DEVELOPMENT AND VALIDATION OF FPGA-BASED AUTORECOVERY LOGIC IN THE CMS HGCAL DATA ACQUISITION SYSTEM

by

Rauf Ozan Akpınar

B.S., Physics, Boğaziçi University, 2022

Submitted to the Institute for Graduate Studies in
Science and Engineering in partial fulfilment of
the requirements for the degree of
Master of Science

Graduate Program in Physics

Boğaziçi University

2025

ACKNOWLEDGEMENTS

I gratefully acknowledge the financial support that made this work possible. This research was funded by:

- Turkish Energy, Nuclear and Mineral Research Agency (TENMAK) grant "HGCAL Dedektöründe Elektronik Sistemlerin Geliştirilmesi, Yumuşak-Kümelenmemiş Enerji Örüntüleri ve HF Çevrimiçi Radyasyon Hasar Sisteminde Çalışmalar";
- The Scientific and Technological Research Council of Türkiye (TÜBİTAK) 2232-A International Fellowship for Outstanding Researchers, project ID 118C223;
- TÜBİTAK 2210/E National MSc/MA Scholarship Program for Former Undergraduate Scholars.

I owe special thanks to Alp Şarkışla for his support and guidance, and to my collaborator Berkay Yazıcı.

ABSTRACT

DEVELOPMENT AND VALIDATION OF FPGA-BASED AUTORECOVERY LOGIC IN THE CMS HGCAL DATA ACQUISITION SYSTEM

This thesis presents the design, implementation and validation of the FPGA-based Autorecovery Block that safeguards the data acquisition (DAQ) chain of the High Granularity Calorimeter against faulty packets and link errors.

Operating at a clock frequency of 40 MHz, the Autorecovery Block monitors errors emitted by the ECON-D data concentrator ASICs. Four independent error counters track timeout, bunch crossing mismatch, event number mismatch, and buffer overflow conditions. When a counter hits its programmable threshold, the Autorecovery Block issues a recovery command.

Many of the block's static parameters were converted into field-configurable ones by expanding the DAQ firmware's existing IPbus infrastructure. The error priority list, thresholds, routing map, and individual ECON-D mute status are made configurable via control registers. Status registers, on the other hand, enable monitoring of a block's internal diagnostics and support testing.

The Autorecovery Block was validated on a Serenity carrier board equipped with an AMD Kintex™ UltraScale+™ KU15P Field Programmable Gate Array. The block's main functions were tested in its isolated form. The block acted on errors injected over 1024 consecutive clock cycles with 100% recovery command trigger rate. It was later embedded in a broader firmware structure and tested with emulated errors, achieving a 100% success rate again.

ÖZET

CMS HGCAL VERİ TOPLAMA SİSTEMİ'NDE FPGA TABANLI AUTORECOVERY MANTIĞININ GELİŞTİRİLMESİ VE GEÇERLENMESİ

Bu tez FPGA tabanlı Autorecovery Bloğu'nun tasarımı, gerçeklenmesi ve geçerlenmesini kapsamaktadır. Autorecovery Bloğu, Yüksek Granülerlikli Kalorimetre veri toplama sisteminde hatalı paketler veya bağlantı problemleriyle karşılaşıldığında harekete geçer.

40 MHz saat frekansında çalışan blok, ECON-D veri yoğunlaştırıcı uygulamaya özel tümdevrelerinden gelen hataların takibini yapar. Dört adet birbirinden bağımsız sayaç ardışık; zaman aşımı, öbek kesişimi numarası uyumsuzluğu, olay numarası uyumsuzluğu ve arabellek taşması hatalarını sayar. Sayaçlardan biri programlanan eşiği geçtiğinde blok ilgili kurtarma komutunu gönderir.

Veri toplama sistemi belleniminin IPbus altyapısı, Autorecovery Bloğu da içine alacak şekilde genişletilmiş ve bloğun sahip olduğu çoğu statik parametre dinamik hale getirilmiştir. Hata önceliği listesi, komut eşikleri, komut yönlendirme haritası ve kurtarma komutu baskılama işlevi IPbus kontrol yazmaçları ile ayarlanabilir hale getirilmiştir. IPbus durum yazmaçları ile de bloğun iç değerlerinin takibi sağlanmıştır.

Blok, Serenity taşıyıcı kartına yerleştirilmiş AMD Kintex™ UltraScale+™ KU15P alanda programlanabilir kapı dizisinde geçerlenmiştir. Bloğun temel fonksiyonları izole haliyle test edilmiştir. 1024 saat döngüsü boyunca enjekte edilen hatalar sonucunda %100 başarı oranı ile kurtarma komutları tetiklemiştir. Blok daha sonra ana sistem yazılımına entegre edilerek, emüle edilmiş hatalarla %100 başarı oranıyla geçerlenmiştir.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF SYMBOLS

n_{eq}	1 MeV equivalent neutron fluence
X_0	Radiation length
β^*	Amplitude-function value at the IP
η	Pseudorapidity
λ	Hadronic interaction length
ϕ	Azimuthal angle in collider coordinates

LIST OF ACRONYMS/ABBREVIATIONS

ADC	Analog-to-Digital Converter
ALEPH	Apparatus for LEP Physics
ALICE	A Large Ion Collider Experiment
ASIC	Application-Specific Integrated Circuit
ATCA	Advanced Telecommunications Computing Architecture
ATLAS	A Toroidal LHC Apparatus
BCMS	Batch Compression Merging and Splitting scheme
BCR	Bunch Counter Reset
BE	Back-End
BU	Builder Unit
BX	Bunch Crossing
CE	Calorimeter Endcap, HGCAL
CE-E	CE Electromagnetic Section
CE-H	CE Hadronic Section
CERN	Conseil Européen pour la Recherche Nucléaire
CKM	Cabibbo-Kobayashi-Maskawa matrix
CMS	Compact Muon Solenoid
CMSSW	CMS Software
CP	Charge Conjugation Parity
CRC	Cyclic Redundancy Check
CSC	Cathode Strip Chamber
DAQ	Data Acquisition
DAQ800	Data Acquisition 800 board
DELPHI	Detector with Lepton, Proton and Hadron Identification
DTH	DAQ and TCDS2 Hub
DTH400	DAQ and TCDS2 Hub 400 board
DT	Drift Tube
EB	Electromagnetic Barrel

ECAL	Electromagnetic Calorimeter
EE	Electromagnetic Endcap
EMP	Extensible, Modular (data) Processor
ES	Electromagnetic Preshower
FASER	Forward Search Experiment
FC	Fast Command
FCC	Fast Command Candidate
FE	Front-End
FED	Front-End Driver
FIFO	First-In, First-Out
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
GBT-SCA	Gigabit Transceiver Slow Control Adapter
GE	GEM Endcap
GEM	Gas Electron Multiplier
GPU	Graphics Processing Unit
HB	Hadron Barrel
HCAL	Hadron Calorimeter
HD	High-Density
HE	Hadron Endcap
HF	Hadron Forward
HGCAL	High Granularity Calorimeter
HL-LHC	High-Luminosity LHC
HLT	High-Level Trigger
HO	Hadron Outer
I2C	Inter-Integrated Circuit
IP	Interaction Point
IPBB	IPbus Builder
IPbus	IPbus protocol
IR	Insertion Region
ISR	Intersecting Storage Rings

L1A	Level-1 Accept
L1T	Level-1 Trigger
L3	L3 experiment
LHC	Large Hadron Collider
LHCb	LHC-beauty
LED	Light-Emitting Diode
LEP	Large Electron-Positron Collider
Linac4	Linear Accelerator 4
lpGBT	low-power GigaBit Transceiver
LS	Long Shutdown
LSS	Long Straight Section
LD	Low-Density
MIP	Minimum Ionizing Particle
MB	Muon Barrel
ME	Muon Endcap
MTD	MIP Timing Detector
OPAL	Omni-Purpose Apparatus for LEP
OCR	Orbit Counter Reset
PF	Particle-Flow
PCB	Printed Circuit Board
PLL	Phase-Locked Loop
PS	Proton Synchrotron
PSB	Proton Synchrotron Booster
pTT	partial Trigger Tower
QGP	Quark-Gluon Plasma
RAM	Random Access Memory
RB	RPC Barrel
RE	RPC Endcap
RF	Radio-Frequency
RPC	Resistive Plate Chamber
RU	Readout Unit

SC	Synchrocyclotron
SiPM	Silicon Photomultiplier
SND@LHC	Scattering and Neutrino Detector at the LHC
SPS	Super Proton Synchrotron
TC	Trigger Cell
TDC	Time-to-Digital Converter
TCDS2	Phase-2 Trigger and Timing Control and Distribution System
ToA	Time-of-Arrival
ToT	Time-over-Threshold
TPG	Trigger Primitive Generator
TS	Technical Stop
TTC	Trigger, Timing and Control
UA-1	Underground Area-1 experiment
UA-2	Underground Area-2 experiment
USC	Underground Service Cavern
VHDL	Very High Speed Integrated Circuit Hardware Description Language
VTRx+	Versatile Transceiver Plus
YETS	Year-End Technical Stop

1. INTRODUCTION

1.1. CERN

After the Second World War, a small number of leading European scientist envisioned a collaboration to revitalize and unite European science that had fallen behind the US and USSR. The European Organization for Nuclear Research, known as CERN (in French, Conseil Européen pour la Recherche Nucléaire), was established in 1954. Having begun with 12 founding states: Belgium, Denmark, France, the Federal Republic of Germany, Greece, Italy, the Netherlands, Norway, Sweden, Switzerland, the United Kingdom, and Yugoslavia; CERN has since grown to include 25 member states, as well as numerous associate and observer states.

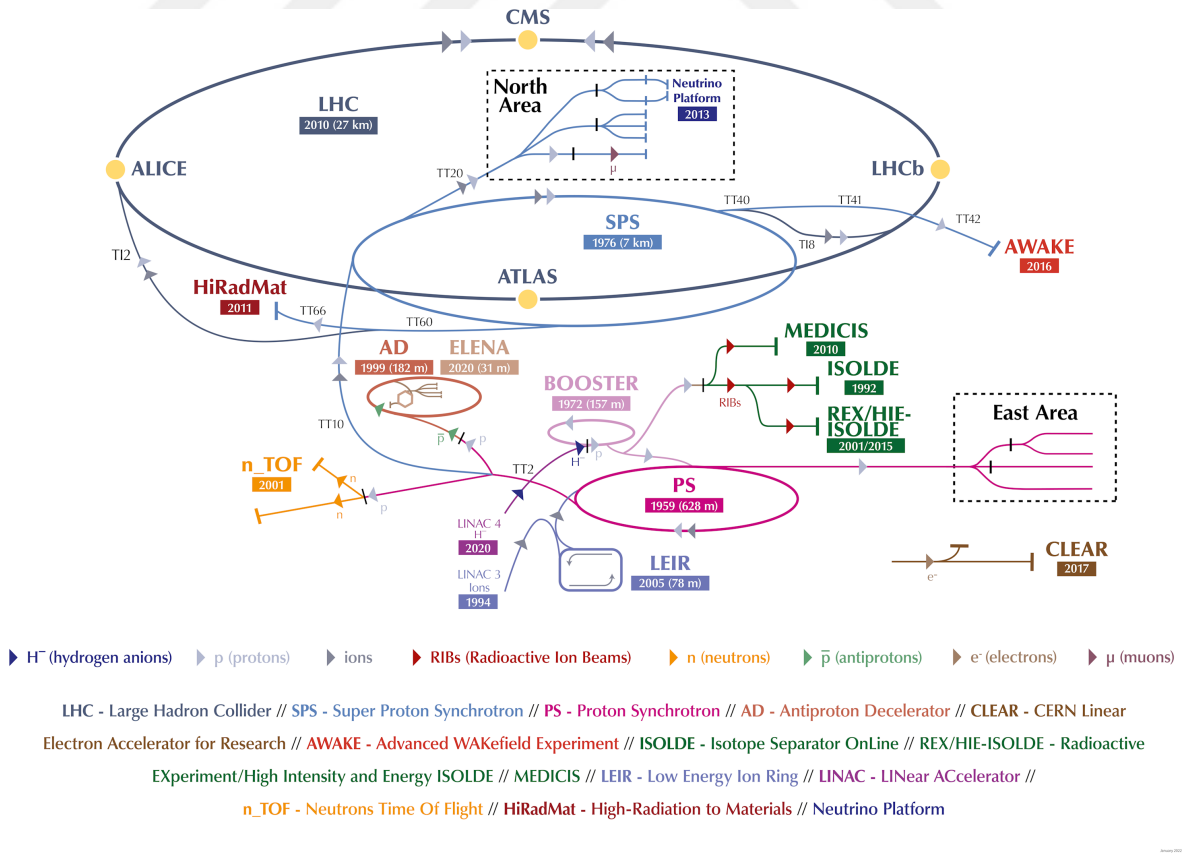

Figure 1.1. CERN accelerator complex. Image from [1] (© 2022-2025 CERN; CC BY 4.0).

1957 marked the year first accelerator came into service at CERN. The Synchro-cyclotron (SC), then third largest in the world, accelerated protons up to 600 MeV. Shortly after beginning operations, the SC demonstrated pion decay to electron and made its first major contribution to nuclear physics. The SC was shut down after 33 years in service.

In 1959, the Proton Synchrotron (PS) entered into service, it was capable of accelerating protons up to an energy of 28 GeV. In its sixth year of operation, the PS beam was targeted at a beryllium target to produce the first antinuclei to be observed, antideuterons [2]. CERN's first synchrotron and briefly the highest energy accelerator in the world, the PS is still operational as a part of the CERN accelerator complex, shown in Figure 1.1. The PS beam was delivered to various facilities including the Antimatter Factory, irradiation test facilities and the Large Hadron Collider (LHC).

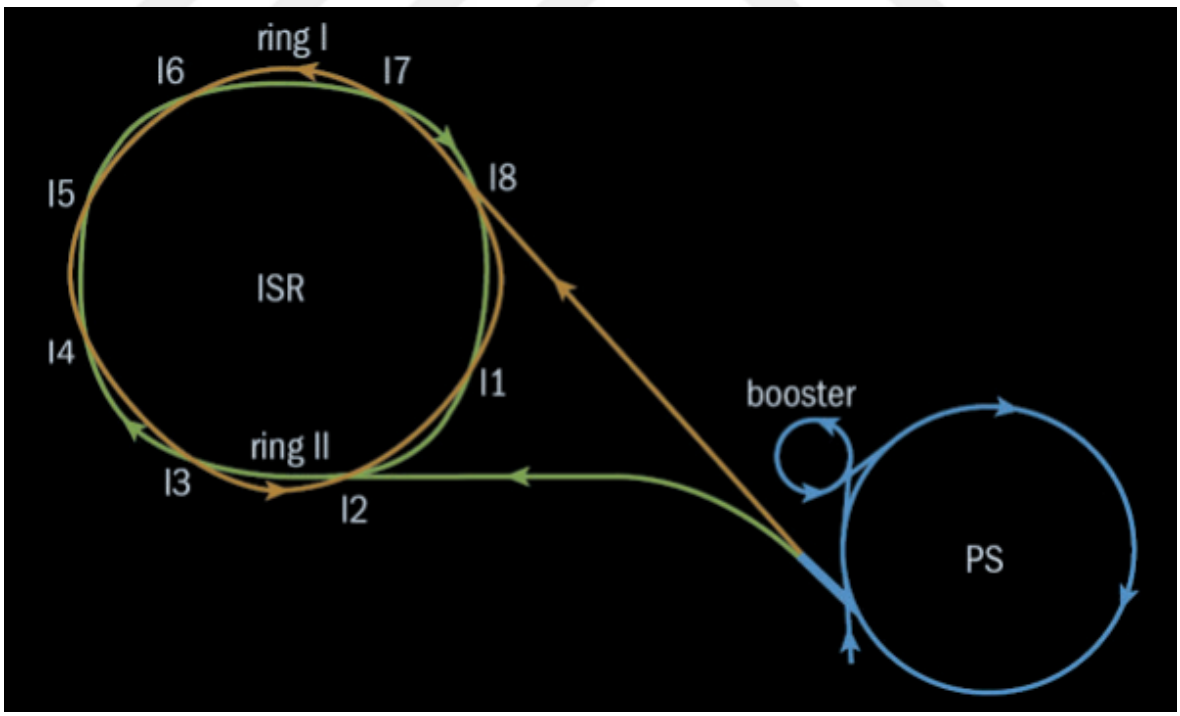

Figure 1.2. Schematic layout of the ISR, its transfer lines and its injector, the PS. Image from [3] (© 2012-2025 CERN; CC BY 3.0).

Fixed target experiments can only use a fraction of beam energy to create new particles, whereas a collider could fully utilize the combined energy of both beams.

In 1965, CERN Council green-lit the Intersecting Storage Rings (ISR) [3], a proton storage ring that would be fed by the PS, shown in Figure 1.2. The ISR was a bold proposition, the first of its kind, a proton-proton collider at a time when only lepton colliders had been built. Delivering its first proton-proton collision in 1971, the ISR remained in operation for 13 years. It held the world record for peak luminosity for another 7 years following end of its service in 1984 [4]. The ISR proved to be a large stepping stone, generating new ideas and identifying challenges in design and theory of future colliders.

