

Low-Latency Track Triggering in High-Energy Physics

Zur Erlangung des akademischen Grades eines
DOKTORS DER INGENIEURWISSENSCHAFTEN (DR.-ING.)
von der KIT-Fakultät für Elektrotechnik und Informationstechnik
des Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

LUIS EDUARDO ARDILA PÉREZ, M.SC.

geboren in Bogotá, Kolumbien

Tag der mündlichen Prüfung: 29.10.2021
Hauptreferent: Prof. Dr. rer. nat. Marc Weber
Korreferent: Prof. Dr.-Ing. Dr. h. c. Jürgen Becker

This work is licensed under a [Creative Commons](#) “Attribution-NonCommercial-ShareAlike 4.0 International” license.



Abstract

The Compact Muon Solenoid (CMS) is a general-purpose experiment at the Large Hadron Collider (LHC) designed to study a wide variety of high-energy physics phenomena. It employs a large silicon tracker within a homogeneous 3.8 T magnetic field, which allows the precise measurement of the trajectories, including transverse momentum (p_T) and vertex position reconstruction of the charged particles emerging from the LHC collisions.

The High Luminosity (HL) upgrade of the LHC in 2025 will increase the simultaneous proton-proton collisions from the current average of 25 to up to 200 every 25 ns. The upgrade will completely replace the silicon tracker with one purposely built to discriminate on-module the charged particles whose p_T is larger than 2 GeV, these hits are called ‘stubs’. The stubs are forwarded to off-detector electronics for real-time track reconstruction under 4 μ s of latency. For the first time in any particle physics experiment, the reconstructed tracker primitives will be included in the first-level trigger with the aim of maintaining the trigger rate of CMS below 750 kHz.

This thesis describes various firmware and hardware developments for a real-time all FPGA-based track finder that employs a regionally segmented and fully time-multiplexed architecture. The Time-multiplexed Track Trigger (TMTT) reconstruction algorithm has four processing stages, two of which were implemented in Hardware Description Language (HDL) by the author and are detailed in this dissertation. Optimizations of such algorithms for increased clock frequency operation and resource utilization optimization are also presented. In addition, the development of specialized hardware utilizing the Advanced Telecommunications Computing Architecture (ATCA) form factor will be presented. The board has sufficient high-speed I/O to be used at the HL CMS tracker off-detector processing system. It implements a novel slow-control solution for ATCA systems by combining the Intelligent Platform Management Controller (IPMC), a Linux slow-control software, and an FPGA for custom slow-control tasks in a single Zynq Ultrascale+ (US+) System-on-Chip (SoC) module.

Zusammenfassung

Das Compact Muon Solenoid (CMS) ist ein Experiment am Large Hadron Collider (LHC), mit dem eine Vielzahl von Phänomenen der Hochenergiephysik untersucht werden soll. Es verwendet einen großen Silizium Spurdetektor innerhalb eines homogenen 3,8 T -Magnetfelds, der die genaue Messung der Trajektorien, des Transversalimpulses (p_T) und die Rekonstruktion der Scheitelposition der geladenen Teilchen aus den LHC-Kollisionen ermöglicht.

Die Aufrüstung des LHC für hohe Luminosität (HL) im Jahr 2025 wird die Zahl der gleichzeitigen Proton-Proton-Kollisionen von derzeit durchschnittlich 25 auf bis zu 200 alle 25 ns erhöhen. Bei der Aufrüstung wird der Silizium Spurdetektor vollständig durch einen ersetzt, der speziell dafür gebaut wurde, die geladenen Teilchen auf den Modulen zu selektieren, deren p_T höher als 2 GeV ist. Die Informationen dieser Treffer werden an eine Off-Detektor-Elektronik zur Echtzeit-Spur-Rekonstruktion mit einer Latenz von 4 μ s weitergeleitet. Die rekonstruierten Tracker-Primitive werden erstmalig in einem Teilchenphysikexperiment in die First-Level-Trigger-Entscheidung einbezogen, mit dem Ziel, die Triggerrate von CMS unter 750 kHz zu halten.

In dieser Doktorarbeit werden verschiedene Firmware- und Hardware-Entwicklungen für einen FPGA-basierten Echtzeit-Spurensucher vorgestellt, der eine räumlich segmentierte und vollständig zeitmultiplexierte Architektur verwendet. Der Time-multiplexed Track Trigger (TMTT)-Rekonstruktionsalgorithmus enthält vier Verarbeitungsstufen, von denen zwei vom Autor in Hardwarebeschreibungssprache (HDL) implementiert wurden und in dieser Dissertation detailliert beschrieben werden. Optimierungen solcher Algorithmen, zur Erhöhung der Taktfrequenz und zur Reduktion der Ressourcennutzung, werden ebenfalls vorgestellt. Darüber hinaus wird die Entwicklung spezieller Hardware auf Basis des ATCA-Formfaktor (Advanced Telecommunications Computing Architecture) vorgestellt. Das Board verfügt über ausreichend Hochgeschwindigkeits-I/O für den Einsatz im HL CMS Tracker Off-Detector Processing System. Es implementiert eine neuartige

Slow-Control-Lösung für ATCA Systeme, die den Intelligent Platform Management Controller (IPMC), eine Linux Slow-Control-Software und einen FPGA für anwendungsspezifische Slow-Control-Aufgaben in einem einzigen Zynq Ultrascale+ (US+) System-on-Chip (SoC)-Modul integriert.

