres must be able to cope with all front-end electronic configuration possibilities. Hence, the aimed Capture Block circuit (the main focus of this paper) was designed to process data from up to 12 ECON-Ds spread across 14 e-links in a configurable manner. As fibres are processed in pairs, each Serenity board is envisaged to host 54 of these units, being able to process data from up to 648 ECON-Ds.

After the processing in the Serenity board, the data are sent to the DTH boards [8] of the central CMS DAQ system via the $O(1100)$ output optical fibres working at 24 Gb/s using the SLink protocol [9].

3 Capture Block architecture

The Capture Block is the elementary DAQ unit of the HGCAL back-end. It was designed to aggregate the data of the same event from all ECON-Ds that are contained in a fibre pair. This process is referred to as event building and is achieved through two major steps: packet detection and packed matching.

To guarantee that the Capture Block can cope with all front-end configuration possibilities that specify the distribution of the data of up to 12 ECON-Ds across 14 e-links in each fibre pair, this circuit was designed to be configurable while minimising resource usage. Figure 2 details the Capture Block architecture and its interfaces with input and output data ports.

Before being transmitted to the Capture Block, a serialiser groups e-links in pairs and serialises the data into a 64-bit data stream transmitted at the operating frequency of the Capture Block (320 MHz) and synchronous with the LHC clock.

After the data are transmitted to the Capture Block, each e-link word is assigned an ID corresponding to an ECON-D. This identification is written in configuration registers accessible to

2025 JINST 20 C03028

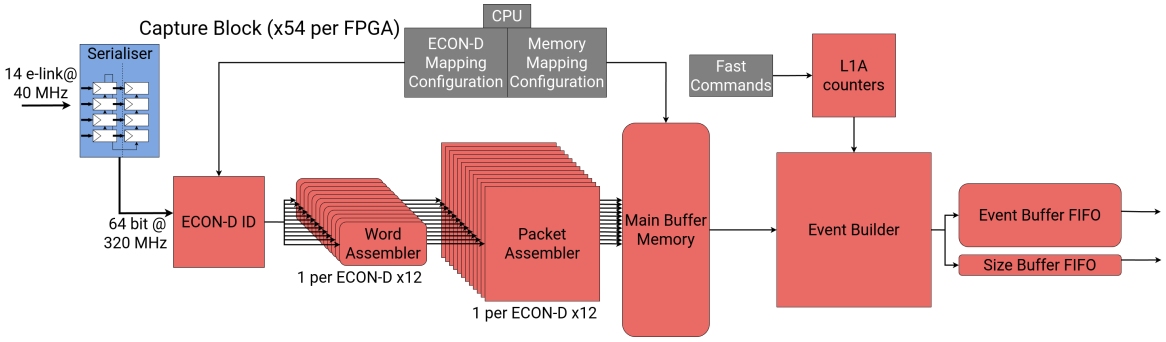

Figure 2. Detailed Capture Block architecture. Capture Block components are highlighted in orange and the connections to the software and fast commands are depicted in grey.

software via the IPBus protocol [10]. Afterwards, the tagged data stream is conveyed to 12 word assemblers, which filter and align the respective data words to only include words from their respective ECON-D. This process creates 12 data streams containing only words from a single ECON-D and ready to undergo the packet detection step of the event building process.

Each of such streams is parsed by a packet assembler locating ECON-D packet headers, while filtering idles and spurious events. This identification is performed by detecting the ECON-D header marker following an idle word in the data stream. To prevent the detection of spurious packets, an additional CRC check is performed in the detected header word to guarantee the detected packet corresponds to valid data. This check involves the 8-bit CRC field embedded in the ECON-D header. If the CRC code calculated in the packet assembler matches the one in the detected header word, the rest of the packet is regarded as valid. Hence, it is read from the data stream, and stored in the main buffer memory, shared by all packet assemblers. Since the ECON-D words are transmitted via 32-bit e-links to the back-end electronics, the packet assembler also inserts a padding word whenever necessary to align the data to 64-bit for further processing.

The main buffer memory is implemented using dense Ultra RAM (URAM) cells and it is shared by all packet assemblers. This allows the optimisation of the memory resources as configuration registers (accessible to software via the IPBus protocol) control the memory allocation for all packet assemblers. It is thus possible to adjust the memory allocation by taking into account the expected data rate of each ECON-D. Furthermore, this also allows to not allocate any memory for unused packet assemblers when the number of processed ECON-Ds in the Capture Block is less than 12.