The second underground ring, after the ISR, was commissioned in 1976: the Super Proton Synchrotron (SPS). The new 400 GeV accelerator was 10 times larger than the PS, which would be used as a proton injector for the SPS. It ran in fixed-target mode until 1981, when it was converted to also serve as a proton-antiproton collider. Two years after the first proton-antiproton collision at CERN's new workhorse, a major discovery was recorded. W and Z bosons were observed for the first time by Underground Area-1 (UA-1) and Underground Area-2 (UA-2) experiments at the SPS. This achievement led to Carlo Rubbia and Simon van der Meer winning the 1984 Nobel Prize in Physics[5] for their substantial contributions to the project.

Situated tens of meters underground, beneath the Franco-Swiss border, the Large Electron-Positron Collider (LEP) accelerator ring was the largest collider in the world with its 27 km circumference. The LEP started operation in 1989 with a center-of-mass energy of 91 GeV and later reached 209 GeV. The LEP is still unsurpassed as the largest and highest-energy electron-positron collider. Four detectors, Apparatus for LEP Physics (ALEPH), Detector with Lepton, Proton and Hadron Identification (DELPHI), Omni-Purpose Apparatus for LEP (OPAL), and L3, observed the collisions at the LEP. Contributions of the LEP experiments include: precise measurements of W and Z bosons, showing there are only three generations of matter and showing exclusion of Higgs mass below 114.4 GeV. In the LEP's final year of service, the ALEPH experiment reported a possible Higgs event with Higgs mass in the range of 114 GeV to 115 GeV. This prediction was off only by a margin, 12 years later, Higgs mass was discovered to

be in the range of 125 GeV to 126 GeV at the LHC [6].

1.2. Large Hadron Collider

The LHC was built in place of the LEP, taking over the title of the world's largest particle accelerator. The LHC began its first physics data-taking period (Run 1) in 2010 colliding protons at 3.5 TeV energy per beam [7]. Despite its main physics program of proton-proton collisions, the LHC is also designed to collide heavy-ions. Products of these collisions are recorded by four main detectors: A Toroidal LHC Apparatus (ATLAS), Compact Muon Solenoid (CMS), LHC-beauty (LHCb) and A Large Ion Collider Experiment (ALICE); along with several smaller experiments, including the two new additions in 2022 [8]: Scattering and Neutrino Detector at the LHC (SND@LHC) and the Forward Search Experiment (FASER).

Protons are accelerated to 6.8 TeV in the LHC, resulting in a total collision energy of 13.6 TeV, the highest ever achieved in a laboratory. Before being injected to the LHC ring, protons are gradually brought to higher energies through a chain of accelerators, as shown in Figure 1.1.

Initially, H^- ions are accelerated to 160 MeV in the Linear accelerator 4 (Linac4). Two electrons are stripped from the ions converting them to protons before their injection into the Proton Synchrotron Booster (PSB), a system of four superimposed synchrotron rings. In Batch Compression Merging and Splitting (BCMS) scheme, protons are accelerated to 2 GeV in the PSB and injected into the PS in 8 bunches (4+4). The bunches are merged, split into 12, and then accelerated to 25 GeV in the PS before getting further split into 48. Five trains of 48 bunches are injected into the SPS, where they are accelerated to 450 GeV. Finally, the protons are injected from the SPS into the LHC beams. See Ref. [9] for more details.

The LHC ring comprises eight Long Straight Sections (LSSs), and eight arcs, as shown in Figure 1.3. The Interaction Points (IPs) are located on the LSS, these points

and their vicinity are where collisions, beam injections, beam dumps, and other beam interactions occur. An IP and its vicinity are called the Insertion Region (IR). The beams are injected near the IP2 and IP8, where the ALICE and LHCb experimental caverns are located, respectively. The particles are accelerated to higher energies at the radiofrequency (RF) cavities located in the IR4. Collimation systems are located in the IR3 and IR7; assisting with control of the background effects at the detectors, protecting the sensitive parts of the collider from beam loss and cleaning the circulating beam off of betatron halo and off-momentum particles. See [10, 11] for further details.

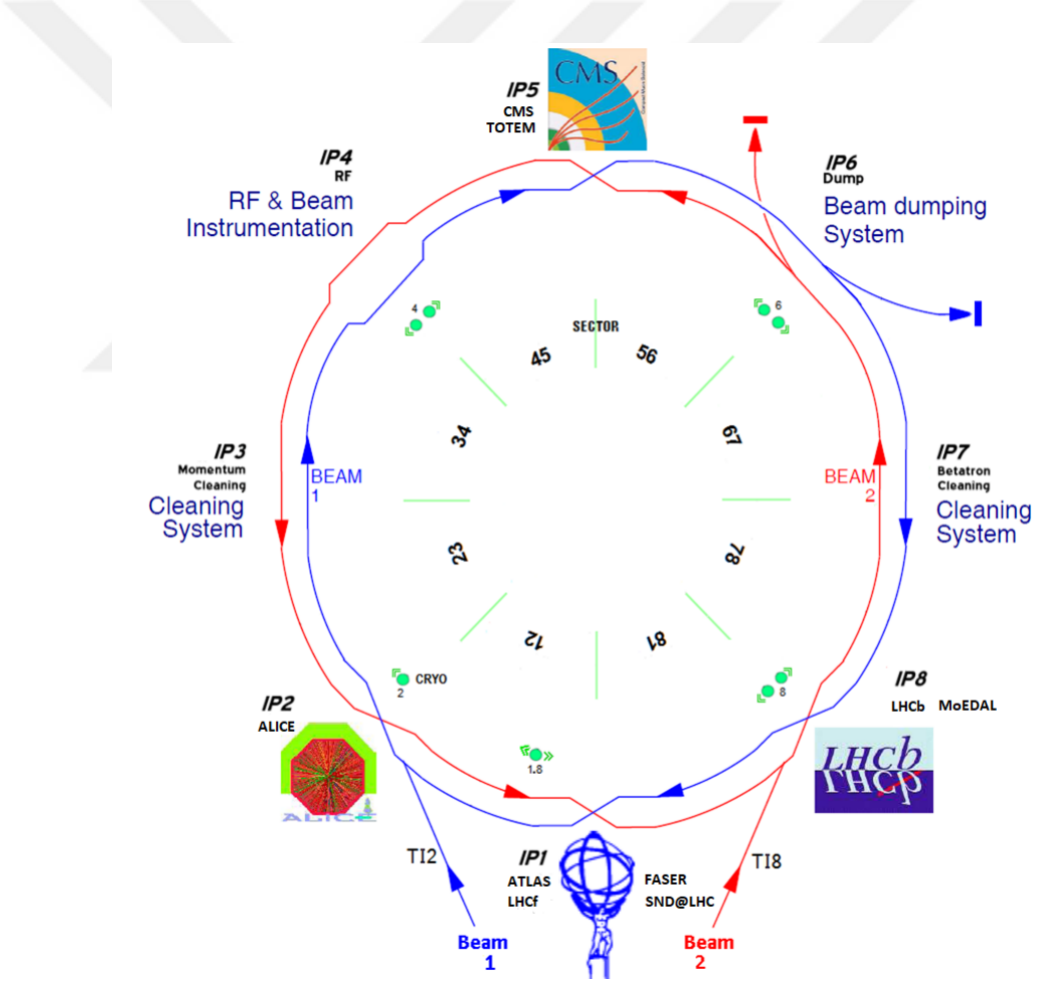

Figure 1.3. Schematic layout of the LHC. The collider is divided into eight sectors with beam interactions occurring in the straight sections of each. Experiments and the LHC systems are positioned at these interaction points. Image from [9] (© 2024 The Author(s); CC BY 4.0)

The LHC's dipole magnets, shown in Figure 1.4, bend the beams while quadrupole magnets form a focusing-defocusing lattice to maintain beam stability and focus. Each section contains 154 twin-aperture superconducting main dipole magnets powered in series. A section is comprised of an arc and its two adjacent half LSS. The magnets are operated at a temperature of 1.9 K and have a length of 14.3 m. Separate beams pass through their twin apertures spaced 194 mm apart. A total of 1232 main dipole magnets and 392 main quadrupole magnets [9] are used in the LHC along with thousands of other magnets of various types.

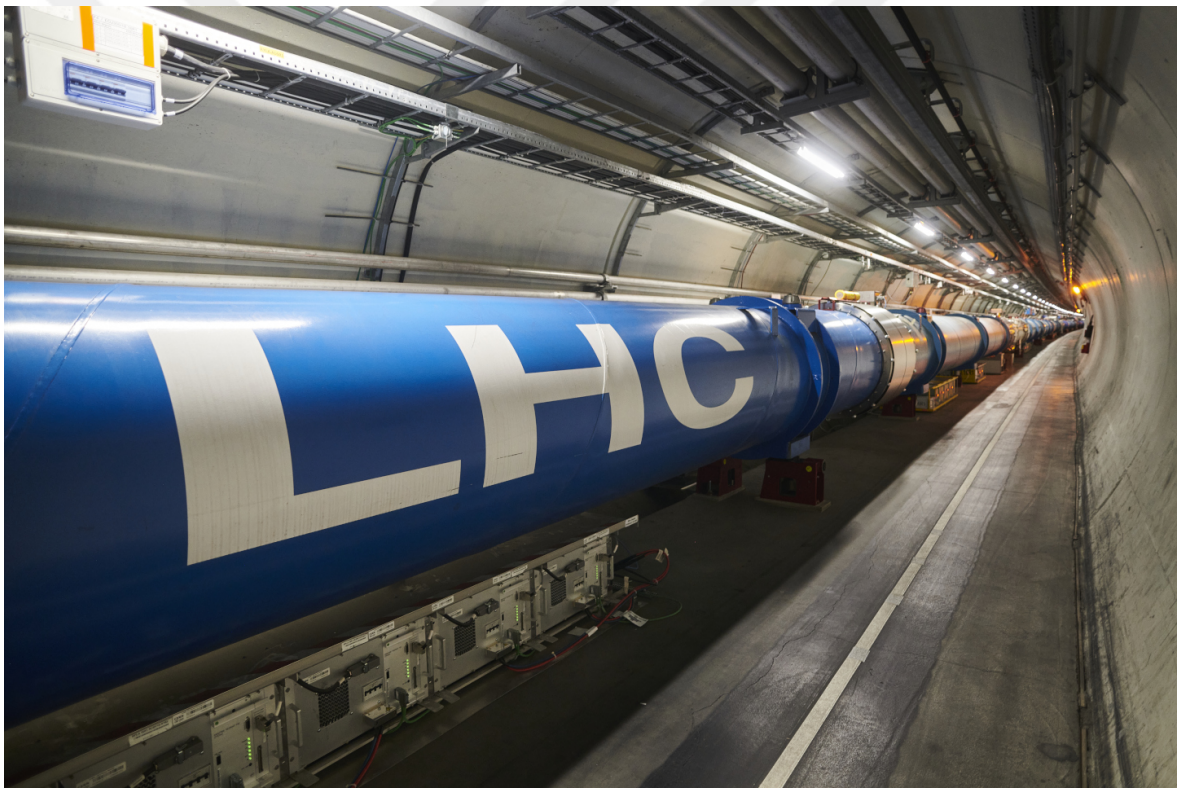

Figure 1.4. LHC dipole magnets in the tunnel. Image from [12] (© 2021-2025 CERN; CC BY 4.0).

To conclude the operation, energy of the beams are absorbed by the beam dump blocks located in the IP6, ensuring a safe and reliable termination of the beams. The operation cycle starting from the injection to the dump is called a fill. Based on inspection of the LHC beam performance parameters from 2024 [13], a fill with a beam energy of 6.8 TeV has a median bunch count of 2352 per beam. Spacing between the

bunches is 25 ns with a median bunch population of 1.59×10^{11} protons. Compared to the number of particles in a bunch, the actual collision rate is very low; in 2024 the average number of interactions per crossing (pileup) at the CMS was 57 [14]. After a crossing, remaining beams keep circulating in the LHC.

Two LHC beams cross in four interaction points housing the main LHC detectors: ATLAS, CMS, LHCb and ALICE. Located at diametrically opposite points on the LHC ring, shown in Figure 1.3, ATLAS and CMS are the two general purpose experiments and are allocated the highest amount of integrated luminosity. These experiments are essential for both the Standard Model measurements and searches for beyond-standard model physics. They enable increasingly precise measurements of the Standard Model physics objects, and rare processes.

The LHCb experiment is specialized in studying the matter-antimatter asymmetry problem in systems involving b- and c-quarks. It performs high-precision measurements of Cabibbo-Kobayash-Maskawa (CKM) quark mixing matrix parameters and of charge conjugation parity (CP) symmetry violation.

The ALICE experiment is dedicated to heavy-ion physics. High-energy lead-lead collisions at the LHC can create quark-gluon plasma (QGP); an extremely hot and dense state of matter believed to have filled the early universe for a few microseconds. ALICE studies the QGP to uncover its properties and to test quantum chromodynamics in the strong-interaction regime.

LHC operates in cycles alternating between data-taking and maintenance periods. A Run refers to a multi-year data-taking phase; so far, Run 1 (2010-2013) and Run 2 (2015-2018) have been completed. The accelerator complex is currently in Run 3, which began in 2022 and scheduled to continue until 2026 [15]. Maintenance and safety tests are carried out during short Technical Stops (TSs) and the Year-End Technical Stops (YETSs). At the end of each run, the complex enters a Long Shutdown (LS) period, during which major consolidation and upgrade projects are undertaken. Following Run

3, LS3 will accommodate the transition to a new era of the LHC: High-Luminosity LHC (HL-LHC).

1.2.1. High-Luminosity LHC

A particle collider is characterized by its luminosity and center-of-mass energy [11]. Luminosity determines the rate of collision events, while center-of-mass energy determines the available energy for the production of new particles.

During the proton operations of Run 1 (2010-2013), the LHC delivered a total integrated luminosity of $29\,\mathrm{fb}^{-1}$ to ATLAS and CMS at center-of-mass energies of 7 to 8 TeV [9]. Following Run 1 and the first Long Shutdown (LS1), Run 2 (2015-2018) began, achieving a center-of-mass energy of 13 TeV and yielding $160\,\mathrm{fb}^{-1}$ of total integrated luminosity at the ATLAS and CMS experiments [9].

During the LS2, LHC injectors were upgraded and various components, such as the superconducting magnets, were consolidated in preparation for Run 3 (2022-2026) [9]. As of the end of 2024, Run 3 is ongoing and the LHC has delivered an integrated luminosity of $195\,\mathrm{fb}^{-1}$ to the ATLAS and CMS experiments [14, 16], as shown in Figure 1.5(a). The total integrated luminosity of the Run 3 is predicted to reach $290\,\mathrm{fb}^{-1}$ by its conclusion in 2026.

The High-Luminosity upgrade entails the replacement of inner triplet quadrupoles. These magnets serve as the final focusing systems before collisions. The new triplet magnets will reduce the transverse size of the beam at the collision point, enabling a low β^*, thereby increasing the luminosity. However, new optics will require a larger crossing angle to maintain sufficient beam separation, which introduces a loss in luminosity. To compensate for this, a novel system, RF crab cavities, will be installed in IR1 and IR5, which will almost completely mitigate this drawback.

