System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC

A thesis submitted for the degree of

Doctor of Philosophy
In
Electronics Science and Engineering

by

PIYUSH KUMAR
(19PHPE04)

Centre for Advanced Studies in Electronics Science and Technology
School of Physics
University of Hyderabad
Hyderabad, Telangana- 500046, India

March-2024

University of Hyderabad
Hyderabad, Telangana- 500046, India

Declaration

I, PIYUSH KUMAR, hereby declare that this thesis entitled "**System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC**" submitted to the University of Hyderabad, Hyderabad, India for the award of degree of Philosophy in Electronics Science and Engineering is a record of original research work carried out by me under the supervision of **Dr. Bhawna Gomber**, Center for Advanced Studies in Electronics Science and Engineering and Technology (CASEST), School of Physics, University of Hyderabad. To the best of my knowledge, the thesis is not submitted for any degree in any University or Institute.

Date: 22 March 2024
Place: Hyderabad, India

Piyush Kumar

Mr. PIYUSH KUMAR
Reg. No. 19PHPE04

University of Hyderabad
Hyderabad, Telangana- 500046, India

Certificate

(For Ph.D. Dissertation)

This is to certify that the thesis entitled "**System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC**" submitted by "**Mr. PIYUSH KUMAR**" bearing the registration number **19PHPE04** in partial fulfillment of the requirements for the award of **Doctor of Philosophy in Electronics Science and Engineering**, is submitted to the Centre for Advanced Studies in Electronics Science and Technology (CASEST), School of Physics, University of Hyderabad, Hyderabad, India is a bonafide work carried out by him under my supervision and guidance.

This thesis is free from plagiarism and has not been submitted previously in part or in full to this or any other University or Institution for the award of any degree or diploma.

Further, the undertaken work has the following publications, talks, and poster presentations before submission of the thesis for adjudication and has produced evidence for the same in the form of an acceptance letter or the reprint in the area of his research.

Journal publications:

1. **P. Kumar** and B. Gomber, "System Design and Prototyping for the CMS Level-1 Trigger at the High-Luminosity LHC," in IEEE Transactions on Nuclear Science, vol. 70, no. 6, pp. 1075-1082, June 2023, doi: 10.1109/TNS.2023.3247510.
2. **P. Kumar** and B. Gomber, "The CMS Level-1 Calorimeter Trigger for the HL-LHC", Instruments, vol. 6, no. 4, p. 64, Oct. 2022, doi: 10.3390/instruments6040064.

Conference proceedings:

1. **Piyush Kumar**. "The Particle Flow Algorithm in the Phase II Upgrade of the CMS Level-1 Trigger". PoS, EPS-HEP2021:771, 202.

2. **Piyush Kumar**, and Bhawna Gomber on Behalf of the CMS Collaboration, "Firmware Development for the Phase-2 Upgrade of the Calorimeter L1 Trigger for the CMS Detector," in Springer Nature (under review).

Talk in Workshops and Conferences:

1. "The Calorimeter Trigger Slice Test" at Level-1 Trigger Workshop @CERN (online), 27-30 April 2022, CERN, Geneva, Switzerland.

2. "System Design and Prototyping for the CMS Level-1 Trigger at the High-Luminosity LHC", 23rd Virtual IEEE Real Time Conference, Online, 1-5, August 2022.

3. "Firmware Development for the Phase-2 Upgrade of the Calorimeter L1 Trigger for the CMS Detector", XXV DAE-BRNS High Energy Physics Symposium 2022, 12-16, December 2022, IISER Mohali, Mohali, India.

4. "System Design and Prototyping for the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC", CMS General Meeting: WGM 589, CERN, Geneva, Switzerland.

5. "Calorimeter Trigger Slice Test Results and Plans", 11-15, September 2023, Level-1 Trigger Workshop @Athens, Athens, Greece.

Poster Presentations:

1. "Calorimeter Level-1 Trigger Upgrade for the CMS Detector at the HL-LHC", XXIV DAE-BRNS High Energy Physics Symposium, 14-28 December 2020, online.

2. "The Particle Flow Algorithm in the Phase II Upgrade of the CMS Level-1 Trigger", The European Physical Society Conference on High Energy Physics 2021, Online, 26-30 July 2021.

3. "The CMS Level-1 Calorimeter Trigger for the HL-LHC", CALOR 2022 - 19th International Conference on Calorimetry in Particle Physics, 16-20 May 2022, University of Sussex, UK.

4. "The Calorimeter L1 Trigger for the HL-LHC", 12th India-Japan Science and Technology Seminar, International Conference on Frontier Areas of Science and Technology (ICFAST-2022), 09-10 September 2022, University of Hyderabad, Hyderabad, India.

5. "System Design and Prototyping for the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC", CMS Upgrade days at CERN, 06-08 February 2023, CERN, Geneva, Switzerland (Won the best poster award).

6. "System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC", 13th LHC students poster session, 27 November 2023, CERN, Geneva, Switzerland.

Further, the student has passed the following courses towards fulfillment of course work requirement for Ph.D.

S. No.	Course Code	Name	Credits	Pass/Fail
1	CS478	Pattern Recognition	3	Pass
2	CS490	Machine Learning	3	Pass
3	PE708	Research Methodology	4	Pass
4	PE716	High Speed VLSI and System On Chip: Design and Implementation	4	Pass

Gomber 22/3/2024

Dr. Bhawna Gomber
Supervisor
CASEST, School of Physics

Dr. Bhawna Gomber
Assistant Professor
CASEST, School of Physics
University of Hyderabad
Hyderabad-500046, Telangana.

M. Shanashyam 22/3/24
Head
CASEST, School of Physics

HEAD, CASEST
(Centre for Advanced Studies in Electronics Science & Technology)
School of Physics
University of Hyderabad-500 046.

22/3/2024
Dean
School of Physics

DEAN
School of Physics
University of Hyderabad
HYDERABAD - 500 046.

ABSTRACT

The High-Luminosity LHC (HL-LHC) project offers a very ambitious physics program that includes high-precision measurements of the Standard Model (SM) and the searches for new physics beyond the SM (BSM). The efficient data collection and precise events reconstruction in the harsh environment of 200 proton-proton interactions per bunch crossing are vital for achieving the success of the HL-LHC program. To fulfill these requirements, the CMS experiment plans to build and install completely new data acquisition (DAQ) and trigger systems during the so-called CMS Phase-2 upgrade. The Phase-2 CMS Level-1 calorimeter trigger system will handle the enormous detector input data bandwidth of 75 Tbps and is desired to complete the single event processing within 12.5 μs. For this purpose, CMS plans to replace the Phase-1 μTCA-based processor boards and crates with an ATCA form factor. Each ATCA board will host Xilinx Ultra-Scale+ family FPGA that supports over a hundred high-speed optical links at 25 Gbps, capable of meeting the high bandwidth and processing requirements of the HL-LHC. Along with the advancement in hardware, the Level-1 trigger system will employ highly modular, flexible, and adequately sophisticated algorithms currently possible only in offline reconstruction, such as a particle-flow algorithm. The modular and flexible architecture will help to address the HL-LHC physics requirements. In this thesis, we will discuss the system design, prototyping, and algorithms being developed for the Phase-2 Level-1 Calorimeter trigger system.

Acknowledgements

First and foremost, I would like to express my deepest gratitude toward my Ph.D. supervisor **Dr. Bhanwa Gomber**. This Ph.D. journey would not have been possible without your immense support and mentorship. Thank you for opening a new world to me by allowing me to work in the CMS collaboration and helping me grow as a better person, personally and professionally. Your passion for the work still inspires me to push myself more. I could not have imagined having a better supervisor and mentor in my Ph.D. journey.

I want to express my sincere gratitude to **Prof. Sridhara Dasu** and **Dr. Alexander (Sascha) Savin** from the University of Wisconsin for providing me with the opportunity to work on the Phase-2 Level-1 calorimeter trigger project. Thank you for providing me with financial and academic support for my one-year stay at CERN. Moreover, I want to thank Sascha for his continuous and tremendous support during my Ph.D. journey. I will always adore the wisdom I have received from him during my Ph.D.

I want to express my special gratitude toward my international collaborators, **Dr. Isobel Ojalvo**, **Dr. Varun Sharma**, **Ales Svetek**, **Luis Moreno**, **Kiran Kumar Das**, and **Dr. Alexander Zabi**. Their academic and technical support made this journey smooth and enjoyable.

I want to thank my Ph.D. review committee members, **Prof. S.L. Sabat**, **Prof. S. V. S. Nageswara Rao**, **Prof. G Rajaram**, and **Dr. Vijaya Shankara Rao**, for their insightful discussions and comments during doctorate review presentations. It gave me a broad research perspective and encouraged me to improve positively.

I want to express my sincere gratitude to Dean, School of Physics **Prof. K.C. James Raju**, and Head of CASEST **Prof. Prof. Ghanashyam Krishna** for providing the necessary facilities and an enjoyable working experience at the University.

I take this opportunity to thank **Dr. Bhawna Gomber**, **Prof. Guruswamy Rajaram**, **Prof. K.C.James Raju**, **Prof. M. Ghanashyam Krishna**, **Prof. S.V.S Nageswara Rao**, **Dr. Vijaya Sankara Rao**, and **Dr. Pratap Kollu** for their teaching and support during my coursework in my Ph.D. and M.Tech.

I want to express my profound appreciation to my colleague **Bisnupriya Sahu** for her constant and pristine support in my Ph.D. journey. Together with her, this Ph.D. journey was extremely enjoyable and memorable.

I am thankful to my colleagues **Shriniketan**, **Shiva**, **Anil**, **Vinay**, **Ravi**, **Gnani**, **Chenna**, and **Rakesh**. They made this journey memorable and enjoyable.

I want to express my special thanks to my mother **Ms. Rekha Gupta**, my father **Mr. Vijay Kumar**, and my maternal grandmother **Ms. Parvati Gupta** for encouraging me in my pursuits and inspiring me to follow my passion. Thank you for teaching me lessons of life that made me the person I am today.

I take this opportunity to especially thank my sister **Sneha Sahu** and brother **Vineet Gupta** for their constant support, love, and care throughout my life.

I take this opportunity to especially thank **Mr. Antoni Jakubowski** and **Dr. Varun Sharma** for their tremendous support and help for my smooth stay at CERN. I learned many new dimensions of life, which eventually helped me personally and professionally. I am also thankful to Kiran for guiding me during my stay at CERN.

I would also very much like to thank all the non-teaching staff of CASEST and the School of Physics, especially **Ravi Babu** sir for guiding me in all the administrative work.

Finally, I would like to thank and acknowledge the support for the fellowship from different funding agencies: the **University of Hyderabad**, the **SERB startup grant**, and the **IoE Grant**. I would also like to thank the University Institute of Eminence office for supporting me in giving a talk at the L1 trigger workshop in Athens, Greece.

Contents

List of Figures

List of Tables

Abbreviations

CERN	European Organization for Nuclear Research
LHC	Large Hadron Collider
CMS	Compact Muon Solenoid
ATLAS	A Toroidal LHC Apparatus
ALICE	A Large Ion Collider Experiment
LHCb	Large Hadron Collider beauty Experiment
HL-LHC	High-Luminosity Large Hadron Collider
ECAL	Electromagnetic Calorimeter
HCAL	Hadronic Calorimeter
HGCAL	High-Granularity Calorimeter
HF	Hadronic Forward
LUTs	LookUp Table
FFs	Flip-flops
PF	Particle Flow
RCT	Regional Calorimeter Trigger
GCT	Global Calorimeter Trigger
APx	Advanced Processor
APd1	Advanced Processor first demonstrator
CT	Correlator Trigger
GT	Global Trigger
TPs	Trigger Primitives
FE	Front-end Electronics
HLT	High-Level Trigger
DAQ	Data Acquisition
ATCA	Advanced Telecommunications Computing Architecture

μTCA	Micro Telecommunications Computing Architecture
MGTs	Multi-Gigabit Transceiver
AXI	Advanced eXtensible Interface
HLS	High-Level Synthesis
RTL	Register transfer level
HDL	Hardware description language
VHDL	Very High-Speed Integrated Circuit Hardware Description Language
SLR	Super Logic Region
SLL	Super Long Line
FPGA	Field Programmable Gate Array
DSP	Digital Signal Processor

Dedicated to my family and my friends. . . .

Chapter 1

Introduction

The Large Hadron Collider (LHC) is the most prominent and powerful particle accelerator and collider [1]. It accelerates two counter-rotating beams of particles in a 27 km long tunnel and collides them at four interaction points: ATLAS, ALICE, CMS, and LHCb. CMS is built to probe the aftermath of collision by reconstructing and identifying the particles. It is installed 100 meters underground near the French commune of Cessy, amidst the mountains of Jura and Lake Geneva [2]. One of the significant achievements of the LHC experiments is the discovery of the Higgs boson particle [3], announced by the CMS and ATLAS experiments in the year 2012. The data mined through the collisions help physicists in their quest to find the answers to some of the most challenging questions of modern-day physics. These include the comprehensive study of the Higgs boson particle, the existence of supersymmetry (SUSY), physics beyond the Standard Model (BSM), dark matter, etc. [4].

The CMS detector helps in the study and measurements of rare occurring physics phenomena, which requires a large amount of data to be delivered by the LHC. Luminosity is a crucial performance indicator of an accelerator: it is proportional to the number of collisions that occur in a given amount of time. A higher luminosity indicates a greater likelihood of particles colliding [5]. The planned High-Luminosity LHC (HL-LHC) will be the upgraded version of the LHC, aiming to deliver a large amount of collision data to the CMS. During collisions, the CMS detector encounters billions of events per second. Storing and processing these events require massive storage and hardware infrastructure capacity. However, many of these events do not carry any rare or significant physics information and could be discarded early in data acquisition. The CMS detector employs the trigger mechanism to perform this rapid rate reduction by efficiently selecting rare and interesting events and rejecting the noise. The CMS trigger system performs this rate reduction in two steps [6]. The first step is the Level-1 trigger system, which is based

on a custom-designed electronics board (hosting FPGAs), which reduces the incoming rate (LHC bunch-crossing) from 40 MHz to 110 kHz (which will increase to 750 kHz during HL-LHC). The second step employs a CPU and GPU-based High-Level Trigger (HLT) to facilitate the rejection from 110 kHz to 2.6 kHz (which will increase to 7.5 kHz during HL-LHC) [7]. The CMS Level-1 trigger system performs this rate reduction by taking the sub-detector information, analyzes the collision data using sophisticated algorithms to build physics objects, such as electrons, photons, jets, tau leptons, muons, and missing transverse energy (MET), and forward this information to the global trigger for the final decision.

The period of CMS and data taking is classified into two phases. The Phase-1 refers to the Run-1, Run-2, and Run-3 data-taking periods during the years 2011-12, 2015-18, and 2022-2025, respectively. The Phase-2 period refers to the HL-LHC operation period, including Run-4, Run-5, and Run-6 operations designated during 2029-32, 2035-38, and 2040-41 years, respectively. The next section discusses the design and drawbacks of the Phase-1 calorimeter trigger system, which demands designing a calorimeter trigger system.

1.1 Literature Review

The CMS Phase-1 Level-1 trigger system is designed and currently operating at the instantaneous luminosity of $2.0 \times 10^{34} m^{-2} s^{-1}$ with a maximum pile-up of 70. With the input bunch-crossing rate of 40 MHz, the Level-1 trigger performs the event selection in 3.8 μs, providing an output rate of 110 kHz. The Phase-1 Level-1 calorimeter trigger takes coarse input from three calorimeter sub-systems: electromagnetic calorimeter (ECAL), hadronic calorimeter (HCAL), and forward hadronic (HF). It is implemented using a two-layer strategy with the calorimeter layer-1 comprising 18 calorimeter trigger processor boards (CTP7) [8] that receive the ECAL and HCAL inputs, perform pre-processing, and provide the output to the layer-2 boards in a time-multiplexing manner. The 10 layer-2 master processor boards (MP7) [8] host various calorimeter trigger algorithms that identify the calorimeter particle candidates and calculate the global energy sums.

The maximum input bandwidth of the calorimeter Layer-1 board is 10 Gbps, which is sufficient to collect the tower information of the ECAL, HCAL, and HF. With an eight times increment in luminosity (reaching up to $7.5 \times 10^{34} m^{-2} s^{-1}$) and, subsequently, an increase in the average pile-up to 200, a more precise triggering strategy is required, mimicking the offline reconstruction methodology at Phase-2. Therefore, it requires full granularity of the detector to be processed at the Phase-2 calorimeter trigger. It involves

a crystal level of information, which increases the input granularity by 25 times compared to the Phase-1 architecture. The processing of crystal information at the Phase-2 calorimeter trigger is made possible using large FPGAs and optical links supporting 25 Gbps of link bandwidth, which was not possible using Phase-1 architecture.

Additionally, the massive influx of radiation in the forward region of the detector requires a new radiation-hardened high-granularity calorimeter (HGCAL) [9] and the corresponding triggering infrastructure at the Level-1 calorimeter trigger, which is currently missing in the Phase-1. In addition, the Phase-1 calorimeter trigger employs Xilinx Virtex-7 FPGAs as the central processing element, which is unable to satisfy the demand of the Phase-2 calorimeter trigger, that requires additional MGTs with high bandwidth, more logic capacity in terms of look-up table (LUT), flip-flops (FF), and digital signal processor (DSP).

On the algorithm side, the Phase-1 system works over the tower's granularity. To process the crystal information, it is necessary to redefine the algorithms and clustering strategies. The Phase-1 algorithm lacks the particle-flow (PF) clustering vital for the PF algorithms at the correlator trigger. Therefore, a completely new calorimeter trigger system is required during HL-LHC.

The Phase-2 Level-1 trigger system, which satisfies the above constraints, receives input from the calorimeter, muon spectrometer, and track finder. The calorimeter trigger processes the barrel information of the electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL) in the regional calorimeter trigger (RCT), also known as the barrel calorimeter trigger (BCT). The global calorimeter trigger (GCT) computes the upstream information from RCT, the forward hadron calorimeter (HF), and the HGCAL. The muon trigger system collects the inputs from various muon spectrometers, such as the resistive plate chambers (RPC), drift tube (DT), cathode strip chambers (CSC), and gas electron multiplier (GEM). In the local sub-division, the muon trigger comprises the muon track finder and the barrel layer-1 processor. The muon track finder processes the input from three distinct regions called BMTF, EMTF, and OMTF for barrel, endcap, and overlap, respectively.

For the first time, the tracker information is included in the Level-1 trigger, which is implemented at the global track trigger (GTT). The GTT sends the tracker output to the correlator trigger and global trigger. This thesis work is entirely dedicated to the calorimeter trigger. The latency budget for the Phase-2 L1 trigger system is proposed to be 12.5 μs with a maximum output rate of 750 kHz. The calorimeter trigger is expected to complete the single event processing within 3 μs for the barrel and 4 μs, including the endcap. One of the main highlights of the Phase-2 Level-1 trigger system is the implementation of the particle-flow algorithm [10], organized in two layers and

called a correlator trigger (CT). It helps to identify higher-level objects such as photons, electrons, muon, taus, jets, and others. The final decision of the Level-1 trigger is executed by the global trigger (GT). Fig. 1.1 represents the architecture of the Phase-2 Level-1 trigger, which includes the calorimeter trigger.

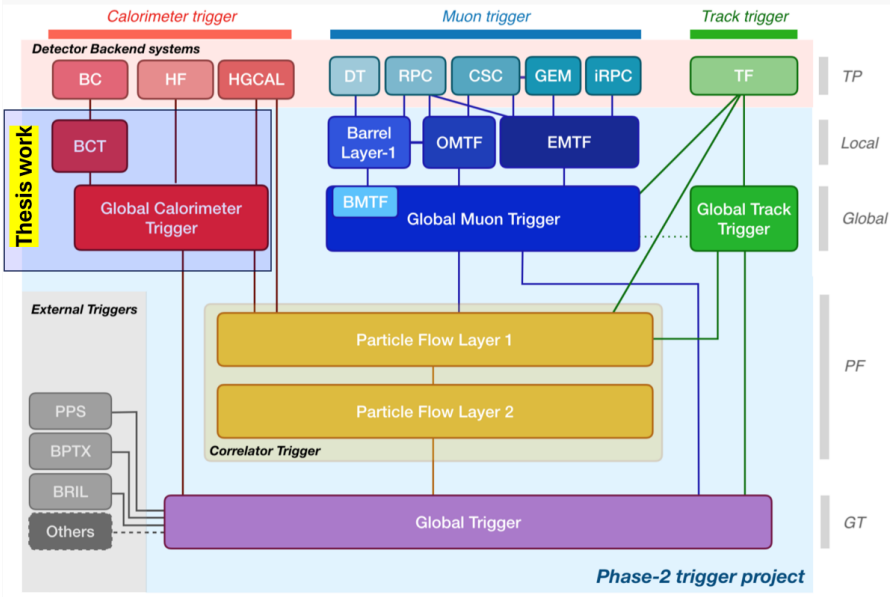

FIGURE 1.1: CMS Phase-2 Level-1 trigger system architecture comprising three trigger sub-systems: calorimeter, muon, and track [6]. The correlator trigger is also included in the L1 trigger architecture. The thesis work entirely focuses on the design and prototyping of the calorimeter trigger highlighted in the box

The next section describe the problem definition, requirements and the associated constraints of the Phase-2 calorimeter trigger system.

1.2 Problem Definition

As mentioned earlier, the FPGA boards are the central processing elements of the Level-1 trigger system, which hosts the trigger algorithms. Therefore, the architecture design is a quest to optimize the board processing region and interconnections to accomplish the required physics strategy with the lowest number of FPGA boards and target latency while maintaining headroom for flexibility and robustness.

The Phase-2 calorimeter trigger plans to handle an enormous bandwidth of 75 Tbps and has access to the information of each crystal in the barrel and tower from the HGCAL. Therefore, the increased granularity of the calorimeter system requires designing new algorithms. Ultimately, designing algorithms with high algorithm clock frequency and bandwidth pose an enormous challenge in the final implementation of the algorithm circuit.

The HL-LHC plans to operate during the year 2029-2041 (13 years). The long year of HL-LHC operation requires a contingency of latency and device logic, which should be reserved for future endeavors. It is necessary to implement calorimeter trigger to fit the current constraint while preserving 50 percent of the logic fabric for any potential future up-scope.

With all the constraints defined above, the proposed system must be prototyped using a minimum number of FPGA boards and employing the calorimeter trigger algorithms to cover all the proposed physics cases. The next section describe the thesis objective.

1.3 Thesis Objective

The main objective of the thesis is to design and implement the Phase-2 calorimeter trigger system, which takes input from four calorimeter sub-detectors.

- Proposing an architecture for the Phase-2 calorimeter trigger by leveraging the unique detector symmetry and identical division of the calorimeter geometry for FPGA processing by optimizing the algorithm and board connectivity with minimum sharing of overlapping regions, utilizing fewer boards and optical links.
- Implementing the barrel calorimeter trigger algorithm and firmware, which includes the information from ECAL and HCAL within the latency budget of 3 μs.
- Implementing the endcap trigger and providing the final output to the global trigger (GT) within the latency budget of 4 μs.
- Implementation of the entire calorimeter trigger with a reserved contingency of 50% of device logic.
- Prototyping and validating the entire system with a handful of FPGA boards and optical connections.

1.4 Proposed Solution and Thesis Contribution

This thesis proposes an architecture of the Phase-2 calorimeter trigger, designed in two levels: regional calorimeter trigger (RCT) and global calorimeter trigger (GCT). The RCT processes the barrel region of the calorimeter using 36 APx boards [6]. In this thesis, an RCT algorithm is developed and validated on the APd1 board, which

prepares the electron/photon clusters and tower information and send it to the GCT barrel board.

The second level is the GCT, which processes the barrel and endcap information using 10 APx boards. The processing of the barrel information from the 36 RCT boards is proposed to be implemented using three GCT barrel boards. In this thesis, a GCT barrel algorithm is developed and validated on the APd1 board. The algorithm includes stitching the RCT electron/photon cluster at the *eta* and *phi* boundaries and performing the clustering of the barrel towers, such as particle flow (PF) clustering, jet, and taus. It also computes the partial MET information (another portion is implemented in the GCT sum algorithm) and time-multiplex the final output for the correlator trigger. The output of the GCT barrel is sent to the correlator (CT) and global trigger (GT).

The endcap part of the detector, i.e., the information from the HGCAL and HF, is planned to be processed using 6 GCT endcap boards. In this thesis, a GCT endcap algorithm is developed and validated on the APd1 board, which demultiplexes the input and performs the jets and taus clustering. The output of the algorithm is sent to the final GCT sum board.

The GCT sum board takes input from the GCT barrel and endcap boards in terms of calorimeter-only objects. In this thesis, a GCT sum algorithm is developed and validated on the APd1 board, which transforms the calorimeter object information as per the requirement of the global trigger, computes the final MET quantity using the CORDIC algorithm, and prepares the output for the global trigger (GT) in a time-multiplex manner.

Infrastructure-wise, the planned architecture foreseen 46 APx boards, 2988 input optical links from barrel calorimeter, HF, and endcap HGCAL, 144 output links to the correlator trigger, and six output links to the global trigger. In the thesis, the overall system is prototyped using only four APd1 board connected using 16 optical links.

Fig 1.2 and Fig 1.3 represent the APd1 board based on APx philosophy and block diagram of the proposed Phase-2 Level-1 calorimeter trigger system.

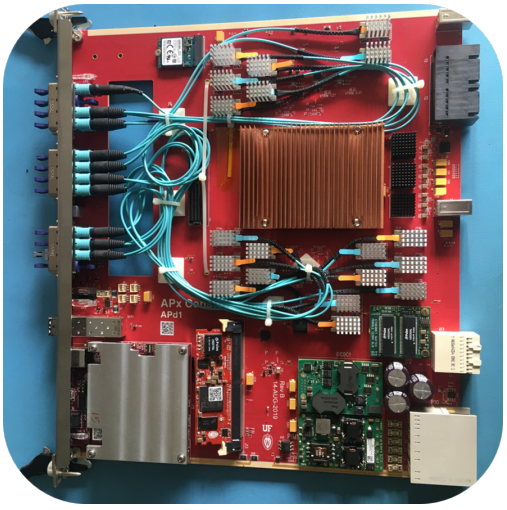

FIGURE 1.2: APd1 board with the XCVU9P FPGA beneath the heatsink. APd1 board is used to prototype the calorimeter trigger system.

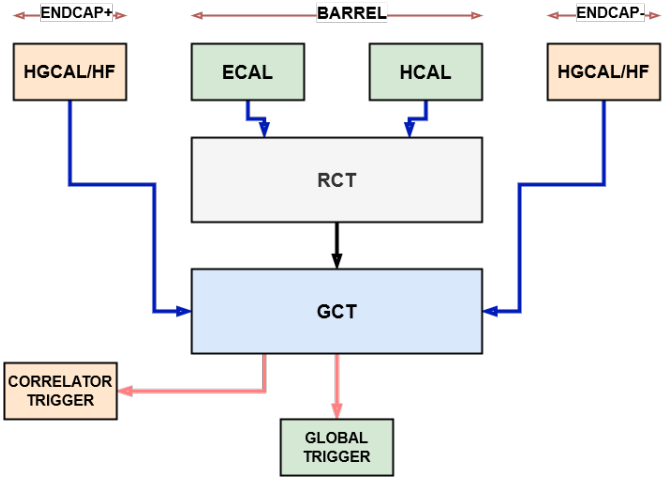

FIGURE 1.3: Two level architecture of the Phase-2 Level-1 calorimeter trigger.

1.5 Thesis Organization

This thesis presents the architectural design, algorithm, and implementation details of the proposed Phase-2 Level-1 Calorimeter Trigger.

1.5.1 Chapter 1: Introduction

This chapter introduces the LHC, CMS, Level-1 trigger, and HL-LHC. It highlights the features of the Phase-1 Level-1 trigger, its drawbacks, and the requirement for a new Phase-2 calorimeter trigger. This chapter also describes the problem definition, proposed solution, thesis contribution, thesis objective, and thesis organization.

1.5.2 Chapter 2: The CMS experiment at the LHC-CERN

This chapter describes LHC design and its accelerator chain. The salient features of the CMS detector and its components, such as the solenoid magnet, tracker, calorimeter, and muon system, are presented. The detector's main electronics components, such as the data acquisition (DAQ) and trigger system, including the hardware-based Level-1 trigger, are also discussed.

1.5.3 Chapter 3: Triggering in High Luminosity LHC

This chapter briefly describes the main highlights and benefits of HL-LHC. It briefly outlines the details of the CMS Phase-2 upgrade, which includes the advancement of pixel, tracker, muon system, calorimeter, data acquisition (DAQ), and Level-1 trigger.

1.5.4 Chapter 4: Phase-2 Level-1 Calorimeter Trigger

This chapter discusses the proposed architecture of the Phase-2 Level-1 calorimeter trigger and its salient feature. It summarizes the hardware infrastructure of the Phase-2 Level-1 calorimeter trigger, such as the details of the ATCA (Advanced Telecommunications Computing Architecture) specification employed for the electronic system of the Phase-2 Level-1 trigger. The details of the APd1 board and its main sub-component are outlined. The design flow, which includes the phases of Monte Carlo simulation, high-level synthesis (HLS), RTL simulation and integration, and final bitfile test, is discussed. Finally, major system design and prototyping constraints are discussed.

1.5.5 Chapter 5: Calorimeter Trigger Algorithm

This chapter describes the various calorimeter objects used and prepared in the Phase-2 Leve-1 calorimeter trigger. It briefly describes the barrel calorimeter geometry and its partitioning for efficient algorithm implementation and prototyping. The regional (RCT) and global calorimeter trigger (GCT) architecture and algorithms are described, along with their performance in terms of latency, device utilization, and F_{MAX} are discussed.

1.5.6 Chapter 6: Calorimeter Trigger Firmware and Prototyping

This chapter describes the calorimeter trigger prototype and test setup. It outlines the test methodology for single-board and multi-board tests. It also discusses the implementation details and single board tests of each calorimeter trigger algorithm, such as

RCT, GCT barrel, GCT endcap, and GCT sum. The multi-board test between the RCT and GCT barrel and the complete prototyping of the calorimeter trigger using only four APd1 boards are also discussed.

1.5.7 Chapter 7: Summary, Conclusions, and Future Work

This chapter summarizes the advantages and details of the overall architecture of the Phase-2 Level-1 calorimeter trigger. The achieved results on algorithms, firmware, and prototyping results of the RCT, GCT (GCT barrel, endcap, and sum), and the complete calorimeter trigger system are outlined.

Chapter 2

The CMS experiment at the LHC-CERN

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is a two-counter-rotating particle accelerator and collider established in the 26.7 km tunnel built between the years 1984 and 1989 for the CERN Large Electron Positron (LEP) machine. The LHC tunnel features 8 arcs and 8 linear sections. It is placed between 45 and 170 m underneath the ground, inclining 1.4% towards the lake Geneva. About 90% of its circumference is in molasse rock, featuring ideal characteristics for this application. The rest 10% are in the limestone beneath the Jura mountain.

The LHC is designed to deliver collisions with the peak energy of $\sqrt{s} = 14$ TeV. During Run-3 (2022-2025), the particles are collided at $\sqrt{s} = 6.8$ TeV. The instantaneous luminosity achieved in Run-3 is $2 \times 10^{34} cm^{-2} s^{-1}$ with the maximum pile up of 70. The LHC also collides heavy ions (Pb-Pb) for a shorter operating duration of approximately one month per year.

In LHC, the two beams orbit the ring inside vacuum tubes guided by the superconducting magnets. Thousands of such magnets supervise the beams in the LHC, including 1232 35-ton and fifteen-meter-long dipole magnets, as shown in Fig. 2.1. These dipole magnets feature a special design called a "2-in-1" configuration, allowing the two beams to orient in opposite directions in the same dipole (with two different beam pipes). The dipoles are configured with a cooling temperature of 1.9 K, with a current flow of 11.7 kA, delivering an 8.4 T magnetic field to support the operation with a 7 TeV energy beam.

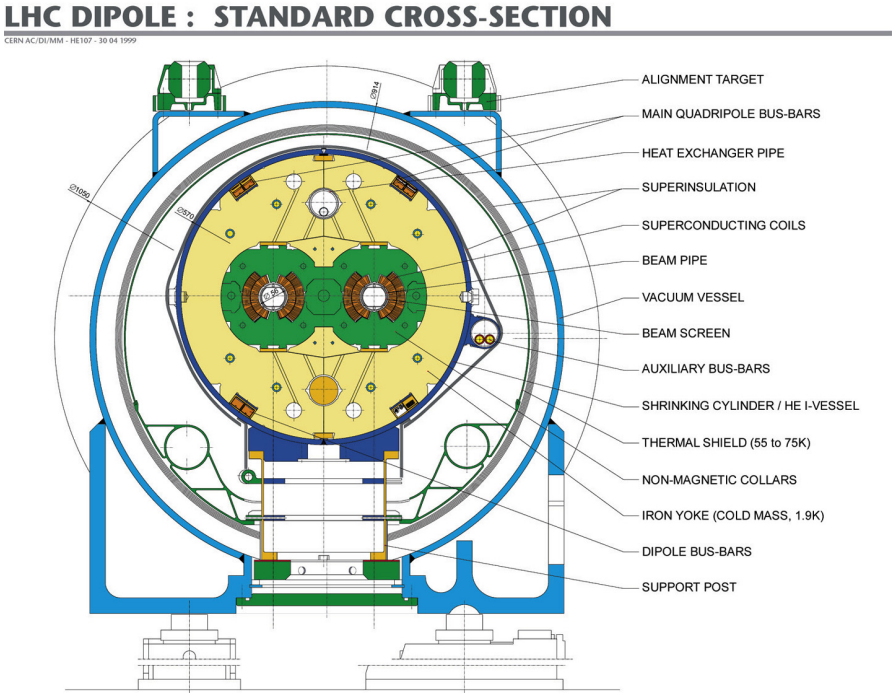

FIGURE 2.1: Cross-section of the LHC dipole [1].

The LHC collides two beams with a maximum of 2808 proton bunches per beam. Each bunch is several centimeters long and retains about 100 billion protons. The probability of a hard collision increases by squeezing the beam as much as feasible at the interaction point of the experiment to a diameter of tens of microns. It is expected to encounter an average of 70 additional proton-proton interactions (pile-up) for these operating design conditions [7].

The CERN injector complex advances the proton beams to 450 GeV (0.45 TeV) of kinetic energy. It accomplishes this through a sequence of accelerators: a linear accelerator (Linac 2) followed by three synchrotrons such as the Proton Synchrotron Booster (PSB) [11], the Proton Synchrotron (PS) [12], and the Super Proton Synchrotron (SPS) [13]. At first, protons with 50 MeV of energy are yielded in the Linac 2 and accelerated to 1.4 GeV in the PSB. The beams are then injected in the PS, gaining energy of 26 GeV and further accelerated to 450 GeV in the SPS. Ultimately, they are injected into the main LHC circle, reaching a maximum energy of 7 TeV. The injector complex also equips the LHC with heavy ion beams, first accelerated via a linear accelerator (Linac 3) and the Low Energy Ion Ring (LEIR) synchrotron and subsequently by the PS and the SPS. Besides delivering beams to the LHC, the CERN injectors also serve several fixed-target experiments, including the CLOUD [14], COMPASS [15], and many others. Fig. 2.2 shows the injector chain of the LHC along with the experiments located on the LHC circumference.

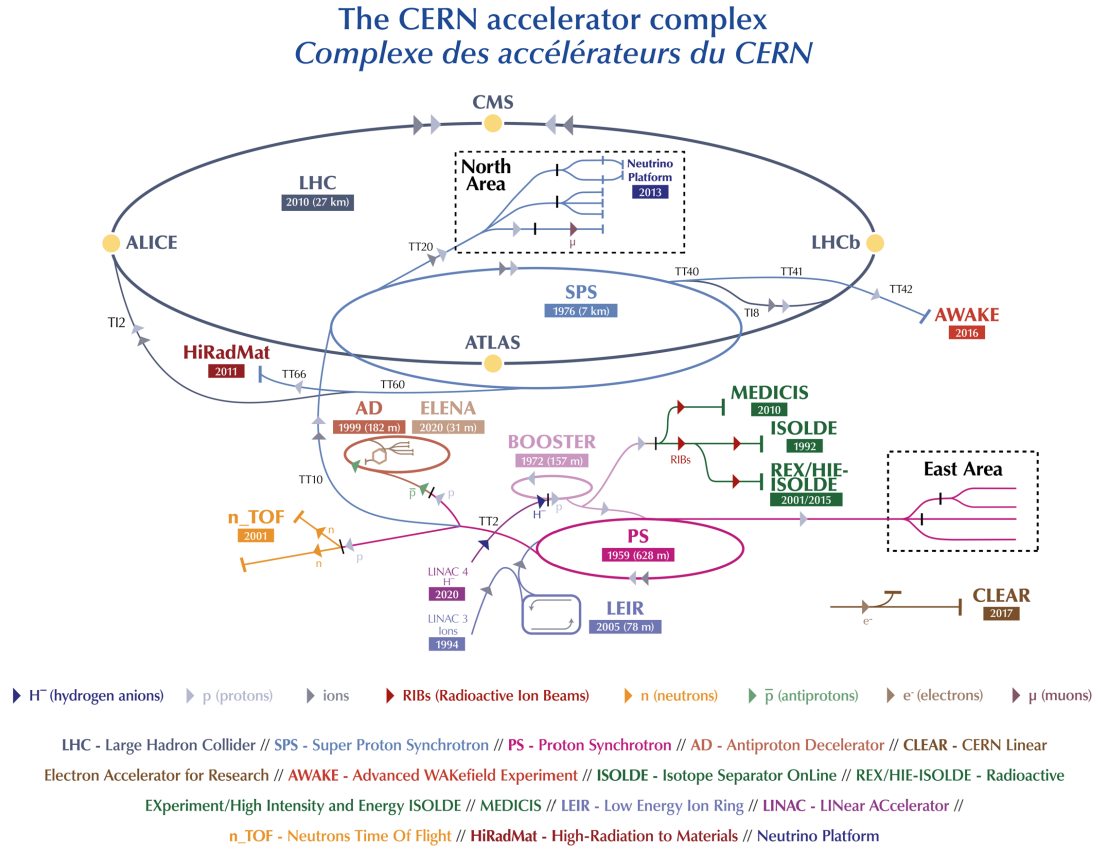

FIGURE 2.2: CERN injector chain and LHC[1].

The LHC hosts nine experiments, including the Compact Muon Solenoid (CMS) detector. The term luminosity is one of the critical parameters of the particle collider. The concept of Luminosity is discussed in the next sub-section.

2.2 The Concept of Luminosity in Colliders

The study and measurements of the rare occurring physics phenomena require a large amount of collision data. The term luminosity estimates the number of collisions that occur in a detector per cm^{-2} and per second. It is the proportional factor between cross-section σ_p and the number of events per second N. Therefore, the unit of luminosity $\mathcal{L}$ is $\mathrm{cm}^{-2}\ \mathrm{s}^{-1}$.

$$\frac{dN}{dt} = \mathcal{L} \cdot \sigma_p \tag{2.1}$$

For two colliding Gaussian beams as shown in Fig. 2.3 the expression of luminosity is:

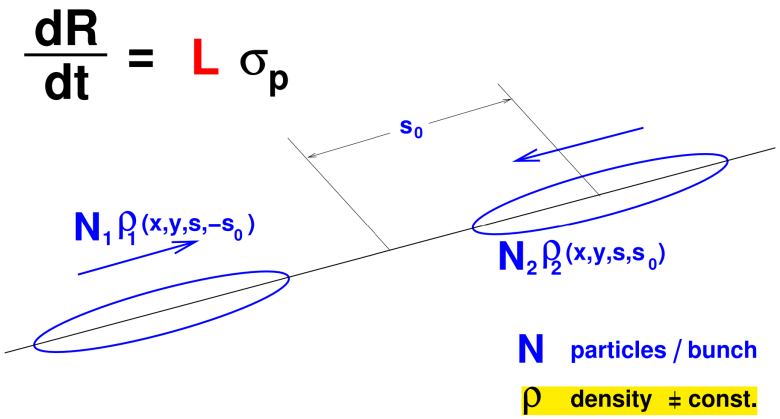

FIGURE 2.3: Graphical representation of two colliding beams [16].

$$\mathcal{L} = \frac{N_1 N_2 f N_b}{4\epsilon\beta^*} \tag{2.2}$$

Where the N_1 and N_2 represent the number of particles in the bunch, f is the revolution frequency, with N_b representing the number of bunches per beam, ϵ represent the emittance and β^* depicting the amplitude function [16].

The LHC recorded luminosity over a given amount of time also known as the integrated luminosity is:

$$L_I \equiv \int_{\Delta t} L dt \tag{2.3}$$

The integrated luminosity directly depends on the instantaneous luminosity and also relies on the machine operating time (Δt) in the collision mode. The enhancement of the operating time is as important and challenging (when the machine run in a highly constrained environment and pushed to the boundaries of its acceptable values) as increasing the instantaneous luminosity. The integrated luminosity is expressed in inverse of cross-section (1/fb or $femtobarn^{-1}$). One inverse femtobarn ($1\ fb^{-1}$) signifies the collisions of approximately 70×10^{12} protons.

LHC has attained the luminosity of $1 \times 10^{34} cm^{-2} s^{-1}$ on 26 June 2016. The peak luminosity of $2 \times 10^{34} cm^{-2} s^{-1}$ was achieved in 2018 due to the reduction in β^* from the nominal value of 55 cm to 30 cm and also due to the small injected beam emittance. This high luminosity and a longer and excellent operational period helped reach the integrated luminosity of 65 fb^{-1} in 2018. The Run-3 agenda is to collect the integrated luminosity of 350 fb^{-1} by the end of 2025.

2.3 The Compact Muon Solenoid (CMS) Experiment

The Compact Muon Solenoid (CMS) detector is a general-purpose detector operating at the Large Hadron Collider (LHC) at CERN. CMS is established about 100 meters below ground, near the French village of Cessy, between the Jura mountains and lake Geneva. CMS is cylindrical in shape, with several concentric layers of sub-detectors installed to precisely measure the tracks, energy, and momentum of the particles emerging from the collisions. The layout and features of the CMS detector are discussed in the next sub-section.

2.3.1 CMS coordinate system

The CMS has adopted the coordinate system with the origin located at the collision point of the detector. The y-axis directs vertically upward, and the x-axis directs radially inward toward the ATLAS detector. The z-axis direct toward the beamline. The x-axis in the x-y plane determines the azimuthal angle ϕ, and the radial position is defined by r. The z-axis in the z-y plane defines the polar angle θ. The Pseudorapidity is described as η = -ln[tan(θ/2)] [17]. The energy and momentum transverse to the beam direction, represented by E_T and p_T, respectively, are estimated from the x and y components. Fig 2.4 represents the CMS coordinate system [2].

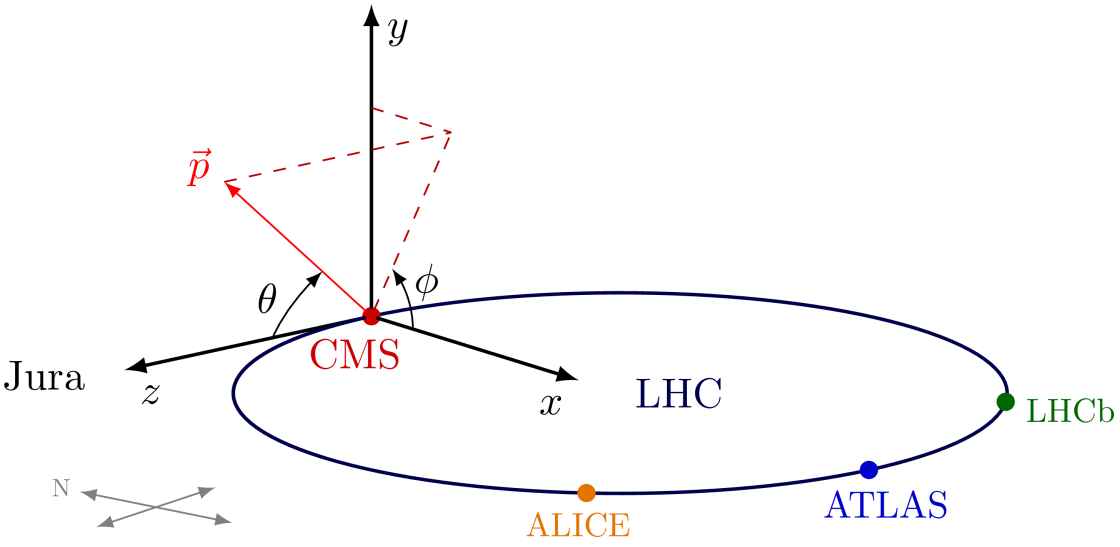

FIGURE 2.4: CMS coordinate system.

2.3.2 Detector Layout and Salient Features

The CMS detector is 21.6 m long and has a diameter of 14.6 m. It has a total weight of 12500 tons. At the core of CMS sits a superconducting solenoid with a dimension of

13m in length, 6m in inner diameter, and 4T of magnetic field providing ample bending power (12 Tm) before the muon system, which calculates the muon bending angle. The return field incorporates four muon stations to ensure full geometric coverage and robustness. The muon system comprises the aluminium-based drift tube (DT) in the barrel region, and the endcap is covered by the cathode strip chambers (CSC). Another muon system complements both regions, known as resistive plate chambers (RPC). Fig. 2.5 represents the layout of the CMS detector highlighting the sub-detector installed at various positions.

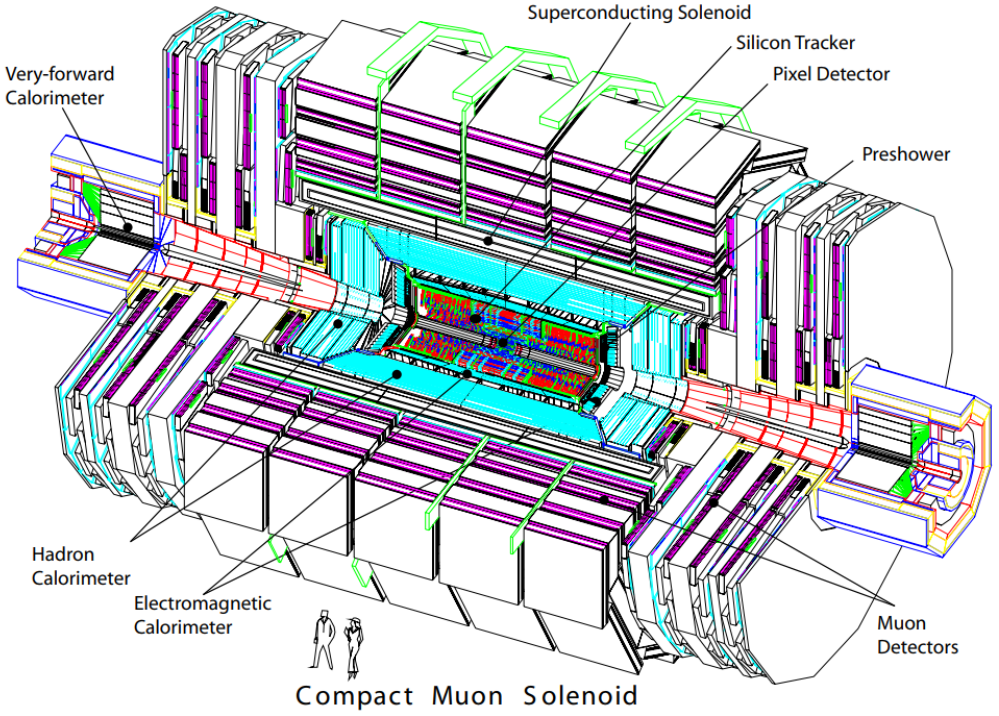

FIGURE 2.5: CMS detector [2].