Acknowledgments

I would like to take this opportunity to express the deepest appreciation to everyone who contributed to the development of this Thesis. I wish to specially thank my supervisors Prof. Marc Weber and Prof. Jürgen Becker for providing me with the opportunity to work on this challenging, yet very rewarding project. I would like to thank my receiving supervisors at RAL Claire Shepherd and Ian Tomalin, for providing an excellent working environment and for keeping the team highly motivated with the enjoyable and educational Friday's pub visit. I want to thank Michele Caselle for his early guidance and introduction to the CMS tracker working group and for his unreserved enthusiasm and stimulus. I want to thank Oliver Sander for the countless hours we spent debugging the system at the beat of the crate-fans at CERN and for the unconditional support and encouragement in almost every aspect. I thank Matthias Balzer for his prompt assistance, specially during the writing of this thesis. I wish to thank Mark Pesaresi for providing me with valuable advises and trusting in my abilities, your encouragement was unquestionably motivating. I thank Greg Iles for the numerous discussions and insights about high-speed optical transceivers, as well as the nitty-gritty detailed suggestions related to FPGA time closure, for always receiving me with a warm welcome and never hesitating in providing what was needed when I visited your lab. I must thank the staff members at KIT Denis, Michael, Uwe and Alex for sharing your valuable expertise and allowing me to learn from you. I want to thank my KIT colleagues Lorenzo, Marvin, Meghana, Nick, Patrick, Timo, Thomas, and Torben for their careful and persistent work effort and for their friendship. I also wish to thank my colleagues at RAL Luigi, Davide, Kostas and Tom for always being diligent when questions arose and for their friendship. I thank Brian and Sven for the interesting discussions which fostered new research directions for the future. Last but not least, I would like to thank Belen for her wholeheartedly loving support and encouragement, specially when things were complicated, and to my family for their immense multivalent guidance which motivates me to go further every day.

This research acknowledges the support by the DFG-funded Doctoral School "Karlsruhe School of Elementary and Astroparticle Physics: Science and Technology"

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement n° [317446] INFIERI "INtelligent Fast Interconnected and Efficient Devices for Frontier Exploitation in Research and Industry"

Declaration

I declare that I have developed and wrote the enclosed thesis entirely by myself, and that I did not use any sources or means that were not explicitly stated in the text. The algorithms and firmware developments in Chapter 4 are collaborative efforts of the community which led to the publications [1], [2] and the results shown in Chapter 6. My specific contributions to the TMTT algorithm and its firmware development are highlighted in Chapter 5. In a similar way Chapter 7 describes the overall community efforts related to the back-end hardware system and my specific contributions to the hardware R&D are highlighted in Chapters 8 and 9.

Contents

Abstract	III
Zusammenfassung	V
1 Introduction	1
2 The CMS experiment at the LHC	3
2.1 The Large Hadron Collider (LHC)	3
2.2 Experimental sites at the LHC	6
2.3 The Compact Muon Solenoid (CMS) Experiment	7
2.3.1 Coordinate System	9
2.3.2 The solenoid	10
2.3.3 Silicon tracker	10
2.3.4 Calorimeter	12
2.3.5 Muon system	16
2.3.6 Trigger and Data Acquisition	17
3 High-luminosity LHC	21
3.1 Physics Motivation	22
3.2 The HL-LHC Machine	23
3.2.1 Insertion region magnets	23
3.2.2 Collimation	24
3.2.3 Crab cavities	25
3.3 The CMS Upgrade	25
3.4 The CMS Tracker Upgrade	29
3.4.1 Requirements	29
3.4.2 The Inner Tacker	31
3.4.3 The Outer Tacker	32
3.5 Track Finding at the Level-1 Trigger	34
4 The Track Finding Algorithms	37
4.1 The Tracklet approach	37

4.1.1	Stub Organization: Layer Router - VM Router	38
4.1.2	Seeding: Tracklet Engine - Tracklet Calculator	38
4.1.3	Projections: Match Engine - Match Calculator	38
4.1.4	Fitting: Linearized χ^2 fit	39
4.1.5	Duplicate Removal: Purge Duplicate	39
4.2	The Time-Multiplexed Track Trigger (TMTT) approach	40
4.2.1	Sector Assignment: Geometric Processor	40
4.2.2	Track Finding: Hough Transform	41
4.2.3	Track Fitting: Kalman Filter	42
4.2.4	Duplicate Removal	43
4.3	The Hybrid approach	43
4.4	Summary	46
5	Track Finder Algorithm Contributions	47
5.1	The TMTT Geometric Processor (GP)	47
5.2	The Duplicate Removal Algorithms for the TMTT approach	56
5.2.1	Pair-Wise Track comparison Duplicate Algorithm	56
5.2.2	A Duplicate Removal Algorithm based on the Hough Transform (HT) parameter space	60
5.3	Summary	65
6	The TMTT Hardware Demonstrator	69
6.1	The Hardware Demonstrator Slice	69
6.2	Software Setup	71
6.3	Track Reconstruction Efficiency	72
6.4	Track Parameter Resolution	75
6.5	Data Rates	78
6.6	Flexibility and Robustness of the System	81
6.7	Latency	83
6.8	Field Programmable Gate Array (FPGA) Resource Usage	84
6.9	Summary	85
7	The CMS Tracker Back-End System	87
7.1	CMS Tracker Rack Configuration	87
7.2	The ATCA Shelf	89
7.2.1	Hardware Platform Management	90
7.2.2	Intelligent Platform Management Controller	90
7.3	The CMS DAQ and TCDS Hub (DTH) Hub Prototype	91
7.4	Tracker Hardware Development Platforms	93
7.4.1	Apollo	94