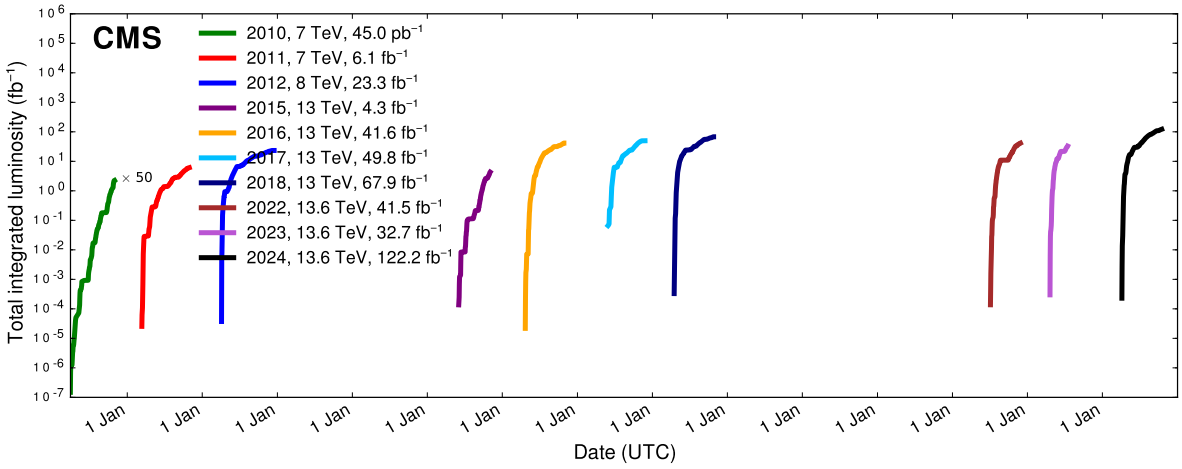

(a)

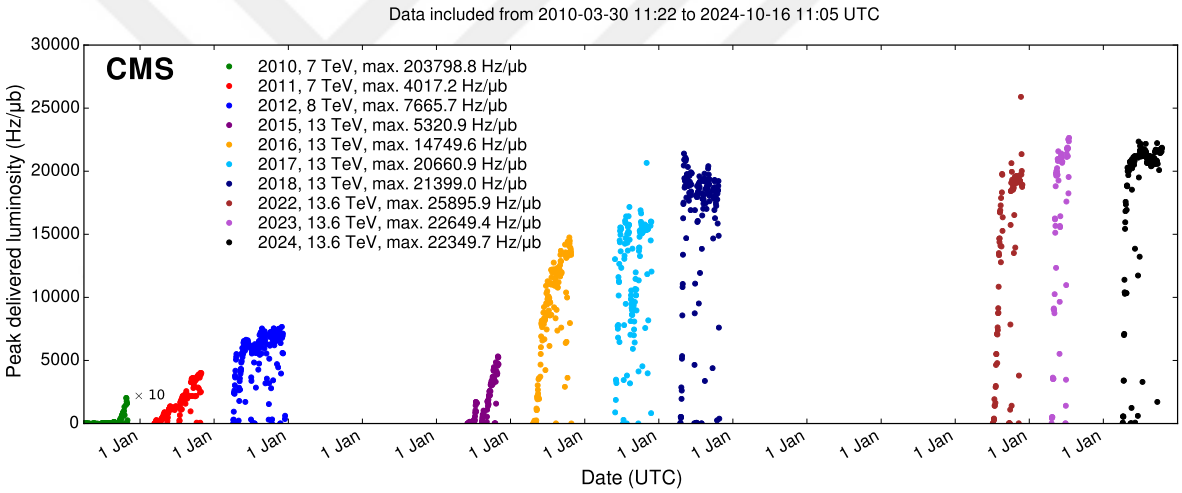

(b)

Figure 1.5. Luminosity delivered to CMS during stable beams for pp collisions at nominal center-of-mass energy: (a) cumulative and (b) peak values. Images from [14] (Public domain).

The HL-LHC project was established in 2010 [10]. Since then, the LHC components have been designed with an overhead. The HL-LHC, through leveled luminosity operation and increased bunch population, will push the machine to its engineering limits towards reaching the ultimate performance goal: a peak leveled luminosity of $75\,000\,\mathrm{Hz}\,\mu\mathrm{b}^{-1}$, resulting in an average pileup of about 200, and a yearly integrated luminosity of $300\,\mathrm{fb}^{-1}$ to $400\,\mathrm{fb}^{-1}$ [10]. These numbers present substantial improvements

compared to the current luminosity rates, shown in Figure 1.5.

1.3. Compact Muon Solenoid

CMS is a 22 m-long, 15 m-wide (diameter), multipurpose detector located in the underground experimental cavern at IP5. The full assembly including the forward calorimeters and their shielding extends to 28.7 m.

A 3.8 T superconducting solenoid of 6 m internal diameter surrounds the inner volume, enclosing the silicon trackers, the electromagnetic calorimeter (ECAL), and the hadron calorimeter (HCAL). Each of these detectors has a barrel section and two endcap sections. A steel flux-return yoke surrounding the solenoid accommodates the CMS muon system between its layers. A diagram of the detector is shown in Figure 1.6. For a more detailed description of the detector see [17].

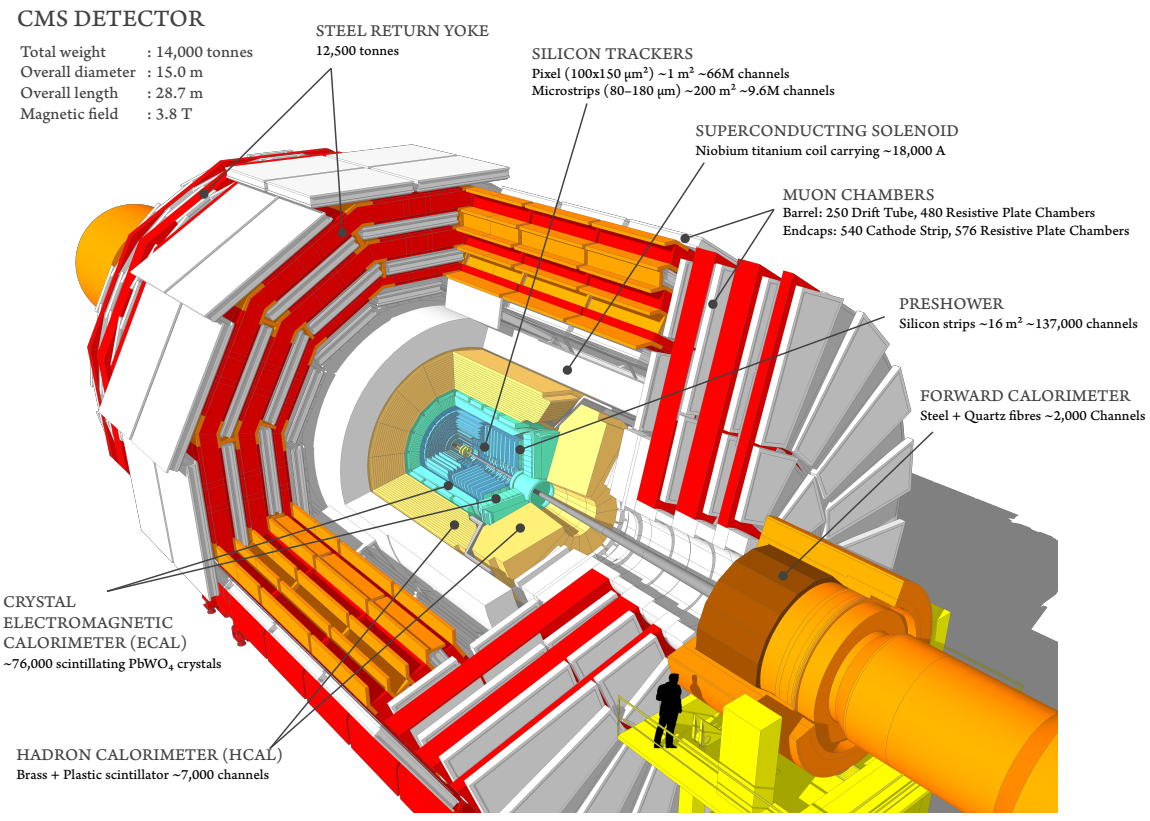

Figure 1.6. Cutaway diagram the CMS detector. Image from [18] (CC BY 3.0).

The CMS detector can detect electrons, photons, muons, and hadrons with its dedicated sub-detectors. The particles interacting with the CMS sub-detectors can be seen in Figure 1.7. The innermost part of the detector is the silicon tracker, which delivers excellent momentum resolution and tracking efficiency. The tracker is divided into two sub-detectors: the inner tracker comprising silicon pixel sensors and the outer tracker comprising silicon strip sensors. The more compact, roughly 1 m long pixel detector sits closest to the IP, where the particle flux is higher. Surrounding it is the strip detector, which extends about 5 m along the beam axis. Both sub-detectors comprise barrel layers and forward disks.

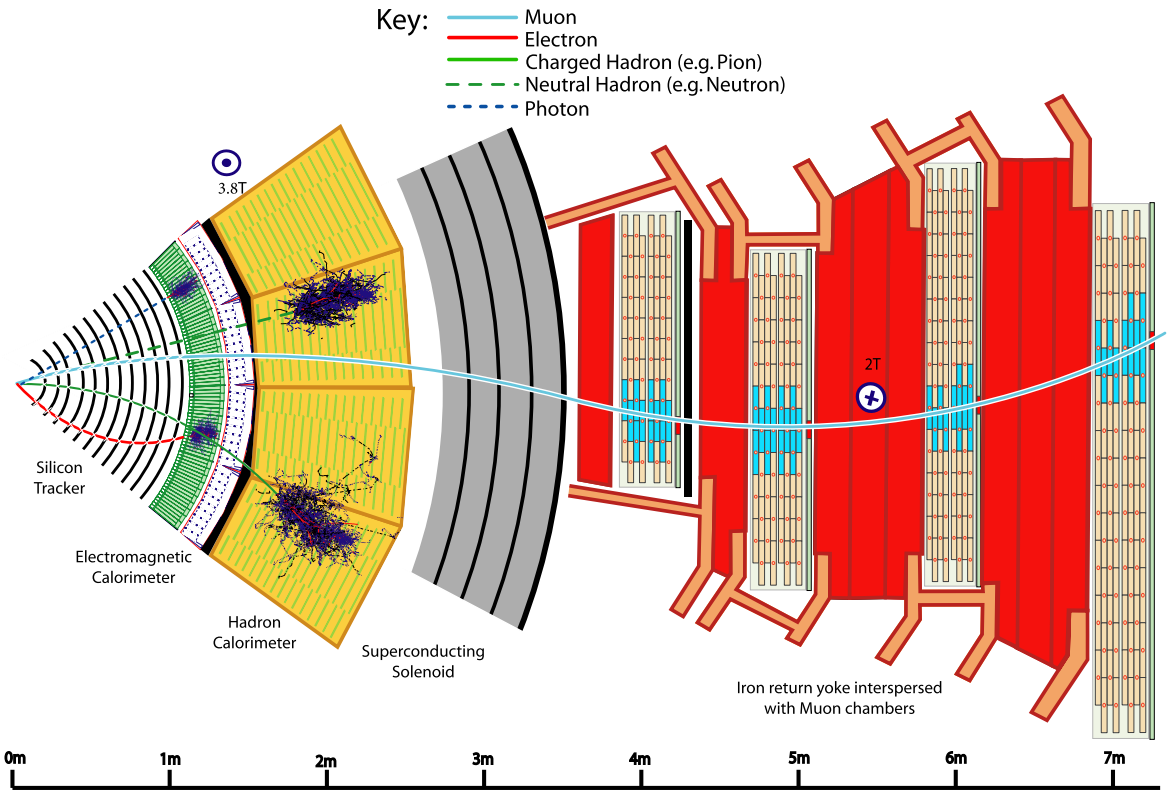

Figure 1.7. Sketch of a transverse slice of the CMS detector barrel section, different particles can be seen interacting with the detector. Image adapted from [19] (© 2017 CERN; CC BY 4.0).

ECAL is installed outside the silicon tracker. It comprises three parts: barrel (EB), endcap (EE) and preshower (ES) detectors, ECAL provides measurements of energy, impact position, and time of arrival for photons and electrons. EB and EE are

collectively made of more than 70 000 scintillating lead tungstate crystals with a short radiation length[1] . ES, installed in front of each EE module, is a sampling calorimeter with two layers: a lead radiator layer and a layer of silicon strip sensors. The main task of ES is to identify neutral pions.

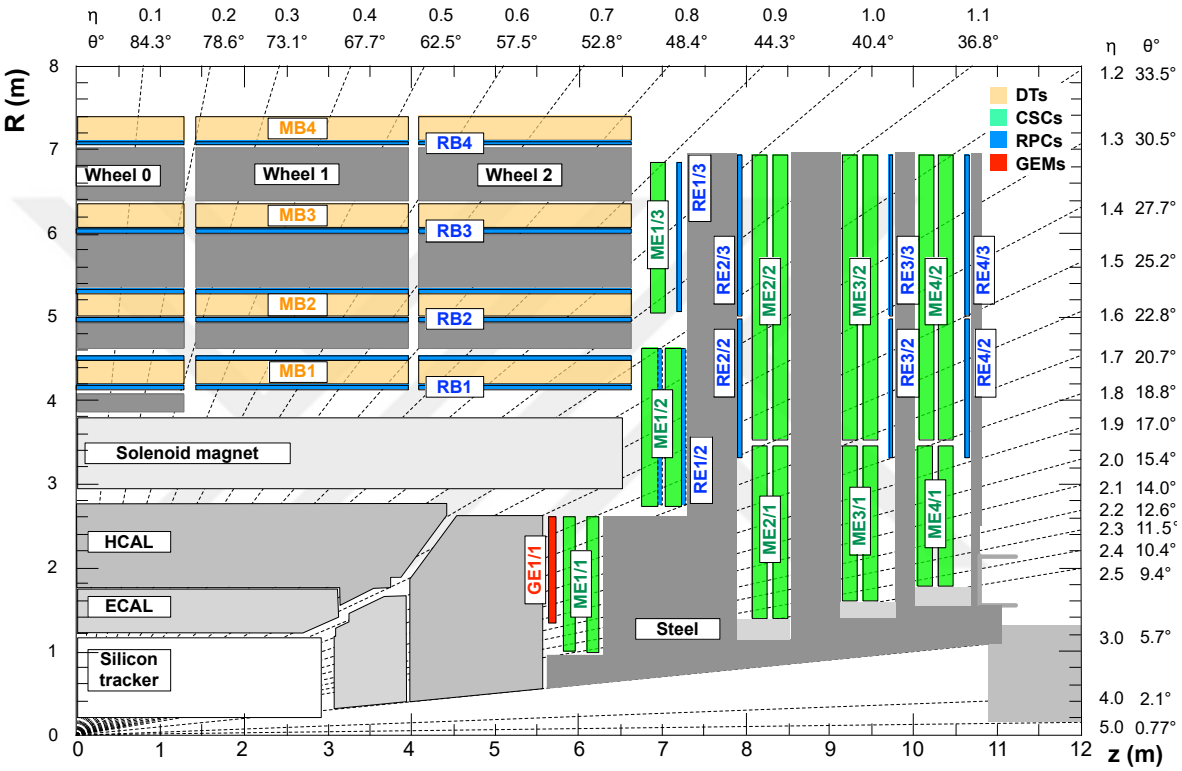

Figure 1.8. Cross section of the CMS detector quadrant. Drift tube stations (DTs) are labeled MB (Muon Barrel), cathode strip chamber (CSCs) are labeled ME (Muon Endcap), resistive plate chambers are labeled RB and RE (barrel and endcap sections), and gas electron multiplier (GEMs) are labeled GE (endcap). Image from [17] (© 2024 CERN for the benefit of the CMS Collaboration; CC BY 4.0).

HCAL measures the energy of hadrons contributing to their identification and measurement of properties. Two of its sub-detectors, hadron barrel (HB) and hadron endcap (HE) calorimeters, are enclosed by the superconducting solenoid. HB and HE are sampling calorimeters, alternating between brass absorbers and plastic scintillating tiles. Several sequential layers are grouped into depth segments. Optical signals from

[1]Mean thickness of a material over which an electron loses $1/e$ of its energy traversing.

scintillators of the same depth, pseudorapidity[2] , and azimuthal angle coordinates are sent to the same silicon photomultiplier (SiPM). HB and HE have a total of 16 and 17 scintillator layers, respectively. To measure the energy of particles reaching beyond HB and to identify late-starting showers, the hadronic outer (HO) calorimeter is installed outside the solenoid. HO is a sub-detector of HCAL with 1 to 2 layers of scintillators. The last HCAL sub-detector is the hadron forward (HF) calorimeter. HF is situated 11.5 m from the Interaction Point on either side of the detector, it covers the pseudorapidity range of $3.0 < |\eta| < 5.2$ [20].

A key component of the CMS detector is its sophisticated muon system, shown in Figure 1.8, which is used for identifying muons and measuring their momenta. Muon detectors are installed outside the solenoid, between steel plates of the return-flux yoke. Due to their outward position it is harder for particles other than muons and neutrinos to reach the muon detectors. The muon system employs four different detector technologies: drift tubes (DTs), cathode strip chambers (CSCs), resistive plate chambers (RPCs), and gas electron multipliers (GEMs).

The DTs are located in the barrel region, a DT chamber houses a matrix of 4 cm drift cells. In the endcap region where particle flux is higher, the CSCs are used. In a CSC, thin anode wires are stretched parallel at a 90° angle to the radially extending cathode strips. The strips divide each chamber into six volumes that are filled with a gas mixture. Traversing charged particles ionize the gas mixture and freed electrons are collected at the wires. CSCs provide spatial measurements with a shorter response time than DTs. The RPCs are interleaved between DT layers in the barrel and CSC layers in the endcap region. An RPC has two gas chambers and a readout strip in the middle for readout. Due to their fast response time, RPCs are able to provide precise timing measurements crucial for bunch crossing identification. The triple GEM subsystem was installed to enhance the muon tracking and identification performance in the endcap region, in preparation for Run 3. The GEM assembly has four gas gaps

[2]Pseudorapitidy, η, is an approximate rapidity measure defined as $\eta \equiv -\ln \tan \frac{\theta}{2}$ where θ is the angle of the particle with respect to beam axis. ($\eta_{0^\circ} = \infty$, $\eta_{5.7^\circ} = 3$)

with drift-transfer-transfer-induction layout. A drift cathode and a readout printed circuit board (PCB) with three GEMs in-between comprise the triple GEM detector.