The inner bore of the solenoid coil houses the tracker and the calorimetry inside. The CMS tracker volume is characterized by a cylinder of size 2.6 m in diameter and 5.8 m in length. CMS utilizes ten silicon microstrip detectors to manage the high-track multiplicities, delivering the required precision and granularity. In expansion, three layers of silicon pixel detectors are established proximate to the interaction region to enhance the measurement of the charged-particle tracks and the position of secondary vertices. CMS electromagnetic calorimeter (ECAL) utilizes lead tungstate ($PbWO_4$) crystals, which cover the regions in pseudorapidity up to $|\eta| < 3.0$. The silicon-based avalanche photodiodes (APDs) in the barrel region and vacuum phototriodes (VPTs) in the endcap, capture and detect the scintillation rays. A preshower system is installed before the endcap ECAL for π^0 rejection [2].

The ECAL is encircled by a brass/scintillator sampling hadron calorimeter (HCAL) with a coverage range of up to $|\eta| < 3.0$. The scintillation light is transformed by wavelength-shifting (WLS) fibers integrated with the scintillator tiles and directed to the photodetectors. The forward calorimeters provide full geometric coverage for estimating the transverse energy in the event. Further, forward coverage is achieved with other diverse sets of dedicated calorimeters (ZDC [2]) and the TOTEM [18] tracking detectors. The following section briefly describes the CMS ECAL and HCAL.

2.3.3 Electromagnetic Calorimeter

The Higgs boson discovery via exploring the 'golden channel' $H \rightarrow \gamma\gamma$ would not be achievable without an excellent and superior electromagnetic calorimeter (ECAL). The ECAL was devised to have fine granularity and excellent energy resolution, which permits the distinction of two photons and, as an outcome, suppresses the background comprised by $\pi^0 \rightarrow \gamma\gamma$ decays.

The ECAL is a homogeneous calorimeter constructed by mounting 61200 lead tungstate ($PbWO_4$) crystals in the central barrel part, sealed by 7324 crystals in each endcap. A preshower detector is positioned in front of the endcaps. As photodetectors, the barrel has employed Avalanche photodiodes (APDs), and the endcap uses vacuum phototriodes (VPTs). The employment of high-density crystals has permitted the design of a fast calorimeter, which has radiation resistance and fine granularity, all essential elements in the LHC environment.

The features of the lead tungstate crystals make them an appropriate option for function at LHC [19]. The high density (8.28 g/cm^3), small radiation length (0.89 cm), and short Molière radius (2.2 cm) result in a fine granularity and a compact calorimeter. The light beam degradation interval of these crystals closely follows the LHC bunch crossing time of 25 ns to emit approximately 80% of the light. The light output of these crystals is relatively low, and changes with temperature (-2.1% $^0C^{-1}$ at 18^0C): at 18^0C, about 4.5 photoelectrons per MeV are accumulated in both VPTs and APDs. Fig. 2.6 represents the lead-tungstate crystal and APD used in the barrel calorimeter (EB). Fig. 2.7 represents the lead-tungstate crystal and VPT used in the endcap calorimeter (EE).

2.3.3.1 ECAL Layout

The pseudorapidity coverage of the ECAL barrel (EB) region is $|\eta| < 1.479$. The EB granularity is 2×85-fold in η, and 360-fold in ϕ resulting in 61200 crystals. The crystals maintain a tapered profile, varying slightly with the position in η. The crystals are

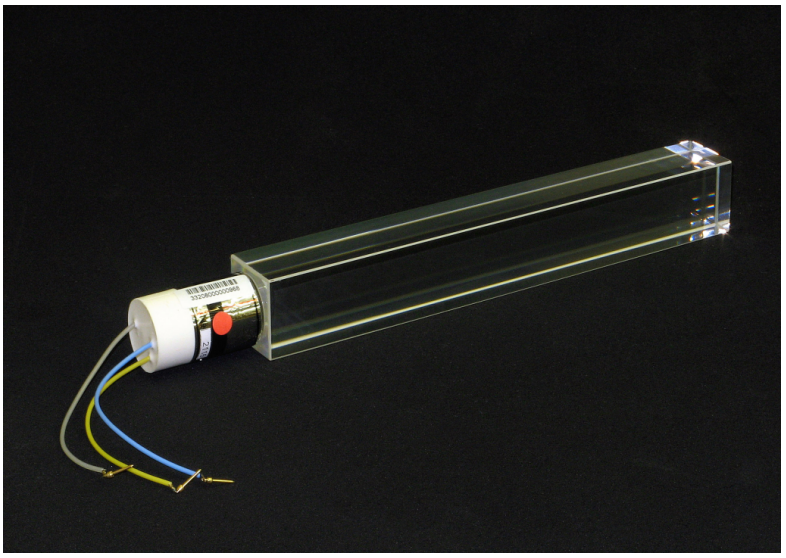

FIGURE 2.6: $PbWO_4$ crystal with APD photodetector attached [2].

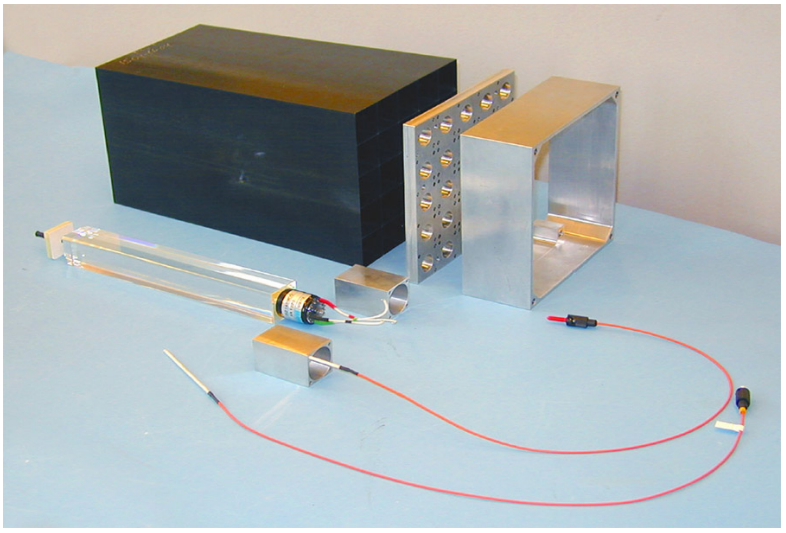

FIGURE 2.7: Alveolar, Interface Plate, Housing, Optical Fibre Insert, Crystal and VPT and End-Stop of the CMS ECAL Endcap (EE) [20].

positioned in a quasi-projective geometry to evade cracks aligned with the trajectories of the particles. Crystal axes make a slight angle (3^0) regarding the vector from the interaction vertex in the η and ϕ projections. The cross-section of the crystal corresponds to nearly 0.0174×0.0174 in η - ϕ or 22 x 22 mm^2 at the front surface of the crystal and 26 x 26 mm^2 at the rear surface. The length of the crystal is 230 mm, corresponds to 25.8 X^0.

The rapidity range of the endcap calorimeter (EE) is $1.479 < |\eta| < 3.0$. The endcap comprises crystals identically organized in the mechanical entity of $5\eta \times 5\phi$ crystals (super-crystals) incorporating a carbon-fiber alveola configuration. Both the endcaps are further partitioned into two identical halves or Dees. Each half maintains 3662

crystals. The crystals and super-crystals are organized in an x-y rectangular grid. The rear face area of crystals is $30 \times 30\ mm^2$. The front face area is $28.62 \times 28.62\ mm^2$ and has a span of 220 mm (24.7 X_0). The layout of the CMS calorimeter is presented in Fig. 2.8.

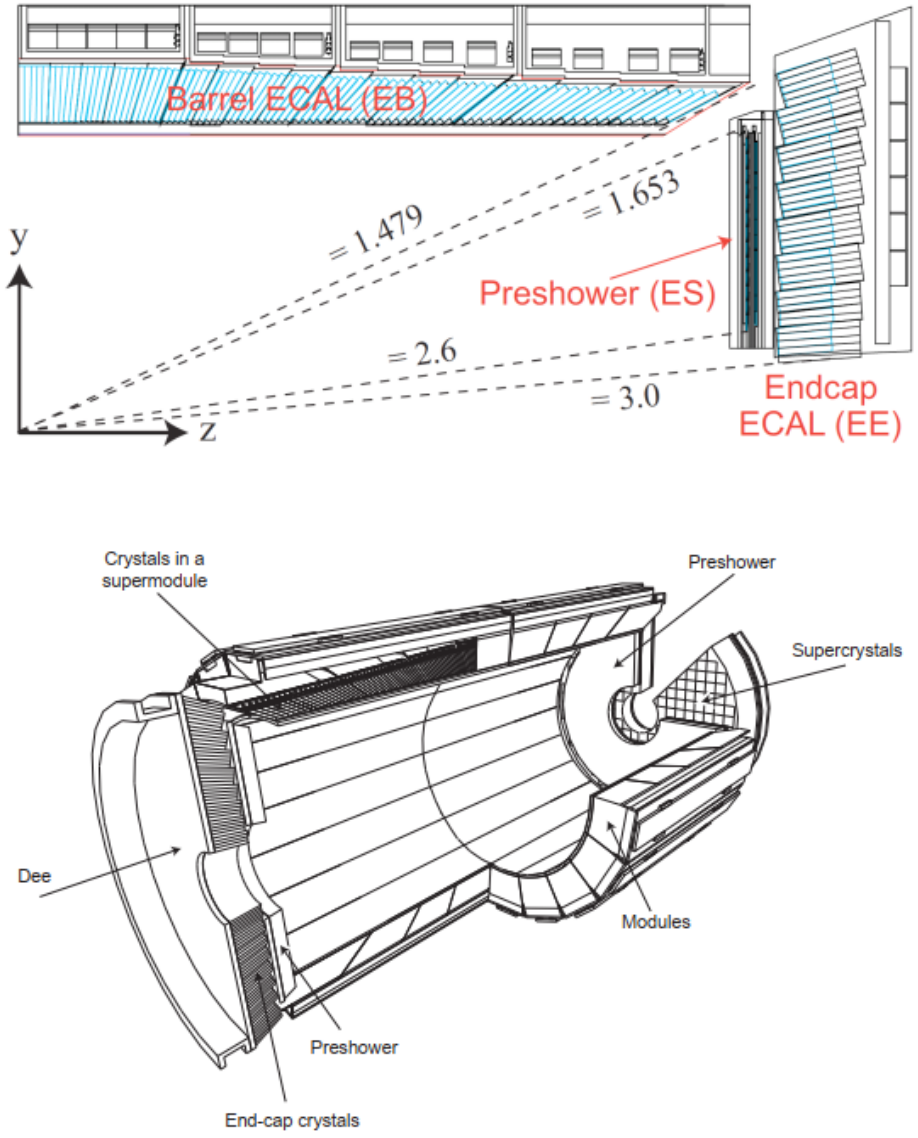

FIGURE 2.8: Longitudinal (top) and 3D (bottom) view of the CMS calorimeter system [2].

As discussed earlier, the essential requirement to suppress the background comprised by $\pi^0 \to \gamma\gamma$ decays, a preshower (ES) detector positioned before the ECAL to prevent such false signals. The ES has a considerably finer and better granularity than the ECAL, with sensor strips 2 mm wide in contrast to the 3 cm wide crystals of ECAL. It comprises two planes of lead (Pb) absorber (of thickness approximately 1 X_0) followed by silicon detectors identical to those employed in the tracker. When a photon transverses

via the lead layer, it generates an electromagnetic shower, detected and measured by the silicon sensors. From this, the photon's energy is estimated while maintaining two detector layers, providing two measurements to accurately locate the particle's position and assist in separating the two photons from a π^0.

2.3.4 Hadronic Calorimeter

The design and placement of the CMS hadron calorimeter (HCAL) were constrained by the volume required to install the detector. The barrel (HB) and endcaps (HE) region of the hadron calorimeter are placed outside the ECAL. The HB is radially confined between the outermost space of the ECAL (R = 1.77 m) and the internal volume of the solenoid coil (R = 2.95 m). This configuration restrains the volume of material that can be placed in to digest the hadronic shower. Further $|\eta| = 3$, the forward hadronic calorimeter (HF) is established at 11.2 m from the collision point, increasing the coverage of the calorimeter till $|\eta| = 5.2$.

As a sampling calorimeter, HB comprises 36 azimuthal wedges in identical shapes covering the pseudorapidity range $\eta < 3$. Each wedge is subdivided into 16 azimuthal plates assembled to eliminate projective dead material. Brass (70% Copper and 30% Zinc) is utilized as the absorber, besides the first and last layers, which are constructed of stainless steel for structural integrity. The coverage of the endcap (HE) region is $1.3 < \eta < 3.0$. The active medium employs the tile and wavelength shifting fiber (WFS) method to carry out the light wave, read out through hybrid photodiodes (HPDs). The HCAL granularity is an integral multiple of the ECAL. Up to $\eta < 1.6$, HCAL towers maintain a dimension of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$, while for $\eta > 1.6$, the size extended to $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$. Fig 2.9 shows the longitudinal view of the CMS detector and highlighting the HCAL and its components.

2.3.5 Solenoid Magnet, Tracker, and Muon System

2.3.5.1 Solenoid Magnet

The characteristics of the magnet are essential for ensuring adequate performance for the physics experiment carried out in the high-energy domain. It requires high bending power to precisely measure the charged particle momenta at a broad range of energies that can be accomplished by employing strong magnetic fields. For a charged particle, the momentum in a uniform magnetic field is given by

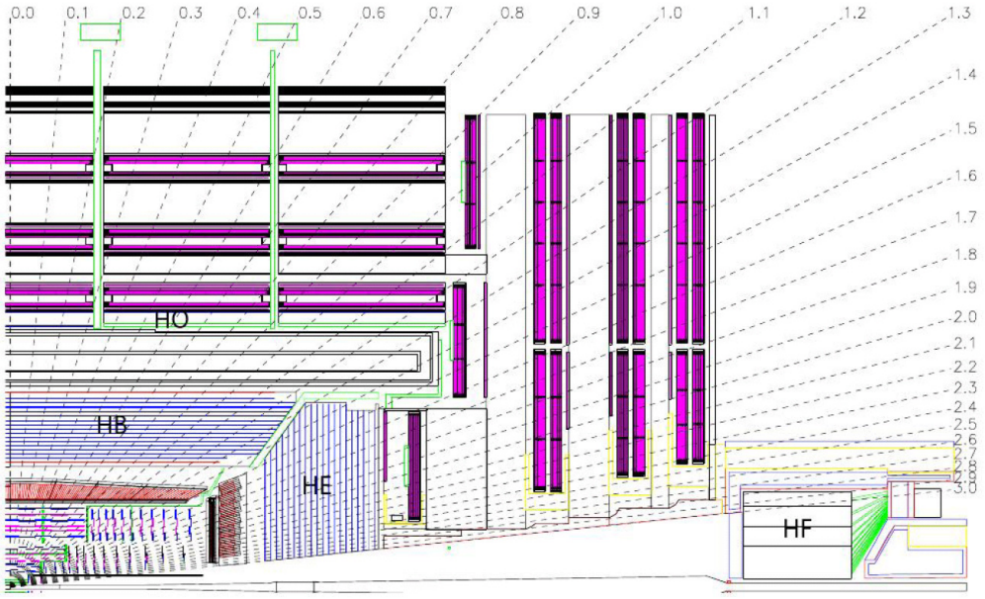

FIGURE 2.9: Schematic diagram of the HCAL elements in the CMS, including the barrel hadronic (HB), the endcap region (HE), the tail catcher beyond the magnet (HO), and the forward hadronic (HF) [2].

$$p = \gamma m v = qBr \tag{2.4}$$

where q is the charge, B is the magnetic field, m is the mass and r is the bending radius. The sagitta of the trajectory is given by

$$s = \frac{L^2}{8r} = \frac{qBL^2}{8p} \tag{2.5}$$

where L is the track length in the magnetic field, and considering that the particle traverses the solenoid from the interaction point, L is equivalent to the solenoid radius. The transverse momentum resolution is proportional to the solenoid radius and magnetic field as

$$\frac{dp}{p} \propto \frac{p}{BL^2} \tag{2.6}$$

Therefore, a high volume and a strong field are required to enhance the resolution. The design of the CMS targets both parameters as it employs a large solenoid of 6 m diameter and an enormous magnetic field of 3.8 T.

2.3.5.2 Tracker

The inner tracking module is designed to equip CMS with the efficient and detailed measurement of the tracks of charged particles arising from the collisions and accurate reconstruction of secondary vertices. It encloses the collision point inside a diameter of 2.5 m and features a length of 5 m. The CMS tracker comprises a pixel detector featuring three barrel layers at radii in the range of 4.4 cm and 10.2 cm and ten barrel detection layers of silicon strip tracker advancing outwards to a radius of 1.1 m. The endcaps provide tracking information in the forward region of the detector. On each the front of the barrel, it comprises 2 plates in the pixel detector and 12 plates in the strip tracker. This topology increases the tracking coverage in the range of $|\eta| < 2.5$. The CMS tracker is the largest silicon tracker ever built, with nearly 200 m^2 of active silicon space. Fig. 2.10 represents the layout of the CMS tracker.

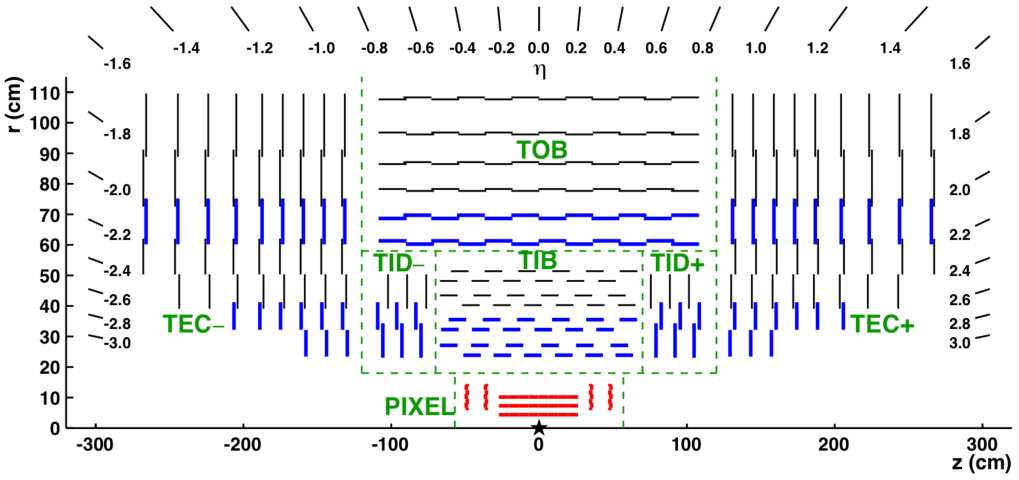

FIGURE 2.10: Cross-section depicting the CMS tracker. The single line defines the detector module, whereas the double lines point to the back-to-back modules [2].

The design and layout of the read-out electronics and silicon sensors must consider and evaluate the radiation damage caused by large particle flux. It must ensure a signal-to-noise ratio (SNR) of 10:1 or higher over the entire lifespan of the detector to ensure robust and accurate track recognition at an adequate fake hit rate. The ionizing particles generate charges in the electronic circuitry of the sensors, which can alter the functioning of memory cells and disrupt the read-out performance. A standard 0.25 μm CMOS technology is used to fabricate the read-out chips of the CMS tracker, which is intrinsically radiation-hard due to unique design rules and thin gate oxide. The deterioration of the silicon sensors due to the high radiation limits the lifespan of the silicon strip tracker. Nuclear interactions of charged particles in the detector material significantly decrease the tracking efficiency for these particles. Furthermore, the sensor material causes the bremsstrahlung and photon conversion, adversely impacting the ECAL measurement

accuracy. Therefore, it is necessary to maintain the optimum quantity of the sensor material.

2.3.5.3 Muon System

One of the key channels in the study and search for the SM Higgs boson is the decay into ZZ or ZZ^*, which decays into four leptons [2]. The most reasonable case is when all four leptons are muons since they are easily detectable and deliver the best mass resolution. Therefore, CMS must identify the muons in the collisions precisely. This task is achieved by the CMS muon system, which must deliver standalone measurements of muon momenta to help in muon identification.

The muon system covers the entire kinematic range of CMS. The muon detector area is nearly 25000 m^2. The CMS employs gaseous detectors to construct the muon system since they are reliable, inexpensive, and robust. Three kinds of muon detectors are installed in the CMS detector.

In the barrel region, drift tube (DT) sections and resistive plate chambers (RPC) are installed, whereas for the endcap sector, Cathode Strip Chambers (CSC) and RPC are established to operate in more dense conditions. The layout of the muon system and its components is depicted in Figure 2.11.

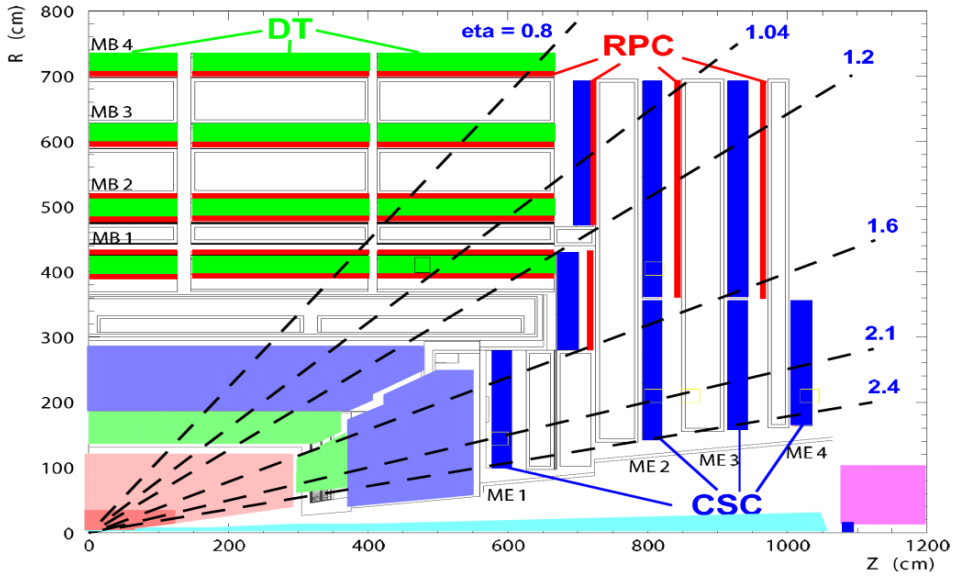

FIGURE 2.11: One quadrant layout of the muon detector in the r - z plane depicting the placement of muon sub-detectors viz. DT, RPC, and CSC. [21]

2.3.6 Triggering at the CMS

During collisions, The CMS detector encounters billions of events per second. Storing and processing these events require massive storage and hardware infrastructure capacity. However, many of these events do not carry any rare or significant physics information and could be discarded early in data acquisition. The CMS detector employs the trigger mechanism to perform rapid rate reduction by efficiently selecting rare and interesting events and rejecting the noise. The CMS trigger system performs this rate reduction in two steps [6]. At first, the Level-1 trigger, based on a custom-designed electronics board (hosts FPGAs), reduces the incoming rate (LHC bunch-crossing) from 40 MHz to 110 kHz. In the second stage, a CPU farm-based High-Level Trigger (HLT) facilitates the rejection from 110 kHz to 2.4 kHz. Fig. 2.12 illustrates the CMS two-level triggering scheme.

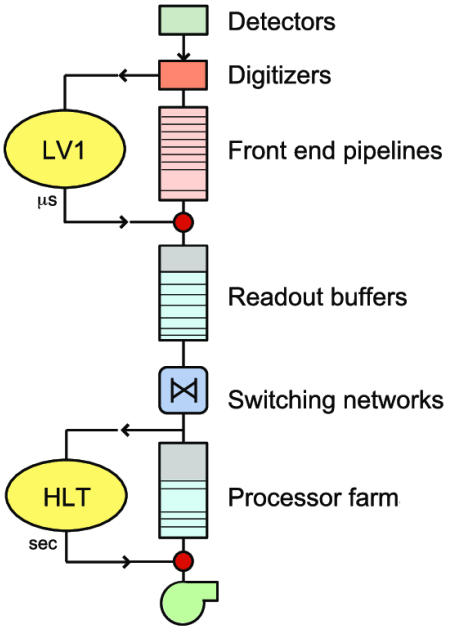

FIGURE 2.12: Schematic of CMS trigger system implemented in two-level: first is the Level-1 trigger (LV1) based on custom-designed FPGA boards, and second is the high-level trigger (HLT) based on processor farm.

2.3.6.1 Level-1 Trigger

The Level-1 trigger is based on a custom-designed electronics board with FPGA as the central processing element. The Phase-1 Level-1 trigger takes coarse (less granular) data from the muon and calorimeter backend system while temporarily retaining the high granular data in the pipelined buffers in the front-end (FE) electronics. The objective of the Level-1 trigger is to achieve adequate reduction from the input bunch-crossing

rate of 40 MHz to deliver an output rate of 110 kHz. The triggering begins from local or regional information in calorimeters and muon sub-detectors. This detector data is encoded in trigger primitives (TPs), essentially hit patterns and energies in the calorimeters and features in muon chambers. The subsequent stage uses the regional triggers that merge data from TPs to create ranked trigger objects in localized parts of the detector. The global calorimeter and muon triggers then sort the regionally constructed candidates and send the highest-ranked ones to the global trigger (GT). The GT incorporates all information from the global calorimeter and muon trigger, and flexibly, it can apply topological conditions and requirements for combinations of trigger objects. The architecture of the Phase-1 Level-1 trigger system is depicted in Figure 2.13.

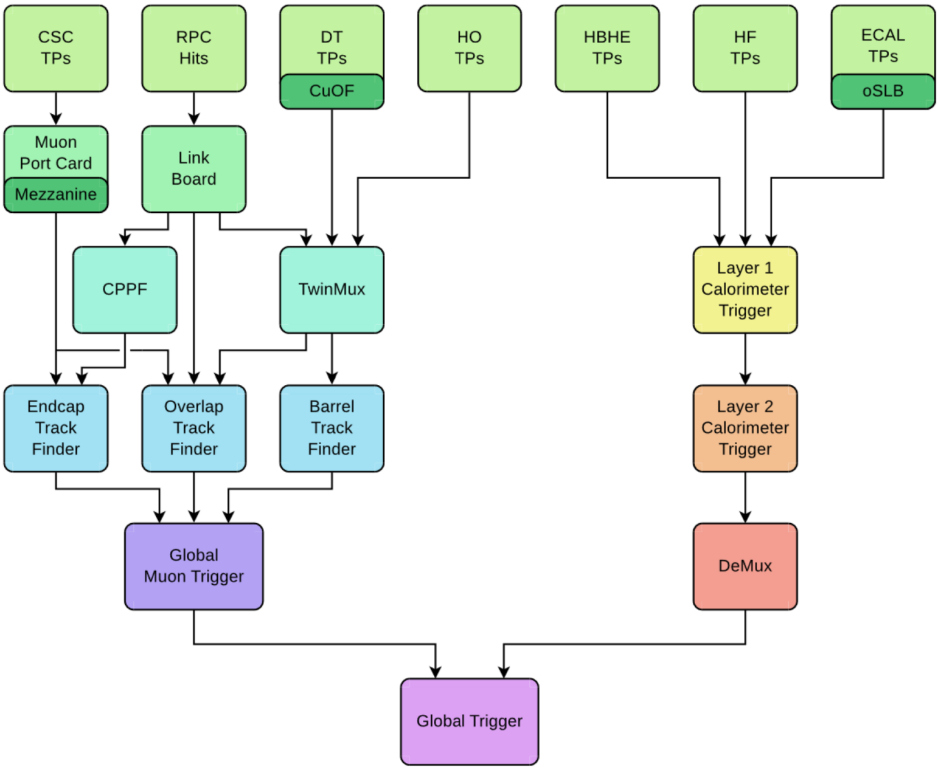

FIGURE 2.13: CMS Phase-1 Level-1 trigger architecture [22].

The Phase-1 calorimeter trigger is implemented into two successive processing layers. The layer-1 comprises 18 Calorimeter Trigger Processor (CTP7) boards that perform pre-processing and data formatting. The algorithms are based on processing the tower level of the ECAL and HCAL, such as the aggregate of ECAL and HCAL energies, energy calibration, and the calculation of the H/E ratio. The 9 layer-2 processor boards (MP7) host calorimeter algorithms that find particle candidates and calculate total global energy. Each MP7 can access an entire event at trigger tower (TT) granularity. The layer-2 employ the time-multiplexing approach which eliminates the requirement of sharing the

of boundary region. The algorithms are completely pipelined and start processing once the lowest portion of data is acquired. The trigger candidates are transmitted to a demultiplexer board, also an MP7, that re-formats the information for the upgraded uGT. The uGT is liable for permitting the event to be further analyzed by the HLT. Fig. 2.14 highlights the two level implementation of the Phase-1 calorimeter trigger.

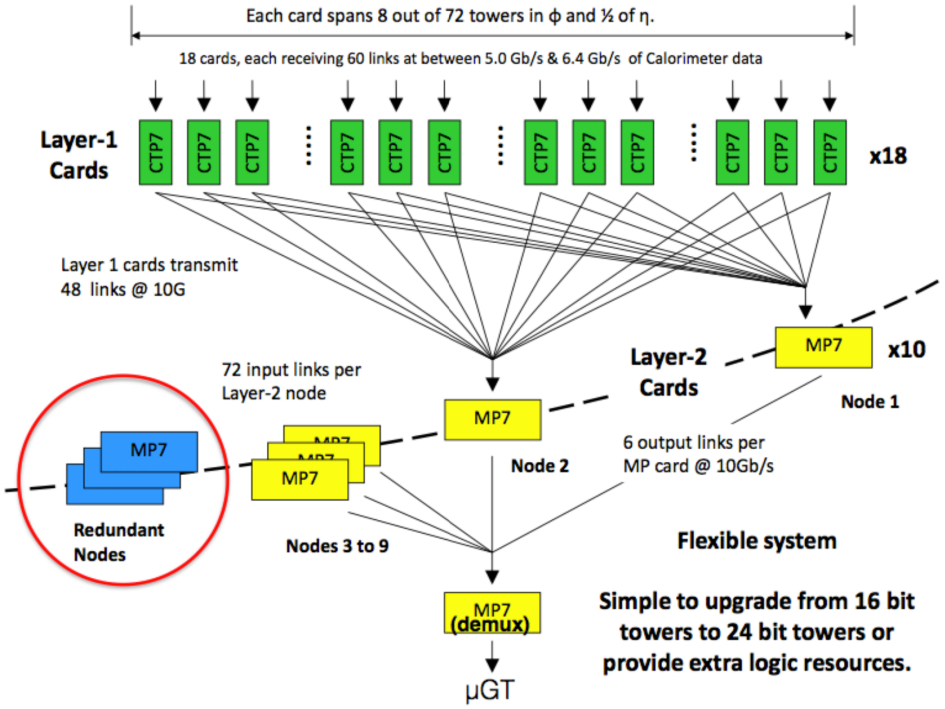

FIGURE 2.14: The Phase-1 Level-1 calorimeter trigger. The boards are organized in two layers. Layer-1 comprises 18 CTP7 boards receiving a total of 1152 input links. Each board sends 48 output links (4 links per bunch-crossing) to the MP7 board (via a patch panel), which are organized in TM9 mode with three redundant or spare boards. Each MP7 board receives 72 input links and sends six output links to the GT [8].

The main highlight of the Phase-1 trigger electronics is the upgrade from VME-based crate system to μTCA standard [23]. The μTCA crate is based on the Vadatech VT892 7U chassis that supports 12 AMC [24] double-width slots. The 12 AMC slots fulfill the requirements of many trigger systems. The double-width cards yield additional front panel space and PCB real estate that can be required, such as routing large numbers of optical links and high-density FPGAs. Fig 2.15 describes the Vadatech VT892 μTCA crate. The crate has dual mTCA Carrier Hubs (MCHs), Power Modules (PMs), and Cooling Units (CUs), allowing complete redundancy for systems. CMS has decided to utilize the second MCH slot to provide the trigger signals, distribute the LHC clock, and optionally consolidate the data via the Data Acquisition (DAQ) system.

The AMC13 is a customized board that provides a feedback mechanism for the Trigger Throttling System (TTS), Trigger Timing and Control signals (TTC), and a high-bandwidth DAQ link for a μTCA crate in CMS [2]. The AMC13 is located in the secondary MCH site with a "dual-star" backplane featuring point-to-point connections to each of the twelve AMC cards in the crate. The AMC13 card corresponds to an MCH form factor. However, it is treated by the main MCH as a "13th AMC". Consequently, the term AMC13. It uses the TTS input to request a change in the state of the CMS trigger throttling system (TTS) output. It utilizes the TTS state response to report the synchronization error conditions or buffer overflow warnings. The response of the AMC13 to these inputs is made programmable.

FIGURE 2.15: The VT892 μTCA crate from Vadatech. The ventilation cavities for the top-bottom cooling are apparently visible. The dual power modules (PMs) are located on the left of the crate. The dual μTCA Carrier Hubs (MCHs) cards are placed in the center. The center top MCH slot is occupied with a custom "AMC13" module.

The main processing components of the Phase-1 trigger is the Calorimeter Trigger Processor (CTP7) board. The CTP7 hosts the Xilinx Virtex-7 FPGA and supports high-bandwidth serial communication capability through MGTs. The Avago technologies provide the transceiver's optical links, which are populated on the front side of the CTP7. The MGTs are placed on a custom backplane designed by Vadatech to facilitate data sharing within the crate. Fig 2.16 shows the CTP7 card used in the Calo Layer-1 featuring the Virtex-7 FPGA without the heatsink.

The CTP7 card feature a dedicated card control system based on Xilinx ZYNQ SoC technology [25]. The ZYNQ supports a Linux operating system that facilitates necessary functions and extensive board monitoring. Special arrangements are made for power supplies and cooling infrastructure to ensure the efficient usage of the logic core for calorimeter trigger applications. The careful design of the clocking infrastructure ensures that numerous speeds can be supported synchronously, asynchronously, or simultaneously with reference to the LHC clock.

Most CTP7 serial links are on the front panel through MiniPOD and Avago CXP optical modules. Three CXP modules handle 36 R_X and 36 T_X links, compatible with the MPO-24 interface. Additionally, there are 31 R_X and 12 T_X links over the Avago MiniPODs on the front panel MPO-48 interface. Furthermore, the backplane supports 14 R_X and 14 T_X links. Out of 48 T_X and 67 R_X optical links, 36 T_X and 40 R_X fibers are linked to the front side of the board, and the remaining 12 T_X and 27 R_X are tied to the backside of the board. On the other hand, There are DAQ and Gigabit-Ethernet (GbE) links over the backplane. The links are developed and validated to function at up to 10 Gbps asynchronously and synchronously with the LHC clock facilitated by the crate.

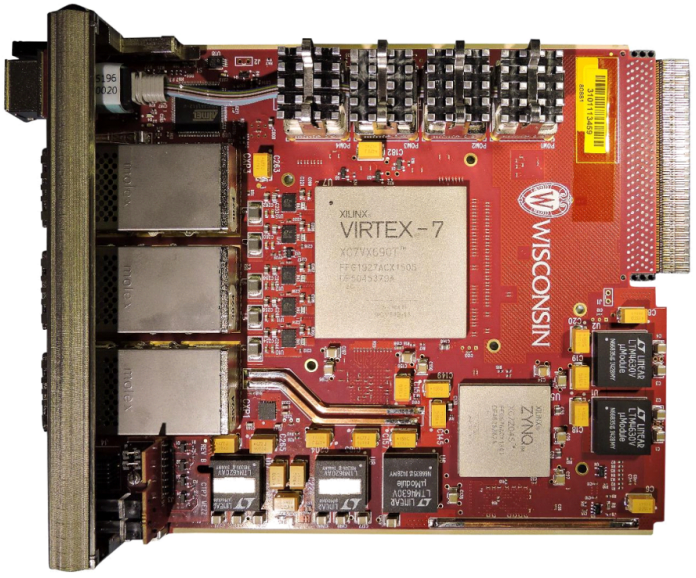

FIGURE 2.16: CTP7 card with Virtex-7 FPGA and ZYNQ SoC.

2.3.6.2 High Level Trigger

The High-Level Trigger (HLT) processes the events using a commercial processor filter farm comprising CPUs and GPUs. The main objective of the HLT is to decrease the incoming event rate from 110 kHz to 2.4 kHz. After receiving the Level-1 accept signal, the events corresponding to the same bunch-crossing must be transmitted from the readout units (RU) to a single processor of the HLT for evaluation. A fast network known as Event Builder (EB) is employed that has a readout (RU) and builder (BU) layer bridging across a high-speed switching system. It constructs one entire event from the data fragments arriving at its readout units. It successively streams all the data of each bunch-crossing to a respective filter unit (FU) for evaluation at HLT. Since HLT processing takes more latency than Level-1, this process is serial, resulting in individual nodes in the farm processing a distinct event. The processing power in the farm must be adequate such that when the last node is populated, the first node is open for the

next event. However, prominent safety factors have to be employed for the possibility of failing nodes or nodes that spend much more time for event processing due to algorithm and events anomalies (i.e., massive number of tracks to be reconstructed).

To optimize the processing time at the HLT, it employs several phases of event rejection. Simple algorithms run first, and essential requirements are applied before executing more complicated algorithms that burn more time. If the event is rejected in the first step, complex algorithms do not execute, and the node is unrestricted to accept the next event. The event accepted by the HLT is finally stored in the CERN TIER0 system. Next chapter describes the High-Luminosity LHC, the drawback of the Phase-1 Level-1 trigger, triggering requirements of HL-LHC, Phase-2 upgrade of the CMS detector, and major design and prototype requirement of the Phase-2 Level-1 calorimeter trigger is discussed in the next chapter.

Chapter 3

High Luminosity LHC and Triggering in Phase-2

High-Luminosity LHC (HL-LHC) will increase the potential of the LHC in terms of precise measurements and physics discoveries. In its ultimate condition, the instantaneous luminosity of the HL-LHC will be $7.5 \times 10^{34} cm^{-2} s^{-1}$, seven times the LHC original design parameter ($1 \times 10^{34} cm^{-2} s^{-1}$) [26]. The goal is to enhance the performance of the LHC after Phase-1 by delivering a large amount of collision data, which helps in solving some of the most challenging questions of modern-day physics, such as the comprehensive study of the Higgs boson particle and its properties, the existence of supersymmetry (SUSY), physics beyond the Standard Model (BSM), and dark matter. These rare physics phenomena have a significantly lower probability of occurrences (the cross-section) and thus require extensive data collection (integrated luminosity). Fig. 3.1 represents the LHC operational timeline, including the HL-LHC, which will initiate the operation during May 2029 after Long-Shutdown 3 (LS3) and planned to operate till year 2041.

FIGURE 3.1: The LHC timeline and its advancement to the HL-LHC (May 2029 to October 2038) after Long-Shutdown 3 [27].

In the ultimate luminosity regime of $7.5 \times 10^{34} cm^{-2} s^{-1}$ and considering thirteen years of operation, the ambitious goal of the HL-LHC is to collect 4000 fb^{-1} of collisions data. The high luminosity of the LHC is associated with the higher number of simultaneous proton-proton interactions, called pileup. A high pileup environment poses more complexity and contaminate the particle reconstruction process. Table 3.1 summarises the Run-2 and HL-LHC expected (instantaneous and integrated) luminosity and the maximum pile-up information.

	Instantaneous luminosity	Pile-up (average)	Integrated luminosity (fb^{-1})
Run-2	$2.1 \times 10^{34} cm^{-2} s^{-1}$	55	160
HL-LHC (baseline)	$5 \times 10^{34} cm^{-2} s^{-1}$	140	3000
HL-LHC (ultimate)	$7.5 \times 10^{34} cm^{-2} s^{-1}$	200	4000

$fb = femtobarn$

TABLE 3.1: Luminosity and pile-up Information of HL-LHC and Run-2.

The next section describes the drawbacks of the Phase-1 system which leads to the upgrade of the CMS Level-1 calorimeter trigger.

3.1 Phase-2 Upgrade of the CMS Detector

The HL-LHC poses the limitation regarding higher pile-up multiplicity and extreme radiation environment. To tackle these challenges, the CMS detector requires significantly upgrading its sub-detectors, front-end and backend readout, data acquisitions (DAQ), and trigger system. The main highlights of the Phase-2 upgrade of the Level-1 trigger are the employment of large FPGAs, high-speed optical links, and the first-time inclusion of the tracker information and PF-based correlator trigger. Fig. 3.2 highlights the major upgrade planned for the CMS detector during Phase-2. The following sections describe the main highlights of the CMS Phase-2 upgrade in detail.

3.1.1 Pixel and Tracker

The CMS has planned to upgrade the entire tracking system by replacing the pixel and strip tracking detector with a smaller pixel sensor that can endure the high radiation environment. The Outer Tracker features micro pixel and strip sensor, extending the overall coverage to $|\eta| = 3.9$ [26]. Stacked strip modules at the outer tracker help lower the hit multiplicity and provide the particle's track information to the Level-1 trigger system. The Outer Tracker performs the identification of charged particles (at 40 MHz of

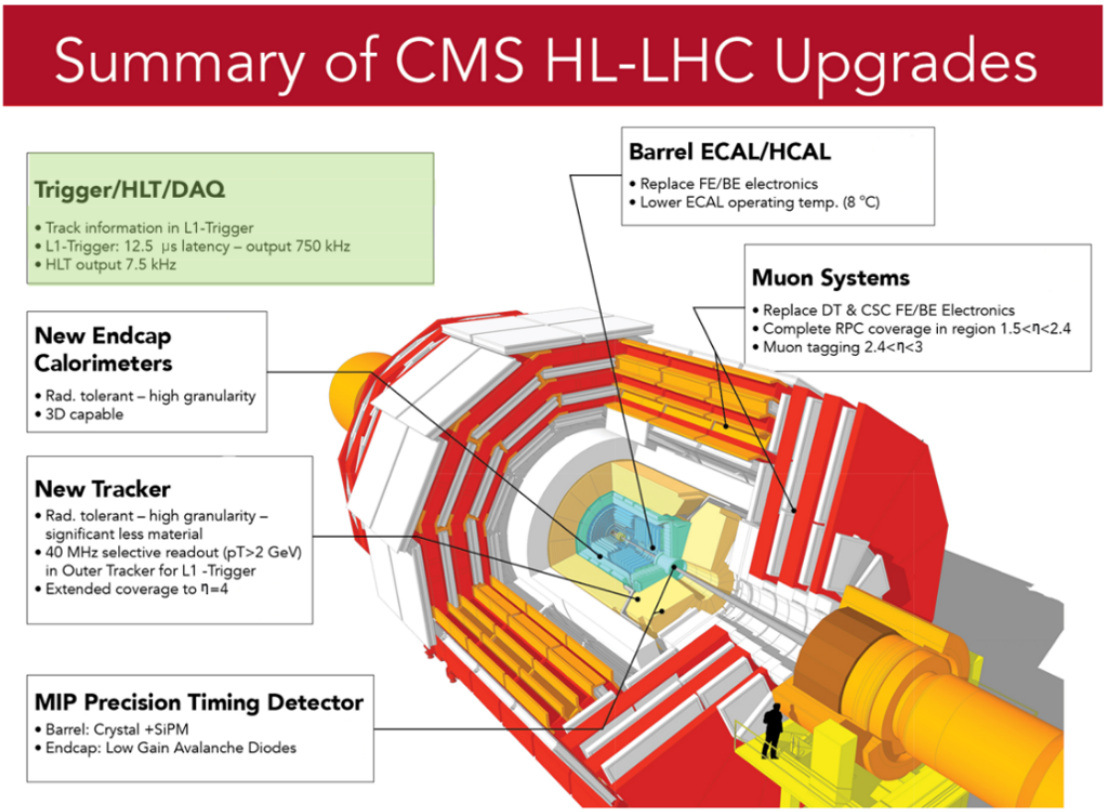

FIGURE 3.2: Summary of the CMS HL-LHC upgrade.

bunch-crossing) through "p_T modules". It uses the p_T threshold of 2 GeV for particles to filter a signal by evaluating the hit correlation between two neighboring silicon sensors. The backend Level-1 track-finding system receives the pairs of correlated tracker hits, also known as "stubs". The track-finding system employs a "hybrid" algorithm [26], and it is implemented on the FPGA-based ATCA platform, extensively exploring parallel processing in geometrical topology and in time.

The implementation strategy is partitioning the outer tracker into nine ϕ sectors, with a single FPGA processing one ϕ sector. It also employs the time-multiplexing (TMUX) scheme [28], leading to the requirement of a 162 FPGA board for the track-finding system.

3.1.2 Muon system

The Phase-2 upgrade of the muon system includes the replacement of the drift tubes (DT) and cathode strip chambers (CSCs) front and backend electronics along with the installation of improved resistive plate chambers (iRPC) and gas electron multiplier (GEM) chambers to increase the coverage upto $|\eta| = 2.4$ and 2.8, respectively. The combined muon barrel trigger primitives generator and the DT DAQ comprise 6 ATCA

crates, each supporting ten Barrel Muon Trigger (BMT) [26] boards to cover 60 DT sectors. Two crates with 12 boards per crate are required to cover 12 DT wedges. Phase-2 upgrade of the RPC backend electronics will utilize 18 Serenity boards with three ATCA crates, one for each RPC barrel, endcap, and iRPC. The Phase-2 upgrade of the CSC includes partially replacing various components in the read-out chain, spanning from off-chamber to on-chamber electronics. The CSC backend system will incorporate two ATCA crates, one per endcap, each housing five APx boards. Fig 3.3 illustrates the barrel muon backend system.

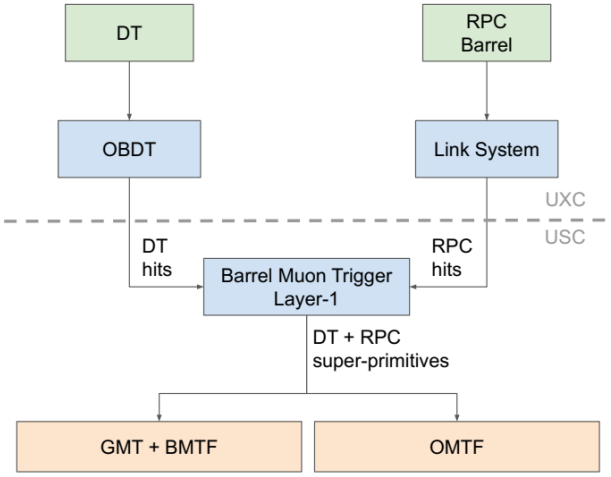

FIGURE 3.3: Barrel muon trigger primitive generation architecture. The DT information is processed via the onboard DT frontend system and ultimately forwarded to the backend barrel muon trigger layer-1 (BMTL1). The barrel RPC hits are transferred via the link system. The BMTL1 performs the clustering of the RPC hits, constructs the DT trigger segments, and merges both collections into super-primitives. The BMTL1 trigger primitive outputs are sent to the Overlap Muon Track Finder (OMTF) and Barrel Muon Track Finder (BMTF) via the high-speed optical links [26].

3.1.3 Calorimeter

CMS collaboration has planned to entirely replace the endcap calorimeter and upgrade the on-detector front end and backend electronics of the barrel calorimeter. The upgrade mainly emphasizes the high radiation tolerance in the forward region, the high granularity information from the endcap and barrel, and the spatial timing information of the individual particles to mitigate the effect of the high pile-up.

The Phase-1 endcap calorimeter will be entirely replaced by a new radiation-tolerant high-granularity 5D imaging capable calorimeter known as the high-granularity calorimeter (HGCAL) [9]. It will provide the energy, positional, and precision-timing information with 30 ps of resolution. Excellent physics performance will be achieved by using radiation-tolerant silicon pad sensors (0.5 or 1 cm^2 of area) and SiPM-on-scintillator

tiles (4 to 32 cm^2) with high transversal and longitudinal segmentation. The HGCAL supports 6.4 million readout channels, almost double the magnitude of channels supported by the Phase-1 endcap system. The electromagnetic section of HGCAL houses 26 active layers of silicon-based sensors interleaved with lead absorbers, copper, and copper-tungsten. The assembling of the hadronic system is planned by using 21 active layers of silicon and scintillator tiles, uncoupled by steel absorbers. Overall, the silicon sensors will occupy 600 m^2 of area (three times the area covered by the tracker), and the SiPMs-on-scintillator tiles will cover 400 m^2 of area.