7.4.2	Serenity-Z	95
7.5	The EMP framework infrastructure firmware	96
7.6	Summary	99
8	Hardware R&D Contributions	101
8.1	FPGA Daughtercard for Serenity-Z	101
8.1.1	PCB Layout High-Speed Differential Lines Tuning	103
8.1.2	High-Speed Optical Transceiver Qualification	104
8.2	Trenz to Serenity Adapter	105
8.2.1	The Unified slow control architecture	105
8.2.2	Hardware components of the Trenz-Serenity adapter	106
8.2.3	IPMC Software Implementations	108
8.3	ATCA ZynqUS+ IPMC Test Board	111
8.4	OpenIPMC-HW DIMM	113
8.5	FMC+ Board for 25 Gb/s Optical Evaluation	116
8.6	Summary	121
9	Serenity-A2577 ATCA Board	123
9.1	Serenity-A2577 Architecture	123
9.2	Integrated Slow Control Management Modules	124
9.2.1	FMC+ version	126
9.2.2	CMX-EXT version	127
9.3	Thermal Analysis	128
9.4	MGT performance	131
9.4.1	25 Gb/s Optical Evaluation with Samtec x12 alpha-v2 parts	132
9.4.2	Bathtub analysis with 25 Gb/s Samtec x12 alpha-v2 parts	134
9.5	Timing Control and Distribution System (TCDS2)	137
9.6	Summary	140
10	Conclusion	141
	Acronyms	143
	List of Figures	154
	List of Tables	156
	Bibliography	171
	List of Publications	173

1 Introduction

The Large Hadron Collider (LHC) located at the European Organization for Nuclear Research (CERN) is the largest particle accelerator in the world. It was commissioned in 2010 delivering proton-proton collisions with a 7 TeV center-of-mass energy. Only two years later, in July 2012 it announced the discovery of the Higgs boson at about 125 GeV by both of its general-purpose experiments A Toroidal LHC Apparatus (ATLAS) [3] and the Compact Muon Solenoid (CMS) [4]. This important discovery was long predicted by the Standard Model (SM) of particle physics [5]–[9], which is currently the best theory to describe the elementary building blocks of the universe.

The LHC has however been designed to go beyond the SM in the search of new physics like the existence, or not, of supersymmetry, the nature of dark matter, and the existence of extra dimensions [10]. To further extend its discovery potential, the LHC is planning a major upgrade [11] to increase its number of collisions per area and time, the so called luminosity, by a factor of five beyond its design value. The integrated luminosity design goal has also increased by a factor of ten to 3000 fb^{-1} [10]. The LHC produces proton-proton (pp) collisions at a rate of 40 MHz, *i.e.* one bunch crossing every 25 ns. The typical event size during Run 1-2 for the ATLAS and CMS experiments was about 1 MB corresponding approximately to a data rate of 40 TB/s [12], for Phase-2 this value will increase to an expected event size of 7.4 MB for a total data rate of 296 TB/s [13], an amount impossible to be stored for later offline analysis.

High-energy Physics (HEP) experiments have a large experience acquiring and consequently reducing very large volumes of data. This is achieved by the use of complex and multiple stages of triggers, which perform different degrees of online data elaboration, with the aim of reducing the final data volume considerably [14]–[16], thus keeping only the events with potentially interesting physical phenomena. The expected trigger rate for the HL operations is 750 kHz which equates to about 50 Tb/s, the HLT is expected to reduce this to an output rate of 7.5 kHz to storage.

Under the challenging conditions of the HL operation, the CMS detector, particularly the silicon tracker, requires a data acquisition system with exceptional

performance to collect the detector data and analyze it in real-time, therefore contributing with track candidates to the Level-1 (L1) trigger. The upgrade of the trigger system will enhance the physics selectivity and maintain the performance necessary throughout the 10 year long HL-LHC program.

This dissertation includes various fundamental developments to demonstrate the feasibility of reconstructing tracks under the tight latency requirements of the HL CMS L1-trigger system. Chapter 5 shows specifically the contributions to the Time-multiplexed Track Trigger (TMTT) algorithm by the author, which were completely implemented in hardware, using last generation FPGAs as a demonstration system, whose tracking results are shown in Chapter 6. Furthermore, this dissertation also includes several hardware developments to contribute to the overall research and development program aimed at designing the outer tracker back-end electronics system, which is based on the Advanced Telecommunications Computing Architecture (ATCA) standard and uses next generation FPGAs and cutting-edge high-speed optical transceivers. The contributions developed under this thesis for the hardware back-end system are specified in Chapter 8 and Chapter 9.

2 The CMS experiment at the LHC

2.1 The Large Hadron Collider (LHC)

The LHC is located at the border between France and Switzerland, nearby Geneva in a 26.7 km long tunnel at 100 m below the surface. It was built on the same tunnel that was previously used by the Large Electron-Positron Collider (LEP) [17]. The Lorentz force, in Equation 2.1, is responsible for particle acceleration and deflection. When charged particles travel in curved paths, synchrotron radiation is emitted, therefore decreasing the energy of the beam. The energy loss is inversely proportional to the mass ($\Delta E \propto \frac{E^4}{m^4}$). For this reason, at the LHC a much greater center-of-mass energy can be achieved by using hadrons instead of electrons or positrons like at the LEP.