Nominal bunch spacing of the proton-proton operation at the LHC is 25 ns corresponding to a bunch crossing rate of 40 MHz. This rate is reduced to 100 kHz or below by the Level-1 Trigger System (L1T), composed of the muon trigger and the calorimeter trigger. The L1T runs on custom hardware located in the underground service cavern (USC), which also houses the back-end (BE) electronics of detector data acquisition (DAQ) systems electronics near the CMS detector. ECAL, HCAL, and HF send trigger primitives to the calorimeter trigger for every bunch crossing. A trigger primitive is a data word containing detector read-out information, a bunch-crossing tag and quality bits, all linked to a single trigger tower, which defines a detector patch with fixed position and geometry. Trigger towers from different detectors are grouped in the calorimeter trigger to form trigger tower sums that are used for reconstruction of physics object candidates.

The four muon sub-systems (DTs, CSCs, RPCs, and GEMs) generate trigger primitives for the muon track finders in the muon trigger. Candidates from the muon trigger and the calorimeter trigger are combined in the global trigger according to the trigger menu, a collection of several hundred algorithms used for physics object identification. The global trigger then issues the level-1 accept (L1A) signal.

The L1A signal triggers full-granularity readout of all sub-detectors. Data fragments are aggregated by the detector front-end drivers (FEDs) and combined into super-fragments by the readout unit (RU) servers. Full events are assembled at the L1T rate by the builder unit (BU) servers from the super-fragments. During 2024, the event building rate of the proton-proton operation was around 110 kHz; the High-Level Trigger (HLT) cuts this rate down to roughly 7 kHz [21]. HLT runs on graphics processing unit (GPU) accelerated computing clusters and reuses many of the reconstruction and selection algorithms used in CMS Offline Software (CMSSW) suite.

The events are then reconstructed into high-level physics objects like jets and photons by the offline[3] computing system which also facilitates detector simulation and physics analyses. The Tier-0 computing center at CERN takes two types of input streams from HLT: the standard streams where every event in the stream is reconstructed within two days, and the parking streams, where the events are buffered for later reconstruction. In practice, every event in the parking streams has been reconstructed during Run 3 [21]. An event visualization can be seen in Figure 1.9.

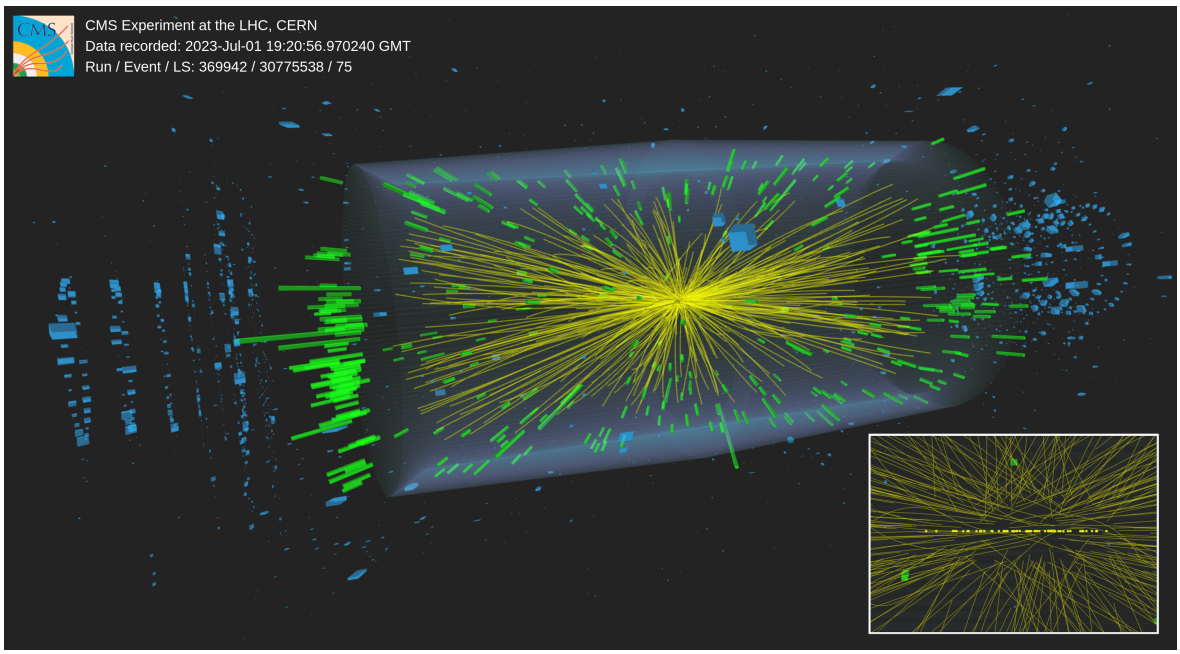

Figure 1.9. Visualization of a reconstructed event with 55 vertices. Tracks (yellow), calorimeter hits (green), and muon chamber hits (blue) can be seen. Image from [22] (© 2023 CERN, for the benefit of the CMS Collaboration).

1.3.1. Phase-2 Upgrade of the CMS Detector

The HL-LHC will deliver bunch crossings with an average pileup of 140 to 200 to CMS, which is a significant increase over 57, the average in 2024. These larger data samples will enable more precise measurements and observation of rarer processes. However, the upgrade presents two major challenges: increased radiation damage on

[3]Offline denotes processing that takes place after the data have left the detector and the DAQ system.

the detector components and the rising pileup.

Particles emerging from proton-proton collisions ionize the detector material and produce hadronic and electromagnetic cascades. These interactions liberate secondary particles, most prominently neutrons, which can travel long distances before capture and accumulate in the largely hermetic detector volume. The resulting neutron fluence significantly degrades silicon sensors, scintillators, and front-end (FE) electronics.

Extra energy deposited in the calorimeters and the additional tracks produced by the increased number of collisions can overwhelm the trigger and reconstruction algorithms. CMS relies on the Particle-Flow (PF) algorithm to reconstruct individual particles. The PF algorithm requires inputs from every sub-detector for optimal results. A finer sensor segmentation is necessary to differentiate between the increased number of incident particles and keep per-channel occupancy low. Coupled with the radiation damage already sustained, many of the CMS sub-systems will be upgraded or completely replaced as part of the Phase-2 Upgrade.

The silicon tracker will be completely rebuilt. The inner tracker will use thinner and smaller pixel sensors. The outer tracker will be equipped with pT modules, capable of transverse momentum filtering [23]. The tracker will contribute to the L1T for the first time.

Between the tracker and the calorimeters a new detector will be installed, the MIP Timing Detector (MTD). It will provide timing measurements with a resolution below 50 ps [24]. Its placement before the calorimeters permits MTD to play a crucial role distinguishing charged particle tracks before they produce showers, as minimum ionizing particles (MIPs) leave only a tiny, well-defined charge deposit.

A new endcap calorimeter, High Granularity Calorimeter (HGCAL), will replace endcap regions of HCAL and ECAL [25]. HGCAL is a sampling calorimeter comprising silicon sensors and plastic scintillators. It can provide precise measurements with

fine granularity and withstand higher radiation levels introduced by the HL-LHC upgrade. Barrel sections of HCAL and ECAL will remain mostly unchanged but new FE electronic boards will allow ECAL to achieve single-crystal granularity at the L1T.

Existing muon chambers will remain but five new stations will be built. Three new triple-GEM based detectors and two detectors with improved RPC (iRPC) technology will extend the coverage in the forward region [26].

The new L1T introduces a correlator layer, resembling HLT by employing PF reconstruction. The maximum L1T rate will be increased to 750 kHz from the current rate around 100 kHz and the latency will be extended to 12.5 µs [27]. These higher rates together with additional data channels, require replacement of most FE and BE electronics across the detector.

2. HIGH GRANULARITY CALORIMETER

As part of the Phase-2 Upgrade, endcap regions of ECAL and HCAL will be replaced with HGCAL. Currently installed endcap detectors were designed to perform efficiently up to only $500\,\mathrm{fb}^{-1}$ of integrated luminosity. HGCAL, spanning the pseudo-rapidity range of $1.5 < |\eta| < 3.0$, will preserve its high resolution in four dimensions beyond $3000\,\mathrm{fb}^{-1}$. It will operate at $-35\,^{\circ}\mathrm{C}$, keeping the noise low at high irradiation levels.

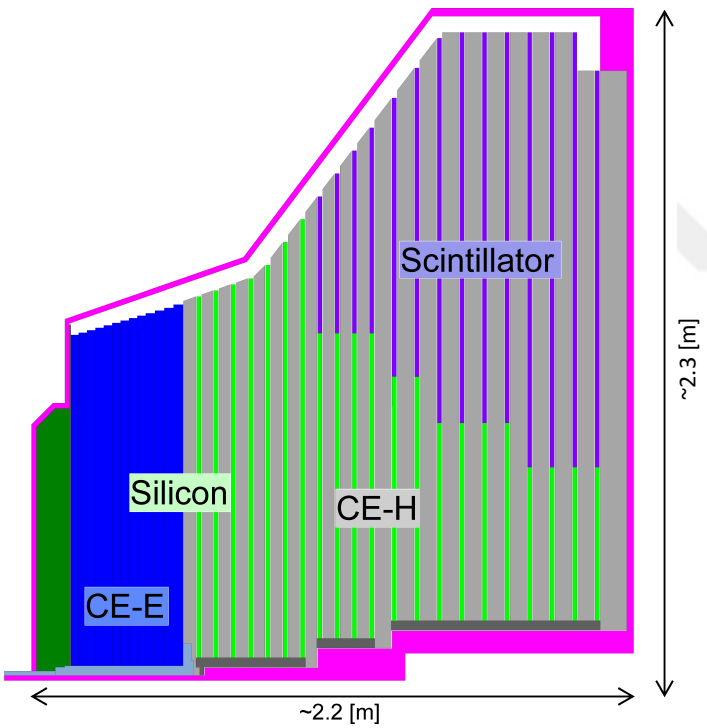

Figure 2.1. HGCAL cross-section schematic with region labels. HGCAL Schematics. Electromagnetic section (CE-E): 26 layers, $27.7X_0$ & $\sim 1.5\lambda$. Hadronic section (CE-H): 21 layers, $\sim 8.5\lambda$. Image from [28], reproduced with permission.

HGCAL is a sampling calorimeter with 47 layers, shown in Figure 2.1. The first 26 layers comprise electromagnetic section (CE-E), which uses silicon sensors between steel clad lead absorbers. The remaining layers comprise the hadronic section (CE-H), the first seven of these layers are silicon sensors only, but the rest utilizes a mixed

structure with both silicon and scintillator sensors. Thicker stainless steel absorber are used in hadronic section. HGCAL features $\sim$620 m^2 silicon sensor coverage and readout through six million channels, and $\sim$370 m^2 scintillator coverage with 240 thousand channels across both endcaps.

Excellent granularity is achieved in the longitudinal direction by the 47 layers and in the transverse direction by silicon sensors as small as $\sim$0.5 cm^2 and highly-segmented scintillator tiles. In addition to the sensor architecture, low-noise FE electronics, cryogenic operation and per-pad calibration enable high resolution measurements of energy and timing.

2.1. On-Detector Electronics

In regions where fluence exceeds 2×10^{14} n_{eq}/cm^2 and reaches up to 1×10^{16} n_{eq}/cm^2, silicon sensors are used. The sensor pads are hexagonal diodes diced from 8 inch circular wafers, approximately 0.5 cm^2 to 1 cm^2 in size and has thickness ranging from 300 µm to 120 µm [25]. Irradiation of the sensors increase the necessary operating voltage for sufficient charge collection. Combined with fluence-driven high leakage current, significant power dissipation is caused in sensors. To overcome this problem, thinner sensor pads are used in regions with higher fluence.

HGCROC is a low-noise, high-speed and high dynamic range FE readout application-specific integrated circuit (ASIC) with up to 72 channels. There are two versions of the chip, one for silicon sensors, capable of measuring charge up to 10 pC and another version for SiPMs, measuring charge up to 300 pC [29]. Each channel of the chip is equipped with a pre-amplifier, at the low-end of the dynamic range, from a few fC to $\sim$200 fC, pre-amplifier operates linearly. The 10-bit 40 MHz successive approximation Analog-to-Digital Converter (ADC) reads the signal through a shaper. In the pre-amplifier's saturation range, charge is measured from Time-over-Threshold (ToT) value, by a discriminator associated to 12-bit Time-to-Digital converter (TDC) with 50 ps resolution. Another TDC provides Time-of-Arrival (ToA) measurement with

25 ps resolution. See Figure 2.2 for a simplified definition of these measurements.

Up to six HGCROC ASICs are installed on an eight-layer hexagonal readout PCB called the Hexaboard. On a full-size high-density (HD) Hexaboard, six HGCROCs read from 432 diode pads while a full-size low-density (LD) Hex-module houses three HGCROCs reading from 192 sensor pads. Silicon sensor pads and a Hexaboard are laminated into a single Hex-module. Silicon sections of HGCAL layers are tiled with the Hex-modules. HD modules are used in higher fluence regions while LD modules used in the rest. Modules with sub-hexagonal geometries are used at the margins. There are 11 Hex-module variants developed for modules with different geometries and layouts.

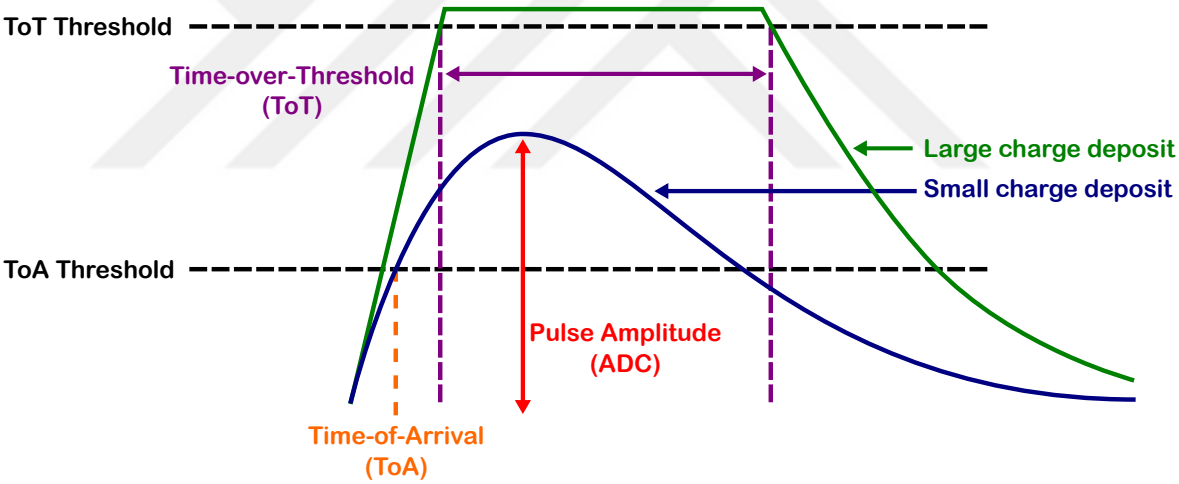

Figure 2.2. Simplified diagram of a single channel analog signal resulting from an incident particle. Image adapted from [30] (© 2025 The Authors; CC BY 4.0).

In LD regions, a concentrator mezzanine board, along with a mezzanine board housing Direct Current-to-Direct Current (DC-DC) converter modules, connects on top of the Hex-module at dedicated interfaces. The concentrator mezzanine holds two concentrator ASICs: ECON-D and ECON-T. Each HGCROC has two 1.28 Gbit s^{-1} e-links to the data path and four 1.28 Gbit s^{-1} e-links to the trigger path. ECON-T applies compression and selection to the data readout in the trigger path at the 40 MHz event rate. It employs one of the five algorithms: Threshold Sum, Super Trigger Cell, Best Choice, Repeater, and Auto-encoder [31]. ECON-D serializes concentrated data

in the data path at the L1A rate of 750 kHz.

A PCB called the Wagon Board connects three to seven Hex-modules to an Engine Board. Wagon Boards can be found many shapes depending on its location on the detector layer. Two LD Wagon Boards connect up to seven LD Hex-modules to an Engine Board, while an HD Wagon Board connects up to four HD Hex-modules to an Engine Board. While LD Wagon Boards are passive, HD Wagon Boards are active boards that house the ECON modules, replacing the concentrator mezzanine in the HD regions. Engine Boards also come in HD and LD variants. An LD Engine Board carries one Versatile Transceiver (VTRx+) [32] and three low-power gigabit transceivers (lpGBTs) [33], the HD variant carries twice as many of each. ECON data are transmitted over Wagon Boards to lpGBT modules on the Engine Board. An lpGBT module is able to serialize ECON inputs from up to seven or six $1.28\,\mathrm{Gbit\,s^{-1}}$ links depending on chosen error correction scheme. Each lpGBT module transmits serialized data to off-detector systems (uplink) at $10.24\,\mathrm{Gbit\,s^{-1}}$ via an optical fiber link driven by a VTRx+ module. One lpGBT module in the data path operates in transceiver mode, receiving clock and control signals at $2.56\,\mathrm{Gbit\,s^{-1}}$ from off-detector systems (downlink). The FE assembly described up to this point is called a train, an HD train schematic is shown in Figure 2.3.