One of the prominent features of the HGCAL is the hexagonal-shaped sensors processed on 8-inch silicon wafers. The hexagonal bee-hive-like structure optimally utilized the circular wafer, making it highly cost-effective. Fig 3.4 illustrates the HGCAL bee-hive shaved trigger cell formed due to the grouping of the sensor cell.

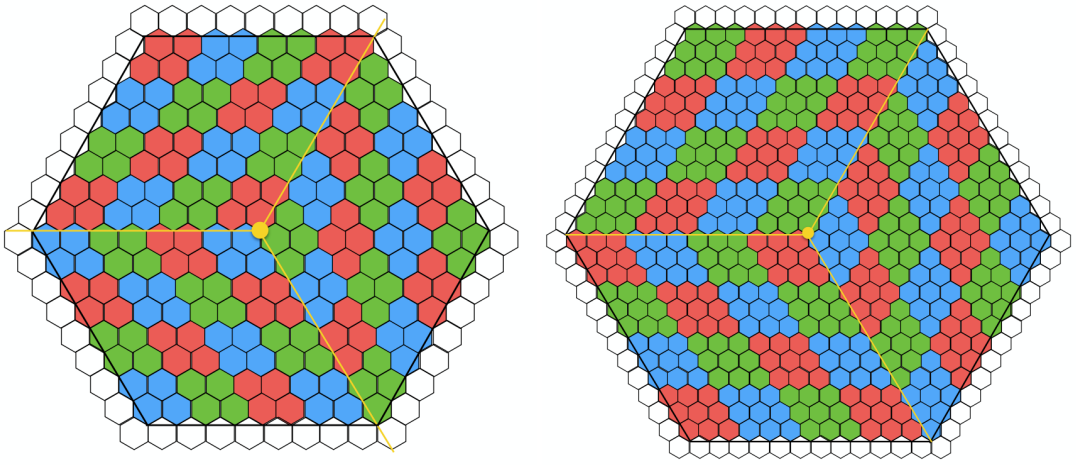

FIGURE 3.4: Illustration of the grouping of sensor cells in three-fold diamond pattern on hexagonal 8-inch silicon wafer. Arrangement of large sensor cell (1.18 cm^2) is shown in left and smaller cell (0.52 cm^2) in the right [9].

The barrel calorimeter (ECAL and HCAL) was crucial in discovering the Higgs boson particle due to its decay in two photons (γ). The Phase-2 upgrade preserves the current ECAL lead-tungstate ($PbWO_4$) crystals and HCAL scintillator. However, the barrel front-end and backend electronics will entirely be replaced to increase the granularity from tower level (five-by-five crystals) to each crystal with further enhancement due to the inclusion of timing information using the MIP timing detector [29]. The timing information helps reduce the pile-ups.

With 61200 crystals in ECAL, the readout of the crystals is planned to be carried out in the group of 25 such crystals, leading to 2448 such groups. It requires four low power GBT (lpGBT) [30] links (10 Gbps) per group, resulting in 9792 links for the whole barrel ECAL. The key technical reason for the barrel calorimeter electronics upgrade is the increase in the L1 trigger latency from 4 μs in the Phase-1 system to 12.5 μs in the

Phase-2, and the increment in L1 trigger output rate from 100 KHz to 750 KHz. Fig. 3.5 represents the proposed front-end readout architecture of the barrel ECAL.

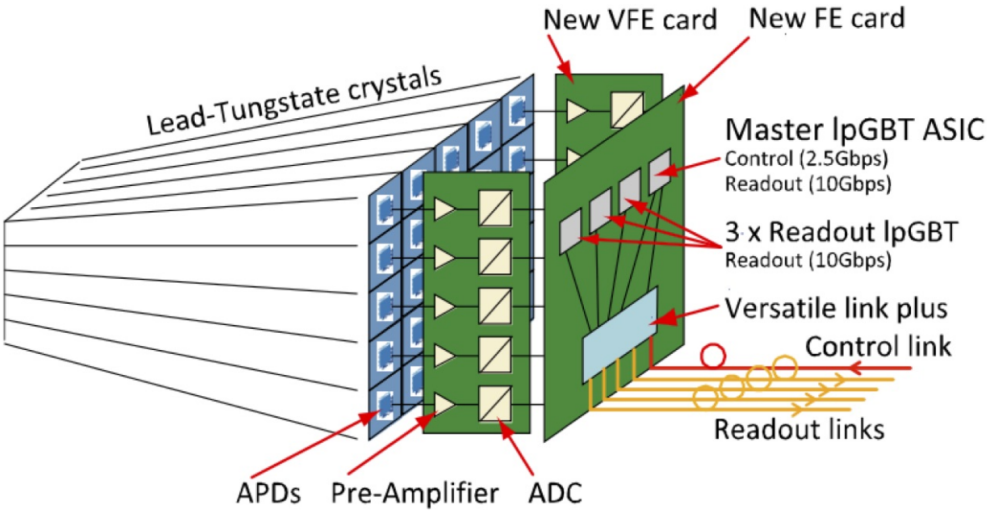

FIGURE 3.5: ECAL Phase-2 front-end readout scheme [31].

The barrel calorimeter signal amplification, shaping, and digitization operation is performed by the Very Front End (VFE) board. These VFE board will be replaced and planned to provide noise filtering and better timing resolution to suppress the APD anomalous signal (known as "spikes") effect at the L1 trigger.

The Phase-2 barrel backend system uses the barrel calorimeter processor (BCP) board, which can process up to 96 lpGBT links [31]. It requires 108 BCP boards housed in 9 ATCA crates (12 cards per crate). The BCP cards perform the operation of data buffering and generate the trigger primitives (TPs) for the Level-1 calorimeter trigger. The mapping and coverage of the BCP crate to process the CMS barrel calorimeter are shown in Fig 3.6.

3.1.4 Data Acquisition (DAQ)

CMS also planned to entirely replace the data acquisition (DAQ) and the trigger system (Level-1 and HLT). The front-end (FE) system is often mounted on the detector or in its vicinity. Conventionally it comprises a combination of an analog and digital module, which performs the task of gathering, processing, and digitizing the signals from the detector. It also acquires additional logic to control these processes and helps monitor the detector parameters. The backend (BE) electronics are localized in the service cavern (USC), which encounters negligible or no stray magnetic field and offers a radiation-free environment. The backend system receives and processes the digitized data from the

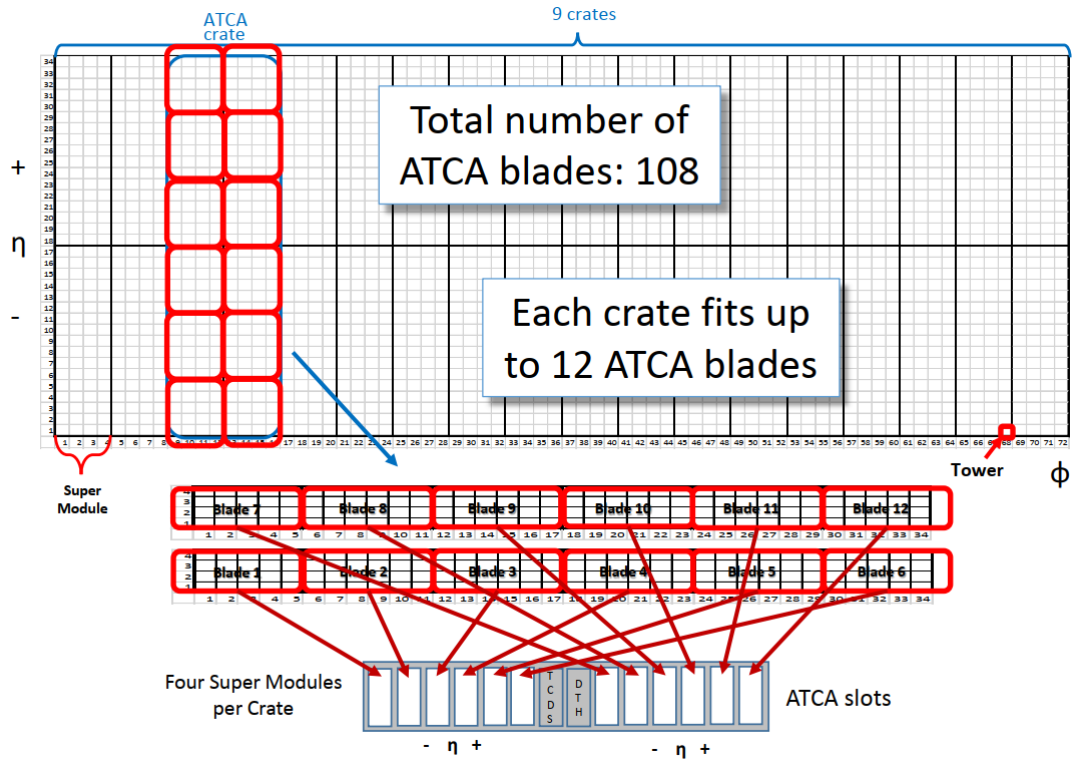

FIGURE 3.6: Mapping of the BCP crate with the barrel ECAL [31].

front-end, ensuring its control and synchronization. Ultimately, it serves as an input to either the data acquisition (DAQ) or L1 trigger.

The detector FE and BE communicate via bidirectional links. The backend uses the downlinks to distribute the L1 accept signal (provided by the L1 trigger), the master clock, and control signals for the FE electronics (mainly reset signals). The FE uses the uplinks to transmit the detector digitized data to the BE electronics for pre-processing before channeling it to the L1 trigger and sending the full granular information to the DAQ system. Several sub-detectors transmit a limited amount of collisions data while buffering the entire information in the FE pipeline buffers (waits for the L1 accept signal), while others transmit the entire bunch-crossing (BX) of the information to the corresponding backend sub-system.

The major role of the CMS DAQ is to facilitate the routing of the detector data and decoupling of the time domain between the synchronous front (backend and hardware trigger system) and the asynchronous HLT and permanent storage. Fig. 3.7 describes the synchronous and asynchronous DAQ front.

The synchronous part of the DAQ system distributes the trigger accept, master clocks, and control signals, such as the Trigger, Timing, and Control (TTC) signal to the BE. It collects the buffer flag signal of each BE board to supervise the application of the L1 accept signal and inhibit the buffer overrun. In CMS, these two functions are merged

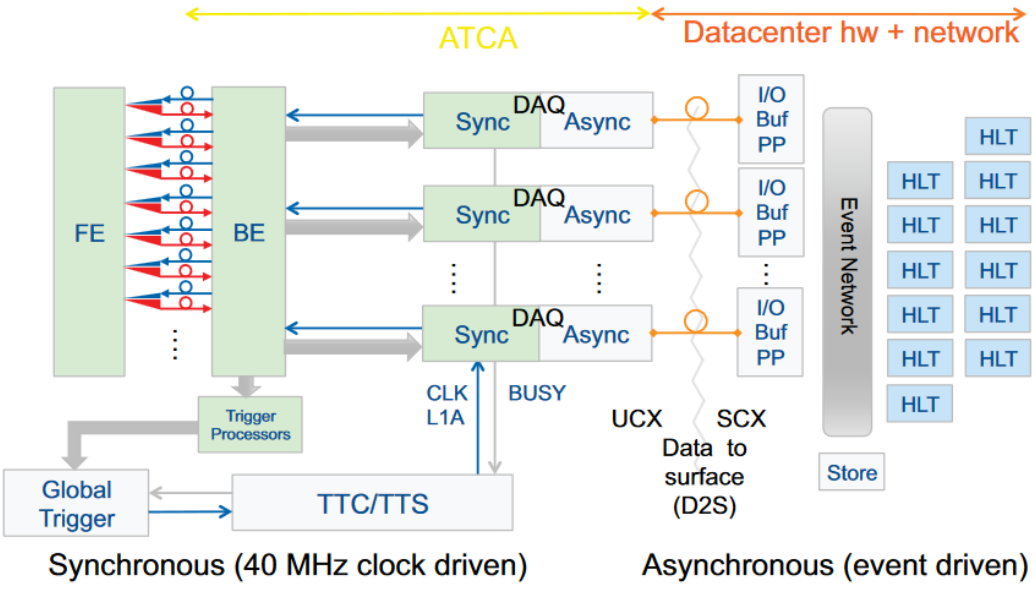

FIGURE 3.7: Asynchronous and synchronous interface of the CMS Phase-2 DAQ system [32].

and performed by the Trigger and Timing Control and Distribution System (TCDS) [33]. The main objective of the DAQ is to collect the L1 accepted events and stabilize the effective deadtime at a nominal level. The DAQ and TCDS Hub (DTH) [26] board is used to accumulate the data from multiple BE boards and merge it to feed the high-speed commercial Data to Surface (D2S) optical links with a bandwidth of 100 Gbps or higher. The D2S links transfer the data to the surface, where the buffering is provided via the set of I/O processors until the event building. Fig 3.8 illustrate the Phase-2 architecture of the CMS DAQ system.

3.1.5 Level-1 Trigger

The Phase-2 upgrade of the CMS Level-1 trigger system focuses on all three fronts, such as the extensive usage of large FPGAs, high bandwidth optical links, and the implementation of sophisticated and complex algorithms, which were only possible at HLT in the Phase-1 system. The Phase-2 trigger system is currently prototyped using Xilinx UltraScale+ XCVU9P FPGA and is planned to use the XCVU13P FPGA in the final system, which can support the maximum bandwidth of 32.75 Gbps. The UltraScale+ FPGA features more than double the amount of LUTs, FFs, and serial bandwidth the Virtex-7 FPGA (Phase-1) can offer. Table 3.2 summarized the specification of the Virtex-7 FPGA used in Phase-1 and UltraScale+ (VU9P and VU13P) FPGA used for prototyping Phase-2.

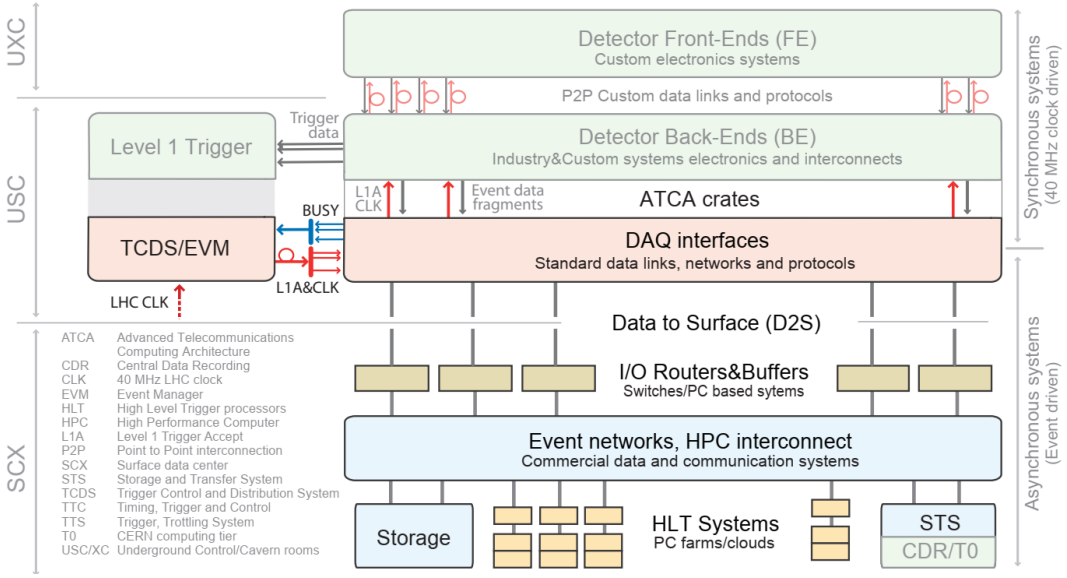

FIGURE 3.8: Graphical representation of Phase-2 DAQ system [33].

	Phase-1 Virtex-7 (XC7VX690T)	Phase-2 UltraScale+ (XCVU9P)	Phase-2 UltraScale+ (XCVU13P)
Technology	28 nm	16 nm	16 nm
LUTs	693,120	1,182,240	1,728,000
FFs	866,400	2,364,480	3,456,300
DSPs	3600	6,840	12,288
Maximum Serial bandwidth (Gbps)	13.1	32.75	32.75
GTs	80	120	128

TABLE 3.2: Logic, and serial bandwidth comparison between the Phase-1 Virtex-7 [34] and Phase-2 UltraScale+ FPGA devices [35].

On the architecture front, the Phase-2 triggering at CMS will preserve its two-level rate reduction strategy while emphasizing the increased reconstruction complexity at the L1 trigger. The latency of the Phase-2 Level-1 trigger system will be 12.5 μs to accommodate more complex algorithms, processing high-granular information and new sub-systems such as correlator trigger (CT) and track trigger. The output rate of the L1 trigger will increase from 100 KHz to 750 KHz while maintaining the same input frequency of 40 MHz of bunch-crossing. The Phase-2 HLT system will employ advanced and high-performance processing systems and is planned to share some of the reconstruction algorithms to compile on GPUs. The input rate for the HLT will be 750 KHz, and it accesses the full granularity of the detector and reduces the rate up to 7.5 KHz within the maximum timing budget of 500 ms.

The system receives input from the calorimeter, muon spectrometer, and track finder. Due to high bandwidth constraints, the computation of the trigger quantities, such as the trigger primitives, is planned to be implemented at the detector backend system. The calorimeter trigger processes the barrel information of electromagnetic calorimeter

(ECAL) and hadronic calorimeter (HCAL) in the regional calorimeter trigger processor cards (RCT). The global calorimeter trigger computes the upstream information from RCT, hadron forward calorimeter (HF), and the HGCAL. The muon trigger system collects the inputs from various muon spectrometers, such as the resistive plate chambers (RPC), drift tube (DT), cathode strip chambers (CSC), and gas electron multiplier (GEM). In the local sub-division, the muon trigger is comprised of the muon track finder and the barrel layer-1 processor. The muon track finder processes the input from three distinct regions called BMTF, EMTF, and OMTF for barrel, endcap, and overlap, respectively.

The Phase-2 Level-1 trigger system employs the particle-flow algorithm, organized in two layers and referred to as correlator trigger (CT). It reconstructs higher-level objects and identifies them into electrons, photons, muons, taus, and jets. The global trigger (GT) makes the final decision, which issues the Level-1 trigger decision on whether to accept the collision data. The Fig. 3.9 represents the architecture of the Phase-2 Level-1 trigger.

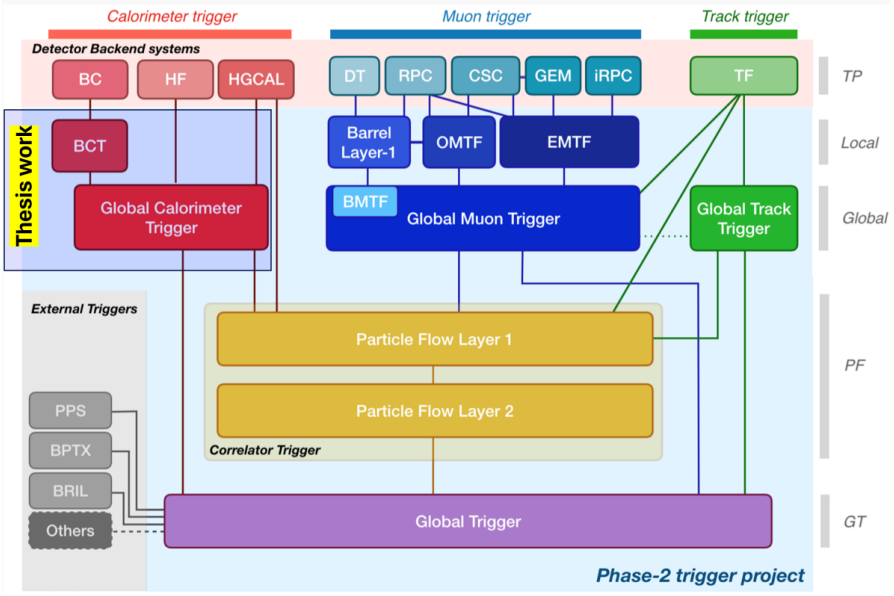

FIGURE 3.9: CMS Phase-2 L1 trigger system comprises three trigger sub-systems: calorimeter, muon, and tracker [26].

This thesis work concentrates on the algorithm and firmware development of Level-1 calorimeter trigger, which requires designing and prototyping the new system. The next chapter of the thesis discusses the proposed architecture of the calorimeter trigger, the hardware components required, and the major design and prototyping constraints associated with the system.

Chapter 4

Phase-2 Level-1 Calorimeter Trigger

The Phase-2 CMS calorimeter system comprises four calorimeter sub-detector systems to strategically cover the barrel and endcap geometry. This includes the electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL) in the barrel part of the detector. A new radiation-hardened high-granular calorimeter (HGCAL) encloses the two end-cap and the forward hadronic (HF) calorimeter in the forward regions of the detector.

4.1 Proposed Solution and Architecture

The proposed architecture of the Phase-2 calorimeter trigger is devised in two levels: regional calorimeter trigger (RCT) and global calorimeter trigger (GCT). The levels are further subdivided into three layers. The first level (layer-1) is the RCT, proposed to be implemented using 36 FPGA boards (each hosting one XCVU9P FPGA). The RCT employs the electron/photon clustering algorithm and processes only the inputs from barrel region of the calorimeter. It sends the barrel calorimeter clusters and tower information to the three GCT barrel boards using 144 optical links.

The second level is the GCT, proposed using ten FPGA boards (each hosting one XCVU9P FPGA). In the second level (layer-2), three GCT barrel boards take the processed input information from the RCT boards, and six GCT endcap boards take the input from the forward region of the calorimeter, viz. HGCAL and HF.

The three GCT barrel boards receive the cluster and tower information of the barrel and perform various hadronic clustering such as particle flow (PF), jets, taus clustering, and computes the partial MET information. It merges the electron/photon clusters

positioned at the RCT board boundaries and sends the final barrel information to the correlator trigger in six-time slices using 144 output links.

The GCT endcap receives the HGCAL and HF information using 234 input links. Six GCT endcap boards cover the HGCAL and HF, with three boards per endcap. The inputs from HGCAL are time-multiplexed and require de-multiplexing at the GCT endcap board before processing in the clustering algorithm, which prepares the jet and taus information. The GCT endcap sends the final output to the GCT sum board using six output links.

In the second level (third layer), the GCT sum board, takes input from all the nine layer-2 boards in the form of calorimeter-based objects from the entire calorimeter (ECAL, HCAL, HGCAL, and HF). It formats the information as per the global trigger requirements, time-multiplex the data, and sends the final output to the global trigger using six output links.

Infrastructure-wise, the planned architecture requires 46 FPGA boards, 2988 input optical links from barrel calorimeter, HF, and endcap HGCAL, 144 output links to the correlator trigger, and six output links to the global trigger. Fig 4.1 shows calorimeter trigger architecture with the required number of FPGA boards and optical links connection.

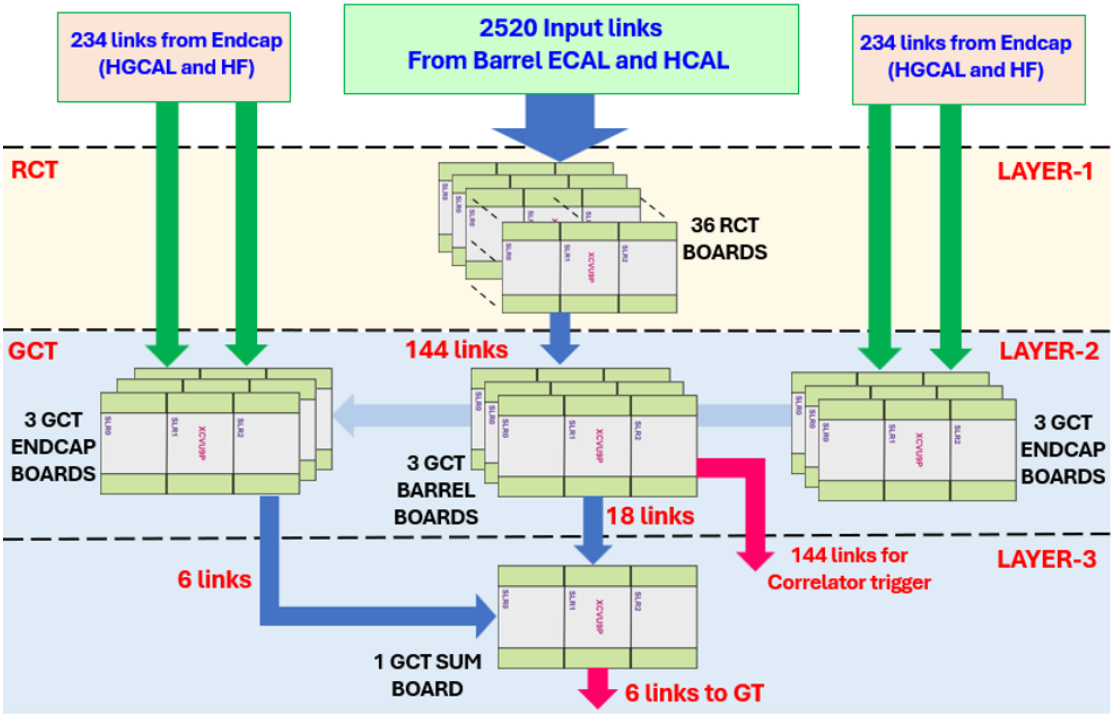

FIGURE 4.1: Two-level and three layers architecture of the calorimeter trigger using APx boards.

4.2 Hardware Infrastructure of Phase-2

The Phase-2 upgrade of the CMS Level-1 trigger system is planned to replace the Phase-1 μTCA-based processors (utilizing in current Phase-1 trigger system) with the Advanced Telecommunications Computing Architecture (ATCA) standard board, crate, and rack system. The ATCA specification [36] promised numerous advantages over the μTCA in terms of more power with associated cooling, high density, and a high-performance backplane compatible with various sets of common serial standards (1/10/100 GbE, PCIe, and others). The CMS Phase-2 ATCA board family features one or two Xilinx large FPGAs (UltraScale/UltraScale+) for increased computational capability for the Level-1 algorithms in Phase-2. The ATCA boards are equipped with high-speed optical links to gather the collision data and thus provide a comprehensive view of the detector for event reconstruction. The next section describes the ATCA specification in details.

4.2.1 Advanced Telecommunication Architecture: ATCA

The ATCA benchmark defines the board, shelf (chassis), frame (rack) form factors, power, core backplane connectivity, management interfaces, cooling, and the electromechanical specifications of the boards. The specification defines the key elements of ATCA as:

- *Boards*: The computing blades/cards that are pluggable and removable in a shelf
- *Shelf*: The 12U (U is rack unit equivalent to 1.75 inches) tall chassis delivers power, backplane connectivity, cooling, and the slots to maintain up to 16 boards.
- *Frame*: A rack (generally 46U high) delivers a rigid framework abiding up to three shelves.

The ATCA board specification defines the form factor of 8U tall (14 in) by 280mm of depth (11.02 in) and 6HP wide (1.2"/30.48mm). The board spans 140 square inches of area, enabling multi-processor strategies or designs to accommodate large chunks of memory. The rear side of the board has alignment keys, power connectors, rear I/O access, and high-speed links to the passive backplane. Fig. 4.2 illustrate the ATCA board and Rear Transition Module (RTM) layout.

Each ATCA crate/shelf incorporates one or two hub slots, determined by slot numbers 1 and 2. Each hub slot is connected to each node slot using a dual-star configuration with four high-quality signal pairs as shown in Fig. 4.3.

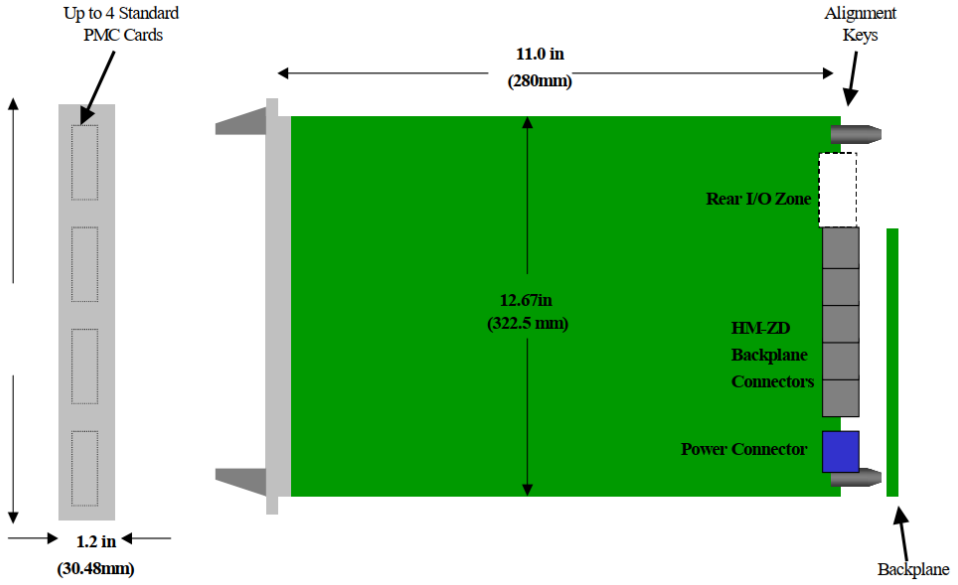

FIGURE 4.2: ATCA board specifications [36].

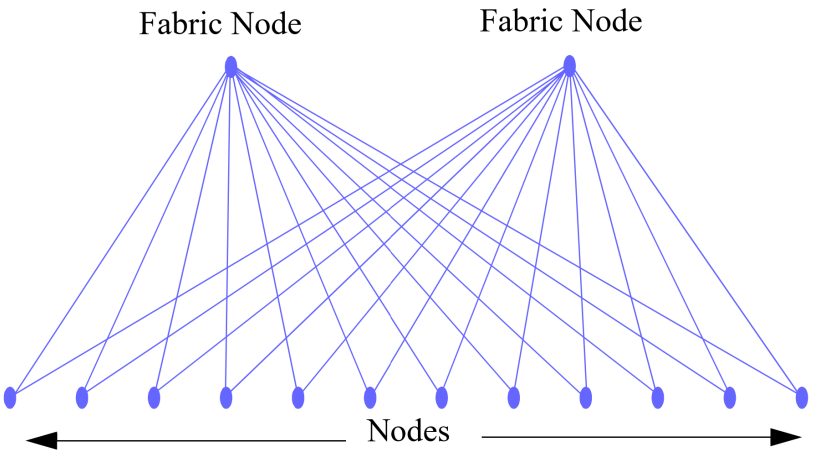

Figure 2. Dual Star Topology

FIGURE 4.3: Dual Star connection topology [36].

The ATCA layout adopted for the Phase-2 upgrades is based on a 14-slot crate, provided with a CERN-endorsed crate manager. For widespread system reasons (availability of spare slots, cooling load, and power) and future requirements, reserving two of the twelve node slots as spares is recommended. Fig 4.4 depicts the initially planned backplane connectivity strategy between the hub and node slots.

CERN-endorsed exterior mains-to-48 V power converters will power the ATCA crates. For installation, the ATCA racks need vertical airflow for cooling in the current electronics rack infrastructure in the service cavern (USC). The cooling system of the USC rack is rated for an ultimate heat load of 10 kW per rack, with an average of 6 kW. Furthermore, the USC Heating, Ventilation, and Air Conditioning (HVAC) system is

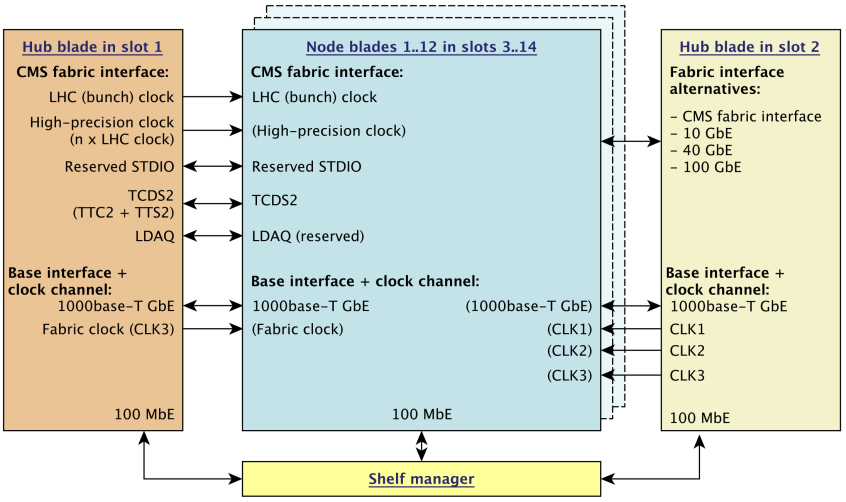

FIGURE 4.4: CMS dual star backplane connection topology. Each of the two hub boards are connected with the node board in "star" convention [26].

designed to manage the maximum heat leakage of 50 W per rack. Therefore, the recommended threshold for the maximum power dissipation for the node blades is 300 W. The next section describe the APd1 board used in prototyping the Phase-2 calorimeter trigger.

4.2.2 First APx Demonstrator Board: APd1

Designed at the University of Wisconsin, the Advanced Processor (APx) philosophy aim is to provide a generic trigger processing board to fulfill the requirements of the Phase-2 trigger system in terms of high-speed optical bandwidth and computation power. It features FPGA as the central trigger processing element. The fundamental design idea is to explore a modular methodology exploiting reusable firmware, circuits, and software. The APx philosophy employs the Xilinx ZYNQ system-on-chip (a blend of programmable logic and processing system on the same silicon) device as the primary embedded Linux-based control site. This SoC-based control site uses integrated programmable logic (FPGA) to assist the Linux system in efficiently handling the optical modules, FPGAs, clock synthesizers, and additional board parts.

The APx board hosts several auxiliary mezzanine cards such as Ethernet Switch Mezzanine (ESM) [26], ATCA Rear Transition Module (RTM) [36], Embedded Linux Mezzanine (ELM) [26], and Intelligent Platform Manager Control (IPMC) [26]. APd1 is the first demonstrator board of the APx philosophy, hosting a single Xilinx XCVU9P FPGA and supporting hundreds of high-speed (25 Gbps) optical links. Fig. 4.5 shows the APd1 board and its elements.

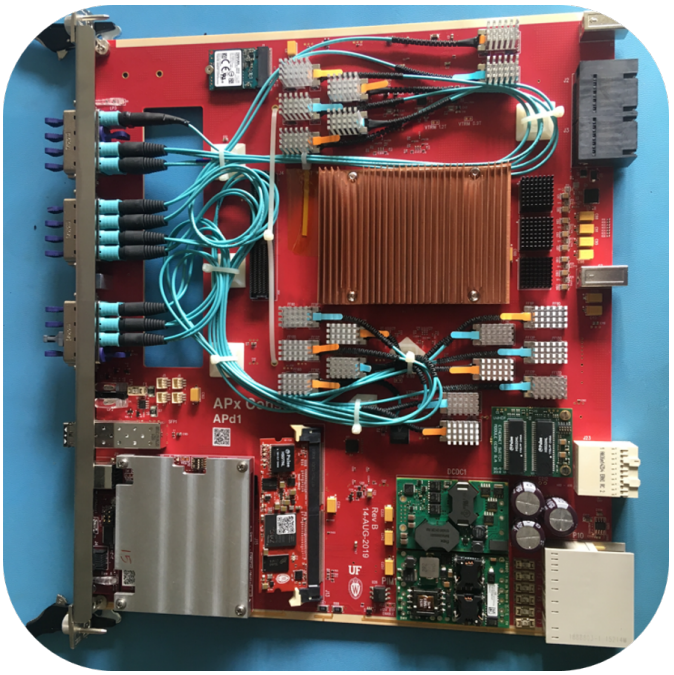

FIGURE 4.5: APd1 board with the XCVU9P FPGA beneath the heatsink.

The IPMC and the ELM provide controls in the APx. The IPMC mezzanine card works with the Shelf Manager (IPM controller) and controls the sensor readouts and crate power on/off. When the ATCA card is plugged in (or the crate is powered up), only the ESM and IPMC cards initially have the power (ESM to equip the IPMC with network connectivity). Upon appropriate negotiation established for crate resources with the Shelf Manager, the IPMC facilitates the power to the ELM and primary board elements (FPGA, optics, etc.). Once booted, the ELM Linux system delivers the operational support and configuration for the platform. It contains the initialization of FPGAs (bitstream loading and register or memory initialization) and the configuration of support devices such as clock switches, Firefly optics, and jitter-cleaner. Communication between the ELM board and APx FPGA(s) is facilitated over an AXI (Advanced eXtensible Interface) bridge which brings the memories and FPGA registers into the address space of the Xilinx ZYNQ CPU at the hardware level. AXI is an interface benchmark advocated by the industry and well merged into the FPGA core catalog and methodology of Xilinx Vivado [37].

4.2.3 CMS Standard Protocol (CSP)

The CMS Standard Protocol (CSP) framework describes the rules and conventions for encoding CMS trigger primitive and processed data onto serial optical links. CSP

is a custom protocol that aims to communicate the physics information at an LHC-synchronous line rate of 25 Gbps over serial links that function at the asynchronous line rate of 25.78125 Gbps utilized in the telecommunication industry.

CSP works by splitting the bandwidth into a data fragment of a determinate size proportionate to the LHC clock (40.079 MHz) and a filler part of indeterminate length, depending on the respective transmitter (T_x) reference frequency. The four main principles of operation of CSP are:

- $Efficiency$: minimize the overhead volume in the LHC bandwidth to expand the amount of physics payload, which can be facilitated in the data stream. CSP fulfills this by accommodating the data integrity checks and synchronization into the filler bandwidth.
- $Flexibility$: Support for the diverse packet lengths required by the various detector subsystems.
- $Reliability$: facilitating means of reliability that diagnose and protect the data stream integrity in the existence of transmission bit errors.
- $Modularity$: allowing various rule-based implementations of CSP to materialize, providing they qualify interoperability tests, thereby widening the area of expertise.

The APx firmware shall incorporate a robust CSP algorithm with extensive reliability and diagnostic capabilities, efficient use of firmware resources, and low latency. The APx implementation of CSP is organized in five processing steps such as:

1. The buffer at the transmitter side (T_x) captures the algorithm's AXI stream data and forwards the stored packets from the buffer to the protocol transmitter module.
2. The protocol transmitter module transforms the captured data from the algorithms to a bit stream according to the CSP conventions for data communication over serial links.
3. An MGT hardware block that incorporates the FPGA transmitter and receiver interfaces.
4. The protocol receiver module receives the serial link data in the CSP structure and transforms it into an AXI stream for the final delivery to the receiver (R_x) side algorithm.
5. A buffer at the receiver side can capture the AXI stream from the CSP protocol receiver and forward the data to the receiver's algorithm.

4.3 DAQ and TCDS Hub: DTH board

In the baseline design planned for the CMS Phase-2 DAQ, sub-detector and sub-system data (i.e., Level-1 trigger systems) are aggregated in the corresponding backend boards. The data is transmitted using mid-board optics over short, point-to-point, and optical links using a custom-developed protocol through the front panel to the DTH. Therefore, the DTH aggregates the data from different backend cards in one backend crate and merges the respective streams for transmission over the data-to-surface (D2S) network for analysis and storage. Hence, the DTH decouples the detector and networking time domains: changes in trigger rate and event size are buffered in the backend boards, and changes in network throughput are buffered in the DTH.

The DTH card is also accountable for circulating the LHC bunch clock, fast commands, Level-1 accepts (L1A), and trigger control directives for synchronization, calibration, etc., to the backend boards, all over dedicated point-to-point backplane connections. On the other hand, the backend cards report their readiness for data-taking to the DTH, where these signals are accumulated for recovery procedures and trigger throttling [33].

The CMS trigger and timing control and distribution system (TCDS) depends on two pieces of information shared between the sub-system electronics and TCDS master. The TCDS master streams the Trigger Timing and Control (TTC) signal to the sub-systems and distributes the beam-synchronous timing commands, LHC bunch clock, and the Level-1 triggers. A Trigger Throttling Stream (TTS) signal notifies the master regarding the readiness of the sub-systems in data-taking. For Phase-2, the throttling (TTS2) and timing (TTC2) streams are established on high-speed serial links (approximately 10 Gbps of line rate) over the backplane and synchronous to the LHC bunch-crossing clock. Fig 4.6 shows the prototype of the DTH board under test.

4.4 Algorithm and Firmware Design Flow

- **Monte-Carlo Simulation**: The algorithms are tested using the Monte Carlo (MC) simulation of the physics background at the HL-LHC condition. Using the MC simulation, we generate the input test vector and reference output for the HL-LHC events. This reference output is used to validate the algorithms during the HLS, RTL, and bitfile testing.

- **HLS Development**: Once the physics scope and the processing region are decided, the algorithm development starts in the high-level synthesis (HLS) tool. The first development step is to perform the HLS simulation and validate the algorithm

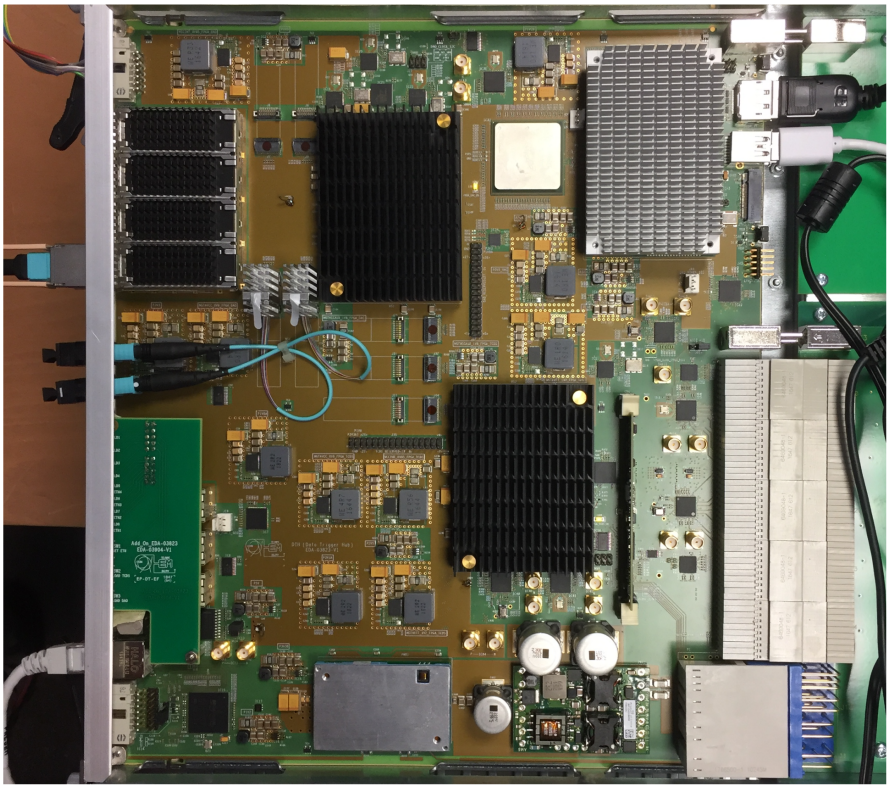

FIGURE 4.6: A first DTH prototype in a test chassis. The left side of the diagram represents the front panel of the DTH board. Two black heatsinks cover the TCDS (bottom) and DAQ (top) FPGAs. Connectivity to central services is supplied via front-panel QSFPs (for the D2S network) and SFPs (for TCDS2). Connections to the node board occur via Firefly optics and backplane connectors, assigning the front-panel optics for sub-system data [33].

using the input test vector and reference output generated in the MC simulation. The HLS tool convert the algorithm written in C/C++ into the RTL, which is further optimized for the area and timing. The timing performance requires designing the algorithm for minimum latency and F_{MAX}. The F_{MAX} defines the robustness of the design, the maximum clock frequency an algorithm can handle without facing any data integrity issue. At the HLS level, the F_{MAX} of the design highlights the maximum clock frequency achieved for a design without providing any input and output pin constraints (standalone mode). Once the algorithm is tuned for the desired latency value, device utilization, and F_{MAX}, the design is exported for the RTL simulation and integration.

- **RTL Integration and Simulation**: In this step, the algorithm is verified in the RTL simulation using the VHDL test bench using the input test vector and reference output generated in the MC simulation. Once the simulation matches the HLS and MC simulated output, the algorithm is integrated with the APx firmware shell, which requires binding the input and output ports of the algorithm

with the device MGTs. The algorithm also required floorplanning in the FPGA in super logic regions (SLR) to meet the F_{MAX} requirement.

- **On Board Bitfile Test**: Once the design passes the timing constraints of the design, the algorithm firmware (bitfile) is then uploaded in the APd1 board and validated using the test vector and reference output generated in the MC simulation. If the output of the bitfile matches the previous validation steps, the system is prototyped successfully. Therefore, successful validation means the output matches at all the following stages mentioned before. Fig. 4.7 represents the design flow of the Phase-2 calorimeter trigger algorithm and firmware.

$$MC == HLS == RTL == Bitfile \tag{4.1}$$

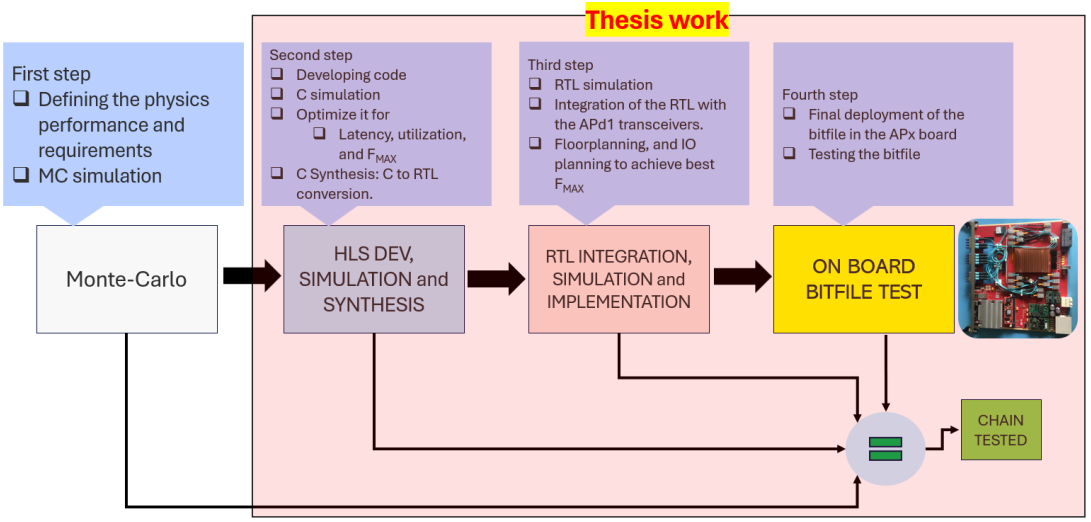

FIGURE 4.7: Design flow of the Phase-2 calorimeter trigger.

4.5 FPGA Specifications

The FPGA used for the calorimeter trigger system design and prototyping is the Xilinx 3D IC XCVU9P device, which is fabricated using SSI technology [38]. In this fabrication technology, several identical FPGA dies, also known as silicon logic regions (SLR), are stacked up to form a larger device using passive silicon interposers with microbumps and through-silicon vias (TSVs). The adjacent SLRs are connected using a high bandwidth and low latency super long line (SLL). The XCVU9P FPGA supports three SLRs where the SLL connects the SLR2 with SLR1 and SLR1 to SLR0 (no direct connection between SLR2 and SLR0). Fig 4.8 represents the side view of the FPGA die fabricated with four SLRs.