$$\dot{\vec{p}} = -q \cdot (\vec{E} + \vec{v} \times \vec{B}) \quad (2.1)$$

Protons at the LHC need to be accelerated in stages using the pre-accelerator complex shown in Figure 2.1. First, protons are extracted from an hydrogen source, ionized, and accelerated using Radio Frequency (RF) cavities to 50 MeV by the LINAC-2 linear accelerator. In 2020, as part of the upgrades for Run 3, LINAC-2 was replaced by LINAC-4 depicted in the Figure 2.1. The protons then enter the Proton Synchrotron Booster raising their energy to 1.4 GeV before they are transferred to the Proton Synchrotron which has a circumference of 628 m and uses conventional electromagnets for bending and focusing the beam. The Proton Synchrotron further raises the energy of the beam to 25 GeV and gives it the required bunch structure. Then, the beam is injected to the Super Proton Synchrotron (SPS), which has a circumference of about 7 km, further raising the energy to 450 GeV before it is transferred to the LHC ring in two opposing beams. At the LHC ring there are 2808 proton packets per particle beam, each of which contains about 10^{11} protons. The protons travel almost at the speed of light and have a collision frequency of 40 MHz with an orbital frequency of around 11 kHz [11].




Figure 2.1: Schematic representation of the LHC with the four detectors: CMS, LHCb, ATLAS, and ALICE. The pre-accelerators complex is also shown containing the Proton Synchrotron Booster, Proton Synchrotron and the Super Proton Synchrotron (SPS) [18].

In order to keep particles with these high energies in the storage ring under control, the LHC is instrumented with 1232 dipole and 392 quadrupole superconducting magnets operating at a temperature of 1.9 K and each producing about 8 T [19]. The LHC ring is composed by eight straight sections, four of them are cross points of the beams equipped with experiments, and eight arcs. The primary job of the magnets is to bend the particle beam at those locations where arcs are present, and to focus the beam into compact packets to increase the probability of collision. The quadrupole magnets, shown in Figure 2.2, are used to focus the beam transversely. The magnets act like a lens on light rays and focus the particle beam. A quadrupole magnet has a focusing effect in one direction, while it has a defocusing effect in the other direction. In order to achieve an even higher focus, several quadrupole magnets one behind the other, in the direction of flight, focus the particle beam, thus reaching a high level of luminosity.




Figure 2.2: a) Model of a superconducting quadrupole magnet for the LHC project [20]. b) Focusing magnet in cross section [21]. The blue arrows indicate the force acting on a positive charge moving into the plane.

The luminosity is an important parameter for the performance of a particle accelerator. It describes the interactions of the particles in a time interval and area. It is specially useful to quantify the event rate of rare phenomena with small production cross sections. There is a direct relationship between the event rate $\dot{N}$, the interaction cross section σ and the luminosity L .

$$\dot{N} = \sigma \cdot L \quad (2.2)$$

A storage ring has a luminosity of:

$$L = \frac{N_1 \cdot N_2 \cdot n \cdot f}{A}, \quad (2.3)$$

where N_1 and N_2 are the number of particles in the two opposing bunches, n is the number of colliding packets, f is the frequency of the collisions and A is the cross-sectional area of the particle bunches. In order to obtain a high luminosity, the minimization of the cross-sectional area of the bunches is of enormous importance [22].

An often used quantity in storage ring experiments is the integrated luminosity.

$$L_{int} = \int L dt \quad (2.4)$$

The LHC was designed for a total center-of-mass energy of $\sqrt{s} = 14$ TeV and a peak instantaneous luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Figure 2.3 shows the total integrated luminosity vs. time for the whole operation of the LHC. During the data taking period Run 1, between 2010 and 2012, the total integrated luminosity reached about 30 fb^{-1} while the center-of-mass energy was $\sqrt{s} = 7$ TeV for the first two years and $\sqrt{s} = 8$ TeV in 2012. With Run 2, between 2015 and 2018, the total integrated luminosity reached almost 200 fb^{-1} . In Run 2 the collider was operated at a center-of-mass of $\sqrt{s} = 13$ TeV.




Figure 2.3: Cumulative delivered and recorded luminosity versus time for 2010-2012 and 2015-2018 during stable beams for pp collisions at nominal center-of-mass energy [23].

2.2 Experimental sites at the LHC

Several experiments are distributed along the ring of the LHC, the four principal ones are located exactly at the crossing points of the beams, two of the biggest experiments are A Toroidal LHC Apparatus (ATLAS) [16] and the Compact Muon Solenoid (CMS), explained in more detail in Section 2.3. Both of them are general-purpose detectors aimed to investigate different fundamental physics questions and complement the findings using measurements from each other. ATLAS is build around a hybrid magnet system which features a central solenoid and several toroids around it. Overall it has an impressive $46 \times 25 \text{ m}^2$ in size. The ATLAS trigger system [14] is arranged in three distinct stages: the Level-1, based on specialized hardware, reduces the detector readout to 75 kHz while its trigger decision needs to reach the front-end detectors under $2.5 \mu\text{s}$; the Level-2, designed to reduce the trigger rate to approximately 3.5 kHz with an event processing time of about 40 ms; and the ‘event filter’ which reduces the event rate to roughly 200 Hz with an event

size of approximately 1.3 MB [16]. The L2 and event filter are based on a large CPU farm, together are known as the High-Level Trigger (HLT).