The HGCAL region where $|\eta| > 2.4$ is effectively silicon-only. While CE-E utilizes only silicon sensors, scintillator tiles of sizes ranging from $4\,\mathrm{cm^2}$ to $32\,\mathrm{cm^2}$ will be used in CE-H regions where ionizing radiation dose is low enough and fluence is below $8 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ [34]. Light from the scintillator tiles will be directly read from radiation-hardened SiPM photodetectors located on tile surface. SiPM-on-Tile technology was developed by CALICE collaboration [35].

Up to five tile modules are laid to make up the basic 10° detector unit. Most tileboards have a single HGCROC chip with two e-links to data path and four e-links to trigger path [36]. A Giga-Bit Transceiver Slow Control Adapter (GBT-SCA) is present on tile modules, the ASIC is responsible for bi-directional data transfer be-

tween HGCROC and motherboard. GBT-SCA is eliminated in silicon section readout chain and its functionality is integrated into lpGBTs. Tile modules in a basic unit is connected to a motherboard via passive intermediary PCBs called Wingboards. Scintillator section differs from the silicon counterparts as ECON ASICs are housed in the motherboard instead of mezzanines. A motherboard houses one ECON-D and two ECON-Ts linked to two lpGBTs, as well as a VTRx+ module.

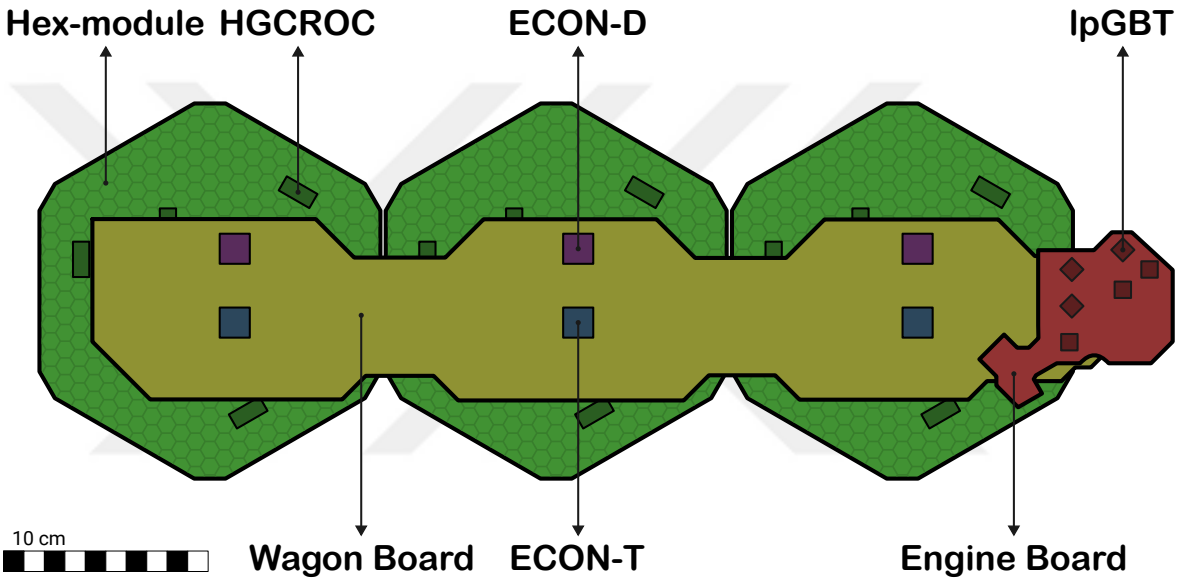

Figure 2.3. Schematic showing an HD train assembly with 3 full-size HD Hex-modules.

Cassettes are wedge-shaped sub-assemblies covering 60 or 30 degrees of detector layers. A cassette is built from Hex-modules, Wingboards, FE readout electronics, a cooling plate and other detector support structures. A CE-E cassette is a 60-degree wedge-shaped assembly whose central copper cooling plate is tiled with Hex-modules on both sides, fitted with peripheral electronics and cover,ed with 2.8 mm thick stainless steel clad lead absorbers. Six cassettes are radially joined to form a disk which corresponds to two active layers of electromagnetic calorimeter. CE-E consists of 13 disks corresponding to 26 silicon module based layers.

Only a single side of the CE-H cassettes are tiled with silicon or scintillator modules. The first seven layers of CE-H are silicon-only; in the remaining 14 layers,

scintillator tile ratio goes up farther away from the IP. CE-H cassettes are covered with a thin protective layer and fitted between 35 mm or 68 mm thick stainless steel absorber disks, which are much thicker than the CE-E absorbers.

2.2. Off-Detector Electronics

2.2.1. Trigger Primitive Generator (TPG) System

3x3 array of individual silicon sensors at the HD region or 2x2 array of larger silicon sensors at the LD region constitute a trigger cell (TC). This grouping ensures that each Hex-module outputs 48 TCs. Each HGCROC sends trigger data to an ECON-T via two e-links through the trigger path. ECON-T applies selection and compression to readout data using mixed algorithms. Threshold Algorithm selects TCs above a certain energy level, as efficient as this method is, additional resources are required to compensate for its variable output rate. Best Choice Algorithm and Super Trigger Cell Algorithm are found to be similar in performance and more feasible when used in parallel [37]. Best Choice Algorithm selects a fixed number most energetic TCs and used in CE-E. Super Trigger Cell Algorithm unifies all or most of the 48 TCs and sends a sum of energies, it is used in CE-H. Repeater algorithm on the other hand is intended for testing. Trigger data is sent through optical links to off-detector trigger systems at $10.24\,\mathrm{Gbit\,s^{-1}}$.

Data is processed initially at the aptly named TPG Stage 1 which is divided into three sections, one for every 120° sector of the detector. Main function of TPG Stage 1 is to repack FE data. TCs are grouped into 84 bins in $r/z \times \phi$ plane and partial trigger towers (pTTs) are built. A tower is a detector constituent of size 0.087×0.087 in $\eta \times \phi$ phase space [27]. TPG Stage 1 logic will be implemented on 14 Serenity boards [38] and TPG Stage 2 logic will be implemented on 18 boards, per 120° sector of the detector [37]. Each TPG Stage 1 FPGA time-multiplexes by bunch crossing and formats data to be sent to two TPG Stage 2 FPGAs [39].

TPG Stage 2 unpacks data from Stage 1 and repacks before sending to the L1T. Partial towers are summed to trigger towers. Cluster seeds are identified from most locally energetic TCs and clusters are built by aggregating adjacent TCs. Trigger towers and clusters are used to build physics objects like electrons and photons or utilized in advanced algorithms. HGCAL trigger primitives, along with primitives from the other calorimeters, the muon systems and the outer tracker, are sent to the L1T. The Global Trigger propagates L1A signal through Phase-2 Trigger and Timing Control and Distribution System (TCDS2) to detector BE and initiates full data readout.

2.2.2. Data Acquisition (DAQ) System

An HD train can have up to 3 full-size and a partial HD Hex-modules, housing a total of 22 HGCROCs and 4 ECON-Ds. Each ECON-D can have up to six $1.28\,\mathrm{Gbit\,s^{-1}}$ output links to two lpGBTs on the Engine Board. These links are spread across a total of 14 $1.28\,\mathrm{Gbit\,s^{-1}}$ input channels on the lpGBTs. Each lpGBT module outputs $10.24\,\mathrm{Gbit\,s^{-1}}$ single-bit stream at 40 MHz, transmitted via an optical fiber link to the DAQ BE system. The DAQ BE system will run on 96 Serenity Advanced Telecommunications Computing Architecture (ATCA) boards, each housing an AMD Virtex™ UltraScale+™ VU13P Field Programmable Gate Array (FPGA) [40]. Readout stream from optical links are processed in pairs by the Capture Blocks. 54 Capture Blocks will be implemented in each FPGA, facilitating readout from 108 fiber-optic input links.

Event fragments from HGCAL enter the CMS central DAQ system through simple point-to-point optical links using SLinkRocket protocol [41] at $25\,\mathrm{Gbit\,s^{-1}}$. Part of the central DAQ system resides in the USC, running on up to 192 ATCA crates [42]. Each crate houses a dual-FPGA DTH400 board and optional DAQ800 boards. A TCDS2 unit and a DAQ unit is implemented on the separate FPGAs of the DTH400 board, hence the name DAQ and TCDS2 Hub (DTH). A TCDS2 "captain" distributes the LHC clock signal, and control and configuration commands including L1A to the DTH boards, which in turn propagate it to FE and other BE electronics of the detector. The DAQ FPGA on the DTH400 aggregates event fragments received over

SLinkRocket links from up to 12 DAQ BE boards at a throughput of $400\,\mathrm{Gbit\,s^{-1}}$ as well as provide buffering. The sub-detector DAQ systems can exert back-pressure through TCDS2 to throttle the trigger. DAQ800 boards, housing another DAQ FPGA instead of a TCDS2 unit, are used to expand the throughput of a DTH by $800\,\mathrm{Gbit\,s^{-1}}$ for each board added. Data is transferred to the online computing center on the surface through $100\,\mathrm{Gbit\,s^{-1}}$ Ethernet streams for event building.

3. CAPTURE BLOCK

Informally, a block is self-contained unit of FPGA firmware that performs well-defined tasks and has specific signal inputs and outputs. The Capture Block is the building block of the HGCAL DAQ BE firmware. A single Capture Block can read from two 10.24 Gbit s^{-1} optical links, and can handle up to 12 ECON-Ds, whose data are spread across 14 input e-links. The Capture Block output are the serialized event data from full-granularity readout. The Capture Block has error handling capabilities as well, constituent blocks implement error identification functions and generate recovery action commands to be propagated downlink. The Capture Block is required to be configured with respect to the detector FE it is linked to.

3.1. Capture Block Logic

A schematic illustrating the Capture Block logic is shown in Figure 3.3. The pair of serial data streams received over two optical links are de-serialized and re-clocked, converting two 256-bit frames at 40 MHz into two 32-bit-wide buses at 320 MHz before entering the Capture Block. The bus width is aligned with the 32-bit ECON-D packet word length.

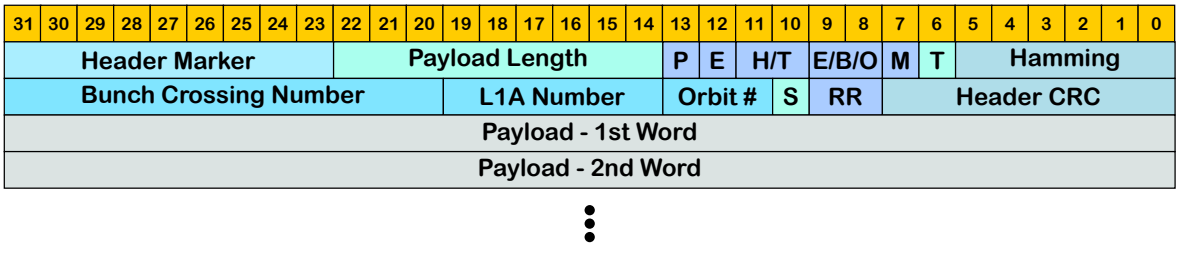

Figure 3.1. Breakdown of a standard ECON-D packet.

An ECON-D packet starts with two header words followed by a payload, a breakdown of the ECON-D packet header is shown in Figure 3.1. Header marker is a programmable 9-bit word. Payload length, truncated (T) flag, and status (S) flag hold information about the packet. Pass-through (P) flag, expected (E) flag, header/trailer

(H/T) flag, event/bunch/orbit (E/B/O) flag, and match (M) flag provide information about ECON-D and HGCROC status, and reconstruction status. Reset request (RR) is sent based on internal conditions. Bunch crossing, L1A, and orbit numbers make up the event timestamp. Hamming and CRC numbers provide error resiliency for the 64-bit header. Payload entails readout data and sub-packet metadata, however it is not examined by the Capture Block.

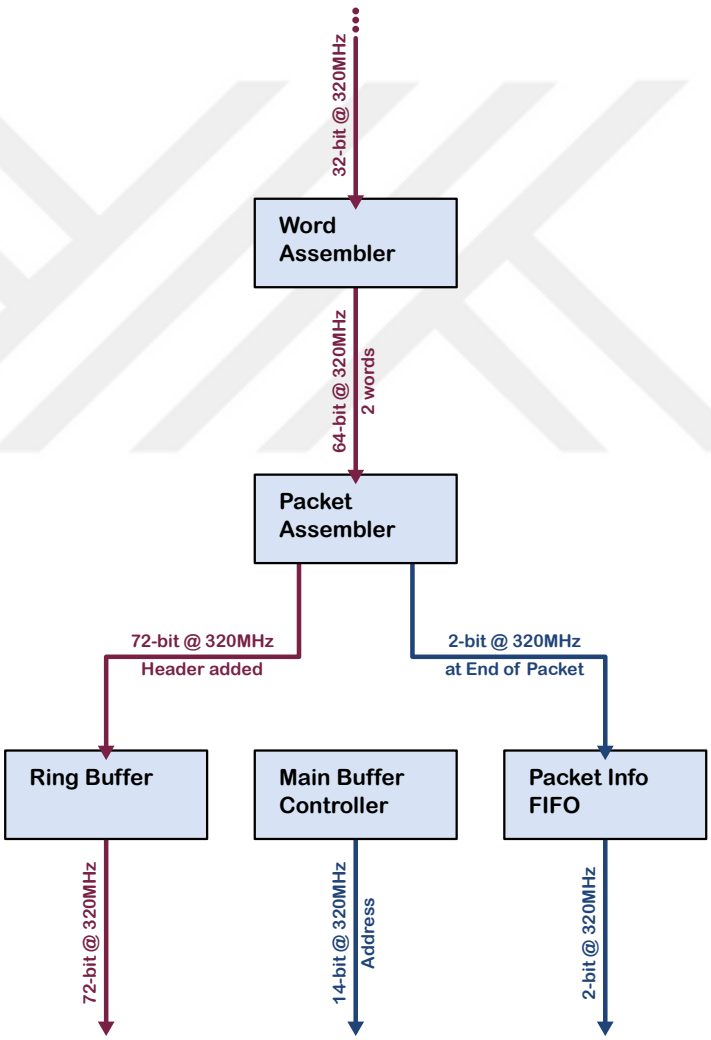

Figure 3.2. Schematic illustrating the Packet Former's internal logic and the flow of data and metadata.

The 32-bit words are assigned ECON-D IDs in the Capture Block according to the initial configuration and sent to the Packet Former.

The Packet Former, instantiated once per ECON-D, packages ECON-D words. A schematic of the Packet Former logic is shown in Figure 3.2. Inside the Packet Former, the Word Assembler reads two lpGBT streams simultaneously and combines two 32-bit words belonging to the same ECON-D into a 64-bit word.

The Packet Assembler searches these words for a 24-bit idle word and a subsequent 9-bit header marker to identify the start of an event readout. The markers are set during the initialization. After a packet start is identified, a cross-check is performed to the 8-bit header Cyclic Redundancy Check (CRC) number, if the check fails an error flag is raised. The Packet Assembler adds an 8-bit header to the 64-bit word and pushes it to the 3-word deep Ring Buffer.

The Main Buffer Controller manages the transfer of data from the Ring Buffer to the Main Buffer. It keeps track of the Main Buffer occupancy and stores the dynamic address for its ECON-D. It asserts back-pressure if the Main Buffer load exceeds a fixed threshold. The Packet Info "First in, First out" (FIFO) latches a 2-bit tag stating if the Main Buffer is overflown every time a packet is finished.

The Main Buffer Multiplexer multiplexes and writes the data from Ring Buffers into the Main Buffer, while communicating with Main Buffer Controller. The Main Buffer is a Xilinx Parameterized Macro for a simple dual-port Random Access Memory (RAM), which implements a 72-bit-wide 2^{14}-deep word buffer with 5 cycle read latency. The Main Buffer stores the data until it is read by the Event Builder.

The Event Builder is a finite state machine (FSM) that groups ECON-D packets of the same event into the event packets. It checks Packet Info FIFOs for new packets and reads them from the Main Buffer. The Event Builder inputs are 64-bit words, 8-bit Packet Assembler headers are discarded at this point.

The Event Builder checks the timestamps of ECON-D packet headers before reading the payload. The headers contain a 12-bit bunch crossing number, a 6-bit

event number, and a 3-bit orbit number. An orbit corresponds to a full revolution of particles around the LHC ring, it entails 3564 bunch crossings. Event number is a stored value incremented with each L1A signal. ECON-Ds read these numbers from HGCROCs and perform a cross-check before sending them to the Capture Block. The Event Builder compares these three numbers in the packet headers with the numbers from the Capture Block L1A counter. The Event Builder raises flags for individual ECON-Ds in case of any timestamp mismatch. The array of error flags arose during the timestamp check is sent to the Autorecovery Block. The array comprises 6-bit vectors corresponding to six error types.