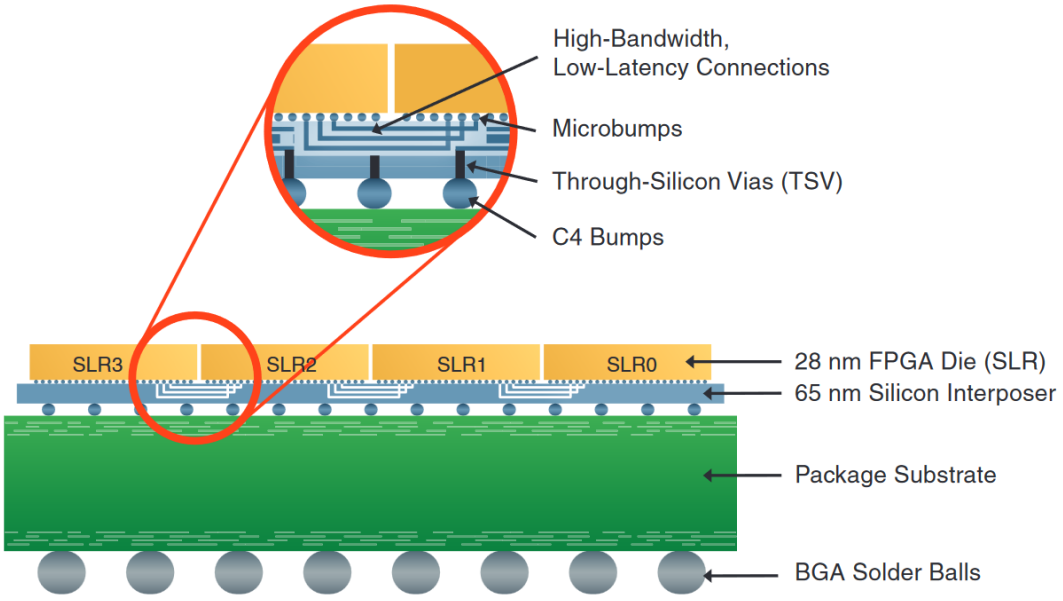

FIGURE 4.8: FPGA fabricated using SSI technology [38].

4.6 Major Design and Prototyping Constraints

Following section defines the major constraints of the design and prototyping.

4.6.1 Very high-throughput

The increased granularity of the input information planned for the Phase-2 upgrade leads to the constraint of very high throughput for the calorimeter trigger system. For the first time, the ECAL front-end system will provide each crystal energy and timing information. Ultimately, it increases the granularity by 25 times compared to the Phase-1 system. Similarly, the granularity at the HCAL and HGCAL increases 4 and 500 times, respectively. For example, the Phase-2 regional calorimeter trigger (RCT) will handle the throughput of 75 Tbps compared to the throughput of 5.18 Tbps of the Phase-1 system. The constraint of high throughput is handled by employing FPGAs with higher MGTs counts and bandwidth. The current Xilinx UltraScale+ XCVU9P FPGA support the transceivers (MGTs) which can operate with the data bandwidth of 32 Gbps having 120 such MGTs block fabricated on the device. Therefore, it can handle the throughput of 3.93 Tbps, approximately 5 times the Virtex-7 FPGA used in the Phase-1 system.

4.6.2 Input and output placement

The SLR division of the FPGA poses significant challenges over the implementation of the trigger algorithms. One the main challenge is the distribution of the transceivers

(MGTs) over the entire FPGA territory. This distribution governs the partition of the trigger algorithm over the FPGA and if the algorithm implemented in their corresponding SLR require to communicate with each other is required to cross the SLR boundary using the SLL. This inter SLR communication impact the timing performance of the algorithm in terms of latency. The major impact of this constraint is discussed in the chapter 5.

4.6.3 Latency

As mentioned earlier, the CMS trigger system should rapidly and selectively take its decision; otherwise, the experiment will lose valuable collision data. Therefore, latency is one of the significant constraints of the trigger system. The Phase-2 calorimeter trigger must process the high throughput data from the calorimeter sub-detectors in the required time budget. The latency budget of the calorimeter trigger (RCT and GCT barrel) to process the barrel part of the detector and provide the output to the correlator trigger is 3 μS. Similarly, the latency budget of the calorimeter trigger to provide the output to the global trigger is 4 μS, which includes the processing of the two endcaps. Therefore, the algorithms must fit the latency budget.

4.6.4 Utilization

The logic utilization of the calorimeter trigger directly impacts the scalability of the system. Lower device utilization increases the scalability of the system. It allows the integration of more complex or logic-intensive algorithms at the FPGA and ensures the overall flexibility of the system throughout the operational period. The device utilization of the trigger algorithm is preferred within the 50% of the device logic.

4.6.5 Input and Output bit-width constraints

The 25 Gbps of links bandwidth and 40 MHz of LHC clock (one bunch-crossing frequency) corresponds to 576 bits of physics payload available for the algorithm in one bunch-crossing. The bit-width constraint on the input ports of the algorithm defines the amount of information it can receive from the detector in each bunch-crossing. Similarly, the output port bid-width constraint defines the amount of data the algorithm can pack for each bunch-crossing. For example, the information of one ECAL crystal is 14-bits wide, and there are 25 crystals in one ECAL tower. Therefore, tower information of 25 crystals occupies 350 bits in the available 576-bit wide input bus. The rest of the bits

can be utilized in the later design stage to pack additional crystal information. Therefore, the number of bits packed in the input and output must be optimum to ensure the flexibility of the system. At the same time, it should not be underutilized, which leads to the utilization of more links. Using more links in the algorithm increases the design compile time from HLS to bitfile and requires additional IO planning to meet the timing constraints.

The next chapter discusses the Phase-2 calorimeter trigger algorithms employed at RCT, GCT barrel, GCT endcap, and GCT sum.

Chapter 5

Level-1 Calorimeter Trigger Algorithms

5.1 Level-1 Calorimeter Trigger Objects

The primitive unit of the electromagnetic calorimeter (ECAL) is the lead-tungstate ($PbWO_4$) crystal. The group of 25 such crystals ($5\eta \times 5\phi$) is one ECAL tower. Similarly, in the hadronic calorimeter, the primitive unit is one HCAL tower. The position of the ECAL crystal, tower, and HCAL tower is geometrically fixed in the CMS detector. The Level-1 calorimeter trigger creates the electron/photon (e/γ), tau lepton, and hadronic jet information by computing cluster energy in various regions of crystals and towers. The preferred region of crystal and tower require to compute the cluster energy depends on the interaction of these particles in the calorimeter detector. For example, a photon created at the interaction point passes through the tracker, penetrates ECAL crystals, and loses all its energy through an electromagnetic shower. If a photon passes through the tracker without interaction, it deposits about 94% (97%) of its energy into $3\eta \times 3\phi$ ($5\eta \times 5\phi$) ECAL crystal region [39]. The calorimeter trigger computes the e/γ cluster energy in the $3\eta \times 5\phi$ crystal region around the peak crystal. The crystal with the highest energy deposition is termed a peak crystal. In addition, the electron might lose its energy due to bremsstrahlung radiation. Hence, it bends further in the *phi* direction due to the presence of the magnetic field. Therefore, the calorimeter trigger computes the cluster energy in the $3\eta \times 5\phi$ crystal region in positive and negative phi direction of the primary e/γ clusters to account for the bremsstrahlung losses. The jets and taus have a larger footprint in the calorimeter as hadronization and the subsequent decays of unstable hadrons can create hundreds of particles close to each other in the CMS detector. The jet and tau energy is computed by accumulating the energy of the

towers in the $9\eta \times 9\phi$ tower geometry [26]. The calorimeter trigger also computes the particle flow (PF) cluster by accumulating the tower energies in $3\eta \times 3\phi$ tower geometry. The correlator trigger uses the PF clusters to identify the hadrons. Fig 5.1 depicts the calorimeter trigger object and the various clustering regions in crystals and towers.

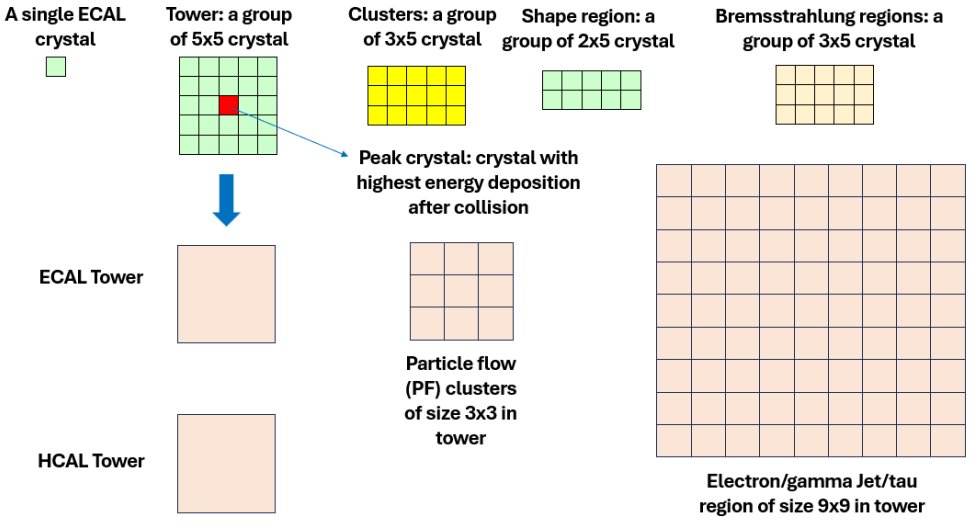

FIGURE 5.1: Illustration of the ECAL crystal, tower, HCAL tower, and the electron/photon clusters. The clustering includes the bremsstrahlung losses in the *phi* direction and jets region to account for the hadronization of quarks and gluons. The particle flow (PF) clustering is considered for the correlator trigger.

5.2 Detector Region Partitioning for Algorithms

The Level-1 calorimeter trigger algorithms are designed to process a distinct detector region and are implemented on the FPGA board. The size of the detector region for the algorithm processing is decided based on the number of multi-gigabit transceivers (MGTs: the input and output 25 Gbps links of the FPGA) the FPGA supports. The XCVU9P FPGA, the central processing element of the APd1 board, supports 100 such transceivers. Therefore, it is necessary to partition the calorimeter detector geometry into smaller regions based on the number of links (MGTs) the FPGA can support.

In addition, the transceivers on this FPGA are spread across the three SLRs, with SLR0, SLR1, and SLR2 supporting 24, 36, and 40 links, respectively. Therefore, the algorithm processing region must partitioned into SLRs in the early stage of development. These transceivers are instantiated using the APx firmware shell, which is floorplanned at the boundary of the XCVU9P FPGA in six sectors. The middle of the SLRs is reserved for the trigger algorithm. Fig. 5.2 shows the floorplan of the XCVU9P FPGA and the sector-wise link allocation in each SLR.

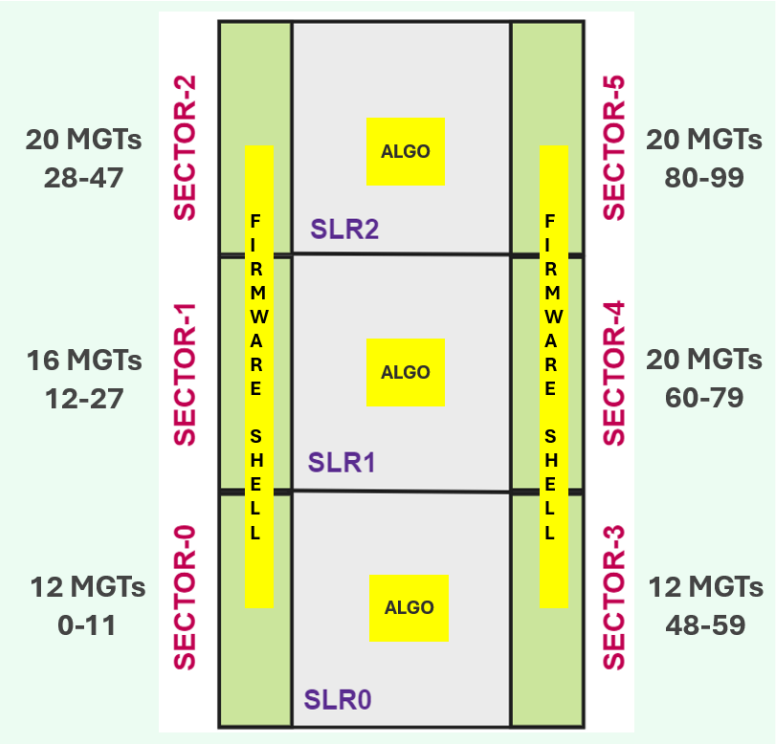

FIGURE 5.2: Floorplan of the XCVU9P FPGA. The device is fabricated by stacking three SLR dies. The MGTs of the FPGA are instantiated using the APx firmware shell, which is implemented in six sectors. The firmware shell is located at the boundary, near the MGTs location [6].

The partitioning of the detector region and the algorithm functionality significantly dictate the architecture of each calorimeter trigger algorithm, which will be discussed in the following sections.

5.3 Layer-1: Regional Calorimeter Trigger (RCT)

As discussed earlier, the implementation of the barrel calorimeter trigger is proposed in two levels. The first level, the Regional Calorimeter Trigger (RCT), takes input from the barrel backend system and processes the entire detector barrel of ECAL and HCAL in 36 APx boards. The division of the detector geometry is identical, and each board processes the tower region of $17\eta \times 4\phi$ of ECAL and $16\eta \times 4\phi$ of HCAL, employing an identical algorithm. The identical division of the barrel region ensures that the prototyping of 36 boards can be readily performed using one board. The processing tower region of $17\eta \times 4\phi$ in ECAL and $16\eta \times 4\phi$ in HCAL corresponds to accommodating 70 (68 from ECAL and 2 from HCAL) input links, which fit into the constraints of 100 links of the APd1 board. Nevertheless, It can also process the region of $17\eta \times 5\phi$ in ECAL and $16\eta \times 5\phi$ in HCAL, corresponding to 88 input links (85 from ECAL and 3 from HCAL). However, this makes the barrel partition non-identical and requires a minimum of two

different algorithms and two boards to prototype the system. Similarly, choosing the tower region of $17\eta \times 6\phi$ in ECAL, which partitions the barrel ECAL into 24 identical regions, can be prototyped using a single algorithm and one board. However, the number of input links associated with this region, 102 for ECAL, made it impossible to implement using the APd1 board. Therefore, it is essential to find a balance between the number of links required and the processing region for a single board, such that it makes the prototyping easier and scalable. Therefore, we choose the tower region of $17\eta \times 4\phi$ in ECAL and $16\eta \times 4\phi$ in HCAL as the processing region for layer-1 RCT board.

The barrel calorimeter geometry highlighting the coverage of one RCT board is depicted in Fig. 5.3. The $x - axis$ represents the integer *phi* ($i\phi$), where each integer represents 5 degrees in *phi*. The $y - axis$ represents the integer *eta* ($i\eta$) or pseudorapidity. This region represents the barrel geometry of ECAL and HCAL. The ECAL comprises 34 towers in *eta* and 72 in *phi* ($34\eta \times 72\phi$). Similarly, the HCAL comprises 32 towers in *eta* and 72 in *phi* ($32\eta \times 72\phi$). Each small square represents one tower of HCAL and ECAL. For ECAL, one tower represents 25 lead-lungstate crystals. The primary function of the RCT is to prepare towers and e/γ cluster information for the GCT barrel. The next sub-sections explain the RCT v1.0 and RCT v2.0 algorithm in details.

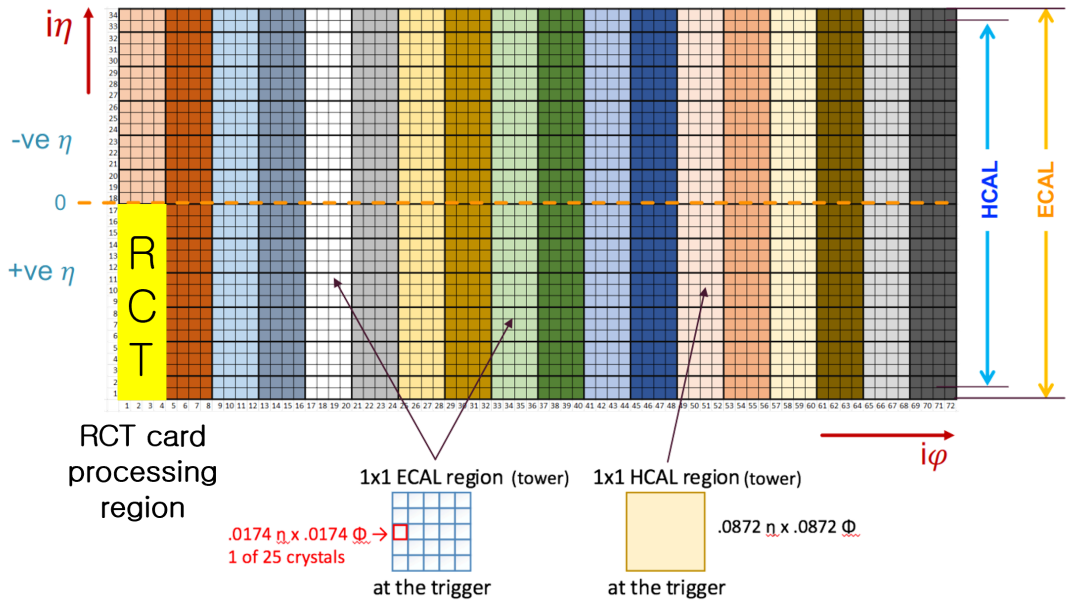

FIGURE 5.3: CMS barrel calorimeter geometry, including the ECAL and HCAL detectors. The processing region of one RCT board, i.e., $17\eta \times 4\phi$ towers in ECAL and $16\eta \times 4\phi$ towers in HCAL, is highlighted in the yellow box [40].

5.3.1 RCT v1.0

The RCT region of $17\eta \times 4\phi$ tower in ECAL corresponds to processing 68 input links, which satisfies the 100 link constraint of the APd1 board. However, in APd1, the maximum number of links it can support in one SLR is not more than 40. Therefore, it is necessary to partition the RCT region of, i.e., $17\eta \times 4\phi$ towers into smaller regions to satisfy the SLR link constraint. In the RCT v1.0 algorithm, the ECAL tower geometry of $17\eta \times 4\phi$ (68 links) is partitioned across the central phi boundary into two identical sub-regions of $17\eta \times 2\phi$ (34 links), proposed to be implemented in SLR2 (40 links) and SLR1 (36 links), respectively [6]. The RCT v1.0 algorithm is proposed in two stages. In the first stage, the RCT17x2 algorithm concurrently processes the two $17\eta \times 4\phi$ ECAL tower regions. The RCT17x2 algorithm receives the ECAL information in 34 input links, each carrying 25 ECAL crystal information. It creates the e/γ cluster and tower information and provides the output to the next stage algorithm called the PhiMerge algorithm. The PhiMerge algorithm combines the e/γ cluster information of both the RCT17x2 and sends the final output to the GCT barrel board. Fig. 5.4 represents the two-stage architecture of the RCT v1.0 algorithm.

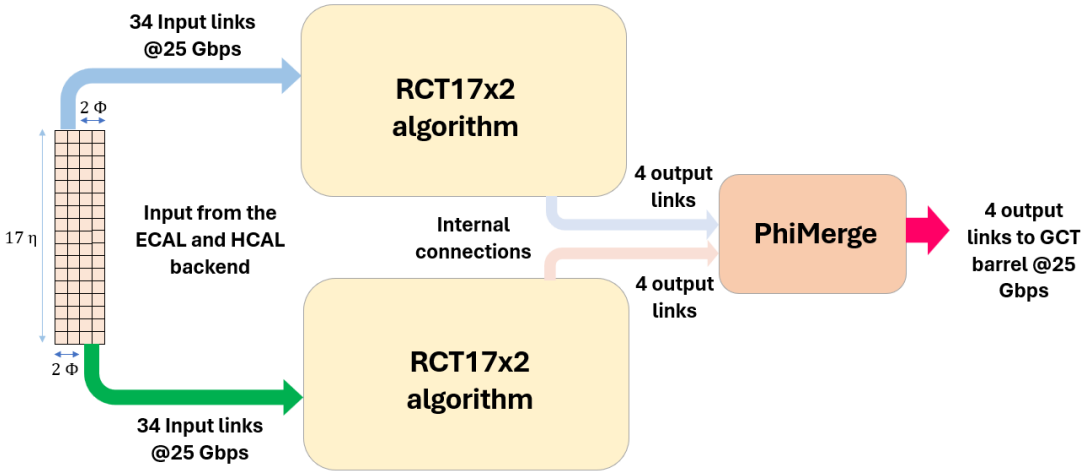

FIGURE 5.4: Two-stage architecture of the RCT v1.0 algorithm. The first level employs two identical RCT17x2 algorithms processing the crystal information. The second stage is implemented using the PhiMerge algorithm, which combines the information from the first level and sends it to the GCT barrel board.

The primary function of the RCT17x2 algorithm is to create the e/γ cluster and tower information. The algorithm first unpacks the input link containing the crystal information with 14 bits allocated for each crystal. Each crystal contains information in terms of 10 bits for crystal energy, 3 bits for timing, and 1 bit for electronic spike in the crystal. The timing and spike information is not included in the RCT v1.0 algorithm. Each input link contains 25 crystals information, occupying 350 bits out of 576 bits available. Table 5.1 summarized the crystal packet details.

Quantity	Number of bit(s)
crystal energy	10
timing	3
spike	1
	Total: 14

TABLE 5.1: Input ECAL crystal information.

The main steps of RCT17x2 algorithm are

- Finding peak crystal in each tower using the energy-weighted sum methodology.
- Computing eγ cluster energy in the $3\eta \times 5\phi$ crystal region.
- Stitching the cluster energy of one tower with another in *eta* and *phi* direction to account for energy loss at the tower boundaries.
- Calculate the total energy of the tower by adding all the 25 crystal energy.
- Final stitching of the two $17\eta \times 2\phi$ tower regions across the *phi* boundary.
- Packing the tower information in 32 bits, which includes 10 bits for cluster energy, 10 bits for tower energy, 3 bits for *phi* position of the peak crystal, 3 bits for eta position of the peak, 3 bits for the timing information, and 3 bits for HCAL to ECAL (H/E) energy ratio. The timing and H/E ratio are not included in the algorithm output as it only processes the ECAL crystal information and uses only the crystal energy information. Table 5.2 summarized the packing of one tower information in a 32-bit word.
- Ultimately, packing the 68 towers (one ϕ slice) information in four output link (17 towers per link).

Quantity	Number of bit(s)
e/γ cluster energy	10
tower energy	10
peak crystal phi	3
peak crystal eta	3
time	3
H/E	3
	Total: 32 bits

TABLE 5.2: Output e/γ cluster information packing in 32 bits.

The position of the peak crystal in a tower is computed using the energy-weighted sum algorithm. This algorithm determines the peak crystal in the two-dimensional

($5\eta \times 5\phi$) crystal matrix. The energy deposit or hit pattern of electrons or photons is used to find the peak crystal in the tower. First, it prepares strips (or arrays) in *eta* (row) and *phi* (column) by adding all the corresponding crystal energy in each of the *eta* and *phi* positions. The tower energy is calculated by summing the energy of either *eta* or *phi* strip elements. The algorithm then computes the weighted sum of the *eta* (*phi*) strip by applying the appropriate weight on each of the elements of the strip and adding them in the end. The difference between the weighted sum and tower energy provides the position of the peak crystal in the *eta* and *phi* direction. Tab 5.3 determines the position of the peak element in the *eta* and *phi* strip based on comparing the weighted sum of the strip energy and tower energy. The position of the peak element in *eta* and *phi* strip determines the peak crystal position in a tower. Fig 5.5 demonstrates the formation of the strip energy, weighted sum, and tower energy.

Comparison equation	Peak element position in *eta* and *phi*
weighted strip energy $\leq$ tower energy	0
weighted strip energy $\leq 2 \times$ tower energy	1
weighted strip energy $\leq 3 \times$ tower energy	2
weighted strip energy $\leq 4 \times$ tower energy	3
others	4

TABLE 5.3: The position of the peak element in the *eta* and *phi* strips is based on comparing the weighted sum and tower energy.

Based on the peak crystal position, the algorithm computes the cluster energy in $3\eta \times 5\phi$ crystal region in each tower. The process of making tower information requires 26 clock cycles. It includes finding the position of the peak crystal, computing cluster energy, and computing the tower energy. One of the main features of the RCT17x2 algorithm is the independent computation of clustering for each of the 34 towers. The algorithm computes the cluster and prepares the tower information for each of the 34 towers concurrently. Thus, reducing the overall latency of the algorithm. Fig 5.6 shows the cluster formation of size $3\eta \times 5\phi$ around the peak crystal in the tower.

The electron and photon deposit their energy arbitrarily in the electromagnetic calorimeter. Therefore, the footprint of the electron and photon energy deposition (clustering region) in ECAL may lie between the boundaries of multiple RCT regions, or between two $17\eta \times 2\phi$ regions, or between two towers. Hence, it is necessary to merge the clusters at the appropriate *eta* and *phi* boundaries. The stitching of two clusters requires satisfying the neighboring condition of the corresponding peak crystals in the *eta* and *phi* direction. In RCT17x2, the stitching logic is applied for two iteration in *eta* to cover two distinct boundary sequences and once in the *phi* direction. The stitching logic is explained below:

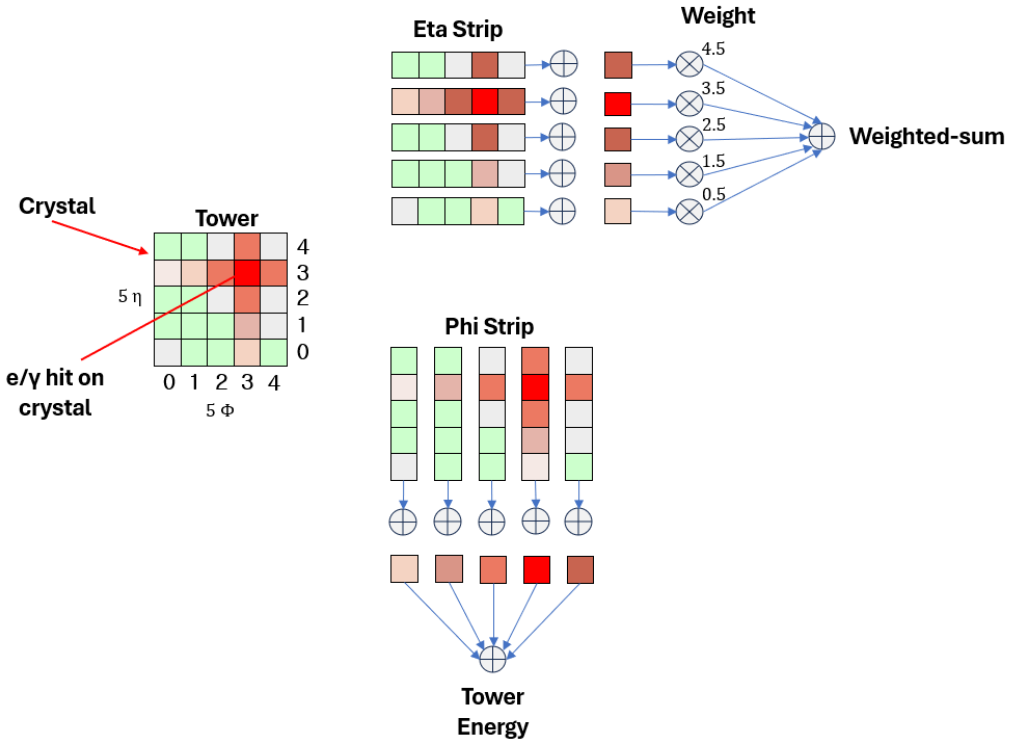

FIGURE 5.5: The illustration of energy-weighted sum algorithm for finding peak in a tower.

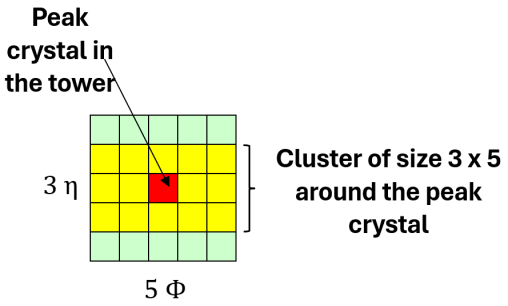

FIGURE 5.6: Cluster formation of size $3\eta \times 5\phi$ around the peak crystal energy in an ECAL tower.

- **Stitching in eta**: For two towers, stitching of the clusters in *eta* is true if the peak crystals of both the clusters have identical *phi* and are neighbors in *eta*.
- **Stitching in phi**: For two towers, stitching of clusters in *phi* is true if the peak crystals of both the clusters have identical *eta* and are neighbors in *phi*.

Fig 5.7 explain the two stitching sequence in *eta* and the required steps for stitching two clusters. Fig 5.8 explain the stitching steps in the *phi* direction.

The latency for stitching clusters in the eta and phi direction is four clock cycles. The stage-2 PhiMerge algorithm stitches the clusters of two $17\eta \times 2\phi$ tower

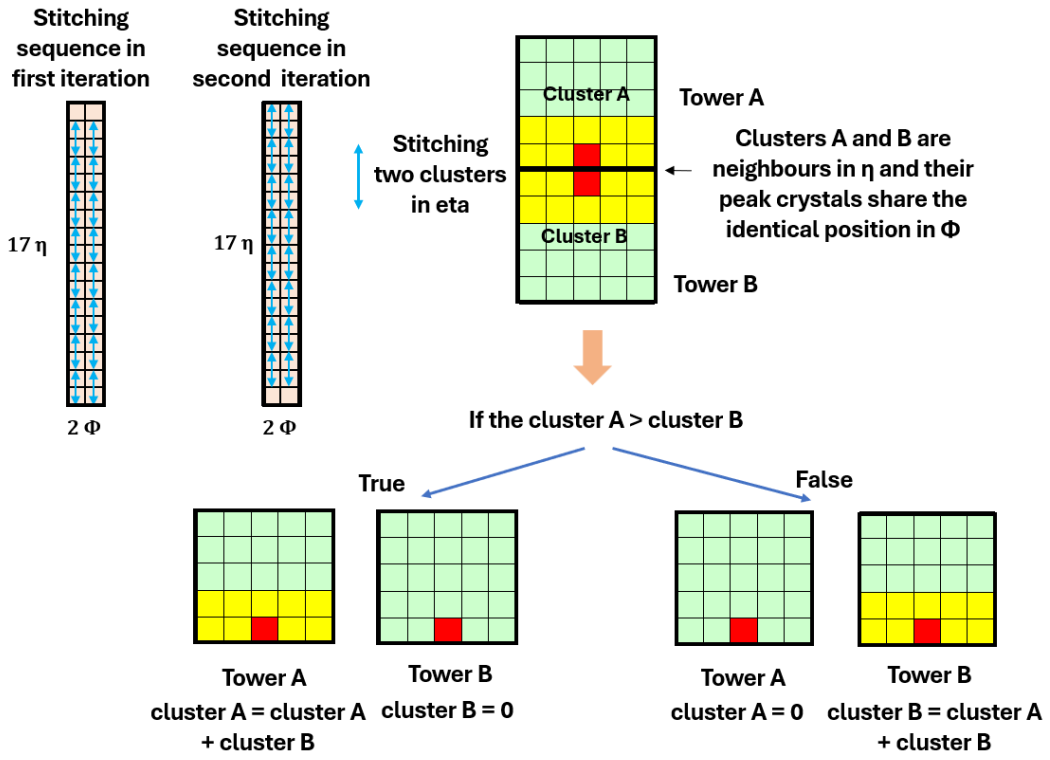

FIGURE 5.7: Stitching logic of clusters in eta. The double-sided arrow in the $17\eta \times 2\phi$ RCT geometry signifies the direction of stitching in *eta* direction.

regions across the *phi* boundary. **The latency of the PhiMerge algorithm is nine clock cycles**. The final *phi* stitching can be implemented at SLR2 or SLR1, which internally receives the output information from the two RCT17x2 algorithms.

The algorithm is synthesized in HLS, and the RTL is verified in Vivado. Fig 5.9 and 5.10 represent the HLS timing and RTL simulation result of the RCT17x2 algorithm. In the RTL, the algorithm latency is measured by taking the difference between the first rising edge of the *algoStart* and the *algoDone* signals.

The RCT17x2 and PhiMerge RTLs are combined using a VHDL wrapper. The wrapper is designed to provide LHC bunch-crossing of input along with the standard control signal such as *start* and *reset*. **The total latency of the RCT v1.0 algorithm, including the VHDL wrapper, is 97 clock cycles**.

RCT v1.0 algorithm handles the latency constraint by employing the flip-flops (FFs) as the primary storage elements for the crystals and using the shift register logic (SRL) for a long buffer chain. Employing FFs and SRLs removes the bottleneck of the read latency associated with BRAMs. Table 5.4 and 5.5 summarizes the latency and utilization results of the RCT v1.0 algorithms, respectively.

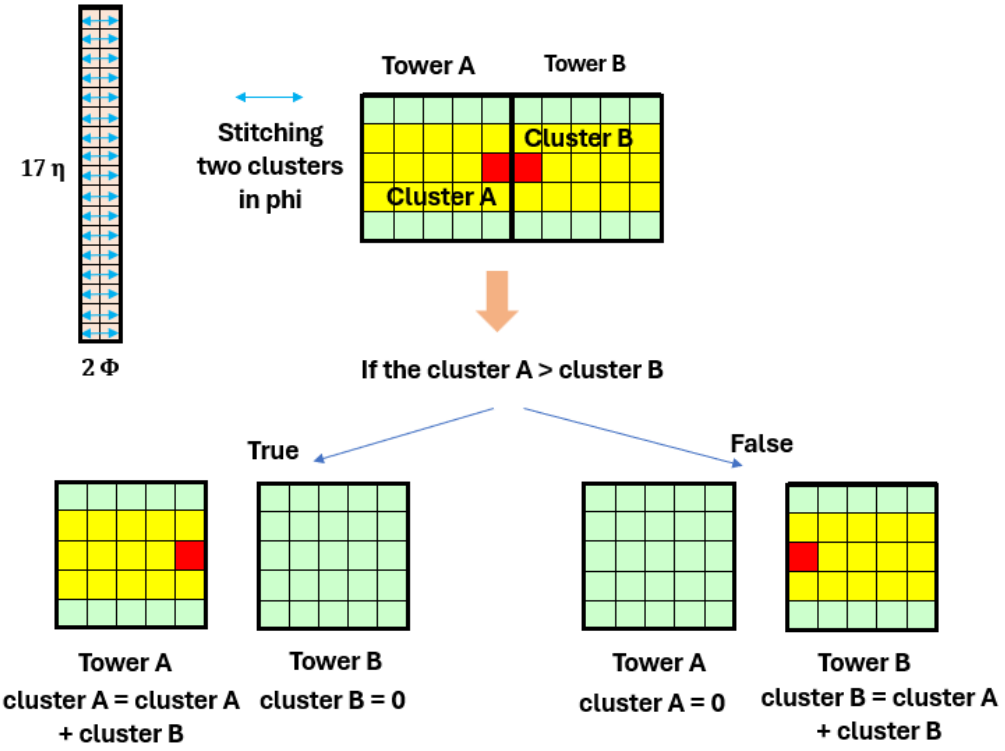

FIGURE 5.8: Stitching logic of clusters in *phi*. The double-sided arrow in the $17\eta \times 2\phi$ RCT geometry signifies the direction of stitching in *phi* direction

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	1.442 ns	1.67 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
50	50	0.139 us	0.139 us	9	9	function

FIGURE 5.9: RCT17x2 HLS latency report.

The implementation details such as floorplaning, validation, and prototyping, are discussed in the next chapter.

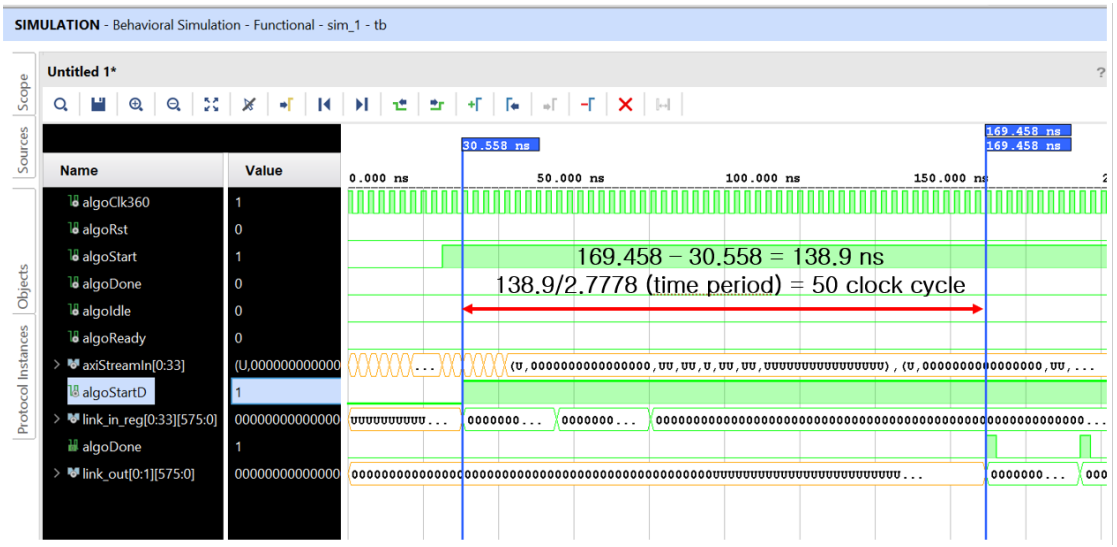

FIGURE 5.10: RCT17x2 RTL simulation validating the latency of 50 clock cycle.

Input to RCT v1.0 algorithm: **ECAL**: a region of $17\eta \times 4\phi$ towers corresponds to $85\eta \times 20\phi$ crystals (1700 crystals) **Input links**: 68 **clock frequency**: 360 MHz	
ECAL region for RCT is sub-divided in two $17\eta \times 2\phi$ tower regions.	
RCT17x2 algorithm	
Creating tower information by unpacking the input links and filling 25 crystals data in each tower	
Processing each towers concurrently	
❶ Finding peak crystal in a tower and calculating the cluster and tower energy	**26 clock cycles**
❷ Stitching clusters in eta direction	**4 clock cycles**
❸ Stitching clusters in phi direction	**4 clock cycles**
PhiMerge algorithm	
Stitching clusters at the phi boundary	**9 clock cycles**
Total latency*	**97 clock cycles**
Output: Output links: 4 Tower: 68 barrel calorimeter tower information Clusters: 8 clusters information	

TABLE 5.4: RCT v1.0 algorithm steps and latency report. * Total latency includes the latency of the wrapper which combines RCT17x2 and PhiMerge algorithm in RTL and it also included the SLR crossing of the internal signals.

	RCT17x2	PhiMerge
Latency (ns)	139	25
F_{MAX} (MHz)	410	422
LUTs	3.6 %	0.2 %
FFs	3.8 %	0.3 %

TABLE 5.5: Timing and utilization summary of the RCT v1.0 algorithm.

5.3.2 RCT v2.0

The processing region of the RCT v2.0 algorithm is $17\eta \times 4\phi$ towers in ECAL and $16\eta \times 4\phi$ towers in HCAL. It requires 68 input links for the ECAL and 2 input links for the HCAL. As discussed earlier, the constraint of a maximum of 40 links available in one SLR of XCVU9P FPGA, the algorithm needs to be partitioned into SLRs. In the RCT v2.0 algorithm, the $17\eta \times 4\phi$ ECAL tower region is partitioned across the *eta* direction, which results in $8\eta \times 4\phi$ and $9\eta \times 4\phi$ tower regions. The two tower regions are subdivided into three smaller $3\eta \times 4\phi$ tower regions (the $8\eta \times 4\phi$ tower region is extended in *eta* by adding a row of $1\eta \times 4\phi$ null towers and then partitioned into three $3\eta \times 4\phi$ tower regions). The partition of the RCT geometry across the *eta* ensures the accuracy of calculating the clusters near the interaction point, which was not considered in the RCT v1.0 algorithm. Fig 5.11 represents the CMS barrel calorimeter geometry highlighting the coverage of one RCT board.

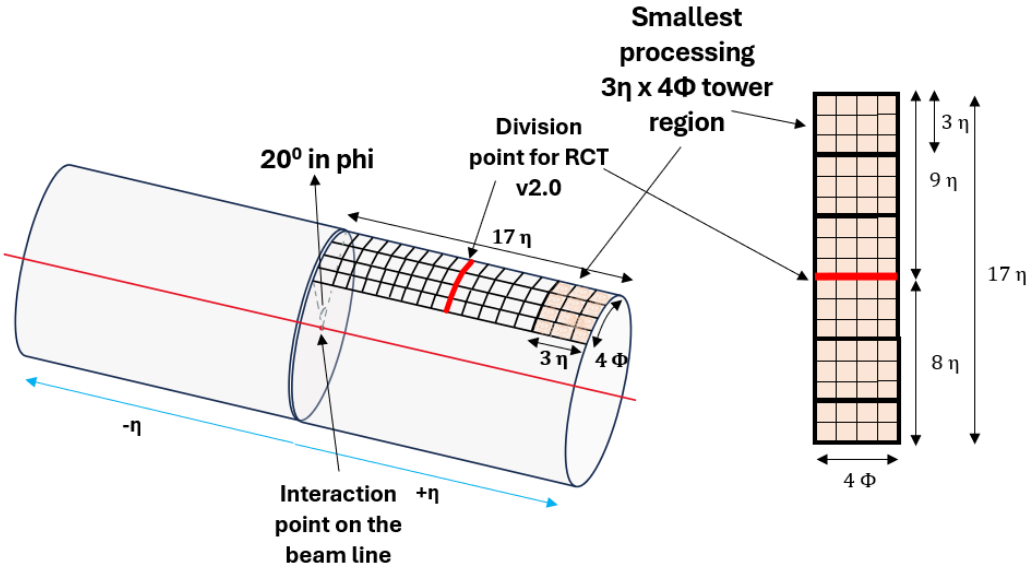

FIGURE 5.11: RCT v2.0 covering region in the CMS ECAL barrel.

As discussed earlier in Chapter 4, the constraint of device utilization and overall timing performance (latency and F_{MAX}), the RCT v2.0 subdivide the algorithm functionality

into three individual IPs. The functionality sub-division ensures the algorithm's performance at the desired frequency of 360 MHz, which was not achieved without this sub-division. Fig 5.12 shows the architecture of the RCT v2.0 algorithm. The tower region of $9\eta \times 4\phi$ and $8\eta \times 4\phi$ takes the input from 36 and 32 ECAL towers via 25 Gbps links and processes them in three distinct algorithms labeled IP1, IP2, and IP3. The HCAL towers energy is added to the corresponding ECAL towers in the final IP3. Therefore, a buffer is required to hold the HCAL information, which is the latency of the arrival time of the ECAL information from IP1 to IP3.

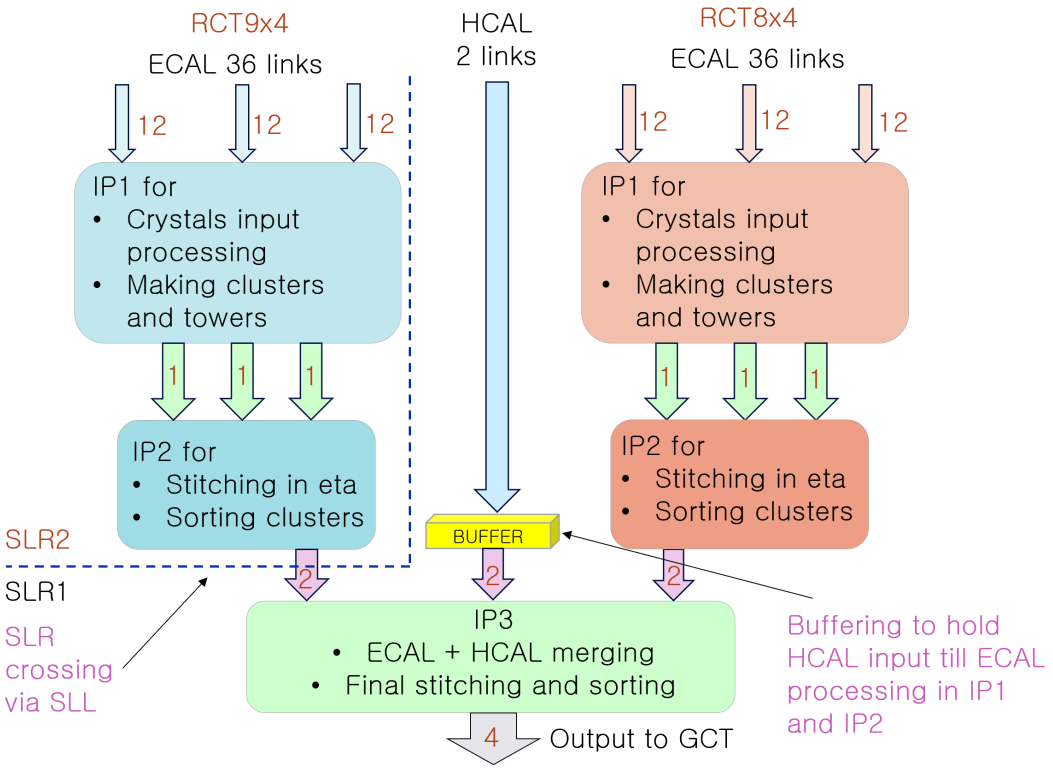

FIGURE 5.12: RCT v2.0 architecture with a regional partition in two tower region of $9\eta \times 4\phi$ and $8\eta \times 4\phi$, and functional division into three IPs or algorithms labeled IP1, IP2, and IP3.

The ECAL tower region of $3\eta \times 4\phi$ is the smallest processing region of RCT v2.0, which corresponds to the $15\eta \times 20\phi$ crystal region. The processing strategy for each IP in RCT v2.0 is as follows.

- **IP1**: Processes the $8\eta \times 4\phi$ and $9\eta \times 4\phi$ tower regions concurrently in two SLR. Following are the algorithm steps:
  - Mapping of the ECAL crystal information into three $3\eta \times 4\phi$ tower regions. Mapping in $3\eta \times 4\phi$ tower region corresponds to filling the crystal information such as energy and position in *eta* and *phi* inside the $15\eta \times 20\phi$ crystal region.
  - Finding peak crystal position in the $15\eta \times 20\phi$ crystal region.

- Computing the central e/γ cluster energy by adding the crystal energies in the $3\eta \times 5\phi$ crystal region around the peak crystal.
- Computing two bremsstrahlung cluster energy by adding the crystal energies in the $3\eta \times 5\phi$ crystal region in the positive and negative *phi* side of the primary cluster.
- Computing two shape cluster energy by adding the crystal energies in the $2\eta \times 5\phi$ crystal region inside the primary cluster.
- Removing the peak crystal from the $15\eta \times 20\phi$ crystals to prepare this region for the next iteration of clustering.
- Repeating the clustering steps for four iterations to build 4 clusters in one $3\eta \times 4\phi$ tower region. Thus preparing 12 clusters in three such regions.
- Computing 12 towers energies in the $3\eta \times 4\phi$ tower region. Thus, creating 36 tower energy in the three such regions.
- Packing three output links for the IP2, each containing 12 towers and 4 clusters information.

- **IP2**: It concurrently processes the output of the $9\eta \times 4\phi$ and $8\eta \times 4\phi$ tower region in two SLRs. Following are the algorithm steps
  - Receives the information of 12 clusters and 36 towers of three $3\eta \times 4\phi$ tower regions from the IP1.
  - Performing the stitching of the clusters at the two eta boundary.
  - Sorting the 12 clusters using *Bitonic sort* algorithm. Select the six highest clusters for IP3. The remaining six clusters are added back with the corresponding tower energy.
  - Send the final output to the IP3 over two output links, each containing 18 towers and 3 clusters information.

- **IP3**: Processes the output of two IP2 of $9\eta \times 4\phi$ and $8\eta \times 4\phi$ tower regions. It also process the HCAL input links. Following are the algorithm steps
  - Receives the e/γ clusters and towers information from the IP2 of $9\eta \times 4\phi$ and $8\eta \times 4\phi$ tower region. It also receives the HCAL tower information in 2 input links. Overall, it receives 12 e/γ clusters, 68 ECAL towers, and 64 HCAL tower information.
  - Perform the stitching of the clusters at the central eta boundary.
  - Sort the 12 e/γ clusters using the *Bitonic Sort* algorithm. The highest eight clusters are the output for the GCT barrel.