The Large Hadron Collider beauty (LHCb) [24] experiment, as opposed to ATLAS and CMS, does not enclose completely the collision point with detector material but uses several layers of detectors to only cover the forward region. The first layer of detectors is located close to the interaction point. The other layers are assembled in only one direction in z with a total length of 20 meters. It is designed to specifically study the slight differences between matter and antimatter by focusing on the study of the beauty quark [25]. Initially, two trigger levels were applied at LHCb: a Level-0 using custom high-speed electronics operating synchronously with the machine clock to reduce the event rate down to 1 MHz; and a HLT trigger which is executed on a processor farm. After the Phase-1 upgrades in 2019-2021, the LHCb experiment is using a triggerless readout and a full software trigger in two stages, the HLT1 and HLT2 [26].

The LHC ring is capable of accelerating heavy ions in addition to protons. In such specialized runs one experiment called A Large Ion Collider Experiment (ALICE) [27] is used to primarily investigate the Quark-Gluon-Plasma (QGP) using lead ion collisions. The QGP is an extreme phase of matter that our universe is thought to have been just a few millionths of a second after the Big Bang, right before quarks and gluons were bound together to form protons and neutrons [28]. ALICE electronics are not pipelined in general and three levels of trigger systems are used to reduced the event rate: a Level-0 with a latency of $1.2\ \mu\text{s}$, a Level-1 at $6.5\ \mu\text{s}$, and a Level-2 after $100\ \mu\text{s}$ implemented once data from the Time Projection Chamber (TPC) are received. Finally, a layer of software-based HLT triggers are run in a CPU farm.

2.3 The Compact Muon Solenoid (CMS) Experiment

The CMS is a particle physics experiment at the LHC. Just like its counterpart, the ATLAS experiment, it is a general-purpose, cylindrical symmetric particle detector. The main design goals of CMS include carrying out precision measurements of the standard model, such as measuring the Higgs self-coupling, which influences the Higgs potential. The range of tasks of CMS include the implementation of Quantum Chromodynamics (QCD) tests, or the search for decay channels with four top quarks in the final state [29]. However, the research is not limited to the standard model. Future measurements will search for supersymmetrical coupled particles [30] with a τ -lepton [31] and other space-time dimensions. The search for dark photons produced by the decay of a Higgs boson in association with a Z boson is a new




Figure 2.4: Modeled cutaway view of the CMS detector [34].

approach to study dark matter [32]. Finally, CMS also contributes to the analysis of heavy ions collisions [33].

The data discussed in this section refers to the current CMS detector. The changes to the CMS detector after its upgrade in the year 2025 are shown in Chapter 3. The structure of the detector, shown in Figure 2.4, is as follows: starting from the center, the silicon tracker is located very close to the Interaction Point (IP). It uses two types of detectors, pixelated or strip detectors with a total sensitive area of 200 m^2 . Next, the crystal electromagnetic calorimeter is found using several PbWO_4 scintillating crystals. Immediately after, the hadron calorimeter is located using another type of scintillator material based on a polymer. The calorimeter detectors are enclosed by the next construction layer of the experiment and one of the fundamental components, the solenoid magnet, which uses a niobium titanium core. Finally, outside the solenoid and enclosed inside the steel return yoke, the muon drift chambers are located.

The purpose of the CMS detector is to measure the decay products resulting from proton-proton collisions. The CMS detector measures mainly photons, jets, hadrons, muons, electrons, and the missing transverse impulse. This is possible by reconstructing tracks with a transverse momentum of over $p_T > 1 \text{ GeV}$ in the

pseudorapidity range of $|\eta| < 2.5$. The total length of the detector is 21.6 m and it has a diameter of 14.6 m. With a total mass of 14 000 t, it is very compact, compared to detectors of similar weight. Another component of the CMS experiment is the trigger system. It is responsible for selecting important events containing relevant physics phenomena, and prompting for its readout from the detectors and storage by the DAQ system. The enormous amount of data that arises from proton-proton collisions can neither be evaluated nor read out in its entirety. For this reason, it is absolutely necessary to select only the interesting events. At the moment it is possible to write to permanent storage with a data transfer rate of 1 Gb/s and a trigger rate of 1 kHz.

2.3.1 Coordinate System

A right-handed coordinate system shown in Figure 2.5 is used on the LHC and therefore CMS. The x -axis points in the direction of the center of the LHC accelerator, the y -axis points vertically upwards and the z -axis runs along the beam axis. The origin of the coordinate system is the center of the CMS detector. As the CMS detector is symmetric around the beam axis, a cylindrical coordinate system is often used. Here φ is the azimuthal angle in the x - y plane. The distance from the beam axis is denoted by r . In addition, the angle θ to the beam axis is introduced. The pseudorapidity, expressed in Equation 2.5, is a dimensionless quantity and serves as a measure for the angle between the flight direction and the jet axis. The number of particles produced per pseudorapidity interval is approximately constant, which is why the pseudorapidity η is preferred to the angle θ .

$$\eta = -\ln \tan \left(\frac{\theta}{2} \right) \quad (2.5)$$




Figure 2.5: CMS coordinate system [35].

2.3.2 The solenoid

The CMS detector is built around a solenoid magnet, with a 6 m in diameter coil formed by superconducting fibers. It is able to generate a homogeneous 3.8 T magnetic field when about 18 000 A are circulating through it, the magnetic field is used to measure the momentum of charged particles. The solenoid is the largest magnet of its type ever constructed. The windings of the solenoid are arranged in four layers of a superconductor made from a Rutherford-type cable using Niobium-Titanium (NbTi) that is co-extruded with Aluminum. Outside the solenoid, the generated magnetic field is enclosed by the iron yoke. In this way, a homogeneous magnetic field of 1.8 T in the central area and an inhomogeneous field of 2.5 T at the ends of the magnet are generated.