An event header containing the timestamp (bunch crossing, event, and orbit numbers) from the Capture Block L1A counter and individual ECON-D status are sent to the Event Buffer. Unless every ECON-D packet is faulty or timed-out, the Event Builder commences to read the packets starting over from the headers. It sends a serialized stream of 64-bit words to the Event Buffer and a 12-bit event length data to the Event Size Buffer. These buffers serve as the final stop before the data are transmitted off the Capture Block.

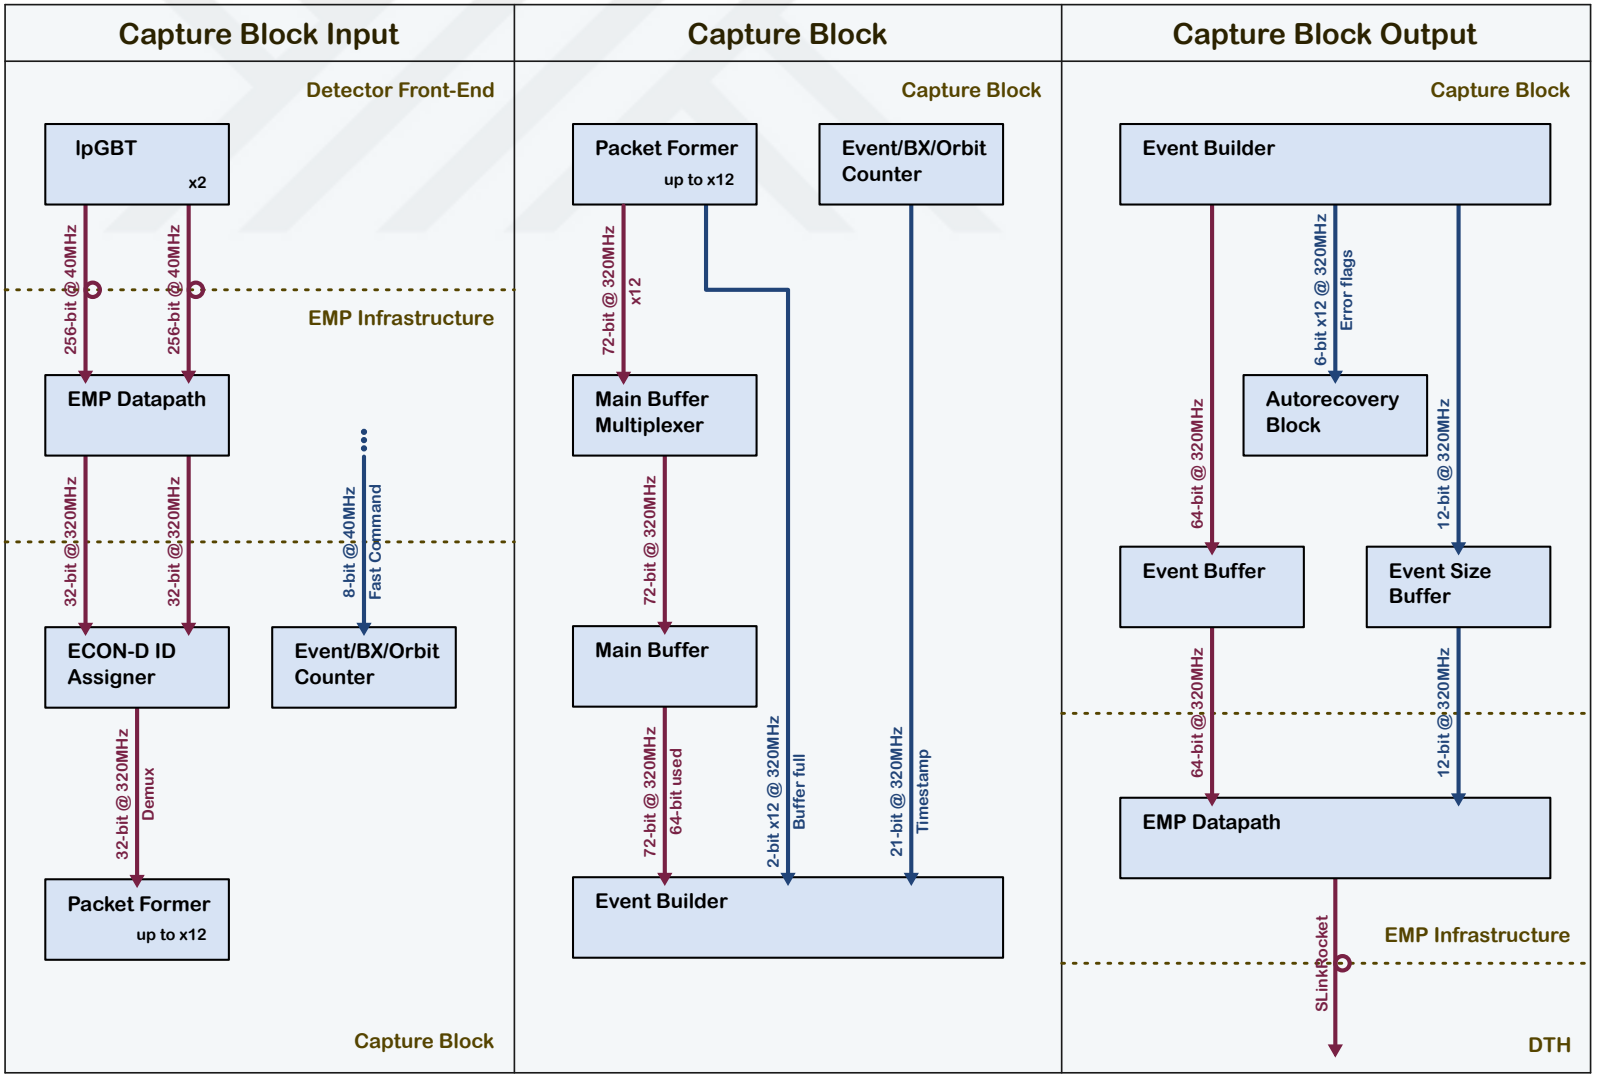

Figure 3.3. Schematic illustrating the Capture Block's internal logic and the flow of data and metadata.

4. AUTORECOVERY BLOCK

The Autorecovery Block resides within the Capture Block, it monitors errors emitted by the ECON-Ds. Four independent error counters track timeout, bunch crossing mismatch, event number mismatch, and buffer overflow conditions. When a counter hits its programmable threshold, the Autorecovery Block issues a recovery fast command (FC), which is fed to the HGCAL control path.

4.1. HGCAL Control Path

A secondary task of the HGCAL DAQ BE system is to distribute clock, and the so-called fast and slow control signals to the FE electronics. FCs come from the TCDS2 at the 40 MHz LHC clock rate. FCs initiate the readout and intervene in it on an event-by-event basis. Clock signal and the FCs comprise the fast control signals. Slow control signals are used for the calibration and configuration of the detector electronics and are generated at the DAQ BE. As HGCAL electronics show great variation in geometry, they need to be configured accordingly.

Clock and control signals are received by the lpGBT acting as the receiver/transmitter (Rx/Tx) in the data path. The control paths of an HD train can be seen in Figure 4.1. Topology for an LD train does not differ substantially, although it is scaled down.

In the SiPM region, there are a few more important differences, e.g., ECON ASICs are moved to the motherboard and GBT-SCA on the tileboards distribute the slow control signals to the HGCROCs. SiPMs are calibrated with an on-board light emitting diode (LED) and ALDO2[43] regulator is used for voltage adjustment.

The lpGBT receives 40 MHz LHC clock from the DAQ BE and multiplies it to 320 MHz with its on-chip Phase-locked Loop (PLL). The clock and fast control signals

are duplicated by a RAFAEL clock and data fan-out ASIC [44] and distributed through the wagon board. Exact configuration of the fast control and clock path varies by the detector region.

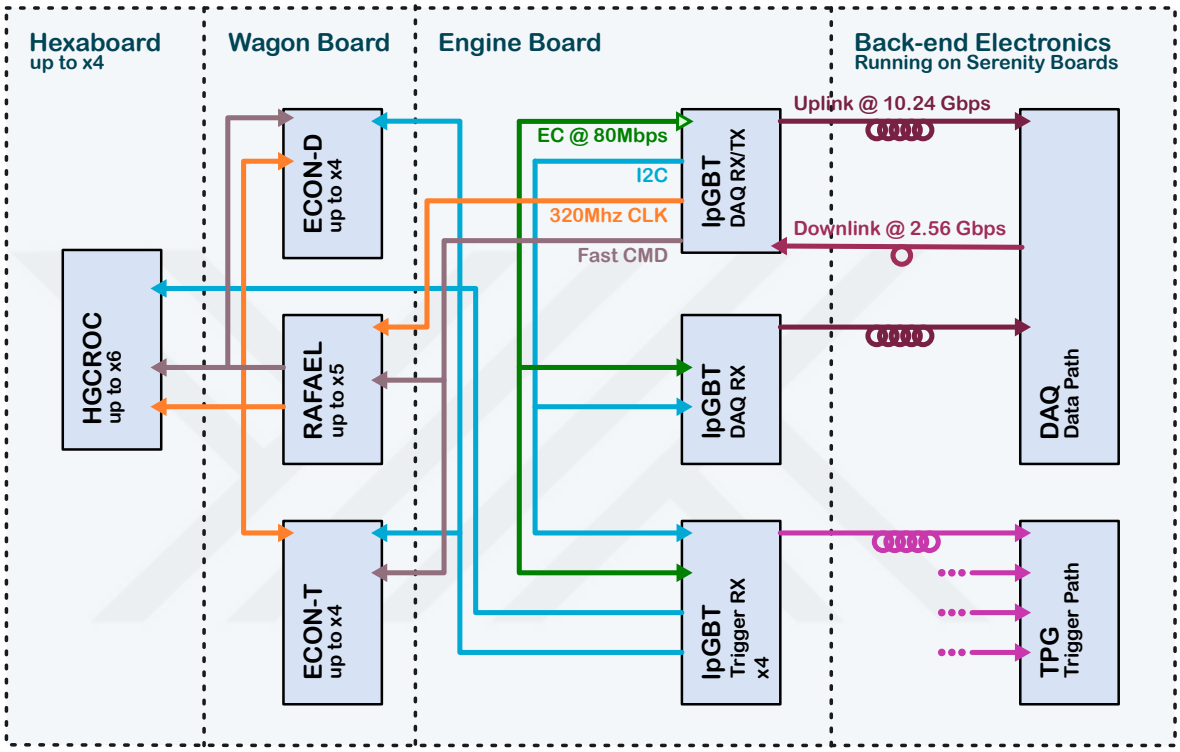

Figure 4.1. Schematic of the HGCAL electronics control path for an HD train.

Each Rx/Tx lpGBT is capable of four independent FC streams. In the HD region, a Hex-module and the ECON ASICs associated with it share one FC stream. In the LD silicon region or the scintillator region multiple Hex-modules can receive the same stream. Each FC stream is $320\,\mathrm{Mbit\,s^{-1}}$, allowing for only one 8-bit FC to be sent at each bunch crossing. FCs can be used to; initiate calibration pulses, reset event/bunch/orbit counters, clear event buffers, resync ECON ASICs with HGCROCs or with the BE electronics, or initiate a more disruptive reset. L1A signal is also an FC, moreover a pre-L1A FC can be sent preemptively to combat loss of the L1A signal.

Slow control signals are distributed to other lpGBTs by the Rx/Tx lpGBT through $80\,\mathrm{Mbit\,s^{-1}}$ bi-directional external control (EC) paths. $\mathrm{I^2C}$ paths between the lpGBTs are present as backup. In the silicon region, $\mathrm{I^2C}$ buses from lpGBTs to the

ECON and HGCROC ASICs, and to the VTRx+ modules are used for configuration. GPIO, ADC, and DAC interfaces of the lpGBTs are likewise integrated into the slow control fabric providing monitoring, voltage setting, and general purpose input/output (I/O) capabilities. Notably, scintillator region differs from silicon region as it utilizes GBT-SCA to deliver slow control signals to the HGCROCs.

4.2. Development Framework

The DAQ BE utilizes the Serenity board, which is an ATCA form-factor carrier board developed for the Phase-2 Upgrade of the CMS detector. It is designed as a generic board for high-speed implementations and intended to be used by many of the CMS sub-systems. The carrier board houses two FPGA-occupied daughter card sites. The firmware runs on these FPGAs. Each FPGA has 16 Firefly optical connectors capable of carrying 192 unidirectional serial links at $25\,\mathrm{Gbit\,s^{-1}}$. The board management, which entails configuration, monitoring and testing is handled by an on-board ComExpress mini CPU running CentOS 7.

The DAQ firmware is authored in Very High Speed Integrated Circuit Hardware Description Language (VHDL). AMD Vivado Design Suite is used for design, synthesis and implementation. The firmware is developed within the Extensible, Modular (data) Processor (EMP) Framework [45], a firmware and software package that provides top-level designs streamlining custom firmware development. The framework connects user-created payload to I/O ports seamlessly. It includes datapath buffers for pattern injection, and an interface for trigger, timing, and control (TTC) signals. EMP framework also provides the fabric for the IPbus protocol [46], which introduces a high-speed interface between the FPGA and the CPU. Slow control signals to the DAQ BE are transmitted via IPbus protocol. Initial configuration, the runtime control tasks and monitoring of the firmware is facilitated over IPbus registers at 31.25 MHz.

4.3. Autorecovery Logic

The Autorecovery Block, shown in Figure 4.2 is a crucial part of the FC generation logic. It counts the errors identified by the Event Builder, and generates FC candidates (FCCs) to initiate a recovery action. FCCs are generated, if only the same error is encountered consequently for a set amount of times. Unlike the rest of the Capture Block, the Autorecovery Block clock domain is 40 MHz of the 320 MHz DAQ clock.

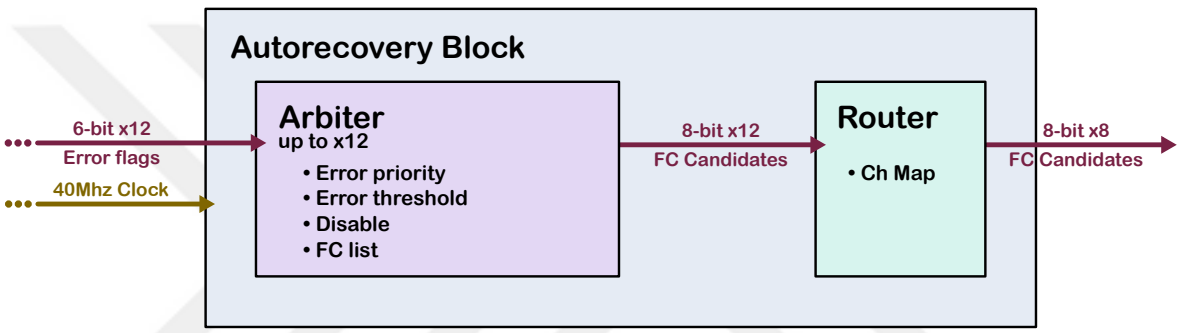

Figure 4.2. Schematic of the Autorecovery Block.

The error flags are received from the Event Builder as an array of 6-bit signals per ECON-D, corresponding to 6 possible error flags. Currently, four types of error checks are implemented in the Event Builder and two remain as redundancy. For the Autorecovery Block to issue an FCC, the Event Builder must assert a single-bit valid signal. The error flags are marked invalid if no ECON-D packet collected for the same L1A is tagged "good", otherwise they are marked valid all at once. The Event Builder logic for "good" packets is shown as a flowchart in Figure 4.3.

Residing in the Autorecovery Block, the Arbiter block is instantiated once per ECON-D. The role of the Arbiter is to keep count of errors of all types and to issue an FCC. The Arbiter runs independent counters for each error type. A counter is incremented if; (1) the corresponding error flag is received, (2) error flags valid signal is received and (3) the ECON-D is not opted out of the Autorecovery logic. The Arbiter expects errors to be consecutive, if the streak is broken, the counter is reset. The Arbiter has an input for excluding individually chosen ECON-Ds from FCC generation

logic.

When an error counter in an Arbiter hits a preset threshold, the corresponding FCC is generated. Multiple FCCs can be triggered in the same cycle, however each Arbiter outputs only one FCC. The predominant FCC is chosen by a static priority list between the error types. It is possible to ignore individual error types by setting their priority to zero.

The Capture Block supports up to 12 ECON-Ds, thus the Autorecovery Block can have up to 12 Arbiter blocks, however there are only 8 FC channels across two lpGBTs connected to a Capture Block. The Router block, residing in the Autorecovery Block, handles the distribution of FCCs to the limited amount of FC channels. It has a static map of ECON-Ds to FC channels. If two or more Arbiter blocks are mapped to the same FC channel, the FCC from the Arbiter with lower index is prioritized. FCCs from the Autorecovery Block are not immediately distributed through the DAQ electronics. A priority logic is to be implemented in the Capture Block, which will decide the final command from FCCs of the Autorecovery Block and other sources.