- The remaining four clusters are added back with the corresponding tower.
- Perform the addition of HCAL and the corresponding ECAL towers. The ratio of the ECAL and HCAL tower energies (H/E) is calculated.
- Packs four output links for the GCT Barrel. Each link carrying 17 towers and 2 clusters information.

The input of the RCT v2.0 is ECAL crystal and HCAL tower. Each ECAL crystal carries 14 bits of information, with 10 bits for crystal energy and 4 for timing information. The HCAL tower information is packed in 16 bits, with 10 bits allocated for the tower energy and 6 bits for the HCAL feature. The HCAL feature bits enable encoding of (a) longitudinal shower shape data for usage in lepton isolation, calibration, and tagging of minimum ionizing particles (MIPs), (b) shower time information created from the 2 bits of the Time-to-Digital-Converter (TDC) data available in the individual constituent channel of the trigger tower. However, in the current RCT v2.0 algorithm, the timing information of ECAL and feature bit information of HCAL are yet to be incorporated. Fig 5.13 shows the bit structure of ECAL crystal and HCAL tower used in the RCT v2.0 algorithm.

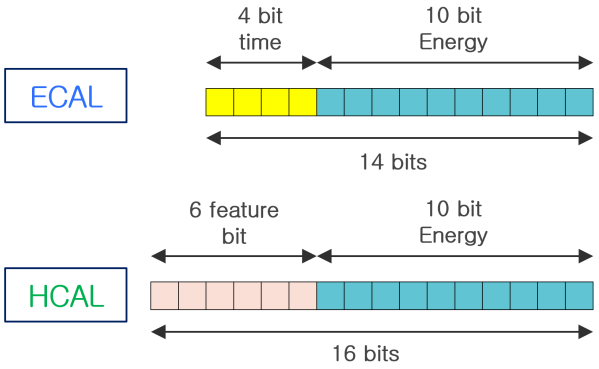

FIGURE 5.13: RCT v2.0 ECAL crystal and HCAL tower input bits pattern.

5.3.2.1 IP1

The IP1 concurrently processes the $9\eta \times 4\phi$ and $8\eta \times 4\phi$ tower regions in two SLR. It partition this region into three $3\eta \times 4\phi$ mutually independent tower regions. In the first step, IP1 maps the ECAL crystal information into three $3\eta \times 4\phi$ tower regions, corresponding to mapping the crystals in three $15\eta \times 20\phi$ crystal regions. During the crystal mapping, the IP1 assigns the positional information to each crystal in terms of its location in the *eta* and *phi* in $15\eta \times 20\phi$ crystal region. Fig 5.14 represents the RCT v2.0 partitioned tower regions and also includes the $15\eta \times 20\phi$ crystal region.

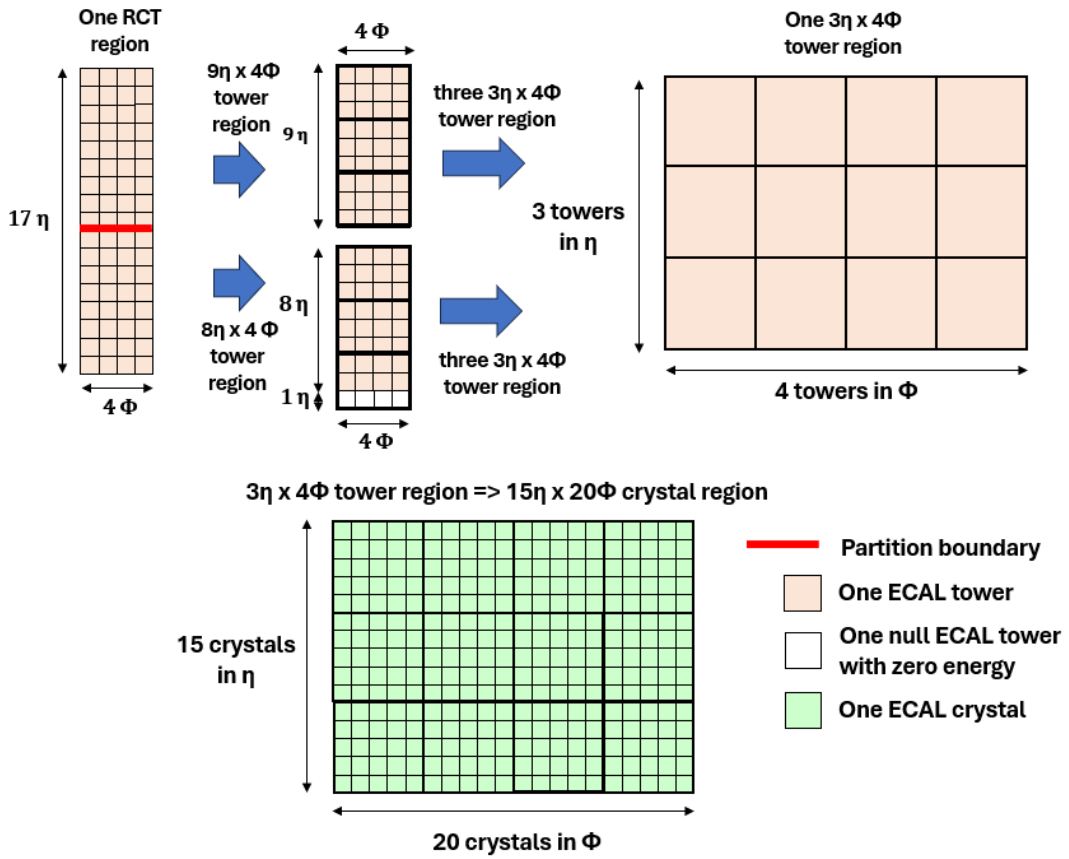

FIGURE 5.14: IP1 ECAL crystal mapping in $15\eta \times 20\phi$ crystal region.

One of the main tasks of the IP1 algorithm is to build the e/γ clusters around the peak crystal, which has the maximum energy deposited by electromagnetic particles such as electrons and photons. In IP1, this process starts with finding the peak crystal in the $15\eta \times 20\phi$ crystal region. The peak finding in this 2-dimensional region is executed in the following two steps.

The $15\eta \times 20\phi$ region is considered 15 strips in *eta*, each consisting of 20 crystals. The process starts with finding the peak crystal in each *eta* strip. Since the strips are mutually independent, finding the peak crystal in each strip is parallel and reduces the overall system latency. A multi-level comparison strategy is utilized to find the peak crystal in the strips. The adjacent crystals are compared, and the superior candidate passes to the next comparison level. Finding the peak in one *eta* bin with 20 crystals requires five levels of comparison. The first element is compared with the second, and the third is compared with the fourth, and so on. The number of comparisons at each subsequent level reduces, and the final level provides the peak crystal. The peak crystal of each of the fifteen eta strips is stored in a one-dimensional array. The final computation now requires finding the peak in the single *phi* strip using the similar strategy. **The latency of finding the peak in the *phi* strip of 20 crystals is**

four clock cycles. Fig 5.15 shows the peak finding algorithm using the multi-level comparison strategy.

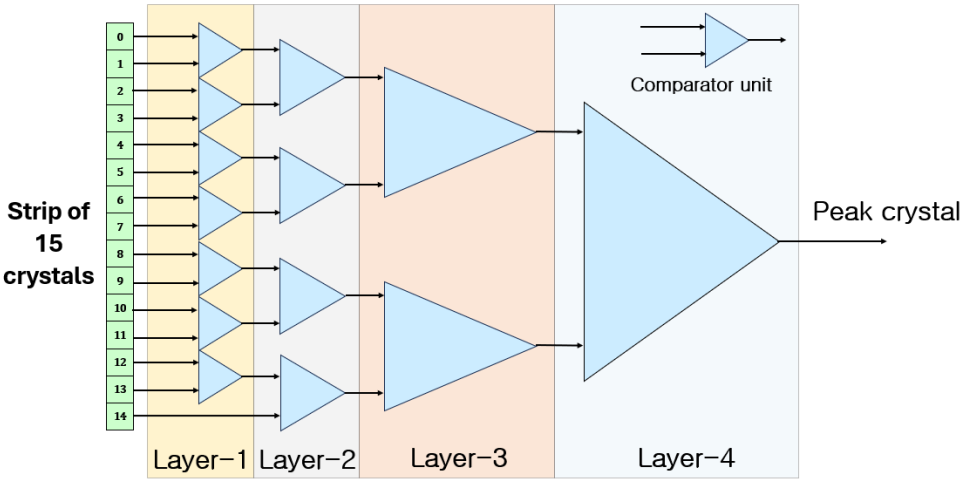

FIGURE 5.15: IP1 peak finding using multi-level comparison strategy.

The peak crystal location in *eta* and *phi* is used to compute the energy of the clusters. The clustering region depends on the peak crystal position in the $15\eta \times 20\phi$ crystal region. For example, the peak crystal position on the first and the last *phi* axis (except the corner) provides the clustering region of $3\eta \times 3\phi$ crystals. Similarly, the peak crystal position on the first and the last *eta* axis (except the corner) provides the clustering region of $2\eta \times 5\phi$ crystals. The crystal at *eta* and *phi* corners corresponds to the clustering region $2\eta \times 3\phi$ crystals. As the primary clustering region is fixed to $3\eta \times 5\phi$ crystals, the various clustering regions are avoided in the algorithm for FPGA implementation by extending the region of crystals in *eta* and *phi* in directions using the null crystals (hypothetical crystal with zero energy deposition). Fig 5.16 shows the clustering scheme for various positions of peak crystal in $15\eta \times 20\phi$ region and the extended region using the null crystals to avoid the multi-dimension clusters.

In IP1, the e/γ clustering starts with computing the primary cluster energy by adding the energy of the crystals in the $3\eta \times 5\phi$ region centered around the peak crystal. It prepares the isolation cluster by adding the energy of the crystals in the $5\eta \times 5\phi$ region centered around the peak crystal. It also computes the two shape clusters by adding the energy of the crystals in the $2\eta \times 5\phi$ region centered around the peak crystal. An identification (shape) aspect is considered to enhance the electrons/photons tagging, which diminishes the potential background contributions [26]. It is established by computing shape clusters within the primary cluster. The final shape energy is considered as the largest of the two shape clusters.

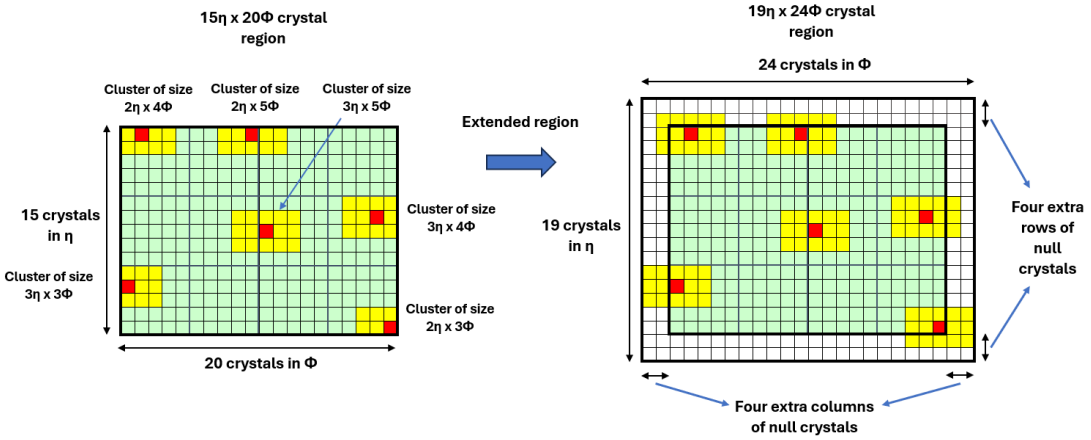

FIGURE 5.16: Various clustering regions in IP1 based on the peak crystal position (left) in the $15\eta \times 20\phi$ crystal geometry. The extension of the crystal geometry from $15\eta \times 20\phi$ to $19\eta \times 24\phi$ by adding null crystals in *eta* and *phi* reduces all the cluster dimensions to a single $3\eta \times 5\phi$ crystal region (right).

In the $15\eta \times 20\phi$ crystal region, the central, isolation, and shape clusters energies are determined by first computing the energy of the five *eta* strips ($1\eta \times 5\phi$ crystals), considering the peak crystal in the central strip. The sum of all the five strips provides the cluster energy of size $5\eta \times 5\phi$, the sum of the three central strips provides the primary cluster of size $3\eta \times 5\phi$, and the addition of the two strips (second, third or third, fourth) out of the three central strips provides the potential shape cluster. **The latency to create the extended crystal region, creating the central, shape, and isolation cluster around the peak crystal, is six clock cycles**. Fig 5.17 shows the central clustering scheme of the RCT v2.0. The next step is building the brems clusters.

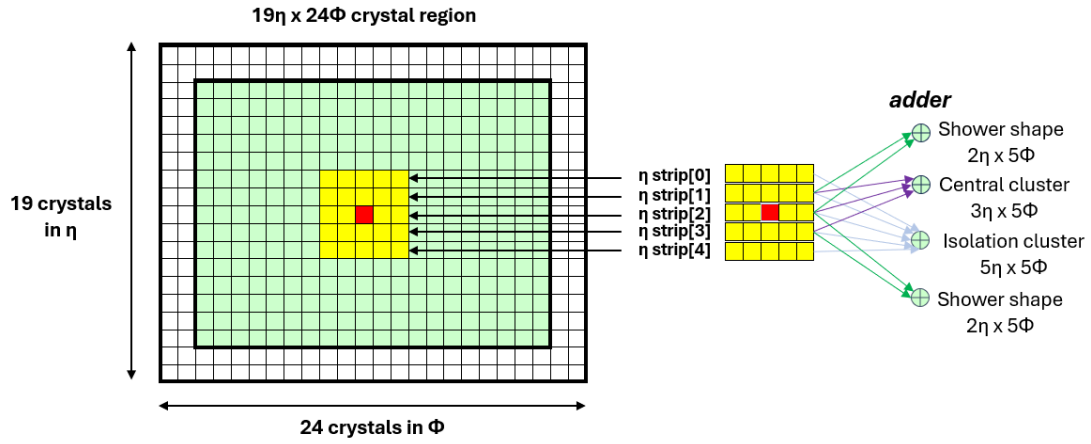

FIGURE 5.17: Computation of the central, shape, and isolation cluster in IP1 in the $19\eta \times 24\phi$ crystal geometry.

An electron might yield considerable energy to bremsstrahlung radiation [39]. It causes the electron footprint in ECAL to be spread in the *phi* direction, which is accounted for in IP1 by computing cluster energy in positive and negative *phi* directions around

the primary cluster. IP creates two arrays to compute the brems energy independently on the positive and negative *phi* sides of the central cluster. The array size for brems calculation is extruded in *eta* and *phi* direction with null crystals to avoid computing brems clusters of different sizes. **The latency of computing the brems clusters is six clock cycles**. Fig 5.18 illustrates the building of a negative brems cluster.

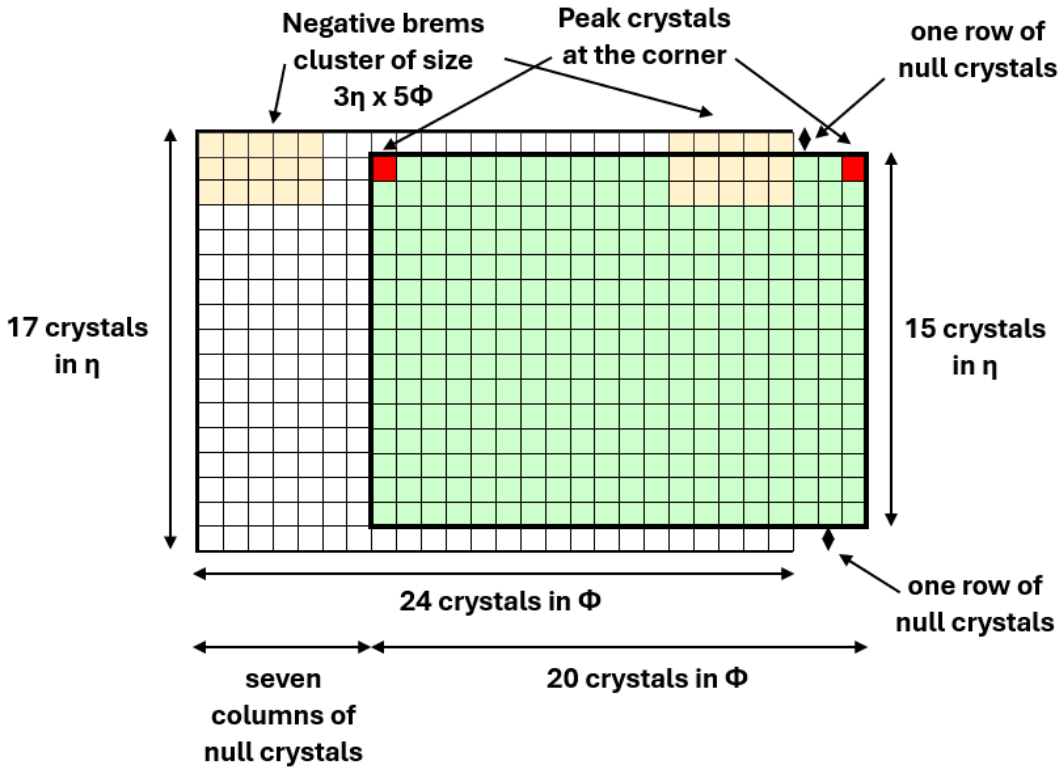

FIGURE 5.18: IP1 negative brem cluster computation with respect to the central cluster. The figure highlights the mapping of the main crystal region of $15\eta \times 20\phi$ to the extruded region of $17\eta \times 24\phi$. The corner cases of making brem clusters are shown with the two peak crystals positioned at the two corners.

The positive and negative brems cluster energy is compared with the central cluster energy. If the negative brems energy is greater than one-eight of the central cluster [26], it is merged with the central cluster energy, and the brem flag is raised to 1. If this condition is false and the positive brem energy is higher than one-eighth of the central cluster, it is merged with the central cluster energy, and the brem flag is raised to 2. Calculating the central, shape, and brems cluster energy completes the clustering scheme of the IP1 algorithm. The clusters are tagged with the location information of the peak in the *eta* and *phi* positions at the tower and crystal level. The tower level determines the peak crystal position within the RCT $17\eta \times 4\phi$ tower geometry, and the crystal level determines the position of the crystal within a single tower. Fig 5.19 shows the peak crystal position within the $3\eta \times 4\phi$ tower geometry with central, shape, and brems clusters.

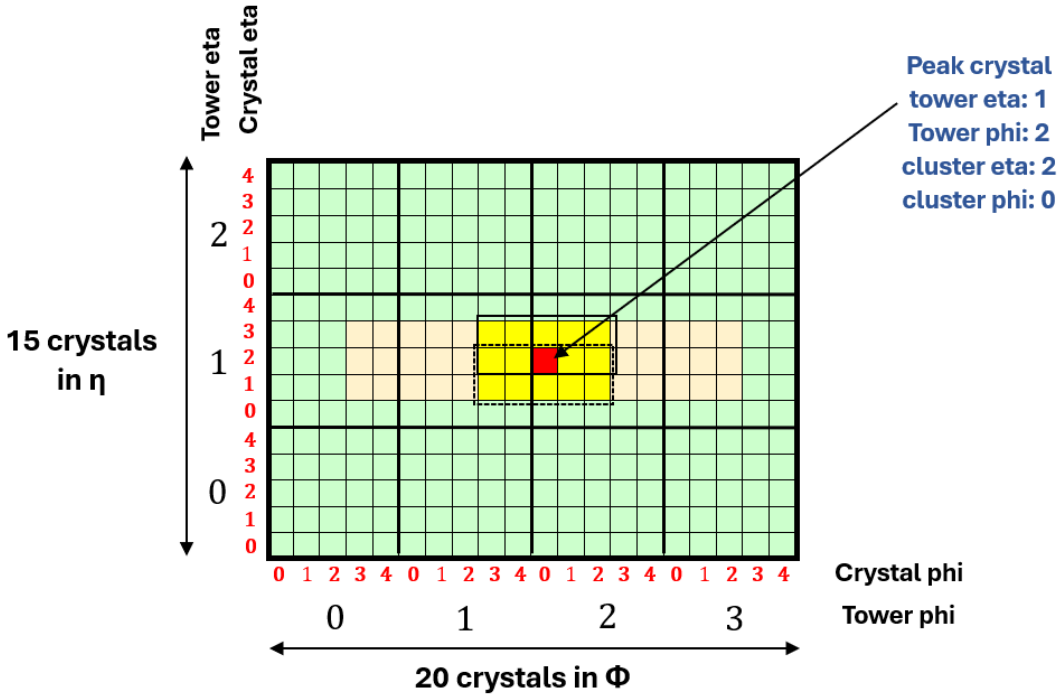

FIGURE 5.19: IP1 clustering scheme depicting the central (yellow), shape (covered in dashed and solid black lines around the peak crystal), and brems clusters. The peak crystal (red) position is encoded in terms of tower *eta*, tower *phi*, crystal *eta*, and crystal *phi*.

Once all clusters are built in the current iteration, the next step is removing all the crystals used to make the central and brem clusters (positive or negative brem cluster). This step is necessary so that the algorithm finds another peak crystal in the next iteration of making clusters and computes the cluster energy around it. **The latency of removing the crystals in the $15\eta \times 20\eta$ crystal region is three clock cycles**. The IP1 algorithm performs the clustering for four iterations for each of the three $15\eta \times 20\eta$ crystal regions and computes twelve clusters (four clusters per $15\eta \times 20\eta$ crystal regions). **The latency of making the cluster information in the IP1 algorithm requires 31 clock cycles in all three regions**.

After clustering, IP1 computes the energy of each tower in the $3\eta \times 4\phi$ tower region. The tower energy is the sum of all the crystals in the $5\eta \times 5\phi$ crystal region. There are 12 towers in a $3\eta \times 4\phi$ tower region; each tower is mutually independent. Therefore, the energy calculation is parallel for each tower and subsequently for the three $3\eta \times 4\phi$ tower regions. The ECAL tower can be represented as a group of five *eta* strips, each containing five crystals. The tower energy is calculated by first computing the energy of each strip and then adding all the five strip energy. It requires six clock cycles of latency to compute the tower energy of one $3\eta \times 4\phi$ region consisting of twelve towers. Fig 5.20 shows the tower region and the algorithmic calculation of the tower energy.

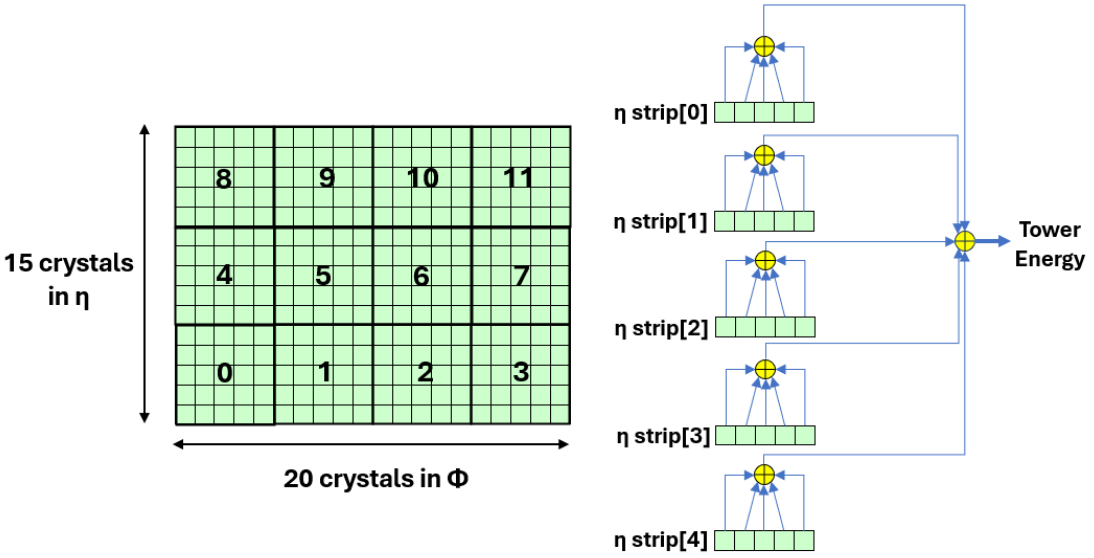

FIGURE 5.20: IP1 tower energy calculation. The figure highlights the twelve towers in the region of $15\eta \times 20\phi$ crystals and demonstrates the tower energy calculation for one of the towers (11).

The IP1 packs the cluster and tower information in 60 bits and 16 bits of data words, respectively. The 60 bits of cluster information consist of 12 bits of central cluster energy, 15 bits each for the isolation, and the shape cluster energy. It also preserves the positional information of the peak crystal in terms of its location in the tower and crystal. For tower position, the algorithm assigns 5 bits for eta and 2 bits for phi. Similarly, for the crystal position, the algorithm assigns 3 bits for each eta and phi. It uses 2 bits for the brem flag and 3 bits for the saturation, which has yet to be implemented. The tower information at the RCT consists of the 12-bit tower energy and the ratio of the HCAL and ECAL tower energies coined as H over E or H/E, packed in 4-bit. The H/E calculation is performed in IP3. Table 5.6 and 5.7 summarized the cluster and tower information packing in the RCT v2.0 algorithm.

The IP1 prepares three output links for IP2, each containing cluster and tower information of one $3\eta \times 4\phi$ tower region. Each link contains information about 4 clusters and 12 towers. **The latency of the IP1 algorithm is 152 clock cycles**. Fig 5.21 represents the HLS timing result of the IP1 algorithm.

The RTL of IP1 is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.8 summarizes the timing and utilization performance. The following subsection describes the IP2 algorithm.

Quantity	Number of bit(s)
central cluster energy	12
tower eta	5
tower phi	3
crystal eta	2
crystal phi	3
saturation	3
isolation cluster energy	15
shape cluster energy	15
brems flag	2
	Total: 60

TABLE 5.6: Output e/γ cluster information packing in 60 bits word.

Quantity	Number of bit(s)
tower energy	12
H/E	4
	Total: 16

TABLE 5.7: Output tower information packing in 16 bits word.

Performance Estimates

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	2.180 ns	1.39 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
152	152	0.422 us	0.422 us	9	9	function

  - **Detail**

FIGURE 5.21: IP1 HLS synthesis report.

	IP1
Latency (ns)	422
F_{MAX} (MHz)	364
LUTs (%)	8
FFs (%)	5

TABLE 5.8: Timing and utilization summary of the IP1 algorithm.

5.3.2.2 IP2

The IP2 receives the cluster and tower information from IP1, and the main task is to stitch the clusters across the two internal eta boundaries in the tower region of $9\eta \times 4\phi$ and $8\eta \times 4\phi$. Fig 5.22 demonstrates the stitching logic of IP2 in the tower region of $9\eta \times 4\phi$ highlighting the two *eta* boundaries. The stitching cases used in IP2 are mentioned below

1. The two clusters at the *eta* boundary are neighbors in *eta*. Therefore,
   - the difference between the *eta* position of the two peak crystals must be one, and
   - Both the peak crystals must have identical *phi*.
2. The two clusters at the *eta* boundary are neighbors in *eta*. Therefore,
   - the difference between the *eta* position of the two peak crystals must be one, and
   - The difference between the *phi* position of the two peak crystals must be one.

The central $3\eta \times 4\phi$ tower region shares the eta boundary with the two outer $3\eta \times 4\phi$ regions, and the clusters in all three regions are built arbitrarily in position and are independent of each other. Therefore, the algorithm concurrently checks the stitching condition at the two eta boundaries. Ultimately, it reduces the latency of the algorithm. **The latency of two clusters at the eta boundary is one clock cycle**.

The IP2 of both $9\eta \times 4\phi$ and $8\eta \times 4\phi$ tower regions receive 12 clusters from the corresponding IP1. However, the IP2 sends only six clusters to the IP3. Therefore, the twelve clusters are sorted based on their energy using the *bitonic sort* algorithm [41]. *Bitonic sort* is a parallel sorting network that compares and sorts the elements in a pre-defined sequence. The basic unit of the *bitonic sort* algorithm is a compare and exchange unit (CAE), which compares two elements. If the inputs are not identical, they are sorted in increasing or decreasing order. A sorting grid is used, which incorporates a series of sorting stages, with each stage containing several CAE blocks that operate in a parallel manner. The latency of a sorting grid is proportional to the level of CAE blocks used. Fig 5.23 shows the high-level implementation of the CAE block. The CAE block, which orders the element in increasing order, is denoted by $\oplus$, and the CAE block ordering in decreasing order is denoted by $\ominus$. Each CAE block is made up of a comparator and two 2:1 multiplexers.

Concatenating an ascending and a descending sequence (or vice versa) creates a single bitonic series. A bitonic sorting grid recursively combines ascending and descending

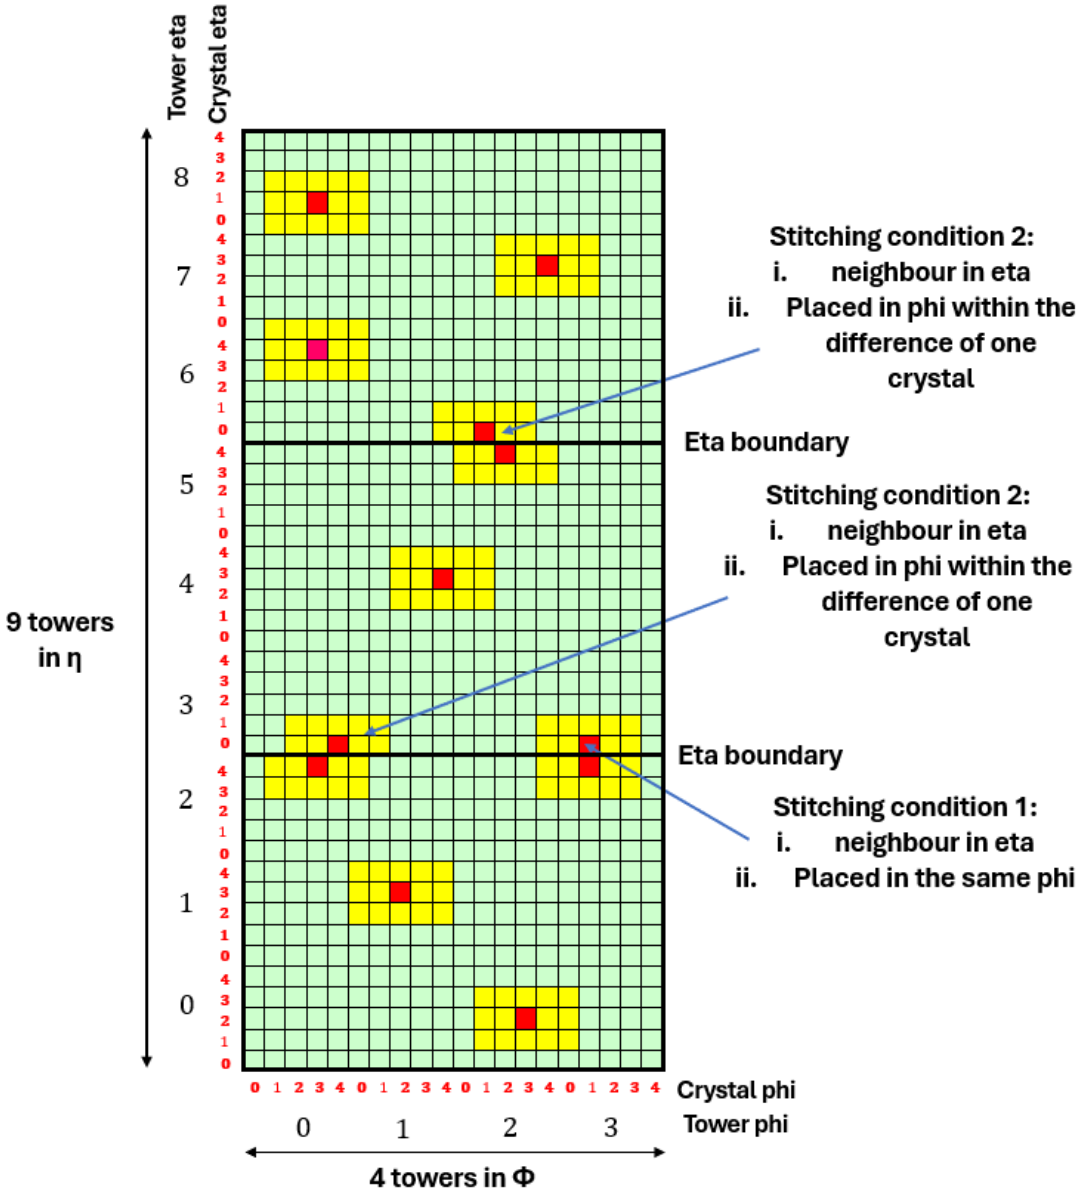

FIGURE 5.22: Stitching of clusters in IP2 at the two eta boundaries.

sequences, each of length N/2, to create a sorted sequence of length N. Each sorting grid comprises several bitonic merging units to integrate bitonic sequences. A M-input bitonic merging unit (designated as BM-M) comprises log_2(M) levels of parallel CAE blocks, where each level comprises M/2 CAE blocks. Thus, a BM-M requires $log_2(M) \times M/2$ CAE blocks. For example, an increasing BM-8 (designated as ⊕BM-8) unit includes three CAE stages and requires $3 \times (8/2) = 12$ CAE blocks. A BM-M unit comprises K/2 parallel CAE blocks and two parallel BM-(K/2) units. The ⊕BM-8 is built from a level of four parallel CAE blocks followed by two parallel BM-4 units incorporating four CAE blocks each. A 16-input *bitonic* sorting unit has eight parallel BM-2 units, four parallel BM-4 units, two parallel BM-8 units, and one BM-16 unit. Therefore, this

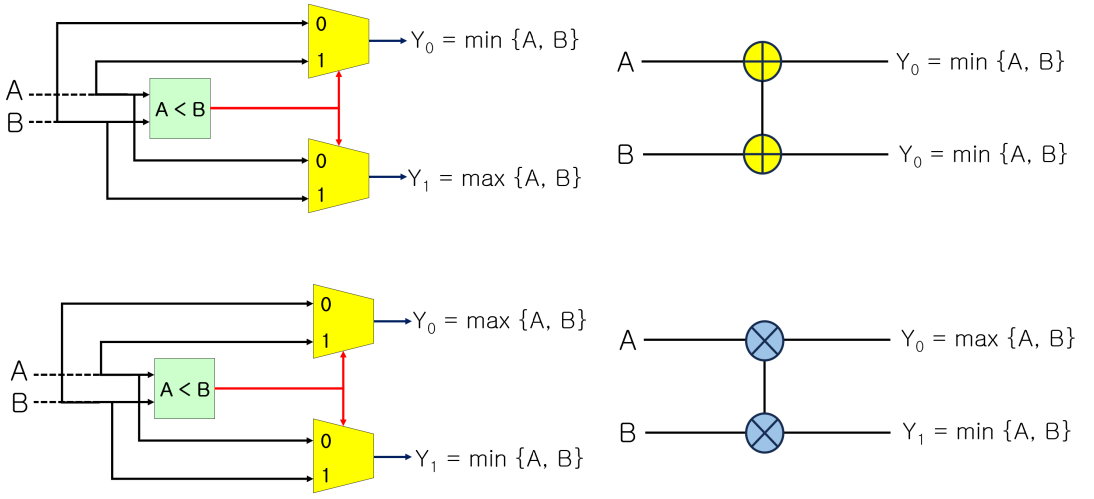

FIGURE 5.23: CAE block high-level implementation with the increasing (up) and decreasing (down) order depicted with the corresponding symbols and equation.

sorting network is comprised of 1+2+3+4 = ten CAE levels. Considering the design is pipelined such that the individual stage takes one clock cycle, it can render the sorted results in 10 cycles and initiate a new sort in each clock cycle. Fig 5.24 shows the 16 input *bitonic sort* grid comprising ten levels of CAE block.

Bitonic sort can only be implemented if the number of inputs is in the order of 2^n where n is a natural number. The IP2 sorts twelve clusters using 16 input bitonic sort algorithm. The sorting algorithm is developed in the HLS tool and implemented in RTL using the standard clock frequency of 360 MHz. **The HLS result shows the latency of 9 clock cycles for the XCVU9P FPGA**. Fig 5.25 shows the HLS synthesis result of the bitonic sort algorithm.

After sorting, the six highest clusters are selected for the IP3. The remaining six clusters must be added back to the towers based on the position of the peak crystal in the existing tower. This operation is crucial as it prepares the tower information for clustering performed in the next level GCT barrel algorithms. As the three $3\eta \times 4\phi$ tower regions are independent, the cluster addition with the tower performs in parallel for all three regions. Fig. 5.26 shows the example of the selected six highest clusters for the IP3 and the addition of one of the un-selected clusters with the tower.

Finally, the IP2 packs the six highest cluster and tower energy in two output links. Each link carries 18 towers and 3 clusters of information, corresponding to 468 out of 576-bit available. The IP2 algorithm is synthesized in HLS. The latency is 62 clock cycles, including the clusters' stitching at two eta boundaries, bitonic sort, and merging the un-selected clusters with the towers. Fig 5.27 shows the HLS latency report.

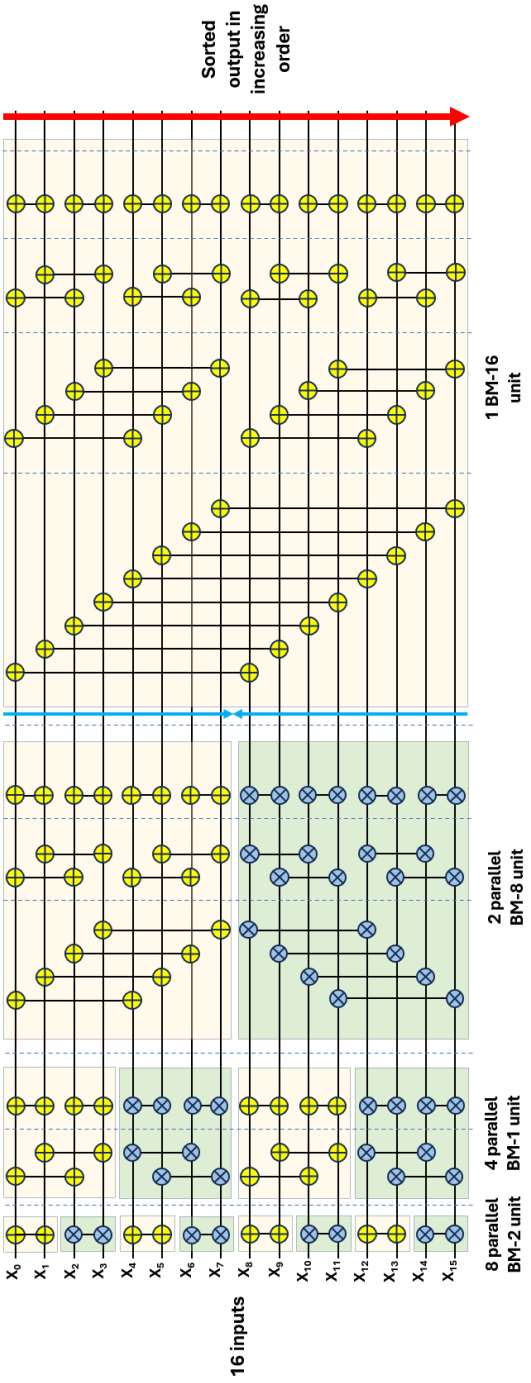

FIGURE 5.24: 16-input bitonic sort grid comprising 10 levels of CAE block. It sorts the input in increasing order.

The RTL of IP2 is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.9 summarizes the timing and utilization performance.

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	1.588 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
9	9	25.002 ns	25.002 ns	9	9	function

FIGURE 5.25: 16-input bitonic sort HLS synthesis report.

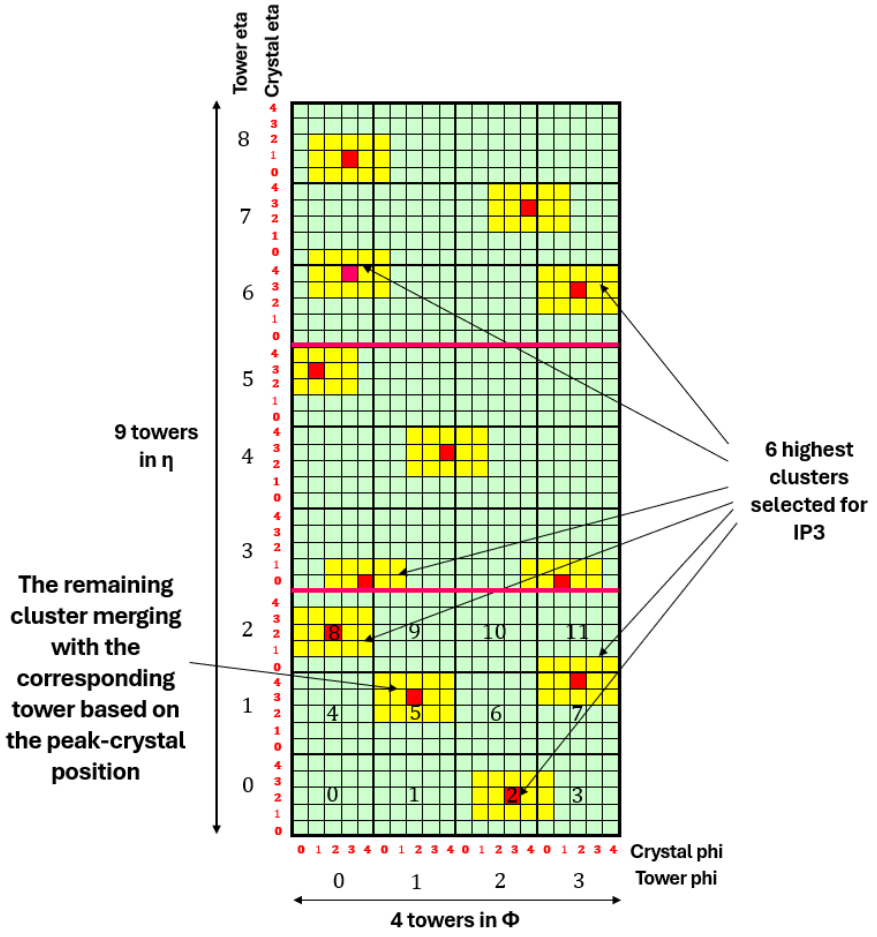

FIGURE 5.26: RCT region of $9\eta \times 4\phi$ highlighting the six highest clusters (after bitonic sort) candidates for the IP3 and the adding back of one of the un-selected clusters with the tower. The highlighted un-selected cluster's peak crystal is inside tower number 5 of the region. Therefore, the cluster energy will be combined with the corresponding tower energy.

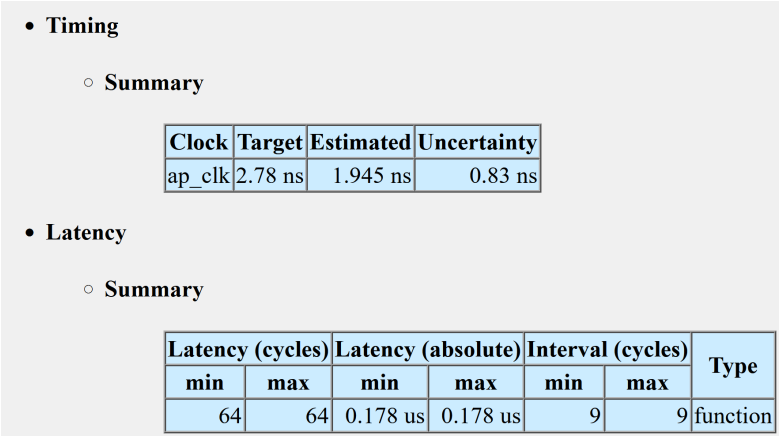

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	1.945 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
64	64	0.178 us	0.178 us	9	9	function

FIGURE 5.27: The IP2 HLS timing reports

	IP2
Latency (ns)	178
F_{MAX} (MHz)	368
LUTs (%)	1
FFs (%)	0.8

TABLE 5.9: Timing and utilization summary of the IP2 algorithm.

5.3.2.3 IP3

IP3 is the final processing stage of the RCT v2.0 algorithm. It takes input from the two IP2 algorithms and the HCAL. The algorithm starts by unpacking the IP2 data and mapping the information for 68 ECAL towers, 12 e/γ clusters, and 64 HCAL towers. The first step of IP3 is to stitch the cluster at the central *eta* boundary. The stitching logic follows the same strategy explained in the case of *eta* stitching in IP2. Fig 5.28 demonstrates the stitching in IP3.

After stitching, IP3 uses the *bitonic sort* algorithm to sort the 12 e/γ clusters based on their energy and select the 8 highest clusters for the GCT barrel board. It uses the 16-input bitonic sort algorithm with 12 inputs from e/γ clusters received from the two IP2, and the remaining four inputs have zero cluster energy. The sorting strategy is similar, as explained in the case of IP2. The eight highest clusters are selected, and the un-selected clusters are added back with the corresponding ECAL towers. The next step of the IP3 algorithm is to add each HCAL tower energy with the corresponding ECAL using a two-input adder. Fig 5.29 shows the addition of ECAL and HCAL towers. The next level GCT barrel algorithm uses the combined tower energy of the ECAL and HCAL to perform PF, jets, and taus clustering.

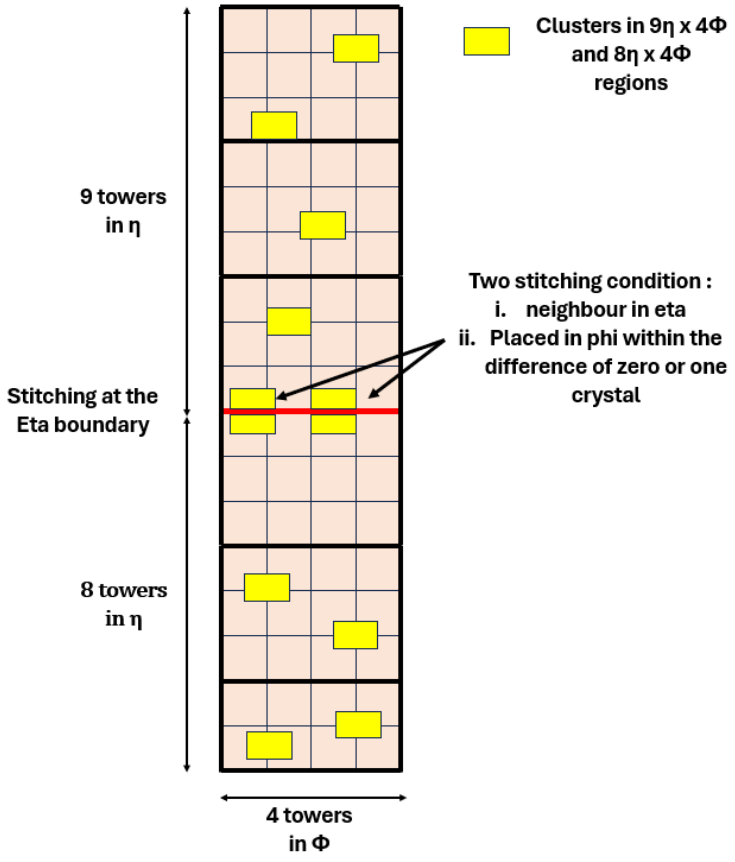

FIGURE 5.28: Stitching of clusters in IP3 at the central eta boundary.