2.3.3 Silicon tracker

The CMS tracker enables the precise measurement of the trajectories of charged particles emerging from the IP. Moreover, it can also determine the momentum with high accuracy by leveraging the homogeneous magnetic field produced by the solenoid and the track curvature as parameters. The tracker modules are constructed using silicon as the sensor material. Silicon sensors enable very fine granularity and fast signal readout, both of which are required to identify and disentangle the tracks of all charged particles produced in the proton-proton collisions. The high particle flux passing through the tracker material requires the use of radiation-hard sensors. This is accomplished by implementing an n+ pixel on n- substrate design, which allows for partial depletion even in close proximity to the beam pipe. Charged particles passing through a semiconducting silicon detector produce electron-hole pairs, resulting in an electric signal that can be readout by the detector front-end electronics. Figure 2.6 depicts an overview of all silicon tracker modules. Overall, the tracker measures 5.8 m in length and 2.2 m in diameter, it is arranged to cover the pseudorapidity range of $|\eta| < 2.5$. The tracking system is divided into two sections: the pixel detector on the inside and the strip detector on the outside. Finally, the entire tracker is run at a temperature of -20°C to help protecting it from the negative effects of radiation damage.

Pixel Detector

The pixel detector is the innermost detector of the CMS experiment. The initial design was made up of three cylindrical layers of silicon pixel modules with radius of 4.4 cm, 7.3 cm, and 10.2 cm that surround the beam pipe. Each layer is composed




Figure 2.6: Layout of the CMS Tracker, showing the $\frac{1}{4}$ view of the r - z plane [36]. Collisions occur at coordinate $(0, 0)$, with the two beams traveling from the left and right of the collision point. Bold lines represent double sided module assemblies.

of 768 pixel modules, where each pixel has an area of $100 \times 150 \mu\text{m}^2$. In addition to the three barrel layers, each side has two endcap disks with silicon pixel modules. They have a 6 cm inner radius and a 15 cm outer radius. The endcap disks are positioned at $z = 34.5$ cm and $z = 46.5$ cm from the center of the detector. Each disk has a geometry with 24 blades that are rotated by 20° . Each blade contains 7 pixel modules, for a total of 672 modules in the endcaps.

Since the pixel detector is the closest to the beam pipe, it received the most radiation damage and was replaced at the end of 2016 with one having four layers in the barrel region located at 2.9 cm, 6.8 cm, 10.9 cm, and 16.0 cm [37]. Similarly, the endcap region of the pixel detector is now a three layer construction located at $z = 29.1$ cm, $z = 39.6$ cm, and $z = 51.6$ cm measured from the center of the detector. The new inner tracker increased the number of pixels from 66 millions to 124 millions. It has a lower material budget and occupies a larger volume. Its innermost layer is situated closer to the interaction point, improving significantly the primary vertex resolution efficiency. Its outermost layer is closer to the strip detector, improving the overall tracking performance of the experiment by reducing the combinatorics during track reconstruction [38].

Strip Detector

The outermost component of the tracking system are the silicon strip detectors. Over 15 000 modules, each with a single-sided silicon strip, are used to construct the system. The strip detector has nearly 10 million readout channels in total. It is

divided into several subsystems, each of which covers a different region of radius and pseudorapidity in the detector layout as seen in Figure 2.6. The overall strip detector has the capability of identifying tracks with pseudorapidities $|\eta| < 2.4$.

According to the Figure 2.6, the outer tracker is comprised by the Tracker Inner Barrel (TIB), the Tracker Inner Disks (TID), the Tracker Outer Barrel (TOB), and the Tracker EndCaps (TEC). The TIB and TID are both located in the region of radii between 20 cm and 55 cm in diameter. The TIB is composed by four layers of modules and has a $|z| < 65$ cm. For the modules in the inner two layers, two modules are always mounted back-to-back with an angle of 100 mrad between their strips, these modules are shown in the Figure 2.6 with bolded lines. This enables dual measurements in both the r - z and the r - x directions at the same time. The TID is made up of three disks on each side, where the dual modules are installed on the inner two rings of each disk. The TID occupy the range $65 \text{ cm} < |z| < 118 \text{ cm}$. The TOB is made up of six concentric layers, the inner two are capable of allowing dual measurements. It is located within $|z| < 118 \text{ cm}$ and radii 55 cm and 116 cm. Finally, the forward regions TEC are covered by nine disks on each side of the detector. It occupies the area between $22.5 < r < 113.5 \text{ cm}$ and $124 < |z| < 282 \text{ cm}$. Dual modules can be found on the two innermost rings and in the fifth ring of TEC.

TID and TEC modules have strips that point towards the beam pipe, whereas the barrel modules (TIB and TOB) have strips that are parallel to the beam pipe direction. Different strip pitches have been used in different parts of the detector. In the first two layers of TIB the pitch size is $80 \mu\text{m}$, resulting in a single point resolution of $23 \mu\text{m}$, in the layers 3 and 4 of TIB the pitch size is $120 \mu\text{m}$ for a resolution of $35 \mu\text{m}$. For TID, the module strip pitch is between $100 \mu\text{m}$ and $141 \mu\text{m}$. In TOB, pitch sizes are between $122 \mu\text{m}$ and $183 \mu\text{m}$ producing a single point resolution of $35 \mu\text{m}$ and $53 \mu\text{m}$ respectively. Finally, TEC can reach a resolution between $230 \mu\text{m}$ and $530 \mu\text{m}$ using modules with strip pitches between $97 \mu\text{m}$ to $184 \mu\text{m}$. The single hit resolutions result in a momentum resolution of 1 to 2 % for the measurement of charged particles with a transverse momentum of $p_T > 100 \text{ GeV}$.