4.4. Error Types Handled by the Autorecovery Block

The Event Builder identifies errors in the ECON-D packets and sends Autorecovery Block an array of errors along with a single-bit valid signal. Criteria for a non-valid signal is, every ECON-D sending a faulty packet or timing-out following an L1A signal. Non-valid signal from the Event Builder disrupts FC generation at the Autorecovery Block and resets its error counters. “Good” packet criteria and the Event Builder error flag logic are shown in Figure 4.3.

4.4.1. Timeout

The Event Builder starts a down-counter when an L1A signal arrives. Size of the counter in clock cycles are set via an IPbus register. When the counter runs out,

ECON-Ds that had not sent a packet are flagged with “timeout” tag. If a packet has a CRC error, it is discarded by the Packet Assembler, thus it never arrives at the Event Builder and causes a timeout error. The corresponding FCC for the timeout error is the bunch counter reset (BCR) signal.

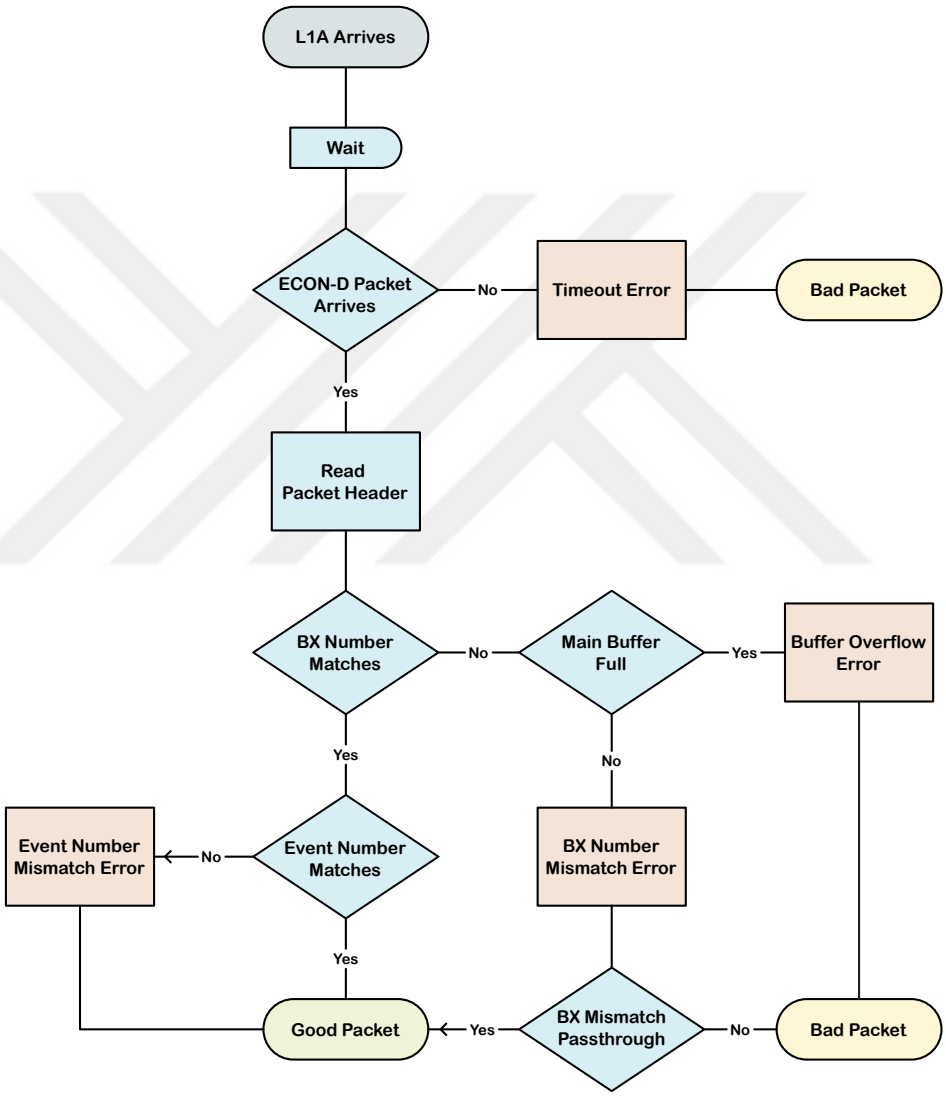

Figure 4.3. The Event Builder error detection and packet validity logic flowchart.

4.4.2. Bunch Crossing (BX) Number Mismatch

The Event Builder compares timestamps from the Capture Block and the ECON-D packet headers. The first step of the check is the bunch crossing (BX) number. During this inspection, the Packet Info FIFOs are also checked for a pos-

sible Main Buffer overflow. If a packet header fails BX number check and the buffer is not overflown, it is flagged with a "BX mismatch" tag. However, an IPbus controlled bypass for the BX number check is implemented. If the bypass is enabled, packet is marked "good" but the BX mismatch error is still conveyed to the Autorecovery Block. The corresponding FCC for the timeout error is the orbit and bunch counter reset (OCR+BCR) signal.

4.4.3. Event Number Mismatch

Event number comparison is performed after the header passes BX number check. The packet is still marked "good" even if it fails the event number check, however the error flag is conveyed to the Autorecovery Block. The corresponding FCC for the timeout error is the chip sync reset (Chip-Sync) signal.

4.4.4. Buffer Overflow

The Event Builder issues the "buffer overflow" tag, if the BX number is incorrect and the concerned Main Buffer section is almost full. The Main Buffer is a fixed-size RAM divided between ECON-Ds. Main Buffer is divided into regions assigned to specific ECON-Ds via IPbus registers at startup in a start/end address basis. A Main Buffer Controller for each ECON-D monitors occupancy of these address ranges. If occupancy reaches a preset threshold (75%) buffer overflow signal is asserted to the Packet Assembler by the Main Buffer Controller. The overflown status is conveyed to the Event Builder by the Packet Info FIFOs. The corresponding FCC for the timeout error is the ECON-D link reset (Link-Reset-ECON-D) signal.

4.5. Improvements to the Autorecovery Block

HGCAL electronics show variety in its sub-configurations, this gives rise to need for a flexible firmware. IPbus infrastructure of the Capture Block is expanded to include the Autorecovery Block, allowing most of the block parameters to be configured via

IPbus protocol. A new IPbus address table for the Autorecovery Block was created to accommodate for five control registers and 30 status registers. The status registers are read-only externally, while control registers are write-enabled. Each IPbus register is 32-bit wide. Schematic of the updated block is shown in Figure 4.4.

Error priorities can now be set via a control register, each of the six types has a 4-bit priority value. If multiple error counters hit the threshold in the same cycle, error with the lower priority value is preferred by the Arbiter block, and no FCC is generated for the other errors while their counters are still reset. It is possible to ignore an error type by setting its priority value to 0. The same priority table is used for every ECON-D.

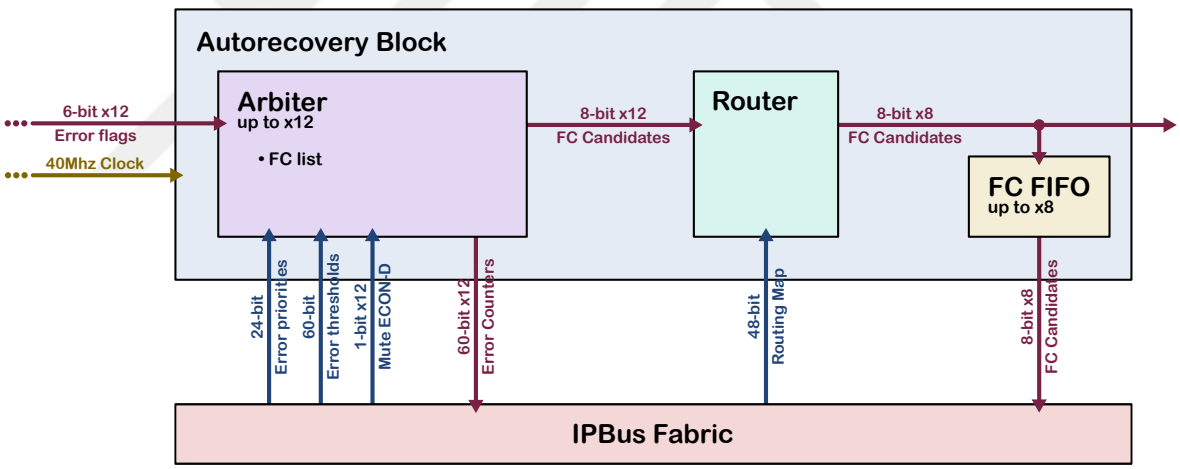

Figure 4.4. Schematic of the updated Autorecovery Block featuring IPbus integration.

The thresholds for error counters can now be set via two control registers. Each of the six error types have a 10-bit threshold value. The same threshold list is used for every ECON-D.

A control register is used to mute errors from individual ECON-Ds. Errors from selected ECON-Ds will not increment any counters.

24 IPbus status registers are implemented to monitor the internal error counters of the Arbiter blocks. A 10-bit counter value for each error type, amounting to 60-

bits per ECON-D, is stored in these registers and updated at the IPbus clock rate of 31.25 MHz.

The routing map is 48-bits split between two control registers. It is used to assign a 4-bit FC channel address to each ECON-D.

One FCC FIFO per FC channel is implemented. They are independent from the existing Autorecovery Block output channels, however they receive the same FCCs. The FIFOs are 8-bit wide, aligned with the FC signal length, and 1024-word deep. They can be read via IPbus registers; a control register is used for read enable signal, while status registers deliver the FCC output. It is possible to monitor the occupancy status of the FIFOs as well.

All of the design, synthesis and implementation for the Autorecovery Block is carried out on Xilinx Vivado Design Suite (Version 2022.2, 64-bit), shown in Figure 4.5, using IPbus Builder (IPBB) [47] tool. IPBB streamlines creating work spaces, managing repositories and building projects. It generates IPbus decoders necessary for IPbus communication, from the address tables.

4.6. Validation of the Autorecovery Block

The validation setup is a Serenity board (version Z1.1) mounted with an AMD Kintex™ UltraScale+™ KU15P FPGA. The assembly is located in Boğaziçi University Kandilli Campus. The Autorecovery Block was validated both in isolation and embedded in the Capture Block. The broader test firmware includes a TCDS2 emulator and two Capture Blocks, each with an Autorecovery Block.

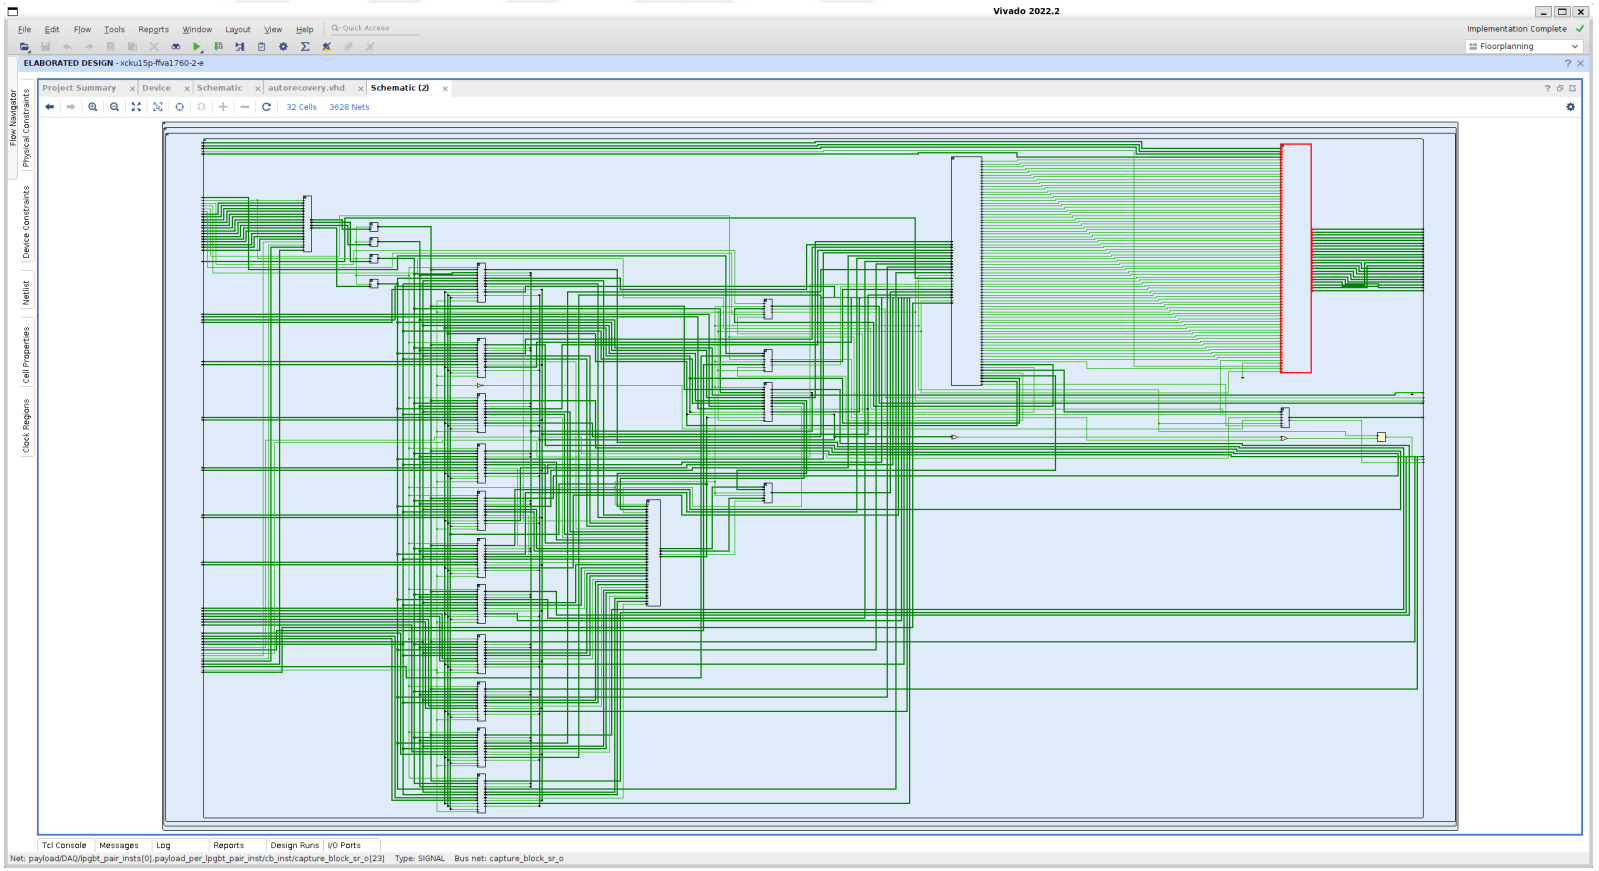

Figure 4.5. Screen capture from Xilinx Vivado Design Suite showing the Autorecovery Block highlighted red in the Capture Block.

It is possible to control, monitor and program the FPGA on the Serenity board, from the CentOS 7 system running on Serenity CPU. After the FPGA design is implemented on Vivado, it is programmed into the FPGA with the Serenitybutler tool, which also provides monitoring functions. The Empbutler tool is used when testing a design, it injects a fabricated input into the FPGA and captures the I/O streams of the FPGA on a frame-by-frame basis. The Empbutler captures data during an orbit and the captures are acquired from the buffers and not streamed in real-time.

Link	050		051	
Frame 0000	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0001	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0002	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0003	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0004	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0005	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0006	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0007	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0008	0101	FF000000ccccff08	0101	FF000000ccccff08
Frame 0009	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0010	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0011	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0012	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0013	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0014	0001	0F000000ccccff08	0001	0F000000ccccff08
Frame 0015	0000	00000000ccccff08	0000	00000000ccccff08
Frame 0016	0101	FF000000aa067080	0101	FF000000aa067080
Frame 0017	0001	0F000000aa067081	0001	0F000000aa067081
Frame 0018	0001	0F000000aa067082	0001	0F000000aa067082
Frame 0019	0001	0F000000aa067083	0001	0F000000aa067083
Frame 0020	0001	0F000000aa067084	0001	0F000000aa067084
Frame 0021	0001	0F000000aa067085	0001	0F000000aa067085
Frame 0022	0001	0F000000aa067086	0001	0F000000aa067086
Frame 0023	0000	00000000aa067087	0000	00000000aa067087
Frame 0024	0101	FF00000001508cc6	0101	FF00000001508cc6
Frame 0025	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0026	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0027	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0028	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0029	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0030	0001	0F00000001e08ce3	0001	0F00000001e08ce3
Frame 0031	0000	0000000001e08ce3	0000	0000000001e08ce3

Figure 4.6. First 32 frames of an injected EMP data pattern. The data are injected from two channels (50 and 51). A 4-bit EMP metadata (red) and a 64-bit data (blue) per channel comprise a frame.

Injected data format is set by the EMP software, a snippet from an injected packet is shown in Figure 4.6. Total data length is usually 1024 frames and it is inserted at the start of every LHC orbit, equivalent to 3564 LHC clock cycles. 64-bit data blocks are injected in parallel to a desired number of channels, along with a 4-bit metadata, which specifies orbit start, packet start/end, and validity.