The final step of the IP3 is to perform the H/E calculation and store the result in a 4 bits word. The LBS bit indicates the possibility that either the ECAL or HCAL energies are zero or HCAL energy is greater than or equal to the corresponding ECAL tower energy. If the above condition is valid, the algorithm sets the LSB bit to zero. If the ECAL tower energy exceeds the HCAL energy, the algorithm sets the LSB bit to one. The next three H/E bits represent the range over which the ECAL or HCAL energies exceed each other. Table 5.10 summarizes the range of the differences between the HCAL and ECAL energies and the corresponding H/E bits. The H/E bits cover a wide range of differences between the HCAL and ECAL energies.

The final clusters and towers information of IP3 are packed in four output links, each carrying 17 towers and two clusters. **The latency of the IP3 algorithm is 125 clock cycles**. It includes the **latency of stitching the clusters, which is 24 clock cycles**. The post-stitched clusters are sorted using the bitonic sort algorithm, which utilizes 9 clock cycles of latency. **The process of merging four un-selected clusters with any of the 68 ECAL towers is 72 clock cycles**. **The final addition of the ECAL and HCAL towers, including the H/E calculation, requires 8 clock cycles**. Fig 5.30 depicts the HLS latency report of IP3.

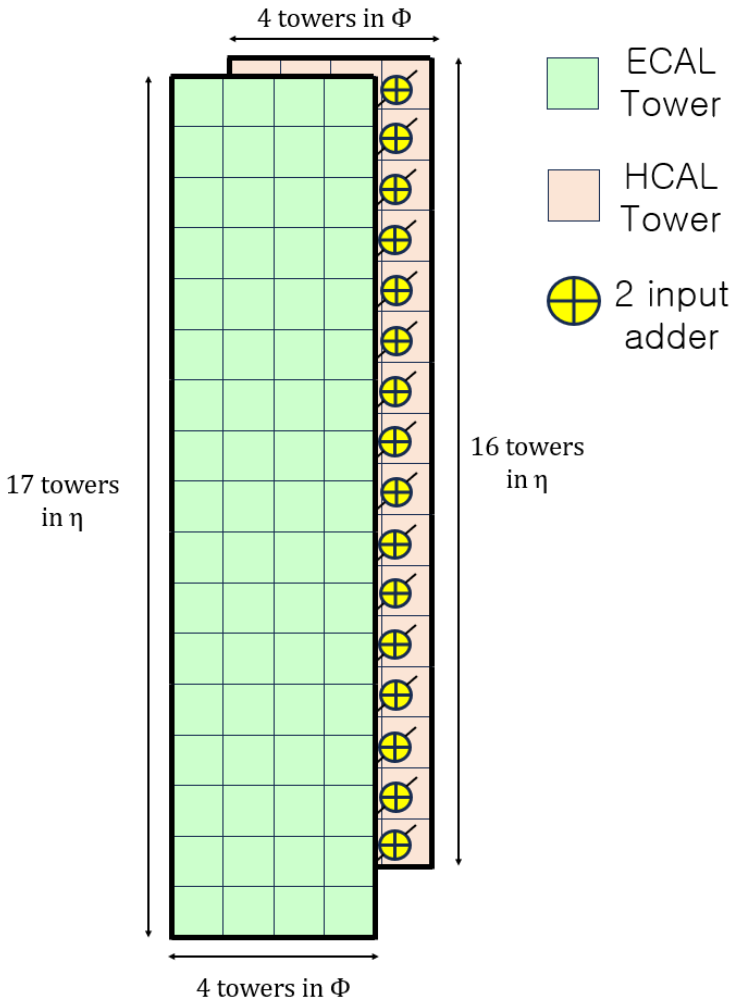

FIGURE 5.29: ECAL and HCAL towers addition.

Performance Estimates

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	1.945 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
125	125	0.347 us	0.347 us	9	9	function

FIGURE 5.30: IP3 HLS timing report highlighting the latency of 125 clock cycle.

The RTL of IP3 is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.11 summarizes the timing and utilization performance.

The IP1, IP2, and IP3 are combined in the RTL using a VHDL wrapper. The wrapper takes nine clock cycles to convert the 64-bit physics payload AXI stream data into the

Range coverage	H/E bits (3 downto 1)
HCAL >ECAL	000
ECAL >HCAL	000
HCAL $> 2 \times ECAL$	001
ECAL $> 2 \times HCAL$	001
HCAL $> 4 \times ECAL$	010
ECAL $> 4 \times HCAL$	010
HCAL $> 8 \times ECAL$	011
ECAL $> 8 \times HCAL$	011
HCAL $> 16 \times ECAL$	100
ECAL $> 16 \times HCAL$	100
HCAL $> 32 \times ECAL$	101
ECAL $> 32 \times HCAL$	101
HCAL $> 64 \times ECAL$	110
ECAL $> 64 \times HCAL$	110
HCAL $> 128 \times ECAL$	111
ECAL $> 128 \times HCAL$	111

TABLE 5.10: H/E 3-bit (from MBS to the LSB+1 bit) calculation based on the difference of energy between HCAL(ECAL) and ECAL(HCAL).

	IP3
Latency (ns)	347
F_{MAX} (MHz)	364
LUTs (%)	1.6
FFs (%)	1.3

TABLE 5.11: Timing and utilization summary of the IP3 algorithm.

576-bit packet information of the LHC bunch-crossing for the algorithm. The latency of the RCT v2.0 algorithm is 1219 ns (1.2 μs). Table 5.12 summarized the algorithm steps and latency of the RCT v2.0 algorithm.

The next section describes the global calorimeter trigger (GCT) algorithms.

Input to RCT:	
ECAL: a region of 85×20 crystals (1700 crystals) **HCAL**: a region of 16×4 towers (64 towers) **Input links**: 70 **Clock frequency**: 360 MHz	
Entire ECAL region for RCT is sub-divided in six $15\eta \times 20\phi$ crystal regions with IP1 serving three such regions.	
IP1:	
Creating three matrix of 15×20 crystals from 36 input links	
Processing the three 15×20 crystals region concurrently	
❶ Finding peak crystal in $15\eta \times 20\phi$ crystal region	**8 clock cycle**
❷ Computing primary cluster of size $3\eta \times 5\phi$, isolation cluster of size $3\eta \times 5\phi$, and shape cluster of size $2\eta \times 5\phi$	**6 clock cycles**
❸ Computing two $3\eta \times 5\phi$ brems clusters	**6 clock cycles**
❹ Removing the used crystals from the region.	**3 clock cycles**
Iterating the steps ❶, ❷, ❸, and ❹ four times to create four clusters.	
Performing twelve towers energy calculation	**6 clock cycles**
IP2:	
Stitching two clusters at the eta boundary	**4 clock cycles**
Sorting clusters using 16-inputs bitonic sorting	**9 clock cycles**
Adding back the unused clusters with the tower	**39 clock cycles**
IP3:	
Stitching clusters at the eta boundary	**23 clock cycles**
Adding back the unused clusters with the tower	**72 clock cycles**
HCAL and ECAL tower energy addition and H/E calculation	**8 clock cycles**
Total latency*	**439 clock cycles**
Output:	
Tower: 68 barrel calorimeter tower information **Clusters**: 8 clusters information **Output links**: 4	

TABLE 5.12: RCT algorithm steps and latency report. * Total latency includes the latency of the wrapper, which combines the two instances of IP1, two instances of IP2, and one IP3 instance of algorithms. It also includes the SLR crossing of the internal signals.

5.4 Layer-2: Global Calorimeter Trigger (GCT)

The global calorimeter trigger (GCT) is the second level of processing for the Phase-2 calorimeter trigger. It is organized in two layers and comprises the GCT barrel, endcap, and sum boards. Fig. 5.31 represents the two-layer connection of the GCT. The following section briefly explains the algorithm employed at the GCT barrel, endcap (HGCAL), and GCT SUM.

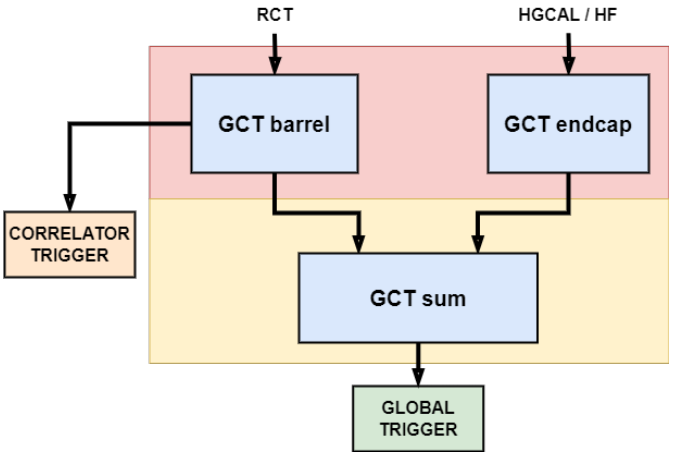

FIGURE 5.31: Two layer organization of GCT comprising GCT barrel, endcap and sum algorithm.

5.4.1 GCT Barrel

The GCT barrel processes the RCT information and prepares the barrel calorimeter information for the GCT sum and correlator trigger. The processing of inputs from 36 RCT boards is subdivided into three GCT barrel boards, each serving the input from 12 unique RCT boards. The three GCT barrel boards require stitching the RCT electron/photon (e/γ) clusters at the partition boundary in *phi* before sending the cluster information to the correlator trigger. This requires the cluster information from the RCT regions, which are located on either side of the two partition boundaries of the GCT barrel region. Thus, each GCT barrel board takes additional inputs from four RCT regions that belong to the processing region of the other two boards. Therefore, each GCT barrel board takes input from 16 RCT boards. The additional RCT regions provide the tower information required for building PF clusters, jets, and taus information. Fig 5.32 represents the CMS barrel geometry, the coverage of one RCT board, and the GCT barrel board. It highlights the three central GCT barrel regions and the additional inputs from 4 RCT regions. Thus, each GCT barrel board takes input from 16 RCT boards.

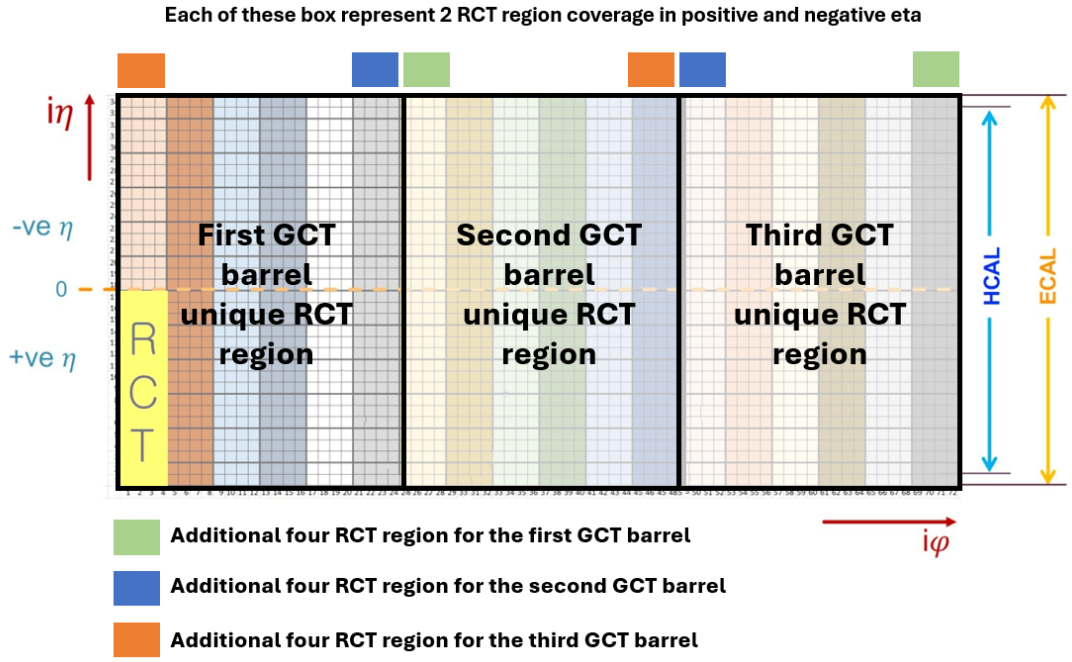

FIGURE 5.32: Calorimeter barrel region segmentation in three GCT barrel board. Each board comprises 12 unique RCT regions and 4 additional RCT regions from either side of *phi*.

One GCT barrel board takes 64 input links from 16 RCT boards (4 links per RCT board). Therefore, the design needs to be partitioned into two SLRs to satisfy the available link constraint of one SLR (40 links). The partition of the GCT barrel region is performed across the central *phi* boundary, which requires processing 8 RCT board inputs in one SLR. As mentioned earlier, partitioning the barrel region requires sharing the boundary information to account for stitching the RCT clusters and performing hadronic clustering at GCT. Therefore, the algorithm in each SLR requires information of two additional RCT regions located at opposite sides of the partition boundary. Thus, the algorithm in each SLR takes input from 10 RCT boards. Fig 5.33 shows the partition of the GCT barrel geometry in two SLRs, proposed to be implemented in SLR1 and SLR2, respectively.

The algorithms employed in both the SLRs are identical. It receives the input from 10 RCT regions, stitches the e/γ clusters in *eta* and *phi* direction, and creates full tower information. The full tower information is required to perform clustering at the GCT level, such as PF clustering, jets, taus, and computing of the partial MET (the remaining portion of the MET algorithm is implemented in the GCT sum board). The stitching of RCT clusters and making full towers are performed in IP1, and PF clustering is performed in IP2. The outputs of IP1 and IP2 are merged in RTL to create 24 output links per SLR and further time-multiplexed in six-time slices (TMUX6) for the correlator trigger. One GCT barrel board provides 48 output links to the correlator trigger. The

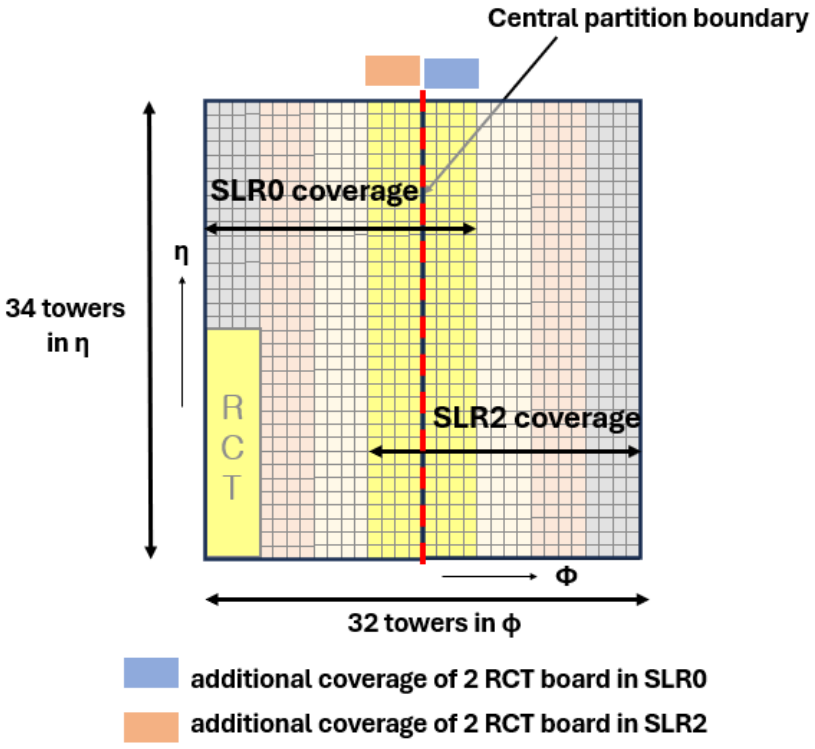

FIGURE 5.33: Partition of one GCT barrel region in two SLR. Each SLR takes input from ten RCT regions.

IP1 in SLR2 and SLR0 sends the full tower information in SLR1, which acts as an input to the jet, taus, and partial MET algorithms, also known as the CaloObject algorithm. The CaloObject algorithm packs six output links for the layer-3 GCT sum board. Fig 5.34 illustrates the architecture of the GCT barrel algorithm.

The processing strategy for each algorithm employed in GCT Barrel is as follows.

- **IP1**:
  - Stitching the e/γ clusters between RCT regions in *eta* and *phi* direction.
  - Computes the cluster information as per the correlator trigger requirements.
  - Create the full tower information by adding back the cluster energy into the corresponding tower. Making full towers in the GCT barrel is crucial because the tower information received from the RCT boards has missing crystal energies, which were used in making the e/γ clusters. Therefore, adding bavk those crystal energies into the corresponding tower is necessary before performing any tower-level clustering in the GCT barrel.
  - Create 24 output links containing e/γ clusters and towers information of the unique RCT region for the correlator trigger.

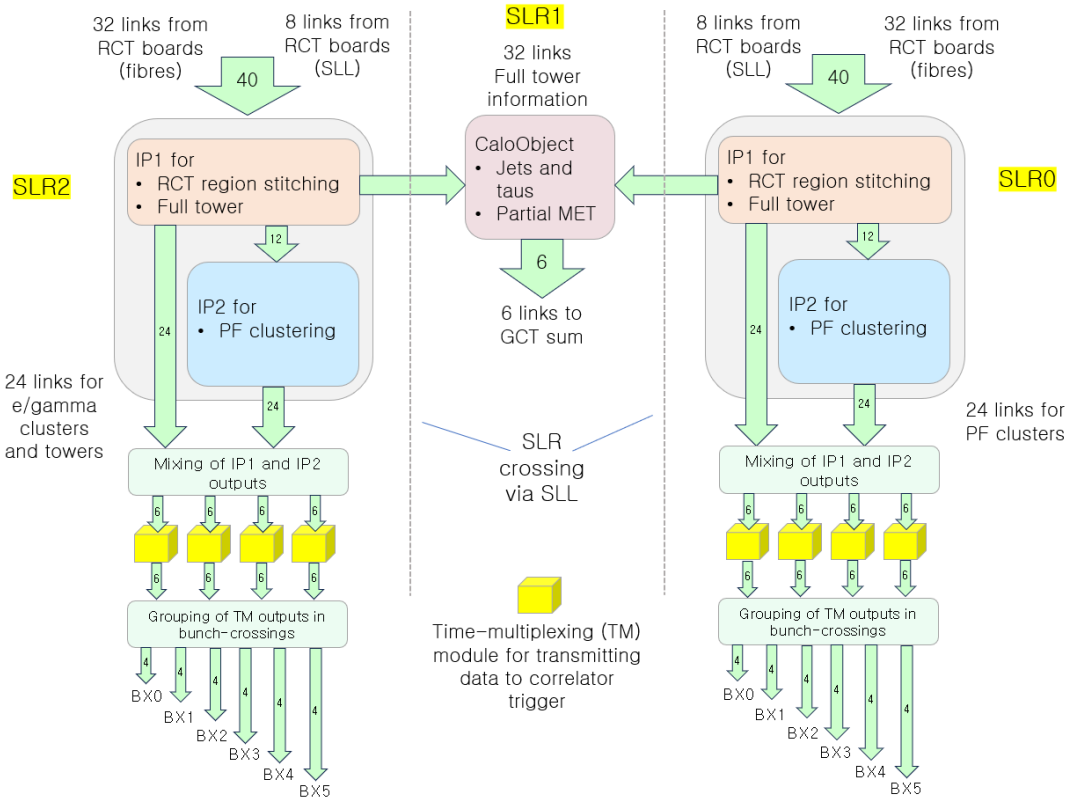

FIGURE 5.34: GCT barrel architecture.

- Prepare 16 output links containing full tower information for the CaloObject algorithm.

- **IP2 (PF clustering)**:
  - It receives the full tower information from the IP1. The main task of the IP2 is to compute 8 PF clusters in each of the central six RCT region by adding the tower energy in $3\eta \times 3\phi$ region around the peak tower. Following are the steps to compute the PF cluster in one RCT region.
  - Calculate the peak tower in the RCT region using the multi-level comparison strategy as discussed in Chapter 5.
  - Compute the PF cluster in $3\eta \times 3\phi$ towers region around the peak tower.
  - Remove the peak tower energy from the region for computing the PF cluster in the next iteration.
  - Iterate eight times to compute eight PF clusters in each RCT region.
  - Pack the PF cluster information in terms of cluster energy, and peak location in *eta* and *phi* in 24 output links.

- **Time-multiplexing**:

- Convert the TM1 information of the IP1 and IP2 algorithm in six-time slices for the correlator trigger.

- **CaloObject**:
  - With the input tower geometry of $34\eta \times 32\phi$, it prepares two (in both the *eta* region) supertower regions of $8\eta \times 12\phi$. Supertower energy is computed by adding the tower energy inside the $3\eta \times 3\phi$ tower region.
  - Find the peak supertower in the supertower region using the multi-level comparison strategy.
  - Compute the jet and tau cluster in $3\eta \times 3\phi$ region around the peak supertower.
  - Remove the peak and corresponding supertower used in the current iteration to compute the jet and taus for the next iteration.
  - If the energy of the peak supertower is higher than 87% of the cluster energy, the cluster is flagged as a tau.
  - The partial MET is computed by adding the tower energy of all the towers in each of the central *phi* columns.

5.4.1.1 IP1 Algorithm

The IP1 algorithm processes the inputs from ten RCT boards. IP1 takes the input from RCT boards in terms of e/γ clusters and towers. The packing of the cluster and tower summarized in Table 5.6 and 5.7, respectively. The task of the IP1 is to stitch the cluster energy at the boundaries of the ten RCT regions. The stitching is performed in three directions, as shown in Fig. 5.35. The stitching algorithm and logic used in IP1 is identical to the one used in the RCT v2.0. **The latency of stitching the RCT regions in each direction is three clock cycles**.

IP1 repacks the e/γ cluster information and makes it compatible with the correlator trigger requirements by adding the cluster position in terms of *eta* and *phi* location of the peak crystal. The cluster information is packed inside a 64 bits word, which includes 12 bits for cluster energy, 8 bits for clusters eta position, 7 bits for phi position, 4 bits for H/E, 2 bits for the H/E flag, 3 bits for isolation, 2 bits for the isolation flag, 6 HCAL feature bits, 5 bits for timing, 2 bits for the shape flag, 2 bits for brems and 11 bits are reserved for future usage. Table 5.13 summarizes the packing of the e/γ cluster for the correlator trigger.

The clusters and towers information of the six unique RCT regions are packed in 24 output links. Each link carries 17 towers (17×16 bits = 272 bits) and 2 e/γ clusters

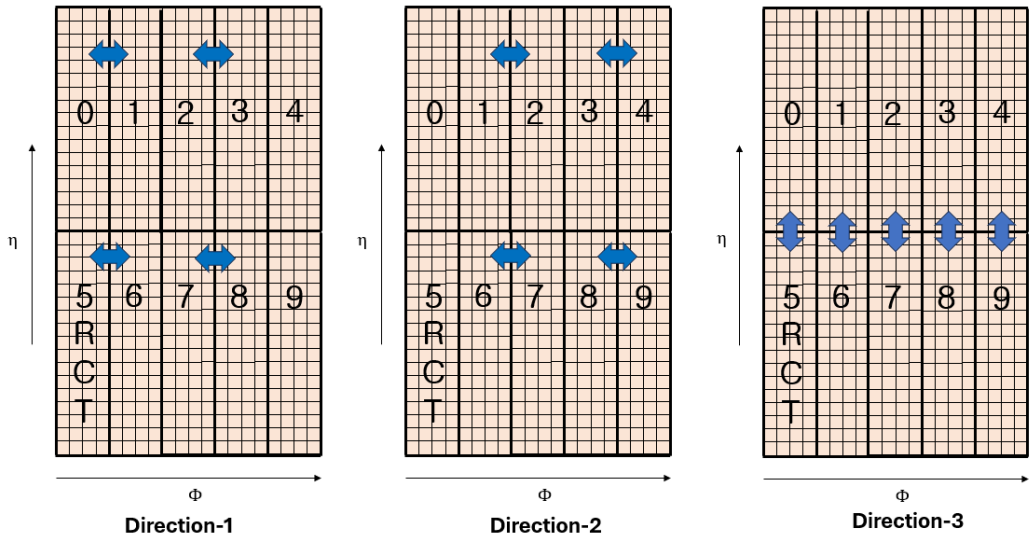

FIGURE 5.35: Stitching of the RCT regions in three directions. Direction-1 and Direction-2 represent the stitching at the phi boundaries and Direction-3 at the eta boundaries.

Quantity	Number of bit(s)
cluster energy	12
peak crystal eta	8
peak crystal phi	7
H/E	4
H/E *flag**	2
*isolation**	3
isolation *flag**	2
HCAL feature	6
*timing**	5
shape *flag**	2
brems	2
Reserved	11
	Total: 64

TABLE 5.13: Output e/γ cluster information packing in 64 bits for the correlator trigger. * *not yet included in the current version of the algorithm.*

information (2×64 bits = 128 bits), occupying (272 + 128 bits) 400 bits out of 576 bits available.

The next step is to create the PF clusters and calorimeter objects, such as jets and taus. Clustering at GCT requires creating full tower information, where cluster energy is added back with the corresponding tower. The full tower computation is concurrent in each RCT region. The computation involves checking the cluster position in tower geometry and adding back its energy with the corresponding tower. **The latency of performing full tower calculation is two clock cycles in one SLR**. The IP1 sends the full tower information to the IP2, which performs the PF clustering.

5.4.1.2 IP2

The main task of IP2 is to compute and send the PF clusters to the correlator trigger from the six central RCT regions. However, if the peak tower lies at the boundary of the RCT region, as shown in Fig. 5.36, the PF clustering is performed by considering the tower energy from the neighboring RCT regions.

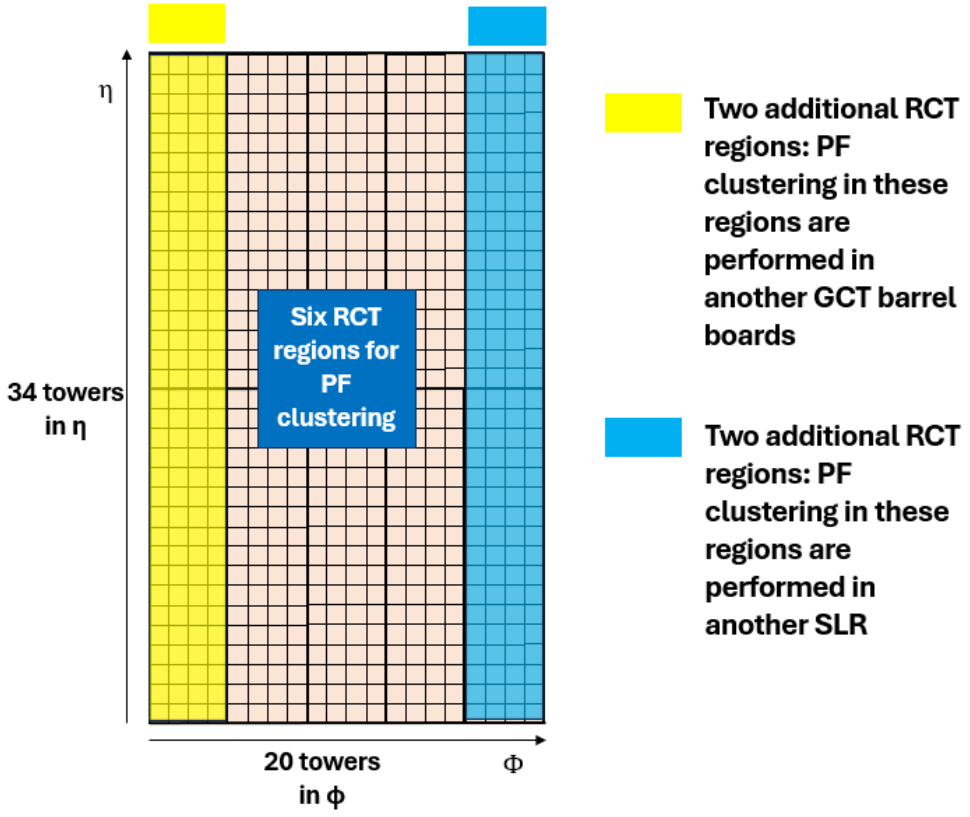

FIGURE 5.36: GCT barrel region highlighting the PF clustering coverage of the RCT regions.

The PF clustering algorithm performs concurrently in all six RCT regions. The first step of the algorithm is to create an extended region for each RCT region to account for the tower energies of the adjacent RCT regions. The extension of the RCT region is performed by including two rows and columns of towers on either side of the eta and phi direction, which increases the processing geometry from $17\eta \times 4\phi$ to $21\eta \times 8\phi$ tower region. PF clustering is an iterative operation that computes eight PF clusters in each RCT region. However, the iterative operation is concurrent in all six regions. The first step in each iteration is to find the position of the peak tower using the multi-level comparison strategy and subsequently compute the $3\eta \times 3\phi$ cluster energy around the peak tower. Once the cluster energy is computed, all the corresponding towers must be discarded to create a new cluster in the next iteration. The algorithm packs the PF

cluster information in 64-bit words, including the geometrical position of the peak tower in *eta* and *phi*. Table 5.14 summarized the 64-bit packing of one PF cluster. Fig 5.37 illustrates the full tower geometry of ten RCT regions, six central RCT regions, and two RCT regions with additional towers. The PF cluster and corresponding peak tower are also highlighted.

Quantity	Number of bit(s)
PF cluster energy	12
peak crystal eta	8
peak crystal phi	7
H/E	4
reserved	33
	Total: 64

TABLE 5.14: Output PF cluster information packed in 64 bits for the correlator trigger.

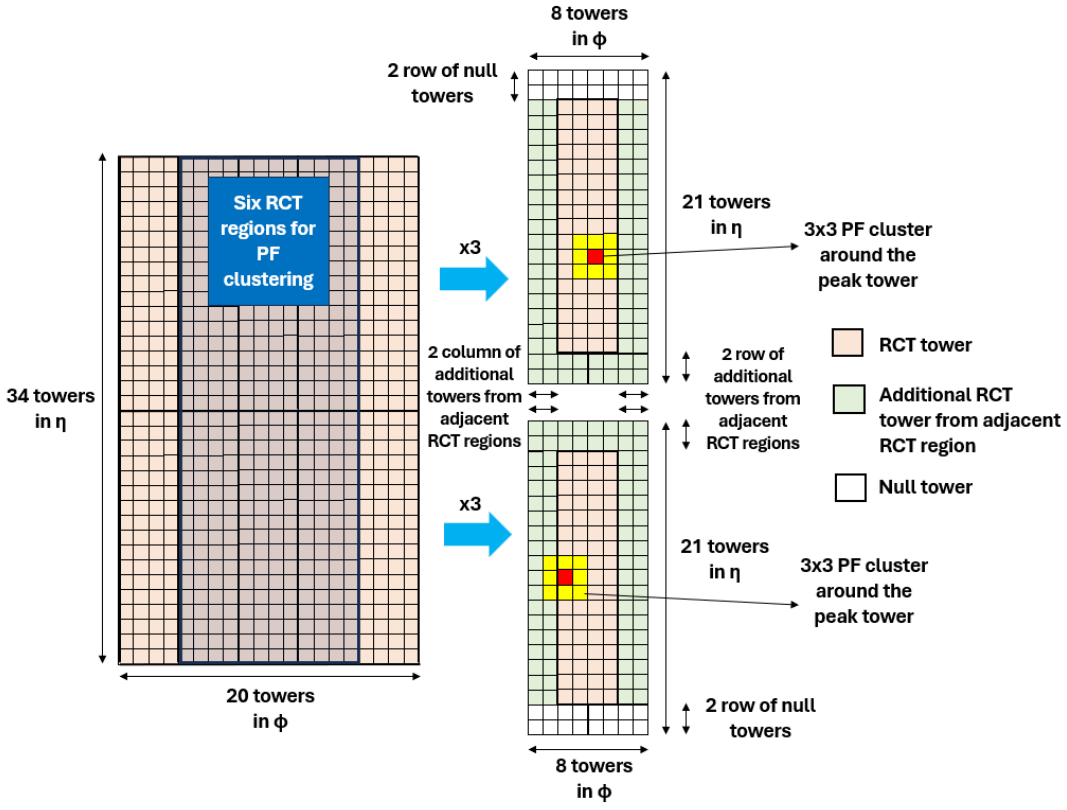

FIGURE 5.37: Full tower geometry of ten RCT regions. The six unique RCT regions are emphasized inside the grey box. One RCT region with additional adjacent towers is highlighted, along with one PF cluster and the corresponding peak tower.

All the 48 PF clusters are packed inside 24 output links with 2 PF clusters per link, which occupies (2×64 bits) 128 bits out of 576 bits available. The IP1 also packs the full tower information of the RCT regions, in 16 output links for the CaloObject algorithm. **The latency of the IP2 algorithm is 80 clock cycles**. Fig 5.38 depicts the HLS latency report of IP1 and IP2 algorithm.

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	2.486 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
117	117	0.325 us	0.325 us	9	9	function

FIGURE 5.38: IP1 and IP2 HLS timing report highlighting the latency of 117 clock cycles.

The RTL of IP1 and IP2 is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.15 summarizes the timing and utilization performance.

	IP1 and IP2
Latency (ns)	325
F_{MAX} (MHz)	360
LUTs (%)	11
FFs (%)	7

TABLE 5.15: Timing and utilization summary of the IP1 and IP2 algorithm.

The 24 output links from IP1 (400 bits per link) and IP2 (128 bits per link) are concatenated in RTL as they can be packed inside the threshold of 576-bit per bunch-crossing. Post concatenation, the number of output links is 24, each carrying 17 towers (17×16 bits = 272 bits), 2 e/γ clusters (2×64 bits = 128 bits), and 2 PF clusters (2×64 bits = 128 bits) information, which amount to (272 + 64 + 64 bits) 528 bits per link. The next step of the GCT barrel is to time-multiplex the output from one-time slice into six-time slices (TM1 to TM6). Time-multiplexing is performed to provide the output to correlator trigger layer-1, which comprises 18 APx boards and employs the time-multiplex architecture (TMUX6). The 18 APx boards are sub-divided into three *phi* regions, with each region corresponds to 6 APx boards taking input from one GCT barrel board. The six APx boards are arranged in a time-multiplexing manner, each serving one bunch-crossing of GCT barrel information. Fig 5.39 shows the connection of one GCT barrel board with 6 APx boards of correlator trigger layer-1.

The GCT barrel sends 48 output links to 6 APx boards; each APx board receives eight links. The time-multiplex algorithm is developed in VHDL, which converts the six input links available in a one-time slice to six output links arranged in six-time slices. To time-multiplex the 24 output links, four instances of the TMUX6 algorithm are used. The

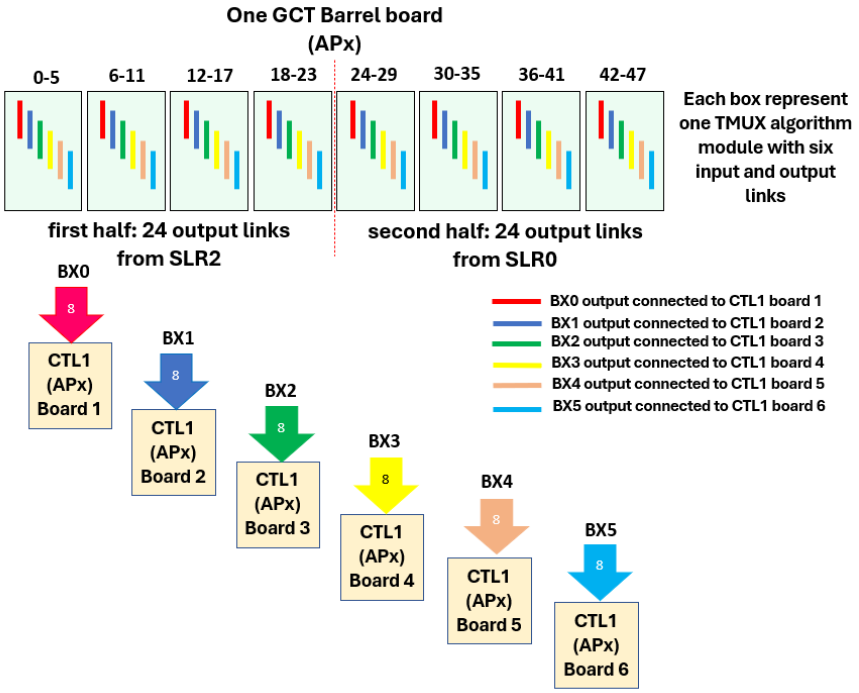

FIGURE 5.39: Connection of one GCT barrel board with six correlator layer-1 APx boards.

RTL of the TMUX6 algorithm is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.16 summarizes the timing and utilization performance.

	TMUX6
Latency (ns)	5.5
F_{MAX} (MHz)	381
LUTs (%)	0.6
FFs (%)	0.9

TABLE 5.16: Timing and utilization summary of the TMUX6 algorithm.

5.4.1.3 CaloObject Algorithm

The CaloObject algorithm takes input from both the GCT barrel algorithm implemented in SLR1 and SLR2. The main task of the CaloObject algorithm is to build jets and taus and compute the partial MET information. CaloJet algorithm processes the $34\eta \times 32\phi$ tower region of the barrel calorimeter. The geometrical coverage of one jet is $9\eta \times 9\phi$ tower region. Fig. 5.40 shows the CaloObject input processing region and the spread of one jet.

The main task of the algorithm is to prepare twelve jets and taus (six in negative and six in positive eta) in one GCT barrel region. The algorithm starts with partitioning

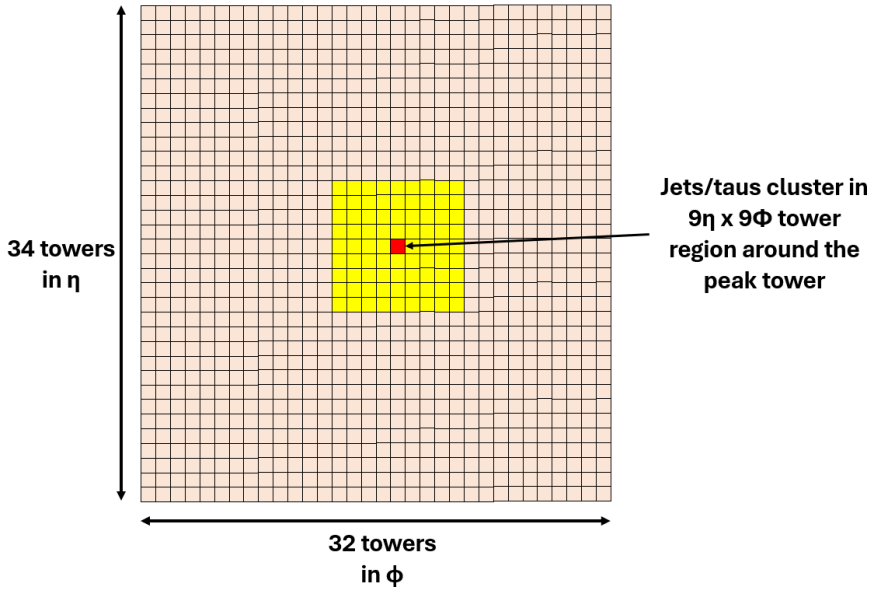

FIGURE 5.40: Input processing region of CaloObject algorithm highlighting the spread of one jet/taus in $9\eta \times 9\phi$ tower region.

the barrel region of $34\eta \times 32\phi$ towers into two eta halves, each comprising 17 unique towers from the corresponding eta region and six additional towers from the other eta region. Therefore, the geometry for each half is $23\eta \times 32\phi$ tower region. The additional towers in eta provide accuracy in making jets and taus at the partition boundary (at the interaction point). With two independent tower regions of $23\eta \times 32\phi$, the task is to build six jets and taus in each region concurrently. For efficient jet and taus clustering, the tower region of $23\eta \times 32\phi$ is reduced by a factor of three using the supertower methodology. In the calorimeter trigger, a supertower is build by adding the tower energy in a $3\eta \times 3\phi$ tower region in a pre-defined sequence. Therefore, to reduce the region of $23\eta \times 32\phi$ into a supertower region, it is decided to add one row of null towers at the extreme edge of the eta and four columns of null towers at the extreme edges in phi (two on either side of phi). This extension creates the towers region of $24\eta \times 36\phi$ on which the supertower algorithm is used, and ultimately, it reduces the region by a factor of three and provides the supertower region of $8\eta \times 12\phi$. Fig 5.41 highlighting the GCT barrel tower region and the corresponding supertower region.

Due to the reduction of the barrel tower region geometry into the supertower region, the area of finding jet and taus in the barrel is now reduced from $9\eta \times 9\phi$ towers to $3\eta \times 3\phi$ supertowers region. The first step in computing jets and taus is to locate the peak supertower in the region of $8\eta \times 12\phi$ supertower using the multi-level comparison strategy. The algorithm then computes the cluster energy by adding the supertower energy in the $3\eta \times 3\phi$ region around the peak supertower. If the energy of the peak supertower is higher than 87% of the cluster energy (to confirm if it is highly collimated), then the

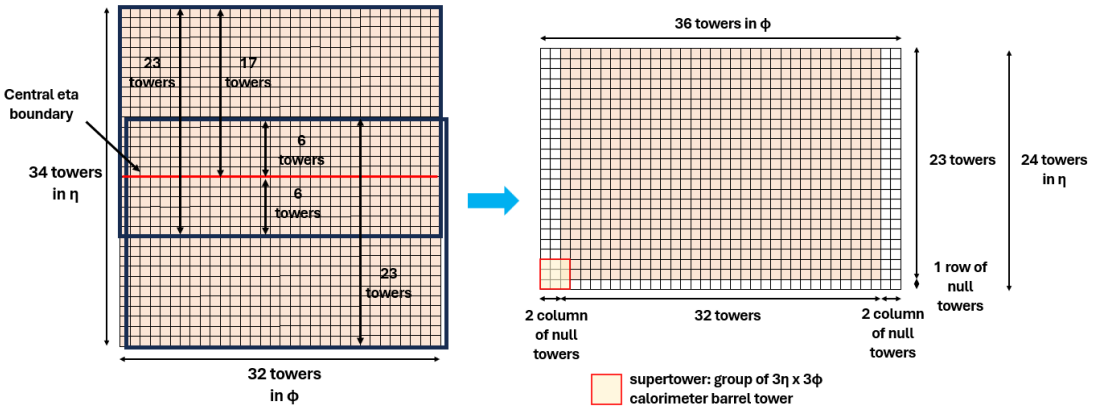

FIGURE 5.41: The conversion of tower geometry of size $32\eta \times 34\phi$ into two $24\eta \times 36\phi$ tower regions. A supertower is highlighted in the red square box depicting the coverage of $3\eta \times 3\phi$ towers.

algorithm flags the cluster as a tau. The algorithm then removes the corresponding supertower involves in computing a jet/taus and prepares the supertower region for the next iteration. Likewise, six such iterations are performed to create six jets and taus in each of the $12\eta \times 8\phi$ supertower regions. **The latency of computing one jet energy is three clock cycles. The overall latency is 99 clock cycles**. Fig 5.42 shows the peak supertower and jet in one $12\eta \times 8\phi$ region.

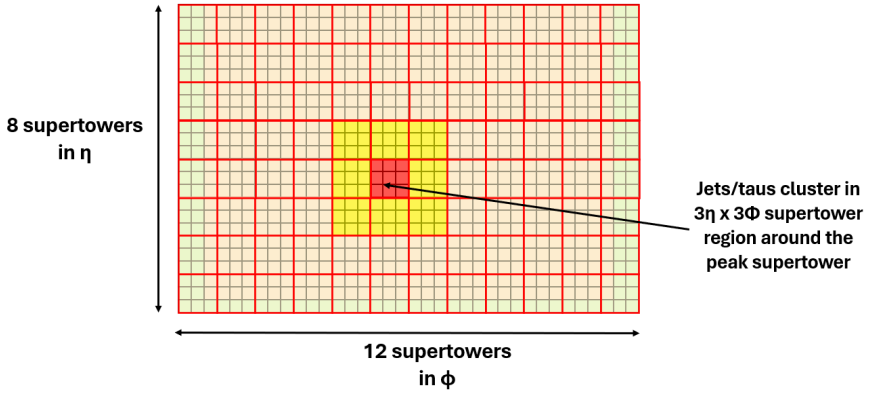

FIGURE 5.42: Jet/taus cluster and peak supertower in the $12\eta \times 8\phi$ supertower region.

The algorithm packs the jets and taus information in a 64-bit wide word, which includes 12 bits for the peak supertower energy, 12 bits for jet/taus energy, 5 bits for peak position in phi, 5 bits for peak position in eta, 1 bit for tagging the tau (to confirm is the cluster computed is a tau) and 29 bits are reserved for future usage. Table 5.17 summarizes the packing of one jet and taus information. The algorithm packs 12 jets and taus information in 2 output links, each carrying 6 jets and taus.

The next task of the CaloJet algorithm is to prepare the partial MET information. The additional portion of the MET algorithm is implemented in the GCT sum algorithm.

The main reason behind this decision is the computation of the sine and cosine of the *phi*, which varies from one barrel region to another. Therefore, to avoid implementing three different firmware in the GCT barrel. This part of the MET algorithm is moved to the GCT sum algorithm, which applies these *phi* dependent values based on the incoming GCT barrel link or board ID. The partial MET computation includes preparing the sum of the tower energy in each central *phi* column. Since there are 34 such towers in each *phi* column, each containing 12 bits of energy, the algorithm stores the sum of 34 towers in 18 bits of words without truncating the sum output. As there are 24 such columns, which are independent of each other, the algorithm performs the sum operation concurrently for each *phi* column. The output of the partial MET algorithm is packed inside a single link; it occupies (24×18 bits) 432 bits in the bit-width of 576 bits available. Fig 5.43 depicts the HLS latency report of the CaloObject algorithm.

Quantity	Number of bit(s)
Peak supertower energy	12
jet/tau energy	12
peak supertower phi	5
peak supertower eta	5
tagging the tau	1
reserve	29
	Total: 64

TABLE 5.17: Output jet and tau information packing in a 64 bits word.

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	2.681 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
99	99	0.275 us	0.275 us	9	9	function

  - **Detail**

FIGURE 5.43: CaloObject HLS timing report highlighting the latency of 99 clock cycles.

The RTL of CaloObject is synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.18 summarizes the timing and utilization performance.

Table 5.19 summarized the algorithm steps and latency of the IP1, IP2, and CaloObject algorithms.

	TMUX6
Latency (ns)	275
F_{MAX} (MHz)	367
LUTs (%)	7
FFs (%)	5

TABLE 5.18: Timing and utilization summary of the CaloObject algorithm.

Inputs:	
Input links: 64 **Clock frequency**: 360 MHz	
Algorithm	
IP1 and IP2	
Stitching RCT regions in *phi*.	**3 clock cycles**
Stitching RCT regions in *eta*.	**3 clock cycles**
Merging the RCT clusters energy into the tower region to create the full towers.	**2 clock cycles**
Making 48 particle-flow (PF) clusters ($3\eta \times 3\phi$ towers size) in six RCT regions	**80 clock cycles**
Jet and taus algorithm	
Making supertower ($3\eta \times 3\phi$ grouping of barrel towers)	**2 clock cycles**
Making six jets and taus in each eta half.	**12 clock cycles**
Time-multiplex algorithm	
Time multiplexing	**2 clock cycles**
Total latency*	**249 clock cycles**
Output:	
For correlator trigger: 48 links, each link carries 17 towers, 2 e/γ clusters, and 2 PF clusters. **For GCT sum**: 2 link carrying 12 jets, and taus information 1 link carrying partial MET information.	

TABLE 5.19: GCT barrel algorithm steps and latency report. * Total latency includes the latency of the wrapper of IP1 and IP2 and CaloObject algorithm. It also includes the SLR crossing of the full tower information from IP1 in SLR2 to the CaloObject algorithm in SLR1.