2.3.4 Calorimeter

In the CMS experiment, there are two calorimeter systems: the Electromagnetic Calorimeter (ECAL) [39] and the Hadronic Calorimeter (HCAL) [40]. The calorimeters are arranged in a symmetrical pattern between the tracker system and the magnet. They provide particle energy measurements and contribute to particle type classification. A scintillator detector is housed within the inner ECAL. Electrons,

positrons, and photons primarily deposit their energy in the scintillator material, whereas hadrons are first stopped in the outer HCAL material.

Electromagnetic calorimeter - ECAL

The ECAL measures the energy of electrons, positrons, and photons. A shower of electromagnetic particles is produced as a result of the subsequent emission of bremsstrahlung in the scintillator material and the formation of conversion pairs. The number of particles in this shower is proportional to the energy of the incoming electron, positron or photon. Shower particles excite molecules in the scintillator material, causing light to be emitted. This light is then detected by photo detectors at the other end of the scintillator. The scintillation light is emitted in the blue-green wavelength range (420-430 nm) with a light yield of 30 photons per MeV. Avalanche Photodiodes (aPDs) in the barrel and Vacuum Phototriodes (vPTs) in the endcaps collect the light and transform it into an electric signal which is then digitized.




Figure 2.7: The electromagnetic calorimeter of the CMS experiment consists of the barrel (EB), the preshower (ES) and the endcap (EE) calorimeters [41].

The ECAL is made up of over 68 000 crystals of lead tungstate (PbWO_4). PbWO_4 is a transparent material with a high density of 8.28 g/cm^3 , a short radiation length of $X_0 = 0.89 \text{ cm}$, and a Molière radius of $R_H = 2.2 \text{ cm}$. As a result, PbWO_4 can function as both an absorber and a scintillator at the same time. It also has a high radiation hardness and a fast scintillation decay time, with approximately 80 % of the light emitted within 25 ns. Using this material, a compact design with fine granularity is possible. The ECAL is made up of two endcaps (EE) and a cylindrical barrel (EB). A preshower detector (SE) is installed in front of each endcap. Figure 2.7 shows the ECAL layout.

The EB has 61 200 PbWO_4 crystals and a pseudorapidity of $|\eta| < 1.479$. Each crystal covers a cross section of $22 \times 22 \text{ mm}^2$ in the η - φ plane. The EB inner radius is 1.29 m and each EB crystal is 230 mm long, corresponding to $25.8 X_0$. All crystals are oriented towards the interaction point. Avalanche photodiodes read the light emitted by high-energy electrons, positrons, or photons passing through the crystal.

The EE disks cover the range $1.479 < |\eta| < 3.0$ in pseudorapidity. Each endcap has two halves. There are 3 662 crystals per half and the crystals are grouped in a 5×5 matrix called supercrystals with a 28.62 mm per side. The crystal has a total length of 220 mm corresponding to $24.7 X_0$.

Preshower detectors identify charged pions and improve electron and photon identification and localization in the pseudorapidity range $1.653 < |\eta| < 2.6$. It is a sampling calorimeter. Electromagnetic showers are created by incoming electrons or photons. These are detected by silicon strips between the lead layers. The preshower detector is 20 cm thick.

Hadronic Calorimeter - HCAL

The HCAL structure differs from the ECAL, since the absorber and scintillator material are arranged alternately, the radiation length is significantly greater than the radiation length of the electromagnetic interaction. A massive 50 mm brass plate, which is embedded in steel plates on both sides, serves as the absorber. Because the HCAL is placed inside the solenoid in a strong magnetic field, the nonmagnetic absorber material brass was chosen, which also has a reasonable hadronic interaction length of $\lambda_I = 16.4 \text{ cm}$. As active material, the absorber is arranged in slices that are interspaced with plastic scintillators. Hadrons passing through the HCAL interact strongly with the absorber material, resulting in showers of secondary hadrons. These showers cause excitations of the active material and the emission of light, which is transported to photo-detectors outside the calorimeter via wavelength shifting fibers. This arrangement alternates 14 times. The scintillator consists of 3.7 mm thick polymer tiles that detect the hadronic shower and are read out via hybrid photodiodes. A total of 70 000 scintillator tiles were installed.

The HCAL, illustrated in Figure 2.8, has a particularly large coverage in order to determine the missing transverse impulse as precisely as possible. The transverse momentum is a conserved quantity. The sum of all transverse impulses before and after the protons collide is approximately zero. The observation of a non-zero transverse momentum therefore provides information about particles that do not interact with matter. Missing transverse momentum can be of crucial importance




Figure 2.8: Cross-section view in the r - z plane of the CMS HCAL. The HCAL barrel (HB) and the HCAL endcap (HE) are located inside the solenoid, the position of the front-end readout electronics (FEE) for HB and HE is indicated. The HCAL forward (HF) is located 11.15 m from the IP. The HCAL outer (HO) is installed outside the CMS solenoid coil [42].

for the detection of new physics. In order to determine the missing transverse momentum as precisely as possible, coverage over a large area is necessary.

The HCAL covers a range of $\eta < 5.2$ thanks to the forward HF detector. The endcaps (HE) cover a range of $1.3 < |\eta| < 3.0$, the barrel (HB) cover the range of $|\eta| < 1.3$. Another part of the calorimeter, HO is located behind the solenoid coil. Due to the high penetration depth of charged hadronic particles, in some cases they are not fully absorbed in the actual calorimeter. To increase the energy resolution, two scintillator plates were installed behind the coil.