The Capture Block also receives the Level-1 Accept (L1A) fast command signal transmitted to the front-end electronics that triggers the acquisition of an event. Dedicated counters corresponding to LHC orbit, Bunch Crossing (BX), and event identification number are kept locally and allow the Capture Block to stay synchronous with the other components, by generating a timestamp identifying the relative time of the event to acquire. Upon the reception of an L1A, the corresponding L1A counter values are stored in a FIFO, later read by the event builder.

The final step of the event building, packet matching, can be started when the event builder, implemented in a state machine, detects both data in the main buffer and an available L1A timestamp. This state machine concatenates the data from all the ECON-Ds corresponding to the same event into a packet, to be transmitted out of the back-end FPGA. For each ECON-D, the local L1A counters are compared with their counterpart, embedded in the ECON-D header. This verification allows the

event builder to guarantee that the data stored in the main buffer corresponds to the L1A timestamp being processed. If the tested timestamps are correct, the event is extracted from the main buffer and stored in the event buffer FIFO and the process is repeated for all ECON-Ds. Hence, when that is not the case, the appropriate metadata flags are added to the Capture Block header. A Capture Block header shall contain metadata regarding the event (namely, the L1A timestamp of the Capture Block) and error information for each ECON-D (if any): timestamp mismatch, timeout, or main buffer overflow flags. This allows the offline software to parse and process the data, after its transmission to the central CMS DAQ system.

The Event buffer FIFO is the output interface of the Capture Block, exposing the processed data to the output channels. Another FIFO, the size buffer FIFO, contains an entry corresponding to the length of each packet in the event buffer FIFO. These data are transmitted to a readout controller, able to merge data from several Capture Blocks and convert them into an SLink stream, which is transmitted out via a 24 Gb/s optical fibre.

4 Performance evaluation

To evaluate the conceived architecture, two Capture Blocks were integrated into a prototyping system that was used at two beamtests performed during the summer of 2024 at CERN SPS North Area. This prototyping system comprises a full HGCAL data chain, extending previous work described in [11], connecting the sensors in the front-end electronics to a Serenity prototype board which is then connected to a DTH board prototype [8]. The used Serenity prototype hosted a VU7P FPGA as the final Serenity prototypes hosting VU13P FPGAs are still under development.

The front-end electronics prototype comprised 6 silicon sensor modules divided into two layers and read out by 6 ECON-Ds. Each layer is connected to the Serenity board via an optical fibre and then connected to a Capture Block. Each Capture Block was configured to receive data from its respective layer's 3 ECON-Ds.

To test the data acquisition at the nominal L1A rate of 750 kHz, the system was stimulated with random L1A triggers at different frequencies, ranging (on average) from 100 kHz to 970 kHz, in different runs. Figure 3 shows the relation of size and number of BXs between consecutive events (instantaneous frequency) for 109 M events comprising two such runs at an average rate of 220 kHz and 970 kHz. The large range of observed packet sizes is correlated with the truncation of the ECON-D packets in the front-end, due to the higher than 750 kHz L1A rates and with the presence (or not) of beam signal among the events.

As the rate increases, the internal buffers of the ECON-Ds progressively fill up, as each packet cannot be transmitted to the back-end in the time interval between L1As. Once the buffers are full, the corresponding ECON-D sends truncated packets to the back-end. This generates the observed horizontal bands in figure 3, each indicating one more ECON-D sent a truncated packet. Inside each band, the variation of packet sizes is explained by the energy deposited by beam particles in the detector. As the SPS beam particles are asynchronous to the data acquisition system, some events have a reduced size when particles arrive far from the signal sampling clock edge.