5.4.2 GCT Endcap

The GCT endcap boards cover the processing of the two endcap calorimeter detectors, which include the new HGCAL and the existing HF. This thesis discusses the algorithm processing the input from the HGCAL detector. The backend of the HGCAL system comprises two identical sub-systems to cover the two endcaps of the detector. The

endcap is partitioned into three identical *phi* sectors, each covering 120 degrees in *phi*. Each sector trigger primitives are prepared using 18 Serenity [26] boards featuring a time-multiplexing (TMUX18) scheme. In this scheme, each Serenity board processes one entire *phi* sector of one endcap for the N_{th} bunch-crossing (where N is from 1 to 18) and provides 4 links containing the HGCAL towers and clusters information. With four links per Serenity board, each sector of the HGCAL endcap provides (18×4) 72 links, amounting to (72×3) 244 links per endcap to the three GCT endcap boards. The algorithm at the GCT endcap and GCT sum are designed to run at TM1 (continuous stream of data); therefore, de-multiplexing of input is required from TM18 to TM1 at the GCT endcap boards. Fig 5.44 shows the schematic of the endcap backend system processing one endcap, providing 72×3 links to the three GCT endcap boards.

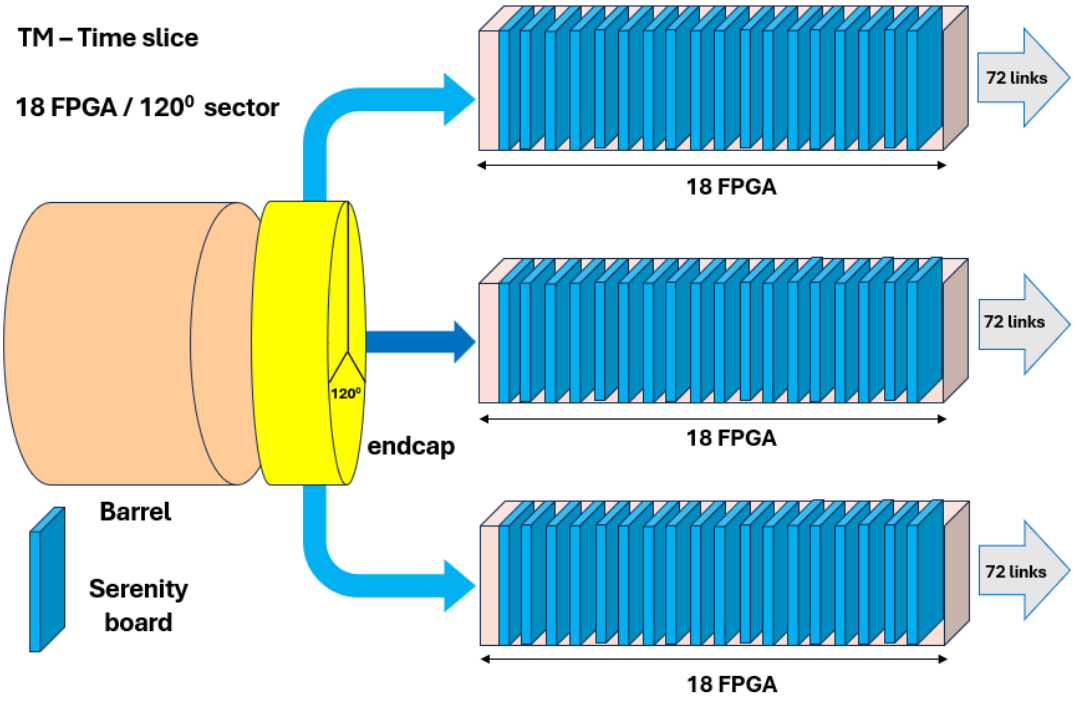

FIGURE 5.44: HGCAL backend (one endcap) connection with three GCT endcap boards.

With 9 physics payload/bunch-crossing (at 360 MHz of the clock), there are 162 payload words in 18 bunch-crossing. Out of 162 words sent by the HGCAL backend system, the first word consists of a header word, the following thirty words consist of HGCAL towers (120 towers), and the rest are clusters. The clusters in HGCAL are ordered in energy (p_T). The GCT endcap algorithm requires all the tower information to build the jets and taus and a few cluster information for e/γ clusters to be sent to the GCT sum card. All the data of the one endcap comprising 54 Serenity boards can be demultiplexed and processed using 3 APx boards. Each board features a VU13P FPGA supporting 4 SLRs,

where the three SLRs process two input time slices. Fig 5.45 shows the architecture of the GCT endcap taking input from one HGCAL endcap.

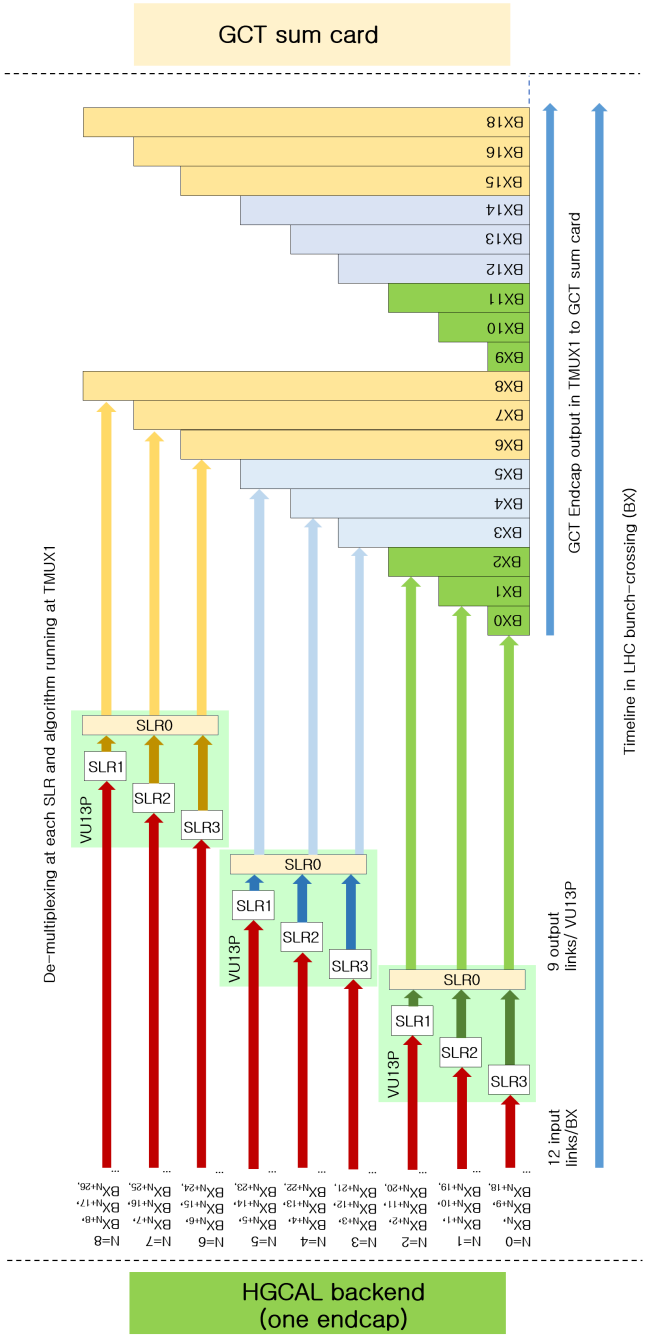

FIGURE 5.45: GCT endcap architecture. The demultiplexing of the entire endcap is performed using three XCVU13P FPGAs. Out of four SLRs, each SLR1-3 receives two bunch-crossing information, which are nine clock cycles apart. The algorithm performs the demultiplexing and provides input in the TMUX1 interval to the jet algorithm. The output from three SLRs is merged in the SLR0 and sent to the GCT sum board at TM1 interval.

To demultiplex the HGCAL input information in the GCT endcap board, a DEMUX

algorithm is developed in VHDL, which captures the incoming tower and cluster information of a one-time slice of all the HGCAL sector in terms of 12 input links (4 links per sector) and provides 108 links to the GCT endcap algorithm, which computes the jet and tau information. The DEMUX algorithm receives the input of a one-time slice arriving in 18 bunch-crossing, where the first header bit is ignored, the following 81 words are stored (the first thirty words contain 120 HGCAL tower information, and the rest are clusters), and the following 81 words (containing only clusters information) are discarded. The algorithm provides the output in 108 internal links, all available in TM1, ready for the GCT endcap jets and taus algorithm. Since half of the incoming data is not crucial at this implementation stage, one DEMUX algorithm can demultiplex two different time-slices, which are nine time-slices apart from each other. This methodology reduces the utilization of the DEMUX algorithm by half. This method leverages the inherent time displacement of two-time slices and uses a single DEMUX algorithm to demultiplex two-time slices. Fig 5.46 shows the DEMUX algorithm processing two-time slices. **The algorithm has a latency of 75 clock cycles**. The DEMUX algorithm output for each time slice is concatenated locally in the SLR0 of each APx device. The final concatenation is planned to be executed in one APx board before sending the final information to the GCT sum board.

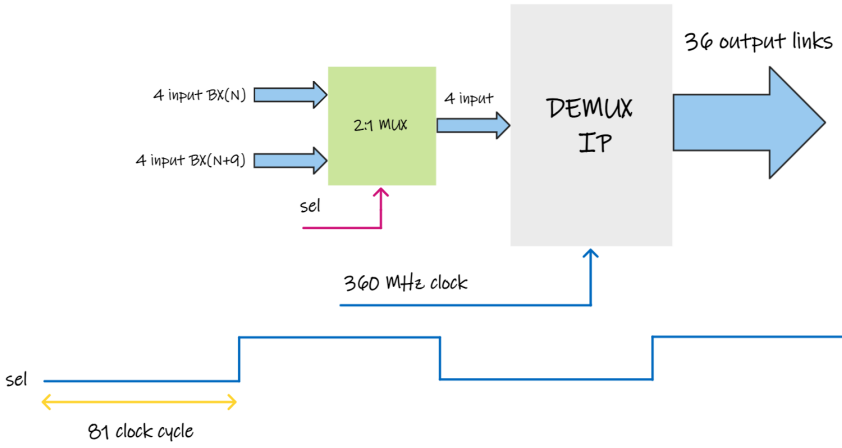

FIGURE 5.46: The demultiplexing scheme of the DEMUX algorithm using a two-input multiplexor. The two inputs have a time difference of 9 bunch-crossings.

The GCT endcap algorithm receives the input from the DEMUX algorithm. The input contains the HGCAL towers and clusters information. The geometrical coverage of the HGCAL is $20\eta \times 72\phi$ tower region. Similar to the GCT barrel, the clustering region of the jet and taus in the HGCAL region is $9\eta \times 9\phi$ towers. The algorithm computes 6 jets and taus in this region. As the operation in each SLR of the 6 APx boards of the GCT endcap is identical, prototyping one SLR can be easily scaled to three SLRs (one APx board) and similarly to six APx boards. The number of input links required

for demultiplexing two-time slices is 24, which can easily fit inside the SLR0 of VU9P FPGA.

To efficiently computes the jets and taus in the HGCAL $20\eta \times 72\phi$ tower region in terms of latency and utilization. The supertower methodology is used to reduce the HGCAL geometry. The algorithm considers two supertower geometry; the first version computes the supertower by adding the towers in the $3\eta \times 3\phi$ HGCAL region. The second version calculates the supertower by adding the towers in the $4\eta \times 3\phi$ geometry to account for the extension of the HGCAL towers in the forward region of the endcap. Post supertower computation, the HGCAL tower geometry reduces from $20\eta \times 72\phi$ to $6\eta \times 24\phi$. **The latency of making the supertower geometry is nine clock cycles**. Fig. 5.47 shows the HGCAL tower, two supertower regions used, and corresponding supertower geometry.

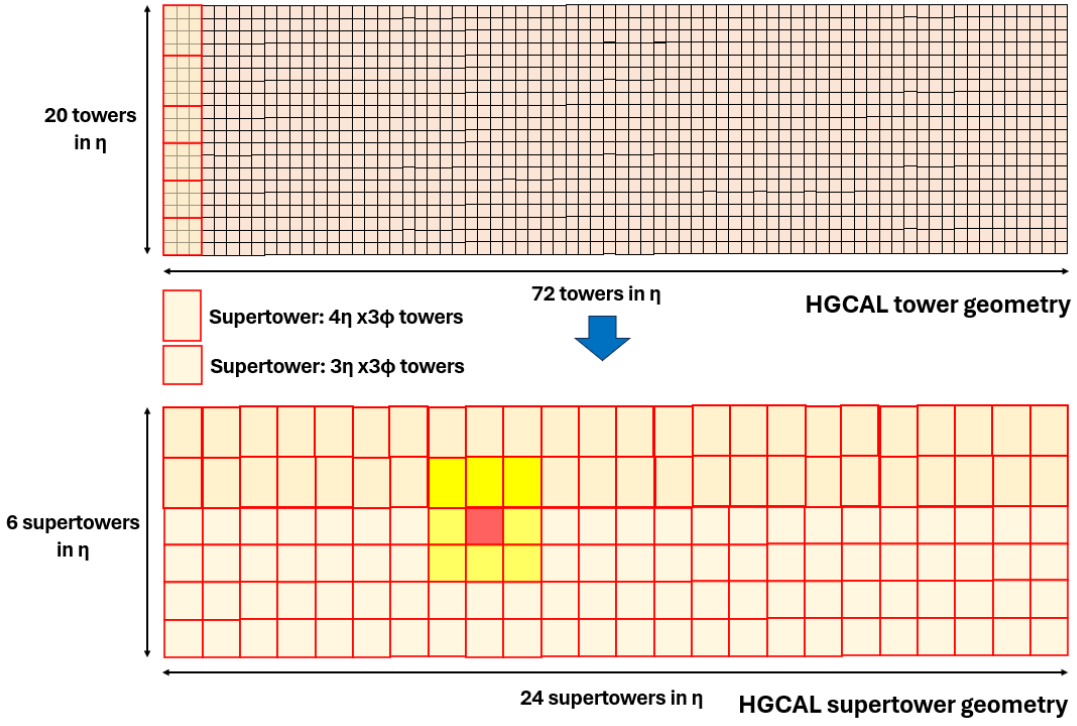

FIGURE 5.47: HGCAL tower geometry and the supertower geometry. It also highlights one peak supertower and the jet/taus of size $3\eta \times 3\phi$ in supertower.

The jets and taus are now computed in the HGCAL supertower region of $6\eta \times 24\phi$. The first step in making a jet is to find the peak supertower using the multilevel comparison strategy. The algorithm then computes the cluster energy by adding the supertower energy in the $3\eta \times 3\phi$ region around the peak supertower. The cluster can be flagged as tau if it is highly collimated such that the central peak supertower energy is higher than the cluster energy by more than 87%. After clustering, all the corresponding supertowers used in making jet and tau in the current iteration are removed to create the region for

next iteration. The output of the jet and tau information is packed in a 64-bit word. It requires only 1 output link ($6 \times 64 = 384$ bits) to carry all the jets and taus information. Similar to the jet clustering described in the GCT barrel, the GCT endcap algorithm also packs the position information of the peak supertower.

The latency of computing one jet and taus in one iteration is 20 clock cycles. The total latency of the algorithm is 151 clock cycles. Fig 5.48 depicts the HLS latency report of the GCT endcap jets and taus algorithm.

- **Timing**
  - **Summary**

Clock	Target	Estimated	Uncertainty
ap_clk	2.78 ns	2.460 ns	0.83 ns

- **Latency**
  - **Summary**

Latency (cycles)		Latency (absolute)		Interval (cycles)		Type
min	max	min	max	min	max	
151	151	0.419 us	0.419 us	9	9	function

  - **Detail**

FIGURE 5.48: GCT endcap jets and taus algorithm latency report.

The RTL of DEMUX and GCT endcap are synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.20 summarizes the timing and utilization performance.

	DEMUX	Endcap jets and taus
Latency (ns)	208	419
F_{MAX} (MHz)	368	361
LUTs (%)	0.5	2
FFs (%)	0.8	3

TABLE 5.20: Timing and utilization summary of the DEMUX and endcap jets and taus algorithms.

Table 5.21 summarizes the algorithm steps and latency of the GCT endcap algorithm.

Input:	
Input link: 24 **Clock frequency**: 360 MHz	
Algorithms:	
Demultiplexing TM18 to TM1.	**75 clock cycles**
Creating $20\eta \times 72\phi$ HGCAL tower region.	**1 clock cycles**
Creating supertower region of size $6\eta \times 24\phi$.	**9 clock cycles**
Creating six jets and taus.	**20 clock cycles**
❶ Finding peak in the 6×24 supertower region.	**8 clock cycles**
❷ Making jets around the peak.	**4 clock cycles**
Iterating the step ❶ and ❷ for six times to create six jets and taus.	
Total latency*	**238 clock cycles**
Output:	
Single link: Six jets and taus to the GCT sum board.	

TABLE 5.21: GCT endcap algorithm steps and latency report. * Total latency includes the latency of the wrapper combining the three instances of DEMUX and the endcap jets and taus algorithm.

5.4.3 Layer-3: GCT Sum

The GCT sum algorithm combines all the calorimeter-only based objects such as isolated and non-isolated e/γ clusters, jets, taus, and sum (H_T) information from the barrel and endcap. It sends the final output to the global trigger (GT). The sum algorithm converts the GCT information of the e/γ cluster, jet and taus as per the requirement of the global trigger. Each calorimeter objects are converted into the same data pattern of 64 bits, which starts with the first bit of valid flag, 16 bits for energy, 13 bits for the position of the calorimeter object in *phi*, 14 bits for position information in *eta*, and the rest 20 bits are reserved. Table 5.22 summarizes the packing of the calorimeter trigger object required by the global trigger.

Quantity	Number of bit(s)
valid	1
phi	13
eta	14
reserve	29
	Total: 64

TABLE 5.22: Calorimeter trigger object output packing in a 64 bits word for the global trigger.

The algorithm also computes the MET information by applying the *sine* and *cosine* values of the *phi* column to the corresponding tower energy information received from the

three GCT barrel boards. The *sin* and *cosine* values are calculated using the coordinate rotation digital computer (CORDIC) algorithm [42] developed in the HLS. The algorithm provides the output for all the *phi* values ranging from 0^0 to 360^0. CORDIC is an iterative algorithm that uses rotations to compute various elementary quantities, such as trigonometric functions. The algorithm is hardware efficient and primarily performs addition, subtractions, and bit shift operations. The *sine* and *cosine* components of the tower energy provide the magnitude of the transverse energy in the horizontal (E_X) and vertical (E_Y) directions. The final MET energy is calculated using the following equations.

$$\begin{aligned} E_{x1} &= E_1 \cos\phi_1 \\ E_{x_2} &= E_2 \cos\phi_2 \\ &\vdots \\ E_{x72} &= E_{72} \cos\phi_{72} \end{aligned} \tag{5.1}$$

$$\begin{aligned} E_{y_1} &= E_1 \sin\phi_1 \\ E_{y_2} &= E_2 \sin\phi_2 \\ &\vdots \\ E_{y72} &= E_{72} \sin\phi_{72} \end{aligned} \tag{5.2}$$

$$\begin{aligned} E_x &= E_{x1} + E_{x2} + \ldots + E_{x72} \\ E_y &= E_{y1} + E_{y2} + \ldots + E_{y72} \\ E_{MET} &= \sqrt{E_x^2 + E_y^2} \end{aligned} \tag{5.3}$$

Where the E_n, n ranges from 1 to 72, denotes the sum of the tower energy in the corresponding *phi* bin, E_{MET} is the final missing transverse energy (MET). For efficient implementation of the MET algorithm, the CORDIC algorithm uses the fixed point representation to store the value of the sine and cosine components. The trigonometric properties of sine and cosine are used to calculate the values in all four quadrants of *phi*, which reduces the latency and logic utilization of the algorithm. To calculate the E_{MET}, the square root operation uses the HLS synthesizable *sqrt* math operator. **The CORDIC and square root operation latency is 72 and 22 ns**, respectively.

Two versions of the GCT sum algorithm are developed. GCT sum v1.0 captures the jet information from the barrel and endcap and provides the output per the specification of the data required by the GT. The v1.0 algorithm is used to prototype the entire calorimeter trigger algorithm, which will be discussed in the next chapter. The MET algorithm is included in the GCT v2.0 algorithm.

The latency of the v1.0 and v2.0 algorithms is 64 and 178 ns, respectively. The RTL of v1.0 and v2.0 endcap are synthesized and implemented on XCVU9P FPGA in standalone mode. Table 5.23 summarizes the timing and utilization performance.

	GCT SUM v1.0	GCT SUM v2.0
Latency (ns)	64	178
F_{MAX} (MHz)	418	365
LUTs (%)	0.1	5
FFs (%)	0.4	3.5
DSPs (%)	0.2	4

TABLE 5.23: Timing and utilization summary of the GCT SUM v1.0 and v2.0 algorithms.

5.5 Conclusions

In this chapter, each layer of the calorimeter trigger algorithm and architecture, such as RCT, GCT barrel, GCT endcap, and GCT sum, is discussed in detail. In the next chapter, the prototyping of each of the algorithms using the APd1 board and the prototyping of the entire calorimeter trigger system will be covered.

Chapter 6

Calorimeter Trigger Firmware, and Prototyping

The unique symmetrical geometry of the CMS detector and the identical division of the calorimeter detector regions for FPGA processing at all the layers of the calorimeter trigger, such as RCT, GCT barrel, GCT endcap, and GCT sum, made the prototyping of the system more simpler and scalable. The prototyping of the system is performed in two steps. In the first step, the firmware of the individual sub-system is developed and tested on the APd1 board using multi bunch-crossing LHC events. Once all four sub-system performance and output are validated independently in a single board, a multi-board test is required to build and prototype the entire system. Therefore, in the upcoming section, the thesis discusses implementing and validating all four calorimeter trigger algorithms. Ultimately, the prototyping and multi-board test of the entire system will be discussed. Before going into the prototyping details, the APd1 device floorplan and link mapping, which heavily determine the firmware implementation of each algorithm, are discussed.

6.1 Calorimeter Trigger Prototype and Test Setup

Prototyping the calorimeter trigger requires designing the algorithms and implementing the firmware for each calorimeter trigger sub-system and layers on XCVU9P FPGA mounted on the APd1 board. The APd1 boards are placed in the ATCA crate operating in the University of Wisconsin lab. Fig 6.1 shows the ATCA crate hosting six APd1 and six APxF boards (based on XCVU13P FPGA).

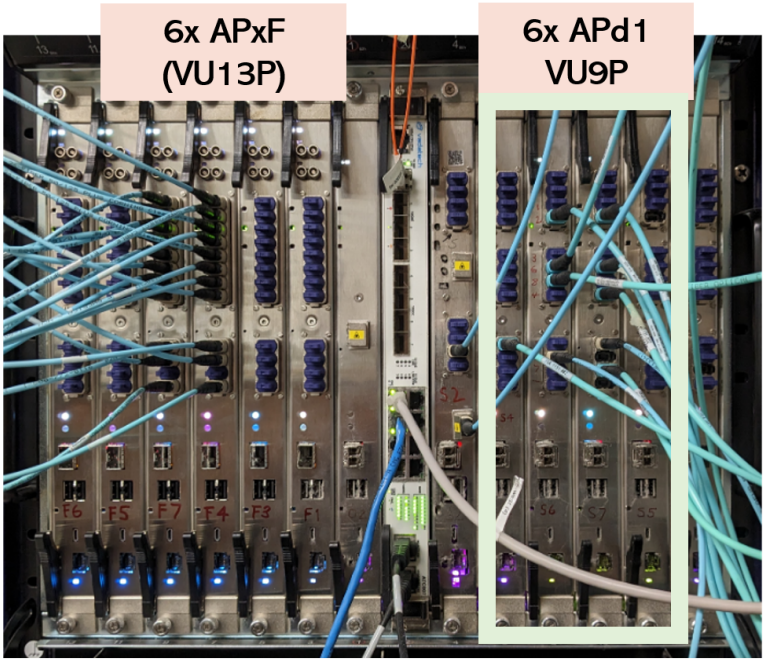

FIGURE 6.1: ATCA crate at the University of Wisconsin lab hosting six APd1 board (XCVU9P) and six APxF board (XCVU13P). Four APd1 boards are used to prototype the Phase-2 Level-1 calorimeter trigger.

Two types of tests are required for prototyping the system. The first method is known as the single board test, which requires uploading the algorithm bitfile on a single APd1 board, and then input(s) is fed to the playback buffer(s) of the firmware shell, which is connected to the input port(s) of the algorithm. The processed output is then captured in the output buffer. If the captured patterns match with the Monte Carlo simulated reference output, it confirms the validity of the prototyping. The second method is known as the multi-board test, which requires connecting multiple boards using the optical links, testing the captured patterns of each board, and verifying it with the MC-generated reference output. The main advantage of performing the multi-board test is the estimation of the latency, which mainly comprises the latency of the algorithm and the serialization and communication of the data from one system to another. Following are the steps involved in the multi-board test.

- Connecting the boards according to the test plan using the number of links required.
- Uploading the input pattern inside the playback/receiver (R_x) buffer of the class-1 board(s).
- The class-1 algorithm processes the input and passes the output to the transmitters (T_x) to send the data to the class-2 board(s) via 25 Gbps optical links.

- The class-2 board capture the class-1 board(s) output in its (R_x) buffer and provide it to the class-2 algorithm.
- The class-2 algorithm processes the input and passes the output to the capture/-transmitter (T_x) buffer.
- Matching the output captured in the class-2 T_x buffer with the MC simulated reference output.

Fig. 6.2 depicts the multi-board connection of three class-1 boards and one class-2 board. The outputs of the class-1 boards are connected to the input of the class-2 board using the 25 Gbps optical links.

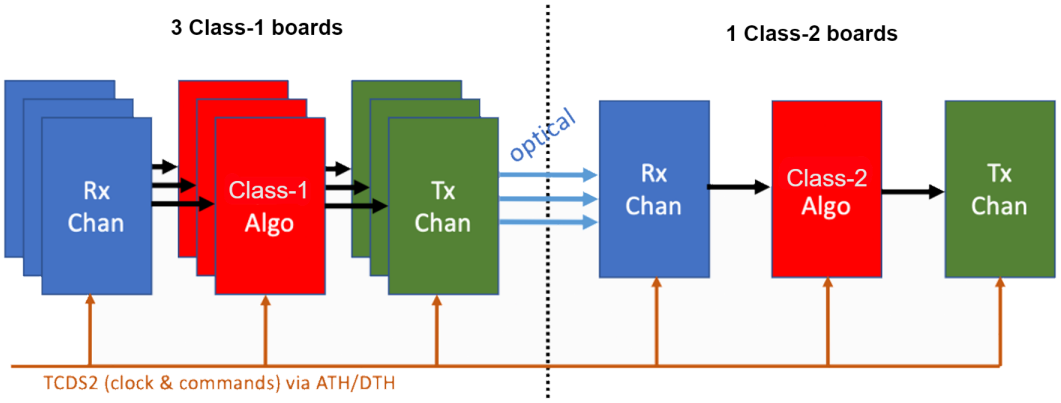

FIGURE 6.2: Multi-boards test setup.

In the next section the prototyping of each of the calorimeter trigger sub-system will be discussed.

6.2 Prototyping Regional Calorimeter Trigger

The next sub-sections describes the implementation of the RCT v1.0 and v2.0 algorithm.

6.2.1 RCT v1.0

The architecture of the RCT v1.0 highlights the partitioning of the $17\eta \times 4\phi$ (68 links) region into two regions of $17\eta \times 2\phi$ (34 links). Both regions are independently implemented in two SLRs and combined in one of the SLRs. The SLR2 (40 links) and SLR1 (36 links) fulfill the requirements of 34 links for the $17\eta \times 2\phi$ region. The SLR0 space is saved for any potential future upgrade. The algorithm clock frequency is 360 MHz; at such a high frequency, it is observed that the data delay between two sequential logic

is more dominated by the net/wire delay than the logic delay. The logic delay is determined by the control and data path logic of the design. However, the net delay is decided by the placement of the logic in the device. Therefore, it is crucial to place the design in an optimum space without causing congestion, which makes the routing more severe. Therefore, the RCT17x2 algorithm is placed inside eight clock regions to reduce the net delay, and the clock regions are selected in the middle of the SLR such that the net delay is equal for all the input links [6]. Fig. 6.3 shows the floorplan of the RCT v1.0 algorithm.

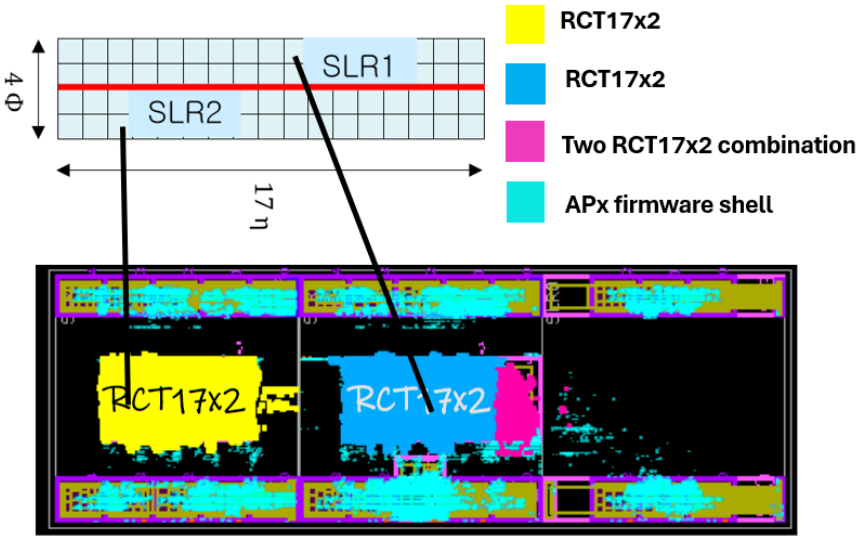

FIGURE 6.3: RCT v1.0 firmware implementation. The two 17x2 regions are implemented in SLR2 and SLR1 (8 clock regions each), respectively. Both the regions are combined in the SLR1 (two clock regions).

The latency and F_{MAX} of the implemented design are 269 ns and 366 MHz, respectively. The implemented design utilizes 14% of LUTs and 15% of Flip-flops. The utilization result includes the LUTs and FFs utilized by the algorithm and firmware shell. The latency and utilization budget is well within the budget of Phase-2. The bitstream of the implemented design is tested and validated successfully on the APd1 board using three distinct Phase-2 Monte Carlo simulated events such as ggHbb (gluon gluon production of Higgs boson decaying to a pair of b quarks), QCD multi-jet, and VBFHToTauTau (vector boson fusion production of Higgs boson decaying to a pair of tau leptons). Table 6.1 summarized the timing and area performance of the RCT v1.0 algorithm.

	RCT v1.0
Latency (ns)	269
F_{MAX} (MHz)	366
LUTs (%)	14
FFs (%)	15

TABLE 6.1: Timing and utilization summary of the RCT v1.0 algorithm implemented on APd1 board (XCVU9P FPGA).

6.2.2 RCT v2.0

RCT v2.0 architecture highlights the partitioning of the $17\eta \times 4\phi$ (68 links) region into two regions of $8\eta \times 4\phi$ (32 links for ECAL) and $9\eta \times 4\phi$ (36 links for ECAL). The partitioning of the region is made in a way such that it satisfies the constraint of the number of links available in each SLR. Apart from the regional partitioning, the algorithm is also partitioned into smaller units (IPs) such that it passes the clock constraint of 360 MHz. Three IPs are used: IP1 and IP2 process the ECAL information, and IP3 merges the ECAL and HCAL information. In the firmware implementation, the region of $9\eta \times 4\phi$, processed using IP1 and IP2, is implemented in SLR2. It requires 36 input links to bring all the crystal information to this region. The output of the IP2 of this region is routed to the SLR1 via SLL. The data crossing between SLR costs in terms of latency, and it is recommended to pipeline this path for a minimum of three clock cycles to reduce the timing-related failure. In the SLR, the region of $8\eta \times 4\phi$, processed using IP1 and IP2, is implemented, which requires 32 input links from ECAL. The IP3 is also implemented in SLR1, which requires input from the IP2 of the $8\eta \times 4\phi$ and $9\eta \times 4\phi$ region, and from the two HCAL links. The IP1 and IP2 of $9\eta \times 4\phi$ region is floorplanned in the middle of SLR2, and the IP1, IP2 of $8\eta \times 4\phi$ and IP3 is floorplanned in SLR1. The HCAL links connected to the IP3 are buffered to hold the HCAL information such that it synchronizes with the ECAL information coming from the IP2 of both regions. IP3 provides the final four output links to the GCT barrel board. Fig 6.4 shows the RCT v2.0 firmware implementation on XCVU9P FPGA. The SLR0 area is saved for the future algorithm implementation.

The latency and F_{MAX} of the implemented design are 1.2 μs and 361 MHz, respectively. The implemented design utilizes 16% of LUTs and 13% of Flip-flops. The latency and utilization budget is well within the budget of Phase-2. The bitstream of the implemented design is tested and validated successfully on the APd1 board using Phase-2 Monte Carlo simulated events. Table 6.2 summarizes the timing and area performance of the RCT v2.0 algorithm.

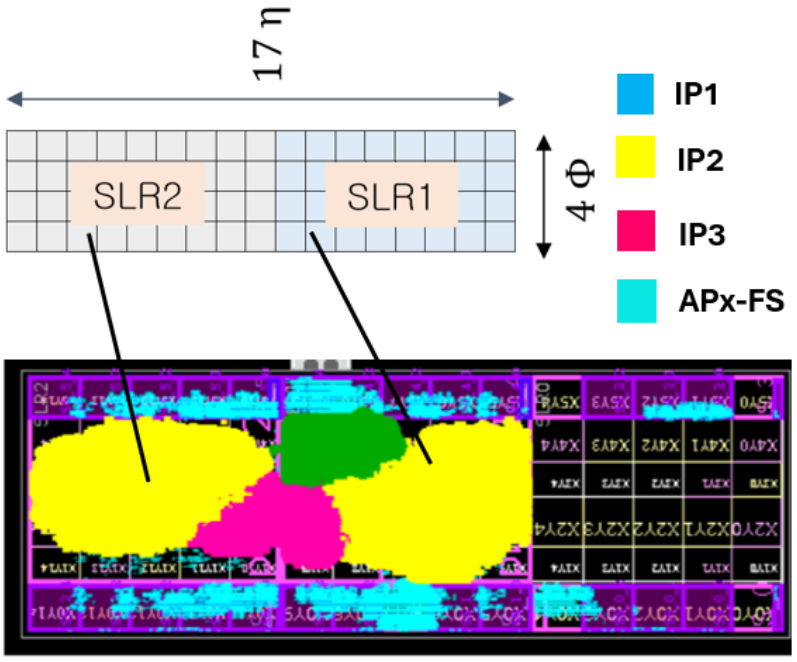

FIGURE 6.4: RCT v2.0 firmware implementation. The two IP1 and two IP2 of $8\eta \times 4\phi$ and $9\eta \times 4\phi$ tower regions are implemented in SLR1 and SLR2, respectively. The IP3 is implemented in SLR1.

	RCT v2.0
Latency (μs)	1.2
F_{MAX} (MHz)	361
LUTs (%)	16
FFs (%)	12

TABLE 6.2: Timing and utilization summary of the RCT v2.0 algorithm implemented on APd1 board (XCVU9P FPGA).

6.3 Prototyping Global Calorimeter Trigger

The following sub-sections discuss the prototyping of the GCT barrel, GCT endcap, and GCT sum algorithms.

6.3.1 Prototyping GCT Barrel

The GCT barrel takes input from 36 RCT boards. It requires three GCT barrel boards to process the input from all the RCT boards. The salient feature of GCT barrel architecture is that prototyping one region in a single SLR can be easily scaled to two SLRs covering the entire GCT barrel region of a single board. Subsequently, it can be scaled to three boards to cover all the RCT boards.

The GCT barrel algorithm covering 10 RCT boards, which comprises IP1 and IP2, is implemented in SLR2. The input of the 10 RCT boards is provided via the 40 input buffers of SLR2. The IP1 stitches the RCT electron/photon clusters and creates the full tower information for IP2 and the jet algorithm. The IP2 creates the 48 PF clusters in the central RCT regions. The output of the IP1 and IP2 are merged in RTL and connected to the 24 output buffers of the SLR2 dedicated for the correlator trigger. The CaloOject algorithm is implemented in SLR1. It receives one-half of the full tower information from the IP1 of SLR2 and another half from 16 input buffers of SLR1. The input from the buffer is pipelined for the time IP1 of SLR2 provides the full tower information so that the CaloObject algorithm receives all the inputs simultaneously. The CaloObject algorithm creates 12 jets and sends the output via two links to the GCT sum algorithm. Fig 6.5 shows the floorplan of the GCT barrel algorithm implemented on XCVU9P FPGA.

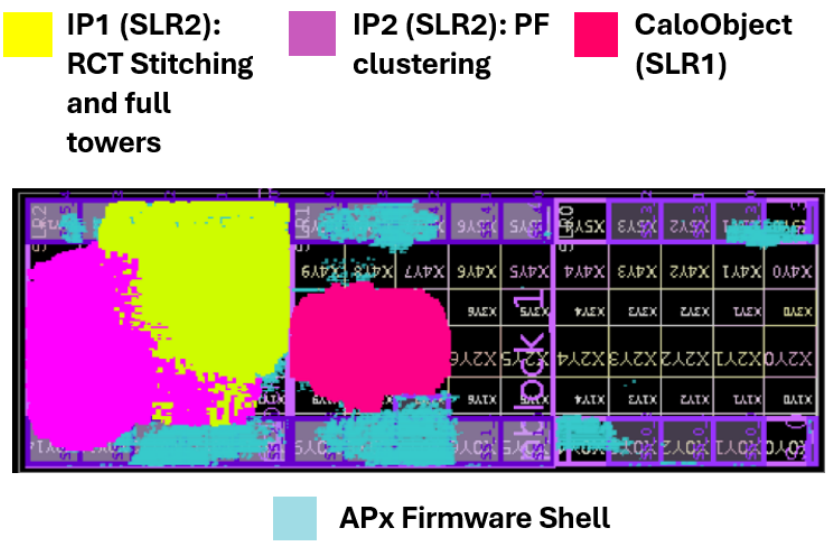

FIGURE 6.5: GCT barrel firmware implementation on XCVU9P FPGA. The IP1 and IP2 is implemented in SLR2 and CaloObject algorithm is implemented in SLR1

The latency and F_{MAX} of the implemented design are 692 ns and 360 MHz, respectively. The implemented design utilizes 22% of LUTs and 18% of Flip-flops. The latency and utilization budget is well within the budget of Phase-2. The bitstream of the implemented design is tested and validated successfully on the APd1 board using the cluster and tower information of the RCT algorithm. Table 6.3 summarized the timing and area performance of the GCT barrel algorithm.

	GCT barrel
Latency (ns)	692
F_{MAX} (MHz)	360
LUTs (%)	22
FFs (%)	18

TABLE 6.3: Timing and utilization summary of the GCT barrel algorithm implemented on APd1 board (XCVU9P FPGA).

6.3.2 Prototyping GCT endcap

The GCT endcap covers the processing of the two endcaps of the CMS detector. It requires six GCT endcap boards to process the input from the two endcaps. Each endcap backend system provides 216 input links in 18 time slices (12 per time slice). As discussed in the architecture of the system, the prototyping of the system in one SLR can be easily scalable to one device and, subsequently, to six boards. The GCT endcap algorithm, which includes the demultiplexing (DEMUX), jets, and taus algorithms, is implemented in SLR0. It requires 24 input (two times slices) links to bring the data of the HGCAL tower and clusters. The output of the algorithm is one link carrying six jets and taus information. The DEMUX algorithm is placed at the top of the SLR0 in two clock regions to demultiplex the 24 links from TM18 to TM1; it creates 108 internal input links for the jets and taus algorithm placed just under the DEMUX algorithm and floorplanned in six clock regions. Constraining the design in eight clock regions reduces the net delay and helps meet the timing constraint of the design. Fig 6.6 shows the floorplan of the GCT endcap implemented on XCVU9P FPGA.

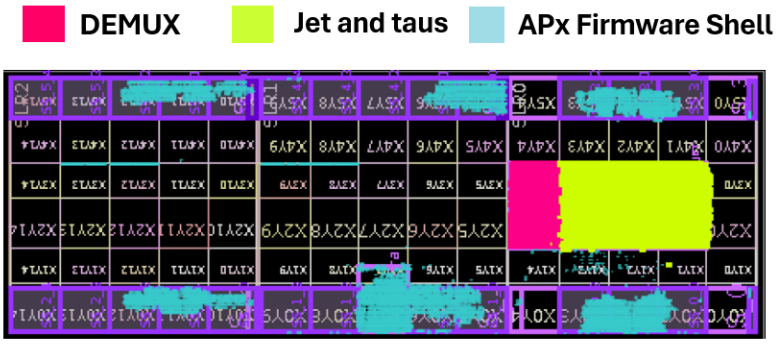

FIGURE 6.6: GCT barrel firmware implementation on XCVU9P FPGA. The IP1 and IP2 is implemented in SLR2 and CaloObject algorithm is implemented in SLR1

The latency and F_{MAX} of the implemented design are 661 ns and 360 MHz, respectively. The implemented design utilizes 6% of LUTs and 7% of Flip-flops. The latency and utilization budget is well within the budget of Phase-2. The bitstream of the implemented

design is tested and validated successfully on the APd1 board using the HGCAL cluster and tower information. Table 6.4 summarized the timing and area performance of the GCT endcap algorithm.

	GCT barrel
Latency (ns)	661
F_{MAX} (MHz)	360
LUTs (%)	6
FFs (%)	7

TABLE 6.4: Timing and utilization summary of the GCT endcap algorithm implemented on APd1 board (XCVU9P FPGA).

6.3.3 Prototyping GCT sum

The GCT sum board combines the calorimeter-only objects from the barrel and endcap, computes the MET quantity for the barrel, reorganizes the data per the GT requirement, and sends it to the global trigger in six-time slices. As discussed in the last chapter, two algorithm versions are prepared. GCT sum v1.0 is developed for prototyping the calorimeter trigger by processing the jet information from the barrel and endcap. The algorithm is implemented in SLR2 utilizing 24 output links, six from the endcap and 18 from the barrel (6 per barrel board). It uses six output links to send the jet information to the GT. Fig. 6.7 shows the implementation of the GCT sum v1.0 algorithm on XCVU9P FPGA. The latency of the algorithm is 105 ns with 4% LUT and 4% FF utilization. Table 6.5 summarized the timing and area performance of the GCT sum v1.0 algorithm.

	GCT sum
Latency (ns)	105
F_{MAX} (MHz)	392
LUTs (%)	4
FFs (%)	4

TABLE 6.5: Timing and utilization summary of the GCT sum v1.0 algorithm implemented on APd1 board (XCVU9P FPGA).

The v2.0 includes the final MET computation using the partial MET information from GCT barrel boards. The algorithm is implemented in standalone mode, utilizes 5% LUTs, 4% of FF, 4% of DSPs, and achieves the F_{MAX} of 365 MHz. Fig 6.8 shows the floorplan of the GCT sum v2.0 algorithm on XCVU9P FPGA. Table 6.6 summarized the timing and area performance of the GCT sum v2.0 algorithm.

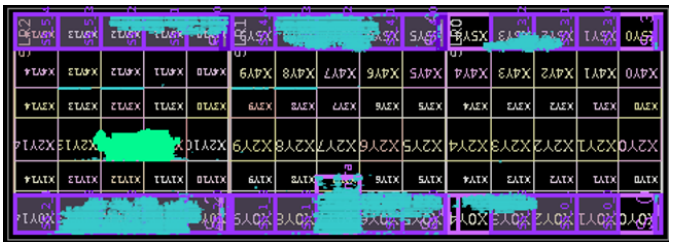

FIGURE 6.7: GCT sum v1.0 algorithm placed in SLR2 (highlighted in green).

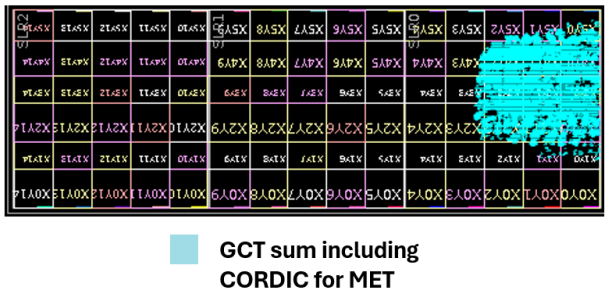

FIGURE 6.8: GCT sum v2.0 algorithm placed in SLR0 (highlighted in green)

	GCT sum
Latency (ns)	178
F_{MAX} (MHz)	365
LUTs (%)	5
FFs (%)	3.5
DSPs (%)	4

TABLE 6.6: Timing and utilization summary of the GCT sum v2.0 algorithm implemented on APd1 board (XCVU9P FPGA).

6.4 Prototyping Barrel Calorimeter Trigger

The barrel calorimeter trigger is prototyped by connecting the RCT and GCT barrel boards and performing multi-board tests. The multi-board test is performed for two RCT and GCT algorithm versions discussed in the following sub-sections.

6.4.1 RCT v1.0 to GCT (MET) algorithm

The first multi-board test is performed to prototype RCT v1.0 and GCT (MET algorithm) using four APd1 boards, with three boards hosting the RCT v1.0 algorithm and one hosting the GCT (MET) algorithm. Fig 6.9 illustrates the RCT and GCT board connection.

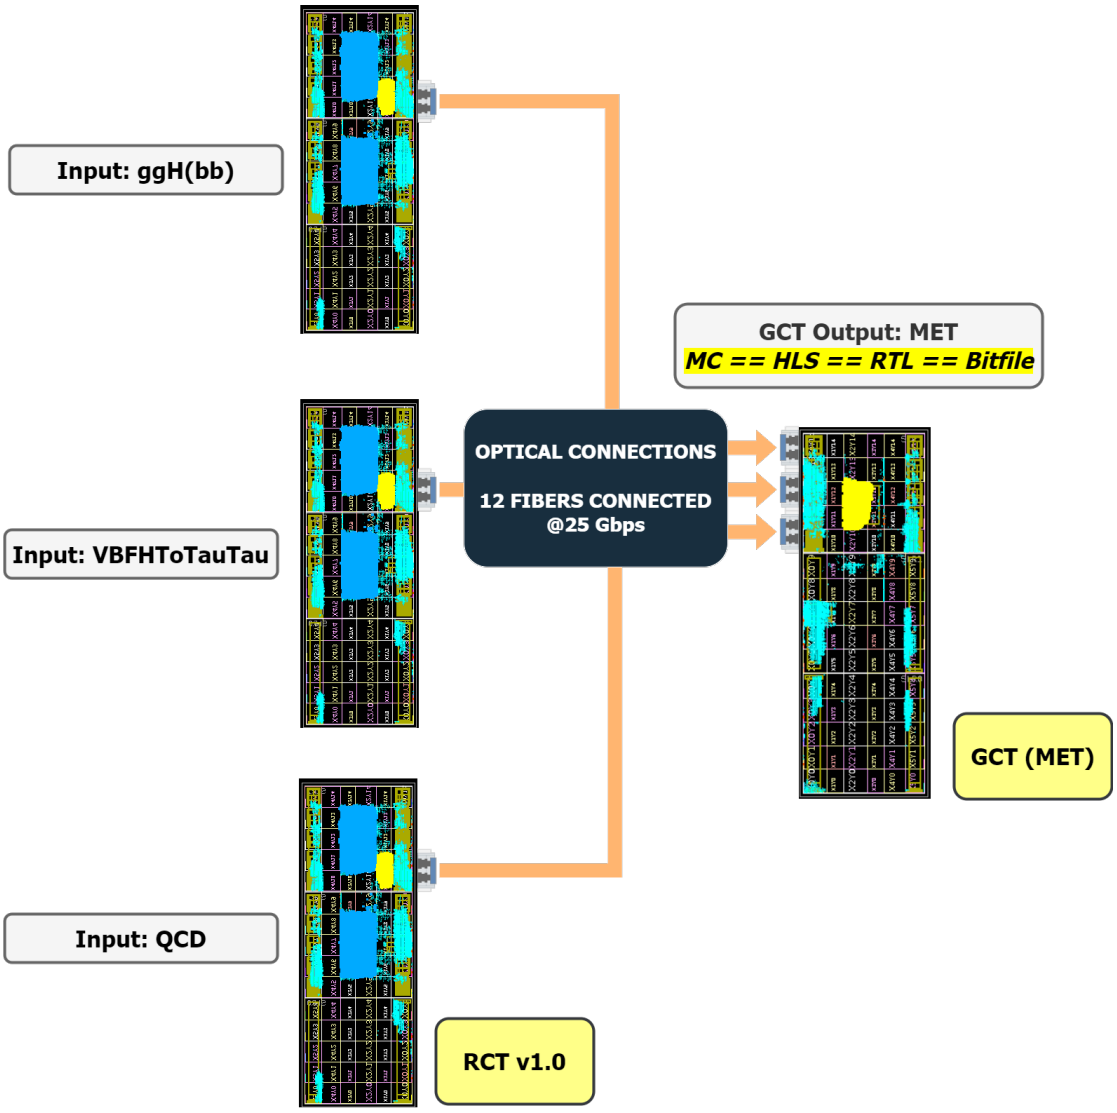

FIGURE 6.9: Four APd1 boards test of barrel calorimeter trigger.