IPbus registers were used for real time monitoring and configuration. Python scripts utilizing "uhal" library were used to read and record these registers.

4.6.1. Standalone Autorecovery Block

The Autorecovery Block is embedded into the EMP framework by itself, 12 input and eight output channels are used. Each input channel is designated as a single ECON-D error flag array input from the Event Builder. Injected input data were formatted as the Event Builder output, it contains six 1-bit error flags, a 1-bit valid, and additionally a 1-bit ECON-D mute signal. The EMP output channels are designated as the Autorecovery Block's 8 FC channels. Output data format is an 8-bit FCC followed by a valid bit.

In this test configuration, the Autorecovery Block clock domain was 360 MHz. Main functions of the Autorecovery Block were tested in its isolated form. The functions and the testing methods are listed below:

- The Autorecovery logic is applied independently to individual ECON-Ds. The routing map was configured as one ECON-D per FC channel via IPbus registers. A data pattern with the error flags were only injected to a single channel. FCC generation was observed only at the expected channel.
- The Autorecovery logic is applied in parallel to the ECON-Ds. The routing map was configured as one ECON-D per FC channel via IPbus registers. Data patterns with a single type of error flags unique to each channel were injected. Correct FCC for different error flags was observed at the corresponding channels.

- FCC generation thresholds can be set individually for the error types. The routing map was configured as one ECON-D per FC channel via IPbus registers. Unique thresholds for error types were set in the firmware, as IPbus registers for the thresholds were not yet implemented. Every error flag was sent in each data frame. It was observed that the period of each FC candidate generation was consistent with the set thresholds.
- Error priorities can be set to establish a hierarchy among error types. The routing map was configured as one ECON-D per FC channel via IPbus registers. Different priorities for error types were set in the firmware, as IPbus registers for the priorities were not yet implemented. Data patterns were created to trigger multiple FCC generation at the same frame. Only the FCC corresponding to the higher priority error type was observed at the expected FC channel.
- ECON-Ds can be muted individually. The routing map was configured as one ECON-D per FC channel via IPbus registers. Mute signal was asserted for selected ECON-Ds via IPbus registers. No FCC generation was observed at the corresponding channels when the error flags were injected.
- The routing map can be configured via IPbus registers. It is possible to independently map ECON-Ds to any of the eight FC channels. It was observed that when multiple FCCs are destined for the same FC channel, the FCC coupled to the ECON-D with the lowest index prevailed.

4.6.2. Integrated Autorecovery Block

Finally, the Autorecovery Block was validated embedded in a setup with two Capture Blocks and a TCDS2 emulator. Seven e-links per channel are set up, each e-link is associated with a single ECON-D. The setup was configured as six ECON-Ds per Capture Block. Idle and header patterns for ECON-D packets are set and Main Buffer address ranges are divided between ECON-Ds via IPbus registers. The Autorecovery Block error priorities and thresholds as well as the routing map are also configured at startup. TCDS2 emulator of the system is configured via IPbus registers. TCDS2 allows injection of L1A and other configured FCs to the Capture Block at each orbit

start.

```
reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_route_00
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_route_00 = 6636321 : 0x654321

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_route_01
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_route_01 = 8947847 : 0x888887

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_err_pris
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_err_pris = 1056768 : 0x102000

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_threshold_00
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_threshold_00 = 1049601 : 0x100401

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_threshold_01
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_threshold_01 = 1049601 : 0x100401

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_disable
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.ct_disable = 0 : 0x0

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_data_00
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_data_00 = 2509608341 : 0x95959595

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_data_01
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_data_01 = 0 : 0x0

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_00
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_00 = 33558528 : 0x2001000

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_01
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_01 = 33558528 : 0x2001000

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_02
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_02 = 16779264 : 0x1000800

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_03
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.fifo_info_03 = 16779264 : 0x1000800

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_00
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_00 = 0 : 0x0

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_01
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_01 = 0 : 0x0

reg_access -d $conid payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_02
payload.daq.per_lpgbt_pair.autorecovery_ctrl_stat.st_flag_02 = 0 : 0x0
```

Figure 4.7. Autorecovery Block's IPbus registers read from a Bash console session during the BX mismatch test. Register names are displayed in red while the data appear in blue.

Two input channels are enabled, corresponding to two lpGBTs. Both Capture Blocks receive the input stream from both channels, except that the channels are

swapped. Input data format is ECON-D packet format, shown in Figure 3.1. FCC outputs were monitored via IPbus registers from the Autorecovery Block FIFOs. A register readout during a BX mismatch test is shown in Figure 4.7. Each of the four implemented error types were tested by injecting fabricated ECON-D packets and FCs.

4.6.2.1. Timeout Error. The error arises when L1A signal is received and no viable packet is received after a timer ran out. To emulate this error, L1A signal is asserted and ECON-D packets constructed with CRC error or with wrong markers are injected. If packet CRC is incorrect it is discarded by the Packet Assembler, thus it never arrives at the Event Builder.

4.6.2.2. BX Mismatch Error. The asserted L1A signal has a BX number tag. The Event Builder compares this number with the BX number of the ECON-D packet. ECON-D packets with different BX numbers were injected to emulate this error. ECON-D packets are otherwise correctly constructed.

4.6.2.3. Event Number Mismatch Error. The event counters of DAQ BE and ECON-Ds start from 1 and are incremented by one with each L1A signal. To emulate this error, L1A is asserted and otherwise faultless ECON-D packets with a fixed L1A number (event number) are injected. The error arose as the Event Builder expects ECON-D packets with increasing L1A numbers.

4.6.2.4. Buffer Overflow Error. When an ECON-D packet is not marked “good”, the Event Builder does not read the payload. Moreover, it does not assert a read enable signal on the Main Buffer Controller and the read address does not shift. The Event Builder will read the same ECON-D packet header in each cycle until it is flagged “good”, this will cause Main Buffer to fill up since it will not be unloading any data. When the Main Buffer count reaches a certain threshold, in combination with BX number mismatch, the Event Builder raises buffer overflow flag. To emulate this error, at least one ECON-D link is configured to send valid packets, preventing invalidation

of all the error flags. Since the Main Buffer regions are allocated via IPbus registers, some ECON-Ds were given much smaller address ranges, leading to the buffer filling up much faster when data patterns with an incorrect BX number were injected. It was also observed that the Main Buffer will fill up faster if packets with smaller payloads are injected.

5. CONCLUSIONS

Testing the Autorecovery Block in its standalone form confirmed that it functions as intended. Every function of the block was validated within a wide range of scenarios. The block acted on errors injected over 1024 frames with 100% FCC generation rate.

IPbus protocol was implemented in the Autorecovery Block as part of my work. The existing IPbus infrastructure of the DAQ BE firmware was successfully expanded into the Autorecovery Block in a fully compatible manner. The new IPbus structure enabled the Autorecovery Block to be configured as desired without re-synthesizing the firmware. Although the core Autorecovery logic remained unchanged, the block's I/O interfaces were overhauled, converting static parameters to dynamic ones. The DAQ BE firmware with the Autorecovery Block successfully met all timing constraints.

5.1. Outlook

Scope of the validation included the core functions of the block. The firmware was configured at the startup, and a single scenario was tested between resets. In the future, the block's robustness against repeatedly changing configuration parameters can be tested. The results of these tests can be used to reduce the cases which require a reset. Although a large variety of data patterns was used, they were cycled for up to only 30 minutes. Long stress tests can be conducted in the future. The next beam test can be a good opportunity to further solidify the Autorecovery Block's functionality. The scope can be expanded into timing, maximum frequency, stress tests and further.

The main logic of the Autorecovery Block can be improved as well. The simple error priority list and shared FC channels cause a loss of the error data. The requirement of consecutive identical errors can lead to some faults getting overlooked, as the total error rate might tell a different story. A new priority logic will be implemented in the DAQ BE firmware, incorporating the TCDS2, the Autorecovery Block and other

components into the FC decision process. The Autorecovery logic must be re-evaluated in light of the global HGCAL synchronization scheme.

REFERENCES

1. Landua, F., "The CERN Accelerator Complex Layout in 2022", 2022, `https://cds.cern.ch/record/2813716`, accessed on August 8, 2025.

2. Massam, T., T. Muller, B. Righini, M. Schneegans and A. Zichichi, "Experimental Observation of Antideuteron Production", *Il Nuovo Cimento A*, Vol. 63, No. 1, pp. 10–14, 1965.

3. Bryant, P.J., "The Impact of the ISR on Accelerator Physics and Technology", *40th Anniversary of the First Proton-Proton Collisions in the CERN Intersecting Storage Rings (ISR)*, Geneva, pp. 15-31, 2011.

4. Hübner, K., et al., "The Largest Accelerators and Colliders of Their Time", in *Particle Physics Reference Library*, Springer, Cham, pp. 585–660, 2020.

5. Nobel Prize Outreach, "The Nobel Prize in Physics 1984", 2025, `https://www.nobelprize.org/prizes/physics/1984/summary/`, accessed on July 1, 2025.

6. Myers, S., "The Greatest Lepton Collider", 2025, `https://cerncourier.com/a/the-greatest-lepton-collider/`, accessed on May 20, 2025.

7. Alemany-Fernandez, R., et al., "Operation and Configuration of the LHC in Run 1", CERN-ACC-NOTE-2013-0041, 2013.

8. CERN, *CERN Annual Report 2023*, CERN, Geneva, 2024.

9. Arduini, G., et al., "LHC Upgrades in Preparation of Run 3", *Journal of Instrumentation*, Vol. 19, No. 5, pp. 1-136, 2024.

10. Aberle, O., et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report*, CERN, Geneva, 2020.

11. Chao, A.W., M. Tigner, H. Weise and F. Zimmermann, *Handbook of Accelerator Physics and Engineering*, Third Edition, World Scientific, Singapore, 2023.

12. Hertzog, S.J., "LHC Dipole Magnets in the Tunnel", 2021, `https://cds.cern.ch/record/2782003`, accessed on August 8, 2025.

13. CERN Beam Performance Tracking - LHC, "LHC SuperTable 2024", `https://bpt.web.cern.ch/lhc/supertable/2024/`, accessed on May 1, 2025.

14. CMS Collaboration, "CMS Luminosity Results", `https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults`, accessed on May 15, 2025.

15. CERN, "New Schedule for CERN's Accelerators", 2024, `https://www.home.cern/news/news/accelerators/new-schedule-cerns-accelerators`, accessed on May 13, 2025.

16. ATLAS Collaboration, "ATLAS Luminosity Results", `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3`, accessed on May 16, 2025.

17. CMS Collaboration, "Development of the CMS Detector for the CERN LHC Run 3", *Journal of Instrumentation*, Vol. 19, No. 5, pp. 1-251, 2024.

18. Sakuma, T. and T. McCauley, "Detector and Event Visualization with SketchUp at the CMS Experiment", *Journal of Physics: Conference Series*, Vol. 513, No. 2, p. 2, 2014.

19. CMS Collaboration, "Particle-Flow Reconstruction and Global Event Description with the CMS Detector", *Journal of Instrumentation*, Vol. 12, No. 10, p. 2, 2017.

20. Acharya, B.S., et al., *The CMS Outer Hadron Calorimeter*, CERN, Geneva, 2006.

21. CMS Collaboration, "Output Rates and Bandwidth of the High-Level Trigger in the 2023 and 2024 Proton-Proton Runs", CMS-DP-2025-015, 2025.

22. CMS Collaboration, "Displays of Events from Run 3 Seen in CMS with Multiple Reconstructed Vertices", 2023, `https://cds.cern.ch/record/2867797`, accessed on August 8, 2025.

23. CMS Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, CERN, Geneva, 2017.

24. CMS Collaboration, *A MIP Timing Detector for the CMS Phase-2 Upgrade*, CERN, Geneva, 2019.

25. CMS Collaboration, *The Phase-2 Upgrade of the CMS Endcap Calorimeter*, CERN, Geneva, 2017.

26. CMS Collaboration, *The Phase-2 Upgrade of the CMS Muon Detectors*, CERN, Geneva, 2017.

27. CMS Collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CERN, Geneva, 2020.

28. Barney, D., "Overview Slide of CE with Main Parameters", 2022, `https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/ShowDocument?docid=13251`, accessed on July 1, 2025.

29. Bouyjou, F., et al., "HGCROC3: The Front-End Readout ASIC for the CMS High Granularity Calorimeter", *Journal of Instrumentation*, Vol. 17, No. 3, pp. 1-6, 2022.

30. Grummer, A., "Overview of the Front-End Electronics of CMS HGCal - Including Readout and Powering", *EPJ Web of Conferences*, Vol. 320, p. 1, 2025.

31. Bergamin, G., et al., "ECON-T and ECON-D: Endcap Concentrator ASICs for the CMS HGCAL", *Journal of Instrumentation*, Vol. 20, No. 3, pp. 1-6, 2025.

32. Troska, J., et al., "The VTRx+, an Optical Link Module for Data Transmission at HL-LHC", *Topical Workshop on Electronics for Particle Physics*, Santa Cruz, pp. 1-4, 2017.

33. Moreira, P., et al., "lpGBT: Low-Power Radiation-Hard Multipurpose High-Speed Transceiver ASIC for High-Energy Physics Experiments", *IEEE Transactions on Nuclear Science*, Vol. 72, No. 1, pp. 24–37, 2025.

34. Hummer, F., on behalf of the CMS Collaboration, "HGCAL SiPM-on-Tile Full-Stack Integration with the Serenity Phase-2 DAQ Hardware", *Journal of Instrumentation*, Vol. 20, No. 1, pp. 1-6, 2025.

35. CALICE Collaboration, "Construction and Commissioning of the CALICE Analog Hadron Calorimeter Prototype", *Journal of Instrumentation*, Vol. 5, No. 5, pp. 1-33, 2010.

36. De Silva, M., on behalf of the CMS Collaboration, "The SiPM-on-Tile Section of the CMS High Granularity Calorimeter", *European Physical Society Conference on High Energy Physics*, Hamburg, pp. 1-6, 2023.

37. Portalès, L., on behalf of the CMS Collaboration, "L1 Triggering on High-Granularity Information at the HL-LHC", *Instruments*, Vol. 6, No. 4, pp. 1-10, 2022.

38. Rose, A., et al., "Serenity: An ATCA Prototyping Platform for CMS Phase-2", *Topical Workshop on Electronics for Particle Physics*, Antwerp, pp. 1-5, 2018.

39. Ehle, I., on behalf of the CMS Collaboration, "Design of the CMS High Granularity Calorimeter Trigger Primitive Generator System", *Journal of Instrumentation*, Vol. 19, No. 2, pp. 1-7, 2024.

40. Rosado, M., et al., on behalf of the CMS Collaboration, "Back-End DAQ System Prototype Testing and Integration on a Full Detector Test System for the CMS HGCAL Detector", *Journal of Instrumentation*, Vol. 20, No. 3, pp. 1-6, 2025.

41. CERN, "SlinkRocket", `https://cms-daq-dth-p1.web.cern.ch/slinkrocket`, accessed on June 16, 2025.

42. Alawieh, J., et al., "Phase-2 CMS DAQ – Growing from Prototype Boards to Demonstrator Systems", *Journal of Instrumentation*, Vol. 20, No. 2, pp. 1-6, 2025.

43. Carniti, P., C. Gotti and G. P. Pessina, "ALDO2, a Multi-Function Rad-Hard Linear Regulator for SiPM-Based HEP Detectors", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Vol. 1039, pp. 1-4, 2022.

44. Guilloux, F., et al., "RAFAEL: A Clock and Data Fan-Out ASIC for CMS HL-LHC Upgrades", *Topical Workshop on Electronics for Particle Physics*, Bergen, p.1, 2022.

45. EMP Collaboration, "EMP Framework Guide", 2019, `https://serenity.web.cern.ch/serenity/emp-fwk/`, accessed on July 1, 2025.

46. Larrea, C.G., et al., "IPbus: A Flexible Ethernet-Based Control System for xTCA Hardware", *Journal of Instrumentation*, Vol. 10, No. 2, pp. 1-9, 2015.

47. IPbus Collaboration, "IPbus Builder Tool (ipbb)", 2025, `https://github.com/ipbus/ipbb`, accessed on July 1, 2025.

APPENDIX A: RIGHTS AND PERMISSIONS

This appendix records permissions for third-party figures reproduced in this thesis.

(i) **Figure 1.9:** Visualization of a reconstructed event with 55 vertices.
Source: See Ref. [22].
Rights holder: CERN, for the benefit of the CMS Collaboration.
Permission: Granted by CERN Copyright committee, on 18 August 2025, for reproduction as-is in this thesis (print + digital) in accordance with Terms of Use for CERN Audiovisual Media.

(ii) **Figure 2.1**: HGCAL cross-section schematic with region labels.
Source: See Ref. [28].
Permission: Granted by David Barney (original author) and Karol Rapacz (revised by), on 11 August 2025, for reproduction in this thesis.