To verify the correct behaviour of the two Capture Blocks under test, it is necessary to check the timestamps on the packet headers and the packet size. Not only must all timestamps match among themselves but also match with the timestamp embedded in the SLink header word [9]. This guarantees the purpose of the Capture Block has been achieved: the correct event building processing

2025 JINST 20 C03028

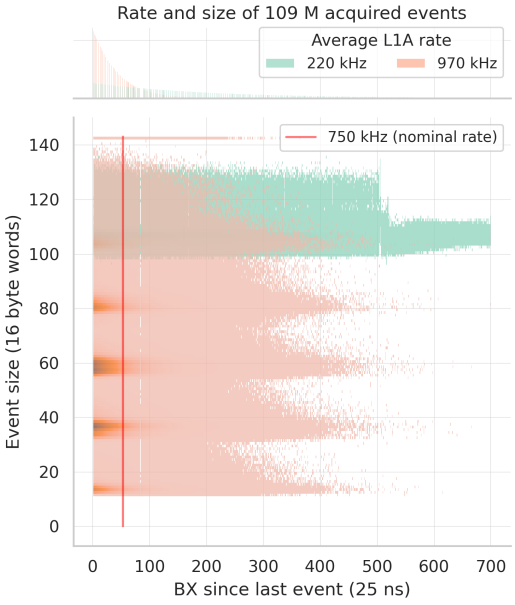

Figure 3. Performance evaluation of the designed Capture Block: results regarding size and relative timing of 109 M events acquired with the system.

of data throughout all the chain. The obtained data underwent offline testing to ensure the correct synchronisation of both Capture Blocks. The performed checks were successful with no error being encountered in the Capture Blocks orbit, BX, or event counters for all the data. These timestamps were also compared with their counterpart embedded in the SLink header word also yielding no errors. Furthermore, the size of the packets was also verified to ensure no missing or extra word anomalies happened in the data. These tests were also successful. Hence, the correct functionality of the Capture Block was demonstrated at average frequencies higher than 750 kHz, the target nominal rate of the HGCAL, with several event sizes and instantaneous frequencies.

5 Conclusions

This work described the architecture of the elementary DAQ unit of the HGCAL back-end DAQ system and discussed its implementation on a prototyping system used in beamtests at CERN. The obtained results allowed us to successfully demonstrate the correct functionality of the Capture Block at the target data acquisition rate of 750 kHz without errors. Future work comprises scaling the beamtest prototyping system up to the final size of 54 Capture Blocks per VU13P FPGA. The main focus of this effort will comprise the load-balanced data transmission across the FPGA to avoid overloading the output data channels.

Acknowledgments

Supported by Fundação para a Ciência e a Tecnologia (Portugal) under project UIDB/50021/2020 and grant PT/BD/154723/2022, as part of a doctoral studentship at CERN.

2025 JINST 20 C03028

References

[1] CMS collaboration, *The CMS Experiment at the CERN LHC*, 2008 *JINST* **3** S08004.

[2] CMS collaboration, *The Phase-2 Upgrade of the CMS Endcap Calorimeter*, CERN-LHCC-2017-023 (2017).

[3] S. Mallios, P. Dauncey, A. David and P. Vichoudis, *Firmware architecture of the back end DAQ system for the CMS high granularity endcap calorimeter detector*, 2022 *JINST* **17** C04007.

[4] A. Rose et al., *Serenity: An ATCA prototyping platform for CMS Phase-2*, *PoS* **TWEPP2018** (2019) 115.

[5] D. Braga et al., *First test results of the HGCAL concentrator ASICs: ECON-T and ECON-D*, 2024 *JINST* **19** C03050.

[6] S. Bonacini, K. Kloukinas and P. Moreira, *E-link: A Radiation-Hard Low-Power Electrical Link for Chip-to-Chip Communication*, in the proceedings of the *Topical Workshop on Electronics for Particle Physics*, Paris, France, 21–25 September 2009 [DOI:10.5170/CERN-2009-006.422].

[7] P. Moreira et al., *lpGBT documentation: release*, (2024), https://cds.cern.ch/record/2809058.

[8] J. Hegeman et al., *Design and development of the DAQ and Timing Hub for CMS Phase-2*, *PoS* **TWEPP2018** (2019) 129.

[9] E. van der Bij, R.A. McLaren, O. Boyle and G. Rubin, *S-LINK, a data link interface specification for the LHC era*, *IEEE Trans. Nucl. Sci.* **44** (1997) 398.

[10] C. Ghabrous Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, 2015 *JINST* **10** C02019.

[11] M. Vojinović, *CMS HGCAL electronics vertical integration system tests*, 2024 *JINST* **19** C02022.

2025 JINST 20 C03028