Each RCT board communicates with the GCT board using four optical links that carry the information at 25 Gbps of bandwidth. Three different test vectors such as *ggHbb* (gluon fusion production of Higgs boson decaying to a pair of b quarks), *QCD* multi-jet, and *VBFHToTauTau* (vector boson fusion production of Higgs boson decaying to a pair of tau leptons) are generated using Phase-2 Monte Carlo (MC) simulation. Each RCT card is fed with one distinct test vector, and the output of the RCT is transmitted and captured at the GCT board, which ultimately serves as the input for the MET algorithm. Theoretically, the GCT algorithm is designed to process the inputs from twelve unique RCT boards, i.e., 48 input links. Therefore, the optical inputs from the three RCT boards are replicated thrice in the RTL wrapper of the GCT algorithm to mimic the processing of 12 RCT boards. The output of the GCT algorithm is captured at the transmitter (T_x) buffer, and it entirely agrees with the reference output generated and validated at Monte-carlo, HLS, and RTL levels.

6.4.2 RCT v2.0 and GCT Barrel algorithm

The second multi-board test is performed by connecting two APd1 boards using four optical links running at 25 Gbps of bandwidth. The first layer-1 board employed the RCT v2.0 algorithm, and the second layer-2 board employed the GCT barrel algorithm (GCTCIP and CaloObject algorithms). The input ECAL crystal and HCAL tower information are fed into the R_X buffer of the RCT board. The RCT algorithm creates the cluster and tower information and sends the output of the GCT barrel board via four optical links. The GCT firmware is designed to process the inputs from 10 RCT boards in one SLR. Therefore, the inputs of the 9 RCT boards (36 input links) are provided using the 36 R_X buffers in the SLR2 of the GCT barrel board. The four inputs from optical links and 36 inputs from the R_X buffers are aligned such that the GCT barrel algorithm processes all 40 input links simultaneously. The GCT barrel provides 24 output links to the correlator trigger and two outputs to the GCT sum board. The T_X buffers capture all the output links. The output of the GCT is matched with the Monte-carlo, HLS, and RTL simulations. Fig 6.10 illustrates the RCT and GCT board connection.

6.5 Prototyping Calorimeter Trigger System

The Prototyping of the entire calorimeter trigger system is performed using four APd1 boards. The boards are connected using sixteen optical links at the University of Wisconsin electronics lab. The connection of the four APd1 boards is shown in Fig 6.1. The input barrel ECAL crystal, HCAL tower, and HGCAL tower information are fed to the RCT and GCT endcap R_X buffers, respectively. The output of RCT in terms of clusters and towers information is sent via four optical links of 25 Gbps bandwidth to the GCT barrel. The GCT barrel replicates this input nine times to emulate receiving the input information from ten RCT boards (4 links via optical fiber and 36 links via R_X buffers. The GCT barrel output for the correlator is captured in the 24-output buffers. GCT barrel sends the jet information to the GCT sum card using two optical links. The algorithm at the GCT endcap demultiplexes the input information, creates the jets, and sends the result to the GCT sum board using one optical link. The final layer-3 sum board receives the input in terms of jets built in the barrel and the endcap. It prepares the data in the format compliance by the GT and saves it in the output buffer. The entire chain is tested successfully with one bunch-crossing of input information. The prototyping results are compared and matched at all stages, such as HLS, RTL, and final bit-file simulation. Fig 6.11 demonstrates the arrangements of the cards connected in three layers.

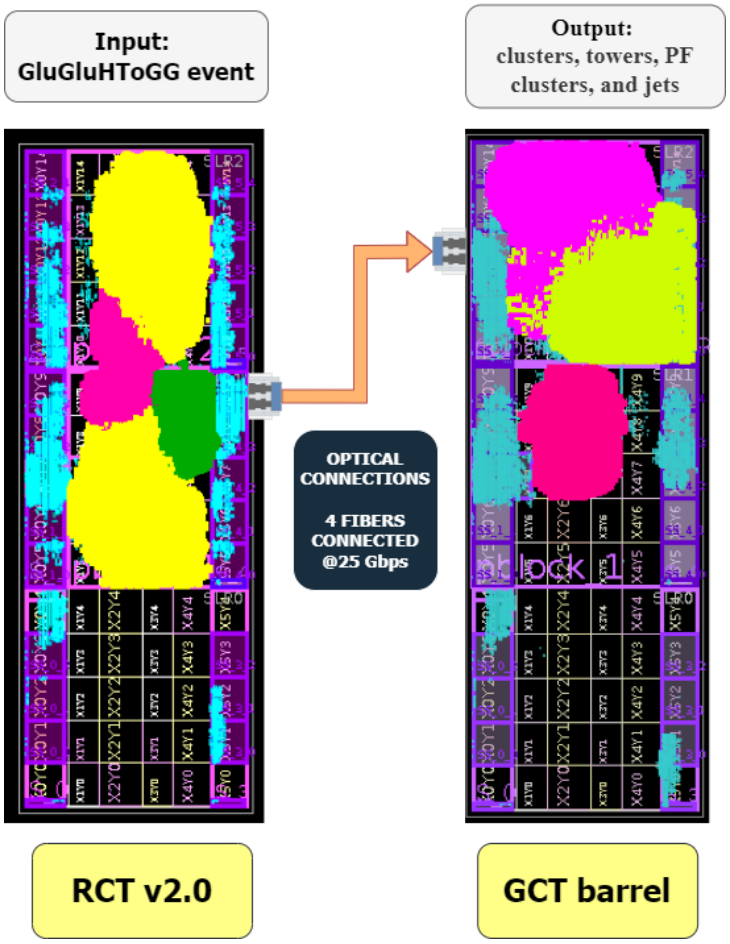

FIGURE 6.10: Two APd1 boards test of barrel calorimeter trigger.

The latency of the system is also measured, which includes the contribution from the trigger algorithms and data transmission latency from one board to another. In this test, the latency is measured for two processing chains. In the first chain, the inputs are provided to the RCT buffers, and the outputs are captured at the GCT barrel buffers for the correlator trigger. The measured latency for this scheme is 1.8 μS, which is well within the budget of 3 μS. In the second chain, the inputs are fed to the RCT and GCT endcap buffers, and the jets are captured at the output of the GCT sum board. The measured latency of this scheme is 2.4 μS, which is well within the budget of 4 μS. Table 6.7 summarizes the measured latency of the Phase-2 Level-1 calorimeter trigger in the multi-board test performed.

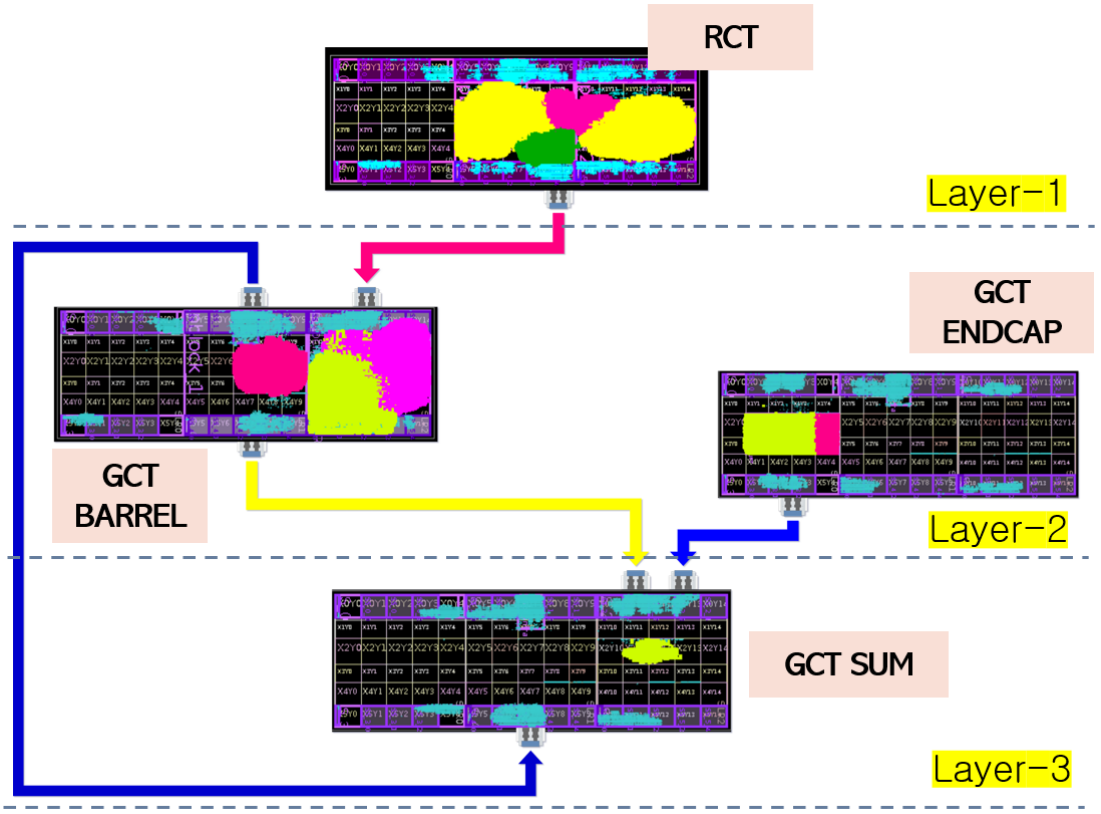

FIGURE 6.11: Connection of the four APd1 boards used for prototyping the entire Phase-2 Level-1 calorimeter trigger. It highlights the firmware used at the RCT, GCT barrel, GCT endcap, and GCT sum together with the optical connection made.

Testing scheme	Latency in μS (budget)
RCT v2.0 to GCT barrel (output to correlator trigger)	1.8 (3)
Calorimeter Trigger (output to global trigger)	2.4 (4)

TABLE 6.7: Latency summary of the multi-board test performed.

Chapter 7

Summary and Conclusions

The inadequacy of the Phase-1 Level-1 calorimeter trigger system in handling the HL-LHC requirements, in terms of the algorithm and hardware, led to the design and prototyping of the Phase-2 level-1 calorimeter trigger. This thesis proposes the architecture and algorithms for the Phase-2 Level-1 calorimeter trigger, which takes input from four sub-detector systems such as ECAL, HCAL, HGCAL, and HF. The proposed architecture processes the entire calorimeter information using 46 APx board.

The architecture is devised in two levels and subsequently in three layers. The first layer hosts 36 RCT boards processing the barrel information of the detector in terms of ECAL crystals and HCAL towers. The second layer hosts the GCT barrel and endcap boards, taking the input from the RCT and the endcap detectors, respectively. The GCT barrel sends the barrel calorimeter information to the correlator trigger using 144 optical links within the latency budget of 3μs. The third layer hosts the GCT sum board, taking the inputs from all the nine layer-2 boards and sending the final output to the global trigger using six optical links within the latency budget of 4μs. The primary constraint of the system is latency, which must be within the budget of HL-LHC. The utilization is preferred to be low and within the threshold of 50% of the device logic.

The Level-1 calorimeter trigger system comprises four trigger module such as RCT, GCT barrel, GCT endcap, and GCT sum. Each of the trigger module hosts an identical algorithm. Therefore, the implementation and validation of each of the subsystems is performed to prototype the complete system. It involves developing a new trigger algorithm for each module and validating it on the APd1 board.

In this thesis, two versions of the RCT algorithms are developed and prototyped, i.e., RCT v1.0 and v2.0. The RCT v1.0 processes the ECAL crystal information, creates the e/γ cluster and tower energy for each of the 68 ECAL towers and sends the final

output to the GCT barrel board. The main drawback of this algorithm is the absence of the HCAL tower processing, comprehensive coverage of the crystals in *eta* and *phi* to create clusters, and account for the bremsstrahlung radiation losses. The algorithm is implemented in XCVU9P FPGA, and the bitfile is validated on APd1 board using the MC simulation of three unique LHC events. Table 7.1 summarizes the timing and area performance of the RCT v1.0 algorithm implemented on the XCVU9P FPGA.

	RCT v1.0
Latency (ns)	269
F_{MAX} (MHz)	366
LUTs (%)	14
FFs (%)	15

TABLE 7.1: Timing and utilization summary of the RCT v1.0 algorithm implemented on XCVU9P FPGA.

The RCT v2.0 algorithm is developed, which covers the missing elements of the RCT v1.0 algorithm, such as the processing of HCAL towers and brems correction. The algorithm creates the e/γ clusters and towers using three distinct IPs such as IP1, IP2, and IP3. All the three IPs are developed independently and merged in RTL. The final RTL is then implemented and validated on the APd1 board. The latency and utilization of the implemented design are well within the budget of HL-LHC requirements. Table 7.2 summarizes the timing and area performance of the RCT v1.0 algorithm implemented on the XCVU9P FPGA.

	RCT v2.0
Latency (μs)	1.2
F_{MAX} (MHz)	361
LUTs (%)	16
FFs (%)	12

TABLE 7.2: Timing and utilization summary of the RCT v2.0 algorithm implemented on XCVU9P FPGA.

In addition, the GCT barrel algorithm is developed, which takes input from 16 RCT boards (12 unique RCT regions and 4 overlapping regions). The architecture of the GCT barrel is proposed in three SLRs, with SLR2 and SLR0 hosting the IP1, IP2, and TMUX algorithm and SLR1 hosting the CaloObject algorithm. The prototyping is performed by implementing the GCTCIP algorithm in SLR2 and the CaloObject algorithm in SLR1. The GCTCIP algorithm stitches the e/γ clusters at the boundaries of the RCT regions and creates PF clusters in the unique RCT regions. The information of the stitched e/γ clusters, towers, and PF clusters are packed into 24 output links for the correlator trigger. It also prepares the full tower information for the CaloObject

algorithm, which performs the jets and taus clustering and computes the partial MET information. The algorithm is implemented on XCVU9P FPGA and validated on the APd1 board. Table 7.3 summarizes the timing and area performance of the RCT v1.0 algorithm implemented on the XCVU9P FPGA.

	GCT barrel
Latency (ns)	692
F_{MAX} (MHz)	360
LUTs (%)	22
FFs (%)	18

TABLE 7.3: Timing and utilization summary of the GCT barrel algorithm implemented on XCVU9P FPGA.

In this thesis, the GCT endcap algorithm is developed, which demultiplexes the input information from the HGCAL and performs the jets and taus clustering on HGCAL towers. The demultiplex algorithm (DEMUX) is developed in VHDL, which converts the 18 time-slice inputs to 1 time-slice and provides 108 internal links containing HGCAL towers and clusters information to the endcap jets and taus algorithm. The DEMUX, jets, and taus algorithms are implemented on XCVU9P FPGA and validated on the APd1 board. The latency and utilization of the implemented design are well within the budget of HL-LHC requirements. Table 7.4 summarizes the timing and area performance of the RCT v1.0 algorithm implemented on the XCVU9P FPGA.

	GCT barrel
Latency (ns)	661
F_{MAX} (MHz)	360
LUTs (%)	6
FFs (%)	7

TABLE 7.4: Timing and utilization summary of the GCT endcap algorithm implemented on XCVU9P FPGA.

In this thesis, Two versions of the GCT sum algorithm are developed, i.e., v1.0 and v2.0. The GCT sum algorithm takes input from the GCT barrel and endcap boards. The v1.0 algorithm combined the jet information from the barrel and endcap, reorganize it as per the global trigger requirements, and send the final output to the global trigger using six output links. The v2.0 is similar to v1.0. However, it includes the missing transverse energy (MET) calculation using the CORDIC algorithm and time-multiplex the six output links for the global trigger. The v1.0 and v2.0 are implemented on XCVU9P, and v1.0 is validated on the APd1 board, which is used in prototyping the complete calorimeter trigger system. The latency and utilization of the implemented design are well within the budget of HL-LHC requirements. Table 7.5 summarizes the timing and

area performance of the GCT v1.0 and v2.0 algorithm implemented on the XCVU9P FPGA.

	GCT sum v1.0	GCT sum v2.0
Latency	105	178
F_{MAX} (MHz)	392	365
LUT (%)	4	5
FF (%)	4	3.5
DSP (%)		4

TABLE 7.5: Latency and utilization report of GCT v1.0 and v2.0 implemented algorithm on XCVU9P FPGA.

Prototyping the calorimeter trigger system also requires connecting one sub-system to another and validating the algorithm output with the MC simulation. The multi-board test also provides the performance of the system in terms of latency, which includes the algorithm and data transmission from one board to another. In this regard, the barrel calorimeter trigger is prototyped by connecting the RCT and GCT barrel boards. In the first scheme, four APd1 boards are used, with three layer-1 boards hosting the RCT v1.0 algorithm and one layer-2 board hosting the GCT (MET) algorithm. The layer-2 GCT board is connected to the three layer-1 boards using 12 optical links (4 links per RCT board) operating at 25 Gbps of bandwidth. The chain is tested successfully using three unique Phase-2 MC events. In the second scheme, two APd1 boards are used, with a layer-1 board hosting the RCT v2.0 algorithm is connected to the layer-2 board using four optical links, which host the GCT barrel algorithm. The chain is tested successfully using one unique LHC Phase-2 MC event.

The entire calorimeter trigger system including the barrel and endcap is prototyped using only 4 APd1 boards and connecting them using 16 optical links. The input barrel and endcap information are fed into the RCT and GCT endcaps buffer, respectively. The 24 correlator output links are captured in the output buffers of the GCT barrel board. The output for the global trigger in terms of jets from the barrel and endcap is captured in the output buffers of the GCT sum board. The complete chain is tested successfully using one LHC event. The measured latency of the system for providing the output to the correlator trigger (1.8 μs) and global trigger (2.4 μs) is well within the budget of the HL-LHC requirements. Table 7.6 summarizes the latency of the Phase-2 Level-1 calorimeter trigger measure in the four boards prototyping.

Testing scheme	Latency in μS (budget)
RCT v2.0 to GCT barrel (output to correlator trigger)	1.8 (3)
Calorimeter Trigger (output to global trigger)	2.4 (4)

TABLE 7.6: Latency summary of the Phase-2 Level-1 calorimeter trigger.

Appendix A

Clock Domain Crossing

In the early development stage, the RCT v1.0 algorithm failed to meet the timing constraints of a 360 MHz clock. Therefore, a clock domain crossing (CDC) scheme is developed to relax the constraint of a 360 MHz clock by moving to a lower clock, such as a 180 MHz. The CDC ensures a lower clock for the algorithm and, at the same time, allows the FPGA to communicate with 25 Gbps of link bandwidth. The 25 Gbps of link bandwidth and LHC clock frequency of 40 MHz (one bunch-crossing frequency) corresponds to 576 bits of physics payload available for the algorithm in one bunch-crossing. The CSP provides the 64 bits of physics payload at 360 MHz clock, which maintains the constant bandwidth ($40MHz \times 576bits = 360MHz \times 64bits$) from 40 MHz to 360 MHz clock domain. Maintaining constant bandwidth is crucial as it prevents any loss of data.

In this thesis, a clock domain crossing scheme is developed, which uses an asynchronous first-in-first-out (FIFO) module to cross the 360 MHz clock domain to the 180 MHz clock. For the FIFO, the write side clock is fixed to 360 MHz with 64 bits of write data width, and for the read side, the clock is 180 MHz with 128 bits of read width. The 128 bits of read data word at 180 MHz ensures the constant bandwidth ($360MHz \times 64bits = 180MHz \times 128bits$) at either side of the FIFO. Though this scheme ensures a constant flow of LHC data, however, on the read side, the two consecutive bunch-crossing of information get mixed after every fifth read occurs. Mixing LHC data prevents the direct injection of FIFO read data into the algorithm wrapper, which prepares the unique LHC data for the algorithm. Fig A.1 demonstrates the FIFO read and write operation.

To tackle the problem associated with mixing of two consecutive LHC bunch-crossing data, an algorithm in VHDL is developed that takes the read output from FIFO in two consecutive bunch-crossing and provides two distinct bunch-crossing of LHC information to the algorithm wrapper. This algorithm is called BunchMaker. The BunchMaker

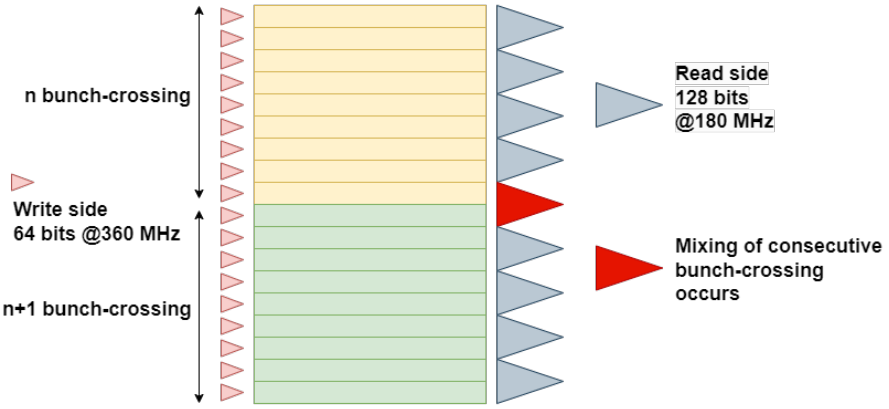

FIGURE A.1: Write and read scheme of the asynchronous FIFO for two LHC bunch-crossing. Write operation in FIFO is performed using 64 bit of data word at the write clock of 360 MHz. It takes 9 clock cycles to complete the write operation of one bunch-crossing. The read operation of FIFO is performed using 128 bits of data word at the read clock of 180 MHz. The mixing of two bunch consecutive bunch-crossing will occur at every fifth read operation.

accumulates nine clock cycles of FIFO read data, which includes the mixed information of two bunch-crossing. The accumulated information amounts to 1152 bits of words. The algorithm redistributes this 1152 bits into three groups. The first group comprises zero data, and the following two groups comprise two LHC bunch-crossing data. Fig A.2 represents the data-flow of the BunchMaker algorithm. The BunchMaker algorithm maintains the constant bandwidth between the input and output.

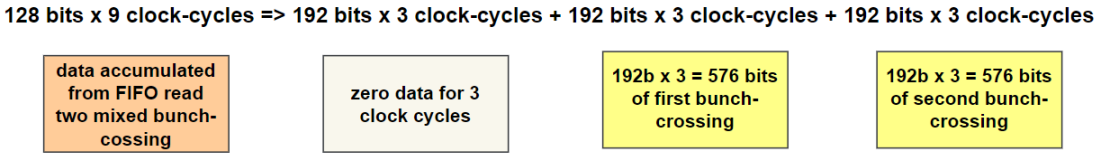

FIGURE A.2: Dataflow of BunchMaker algorithm.

The BunchMaker provides the unique LHC events to the algorithm wrapper in three clock cycles of 180 MHz. This mandates the HLS algorithm to operate at a 180 MHz clock with a pipeline interval of three. Similarly, at the output side, a BunchBreaker algorithm is developed to re-arrange the output of the algorithm and provide two bunch-crossing of LHC information to the output FIFO. Fig. A.3 represents the dataflow of the BunchBreaker algorithm.

The output FIFO provides the clock domain crossing from 180 MHz of the algorithm clock to the 360 MHz clock. The write side consists of 128 bits of write word and 180 MHz clock and the read side consists of 64 bits of read word and 360 MHz of read clock.

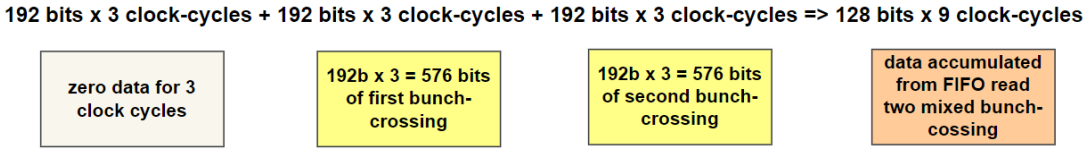

FIGURE A.3: Dataflow of BunchBreaker algorithm.

Fig A.4 shows the architecture of the CDC scheme designed for the Phase-2 calorimeter trigger.

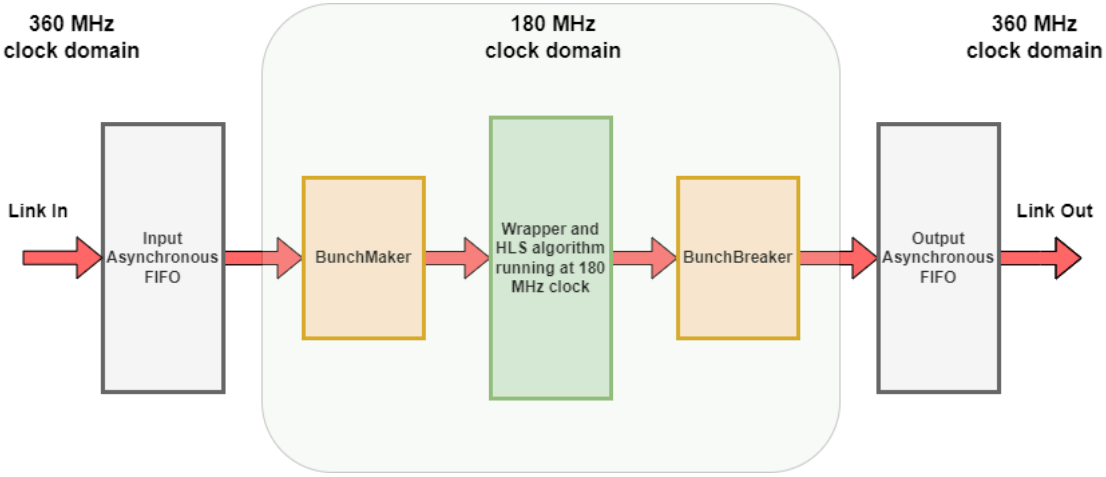

FIGURE A.4: Architecture of the CDC scheme. It consists of the input FIFO, BunchMaker, BunchBreaker, and Output FIFO to support the input and output integrity for the central algorithm block

As mentioned earlier, the RCT v1.0 algorithm failed to meet the 360 MHz clock constraint, so the CDC scheme is used to tackle this problem. As the input FIFO defined the 180 MHz clock for the RCT v1.0 algorithm, the BunchMaker defined the pipeline interval of three for the algorithm. The algorithm is synthesized in HLS and implemented in XCVU9P FPGA with the input and output link constraint discussed in Chapter 6. The additional circuitry of FIFO, BunchMaker, and BunchBreaker modules require additional floorplanning to meet the timing constraint. The input and output FIFO requires floorplanning near the corresponding input and output links. The implemented design tested successfully on the APd1 board with multiple LHC events of test vectors. The latency of the RCT v1.0 algorithm synthesized at 360 (pipeline interval of 9) and 180 MHz (pipeline interval of 3) clocks are 150 ns and 106 ns, respectively. The additional circuitry of FIFO, BunchMaker, and BunchBreaker associated with the CDC requires more logic elements such as BRAM, LUTs, and Flip-flops. Table A.1 compares the utilization results of the RCT v1.0 algorithm implemented on the XCVU9P FPGA using 360 MHz and 180 MHz clock.

	RCT v1.0 360 MHz design	RCT v1.0 180 MHz design
HLS Latency (ns)	150	106
LUT (%)	16	18
FF (%)	17	26
BRAM (%)	27	35

Table A.1: Latency and utilization report of RCT v1.0 algorithm designed using 360 and 180 MHz clock.

CDC provides a solution in terms of lowering the algorithm clock from 360 MHz to 180 MHz. However, the additional circuitry and lower pipeline interval (from 9 to 3) results in higher utilization of BRAM, LUTs, and flip-flops. Moreover, an additional floorplan of the input and output FIFO is also required. Therefore, for the calorimeter trigger algorithm design, a new strategy was developed to partition the functionality of the algorithm into multiple smaller IPs, with each IP being developed independently in HLS and merged together in RTL. This scheme is used to develop the RCT v2.0 and GCT barrel algorithms.

Bibliography

[1] LHC Machine. *JINST*, 3:S08001, 2008. doi: 10.1088/1748-0221/3/08/S08001.

[2] S. Chatrchyan et al. The CMS Experiment at the CERN LHC. *JINST*, 3:S08004, 2008. doi: 10.1088/1748-0221/3/08/S08004.

[3] S. Chatrchyan et al. Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc. *Physics Letters B*, 716(1):30–61, September 2012. ISSN 0370-2693. doi: 10.1016/j.physletb.2012.08.021. URL `http://dx.doi.org/10.1016/j.physletb.2012.08.021`.

[4] Andrea Dainese, Michelangelo Mangano, Andreas B Meyer, Aleandro Nisati, Gavin Salam, and Mika Anton Vesterinen. Report on the Physics at the HL-LHC, and Perspectives for the HE-LHC. Technical report, Geneva, Switzerland, 2019. URL `https://cds.cern.ch/record/2703572`.

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

[6] P. Kumar and B. Gomber. System design and prototyping for the cms level-1 trigger at the high-luminosity lhc. *IEEE Transactions on Nuclear Science*, 70(6): 1075–1082, 2023. doi: 10.1109/TNS.2023.3247510.

[7] CMS Collaboration. Development of the cms detector for the cern lhc run 3, 2023.

[8] Alexandre Zabi, Florian Beaudette, Luca Cadamuro, Olivier Davignon, Thierry Romanteau, Thomas Strebler, Maria Luisa Cepeda, Jean-baptiste Sauvan, Nicholas Wardle, Robin Cameron Aggleton, Fionn Amhairghen Ball, James John Brooke, David Newbold, Sudarshan Paramesvaran, D Smith, Joseph Ross Taylor, Konstantinos Fountas, Mark David John Baber, Aaron Bundock, Shane Davy Breeze, Matthew Citron, Adam Christopher Elwood, Geoffrey Hall, Gregory Michiel Iles, Christian Laner Ogilvy, Bjorn Penning, A Rose, Antoni Shtipliyski, Alexander Tapper, Timothy John Durkin, Kristian Harder, Sam Harper, Claire Shepherd-Themistocleous, Alessandro Thea, Thomas Stephen Williams, Sridhara Rao Dasu,

Laura Margaret Dodd, Pamela Renee Klabbers, Aaron Levine, Isabel Rose Ojalvo, Tyler Henry Ruggles, Nicholas Charles Smith, Wesley Smith, Ales Svetek, R Forbes, Jesra Lilah Tikalsky, and Marcelo Vicente. The CMS Level-1 Calorimeter Trigger for LHC Run II. Technical report, CERN, Geneva, 2017. URL `https://cds.cern.ch/record/2233176`.

[9] The Phase-2 Upgrade of the CMS Endcap Calorimeter. Technical report, CERN, Geneva, 2017. URL `https://cds.cern.ch/record/2293646`.

[10] Piyush Kumar. The Particle Flow Algorithm in the Phase II Upgrade of the CMS Level-1 Trigger. *PoS*, EPS-HEP2021:771, 2022. doi: 10.22323/1.398.0771. URL `https://cds.cern.ch/record/2827244`.

[11] Lorenzo Teofili, Jose Briz Monago, Marco Calviani, Nicolas Chritin, Jaakko Esala, Simone Gilardoni, Inigo Lamas Garcia, Jorge Maestre, Mauro Migliorati, Tobias Polzin, and Tatiana Rijoff. Design of the New Proton Synchrotron Booster Absorber Scraper (PSBAS) in the Framework of the Large Hadron Collider Injection Upgrade (LIU) Project. *J. Phys.: Conf. Ser.*, 1067(8):THPAK091, 2018. doi: 10.18429/JACoW-IPAC2018-THPAK091. URL `https://cds.cern.ch/record/2671912`.

[12] Daniel Jean Simon. The CERN PS complex: a versatile particle factory. 1996. URL `https://cds.cern.ch/record/305985`.

[13] P. Collier and B. Goddard. The sps as lhc injector. In *PACS2001. Proceedings of the 2001 Particle Accelerator Conference (Cat. No.01CH37268)*, volume 4, pages 3150–3152 vol.4, 2001. doi: 10.1109/PAC.2001.988034.

[14] Jasper Kirkby, Antonio Amorim, Urs Baltensperger, Kenneth Carslaw, Theodoros Christoudias, Joachim Curtius, Neil Donahue, Imad Haddad, Richard Flagan, Hamish Gordon, Armin Hansel, Hartwig Harder, Heikki Junninen, Markku Kulmala, Andreas Kürten, Ari Laaksonen, Katrianne Lehtipalo, Jos Lelieveld, Ottmar Möhler, and Douglas Worsnop. Atmospheric new particle formation from the cern cloud experiment. *Nature Geoscience*, 16:948–957, 11 2023. doi: 10.1038/s41561-023-01305-0.

[15] P. Abbon et al. The COMPASS experiment at CERN. *Nucl. Instrum. Meth. A*, 577:455–518, 2007. doi: 10.1016/j.nima.2007.03.026.

[16] Werner Herr and B Muratori. Concept of luminosity. 2006. doi: 10.5170/CERN-2006-002.361. URL `https://cds.cern.ch/record/941318`.

[17] C. Y. Wong. *Introduction to high-energy heavy ion collisions*. 1995. ISBN 978-981-02-0263-7.

[18] G. Antchev et al. The TOTEM detector at LHC. *Nucl. Instrum. Meth. A*, 617: 62–66, 2010. doi: 10.1016/j.nima.2009.08.083.

[19] P. Lecomte. Large lead tungstate crystals for calorimetry. *IEEE Transactions on Nuclear Science*, 49(2):583–587, 2002. doi: 10.1109/TNS.2002.1003678.

[20] Images of CMS ECAL Endcap (EE), 2008. URL `https://cds.cern.ch/record/1431479`. CMS Collection.

[21] A. Sirunyan et al. Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ TeV. *JINST*, 16:P07001, 2021. doi: 10.1088/1748-0221/16/07/P07001. URL `https://cds.cern.ch/record/2751483`. Replaced with the published version. Added the journal reference and the DOI. All the figures and tables can be found at http://cms-results.web.cern.ch/cms-results/public-results/publications/MUO-19-001 (CMS Public Pages).

[22] Hyejin Kwon. Performance of the cms level-1 trigger during run 2. page 806, 02 2021. doi: 10.22323/1.390.0806.

[23] Micro telecommunications computing architecture short form specification. `https://www.picmg.org/wp-content/uploads/MicroTCA_Short_Form_Sept_2006.pdf`, . [Accessed 19-03-2024].

[24] Vadatech microtca overview. `https://www.vadatech.com/media/pdf_MicroTCA_Overview.pdf`, . [Accessed 19-03-2024].

[25] Ultrascale architecture and product data sheet: Overview. `https://docs.amd.com/v/u/en-US/ds890-ultrascale-overview`, . [Accessed 19-03-2024].

[26] The Phase-2 Upgrade of the CMS Level-1 Trigger. Technical report, CERN, Geneva, 2020. URL `https://cds.cern.ch/record/2714892`. Final version.

[27] Longer term lhc schedule. `http://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm`, . [Accessed 19-03-2024].

[28] Ricky Frazier, S Fayer, G Hall, C Hunt, Gregory Iles, D Newbold, and Andrew Rose. A demonstration of a time multiplexed trigger for the cms experiment. *Journal of Instrumentation*, 7:C01060, 01 2012. doi: 10.1088/1748-0221/7/01/C01060.

[29] R Tramontano. Precision Timing with the CMS MTD Barrel Timing Layer for HL-LHC. *Nuovo Cimento C*, 44(2-3):80, 2021. doi: 10.1393/ncc/i2021-21080-y. URL `https://cds.cern.ch/record/2791322`.

[30] N. Guettouche, S. Baron, S. Biereigel, D. Hernandez Montesinos, S. Kulis, P. Vicente Leitao, J. Mendez, P. Moreira, D. Porret, and K. Wyllie. The lpGBT production testing system. *JINST*, 17(03):C03040, 2022. doi: 10.1088/1748-0221/17/03/C03040.

[31] The Phase-2 Upgrade of the CMS Barrel Calorimeters. Technical report, CERN, Geneva, 2017. URL `https://cds.cern.ch/record/2283187`. This is the final version, approved by the LHCC.

[32] J-M Andre, U Behrens, J Branson, P Brummer, O Chaze, S Cittolin, C Contescu, B G Craigs, G-L Darlea, C Deldicque, Z Demiragli, M Dobson, N Doualot, S Erhan, J F Fulcher, D Gigi, M Gładki, F Glege, G Gomez-Ceballos, J Hegeman, A Holzner, M Janulis, R Jimenez-Estupiñán, L Masetti, F Meijers, E Meschi, R K Mommsen, S Morovic, V O'Dell, L Orsini, C Paus, P Petrova, M Pieri, A Racz, T Reis, H Sakulin, C Schwick, D Simelevicius, and P Zejdl. The cms data acquisition - architectures for the phase-2 upgrade. *Journal of Physics: Conference Series*, 898 (3):032019, oct 2017. doi: 10.1088/1742-6596/898/3/032019. URL `https://dx.doi.org/10.1088/1742-6596/898/3/032019`.

[33] CMS Collaboration. The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger. Technical report, CERN, Geneva, 2021. URL `https://cds.cern.ch/record/2759072`. This is the final version of the document, approved by the LHCC.

[34] 7 series fpgas data sheet: Overview. `https://docs.amd.com/v/u/en-US/ds180_7Series_Overview`, . [Accessed 19-03-2024].

[35] Ultrascale architecture and product data sheet: Overview. `https://docs.amd.com/v/u/en-US/ds890-ultrascale-overview`, . [Accessed 19-03-2024].

[36] Picmg 3 . 2 advanced telecommunications and computing architecture infiniband technology delivers performance , scalability , and fault tolerance for next generation platforms. . URL `https://api.semanticscholar.org/CorpusID:13862443`.

[37] Axi protocol overview. `https://developer.arm.com/documentation/102202/0300/AXI-protocol-overview`, . [Accessed 18-03-2024].

[38] Xilinx stacked silicon interconnect technology delivers breakthrough fpga capacity, bandwidth, and power efficiency. `https://docs.amd.com/v/u/en-US/wp380_Stacked_Silicon_Interconnect_Technology`, . [Accessed 19-03-2024].

[39] Stephanie Baffioni, Claude Charlot, Federico Ferri, David Futyan, Paolo Meridiani, Ivica Puljak, Chiara Rovelli, Roberto Salerno, and Yves Sirois. Electron

reconstruction in CMS. Technical report, CERN, Geneva, 2006. URL `https://cds.cern.ch/record/934070`.

[40] Piyush Kumar and Bhawna Gomber. The cms level-1 calorimeter trigger for the hl-lhc. *Instruments*, 6(4), 2022. ISSN 2410-390X. doi: 10.3390/instruments6040064. URL `https://www.mdpi.com/2410-390X/6/4/64`.

[41] Qi Mu, Liqing Cui, and Yufei Song. The implementation and optimization of bitonic sort algorithm based on cuda. 06 2015.

[42] Jack E. Volder. The cordic trigonometric computing technique. *IRE Transactions on Electronic Computers*, EC-8(3):330–334, 1959. doi: 10.1109/TEC.1959.5222693.

System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC

by Piyush kumar

Submission date: 21-Mar-2024 05:25PM (UTC+0530)
Submission ID: 2326735862
File name: Piyush_Kumar.pdf (25.64M)
Word count: 33855
Character count: 170530

System Design and Prototyping of the CMS Level-1 Calorimeter Trigger at the High-Luminosity LHC

ORIGINALITY REPORT

13%	8%	12%	2%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	P. Kumar, B. Gomber. "System Design and Prototyping for the CMS Level-1 Trigger at the High-Luminosity LHC", IEEE Transactions on Nuclear Science, 2023 Publication	3%
2	Michail Bachtis. "Chapter 3 Experimental Setup", Springer Science and Business Media LLC, 2014 Publication	1%
3	mdpi-res.com Internet Source	1%
4	www.pure.ed.ac.uk Internet Source	1%
5	Alexandre Zabi. "Design of the CMS calorimeter trigger upgrade from Phase I to Phase II of the LHC", Journal of Physics: Conference Series, 2019 Publication	<1%
6	research-information.bris.ac.uk Internet Source	<1%

Dr. Bhawna Gomber
Assistant Professor
CASEST, School of Physics
University of Hyderabad
Hyderabad-500046, Telangana.

This is from Student Publication
Gomber 22/3/2024

M. Thamarasan 22/3/24
(Centre for Advanced Studies in Electronics Science & Technology)
School of Physics

7	Svetek, A., M. Blake, M. Cepeda Hermida, S. Dasu, L. Dodd, R. Fobes, B. Gomber, T. Gorski, Z. Guo, P. Klabbers, A. Levine, I. Ojalvo, T. Ruggles, N. Smith, W.H. Smith, J. Tikalsky, M. Vicente, and N. Woods. "The Calorimeter Trigger Processor Card: the next generation of high speed algorithmic data processing at CMS", Journal of Instrumentation, 2016. Publication	<1 %
8	Cristina Martin Perez. "Search for the Higgs Boson Produced in Association with Top Quarks with the CMS Detector at the LHC", Springer Science and Business Media LLC, 2022 Publication	<1 %
9	digital.library.wisc.edu Internet Source	<1 %
10	Csaki, Csaba, and Scott Dodelson. "INTRODUCTION TO THE LHC EXPERIMENTS", Physics of the Large and the Small, 2011. Publication	<1 %
11	Submitted to Kyungpook National University Student Paper	<1 %
12	olympias.lib.uoi.gr Internet Source	<1 %
	moam.info	

13	Internet Source	<1 %
14	hdl.handle.net Internet Source	<1 %
15	Piyush Kumar, Bhawna Gomber. "The CMS Level-1 Calorimeter Trigger for the HL-LHC", Instruments, 2022 Publication	<1 %
16	cerncourier.com Internet Source	<1 %
17	othes.univie.ac.at Internet Source	<1 %
18	thesis.library.caltech.edu Internet Source	<1 %
19	charma.uprm.edu Internet Source	<1 %
20	cds.cern.ch Internet Source	<1 %
21	theses.hal.science Internet Source	<1 %
22	docplayer.net Internet Source	<1 %
23	etd.adm.unipi.it Internet Source	<1 %

Dr. Bhawna Gomber
Assistant Professor
CASEST, School of Physics
University of Hyderabad
Hyderabad-500046, Telangana.

B Gomber 2/3/2024

This is from student publications

24	Submitted to University of Hyderabad, Hyderabad Student Paper	<1 %
25	ufdcimages.uflib.ufl.edu Internet Source	<1 %
26	"Search for Contact Interactions in the Dimuon Channel in pp Collisions at $\sqrt{s} = 7$ TeV at CMS", 'Office of Scientific and Technical Information (OSTI)' Internet Source	<1 %
27	Vasileios Amoiridis, Ulf Behrens, Andrea Bocci, James Branson et al. "Progress in Design and Testing of the DAQ and Data-flow Control for the Phase-2 Upgrade of the CMS Experiment", IEEE Transactions on Nuclear Science, 2023 Publication	<1 %
28	Wu. "Status of CMS", Springer Proceedings Physics, 2006 Publication	<1 %
29	repositorium.uminho.pt Internet Source	<1 %
30	Submitted to HTM (Haridus- ja Teadusministeerium) Student Paper	<1 %
31	iihe.ac.be Internet Source	<1 %

32	publikationen.bibliothek.kit.edu Internet Source	<1 %
33	siba.unipv.it Internet Source	<1 %
34	www2.mdpi.com Internet Source	<1 %
35	A. Zabi, F. Beaudette, L. Cadamuro, O. Davignon et al. "The CMS Level-1 Calorimeter Trigger for the LHC Run II", Journal of Instrumentation, 2017 Publication	<1 %
36	Antel, Claire. "Enhancing low mass dark matter mediator resonance searches with improved triggering in the ATLAS detector", 'American Association for Cancer Research (AACR)', 2019 Internet Source	<1 %
37	Farmahini-Farahani, Amin, Henry J. Duwe III, Michael J. Schulte, and Katherine Compton. "Modular Design of High-throughput, Low-latency Sorting Units", IEEE Transactions on Computers, 2012. Publication	<1 %
38	indico.cern.ch Internet Source	<1 %
39	m.moam.info Internet Source	

		<1 %
40	Luca Cadamuro. "Search for Higgs Boson Pair Production in the $b\bar{b}$ τ+ τ- Decay Channel", Springer Science and Business Media LLC, 2018 Publication	<1 %
41	dokumen.pub Internet Source	<1 %
42	web.itu.edu.tr Internet Source	<1 %
43	www.universetoday.com Internet Source	<1 %
44	Jonathan Fulcher, Joschka Lingemann, Dinyar Rabady, Thomas Reis, Hannes Sakulin. "The New Global Muon Trigger of the CMS Experiment", IEEE Transactions on Nuclear Science, 2017 Publication	<1 %
45	Kar, Tamasi Rameshchandra. "A Triplet Track Trigger for Future High Rate Collider Experiments", 'American Association for Cancer Research (AACR)', 2020 Internet Source	<1 %
46	scholar.colorado.edu Internet Source	<1 %

47	Submitted to Cranfield University Student Paper	<1 %
48	repozitorij.pmfst.unist.hr Internet Source	<1 %
49	"Proceedings of International Conference on Technology and Instrumentation in Particle Physics 2017", Springer Science and Business Media LLC, 2018 Publication	<1 %
50	Anders Ryd, Louise Skinnari. "Tracking Triggers for the HL-LHC", Annual Review of Nuclear and Particle Science, 2020 Publication	<1 %
51	D Futyan. "Search for the standard model Higgs boson in the two-electron and two-muon final state with the CMS detector", Journal of Physics G Nuclear and Particle Physics, 09/01/2007 Publication	<1 %
52	Submitted to Higher Education Commission Pakistan Student Paper	<1 %
53	Submitted to Indian Institute of Technology, Kharagpure Student Paper	<1 %
54	digitalcommons.unl.edu Internet Source	<1 %

55	web.physik.rwth-aachen.de Internet Source	<1 %
56	www.jstage.jst.go.jp Internet Source	<1 %
57	acikbilim.yok.gov.tr Internet Source	<1 %
58	Fasanella, Daniele <1981>(Rossi, Antonio Maria). "Study of the X(3872) state with the CMS experiment at LHC", Alma Mater Studiorum - Università di Bologna, 2012. Publication	<1 %
59	Springer Theses, 2016. Publication	<1 %
60	www.artscroll.com Internet Source	<1 %
61	s3.cern.ch Internet Source	<1 %
62	Biplob Bhattacherjee, Swagata Mukherjee, Rhitaja Sengupta, Prabhat Solanki. "Triggering long-lived particles in HL-LHC and the challenges in the first stage of the trigger system", Journal of High Energy Physics, 2020 Publication	<1 %
63	Taek-Jun Kwon, J. Sondeen, J. Draper. "Design Trade-Offs in Floating-Point Unit Implementation for Embedded and	<1 %

Processing-In-Memory Systems", 2005 IEEE International Symposium on Circuits and Systems, 2005
Publication

64	etheses.bham.ac.uk Internet Source	<1 %
65	inspirehep.net Internet Source	<1 %
66	pure.uva.nl Internet Source	<1 %
67	www.zora.uzh.ch Internet Source	<1 %

Exclude quotes	On	Exclude matches	< 14 words
Exclude bibliography	On		