The HB is made up of 36 identical wedges that form two half-barrels. The absorber plates in the HB are parallel to the beam axis. The HB wedges are divided into towers in the η - φ plane. The HB has an inner radius of 1.78 m and an outer radius of 2.88 m. The thickness of the total amount of absorber material in the HB corresponds to $5.8 \lambda_I$ for the central towers with $|\eta| < 0.087$ and up to $10.6 \lambda_I$ at $|\eta| < 1.3$. The HB is followed by the HO, which increases the effective thickness of the absorber in the barrel region. The HO is divided into five rings based on the geometry of the muon chambers in the outermost part of the detector. The outer four rings each have a single scintillator slice with the solenoid material acting as an absorber. In the

central ring, an additional lead plate serves as an absorber between two active layers. This results in an absorber thickness of at least $10 \lambda_I$ over the entire barrel region, including about $1 \lambda_I$ from the ECAL. The HE is made up of 18 segments arranged in a radial symmetry around the beam pipe. The brass absorber in the HE has a total length of about $10 \lambda_I$. The HF system is located outside the solenoid and is used to detect hadron showers in the very forward region of the CMS detector. The HF has a cylindrical structure that surrounds the beam pipe and has an inner radius of 12.5 cm and an outer radius of 130 cm. The HF must withstand high particle fluxes that occur in the forward region of high-energetic particle collisions. As a result, the HF is made up of steel absorber plates and quartz fibers as active material. When compared to plastic scintillators, quartz fibers are more radiation-hard. Cherenkov light produced in the quartz fibers by hadron showers is detected by photomultipliers.

2.3.5 Muon system

Figure 2.9 shows in the right side the muon system that forms the outermost layer of the CMS detector. The muon system can only be reached by neutrinos, muons and leaking hadron showers. It is responsible for the detection of muons. The detection of muons proves to be particularly helpful in order to later reconstruct decays in which at least one muon is generated. The muon system can provide in real time the trigger with reconstructed muons. The reason for the easy penetration of the muons through the detector material lies in their large mass $m_\mu = 105.6 \text{ MeV}$ and the fact that the radiated energy in the form of bremsstrahlung is proportional to the inverse of the square of the mass.

The detection of muons takes place via ionization chambers. A detection area of $25\,000 \text{ m}^2$ must be covered with these chambers. Since the particle density is lower in the outer detector layers, there is no need for a high granularity in the muon system as in the inner area. Smaller drift cells are used in the cylindrical area, which ultimately form the drift chambers (DT). Drift chambers show the trajectory of charged particles. The gas in the ionization chambers is ionized as the particle passes through it. Parallel wires with a high potential difference ensure that a current flows between the wire and the cathode. Using the current pulses and the drift time, it is possible to reconstruct the particle trajectories. The drift chambers are embedded in four groups in the magnet yoke. The first three groups consist of layers of eight muon chambers each. They are used to determine the position in the r - φ plane. Four further drift chambers make it possible to determine the z coordinate. The spatial resolution in the muon detector is $250 \mu\text{m}$. In contrast, in the endcaps a different technology has to be used. A smaller signal decay time is important because the




Figure 2.9: Cross-section through a slice of the CMS detector. The drift chambers of the muon spectrometer in the iron return yoke are shown in the outer area [43].

particle density is higher than in the radial area. For this reason, Cathode Strip Chambers (CSC) are used in the endcaps. The cathode strips are perpendicular to the beam axis and the anode wires are perpendicular to the cathode strips. The resolution in the radial direction is 6 mm. By forming the center-of-mass, the spatial resolution in the r - φ plane can be improved to 150 μm .

Resistance plates (RPC) are also used to detect muons. There is no center-of-mass formation because the spatial resolution is not high. The RPCs offer the advantage of a good time resolution of 3 ns and help in the reconstruction of the muons. Due to the large distance to the beam axis, the muon system helps to increase the impulse resolution in the high impulse range. By combining information from the muon chambers with tracking information from the inner tracker, the momentum resolution can be significantly improved.

2.3.6 Trigger and Data Acquisition

The LHC produces proton-proton collisions at a relatively high rate of 40 MHz, *i.e.* of one every 25 ns. At the designed instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the accelerator is able to produce in average about 25 simultaneous collisions with a




Figure 2.10: Architecture of the L1 trigger. The final decision is made when all hierarchies have been run through in the calorimeter trigger and the muon trigger. A decision is made for each of the local subsystems, on which the overall decision of the L1 trigger later depends [15].

massive number of about 1 000 particles penetrating the detector each bunch crossing. This massive number of events corresponds to a significant amount of data that cannot be stored without a filtering stage. As a result, the trigger system of CMS [15], [44] reduces the amount of data collected on a short time scale by deciding whether an event should be stored or rejected. The trigger system consists of two subsystems called Level-1 (L1) trigger and High-Level Trigger (HLT). The data collection rate is reduced to a frequency of several hundred hertz, corresponding to a reduction factor of about 10^7 . The signature of processes of physical interest is primarily used to determine which events are stored. It is critical, especially for the search for rare processes, that almost all events of an specific process are saved to permanent storage for offline analysis. As a result, the trigger system ensures that information about events involving one or more charged leptons or photons with high momentum, large missing transverse energy, or a high multiplicity of high-energetic hadron jets are identified within a low latency and saved on a relatively short time scale.

The original CMS L1 trigger design, represented in Figure 2.10, is a hardware-based trigger system with components at the local, regional, and global levels. All components are programmable electronic devices that only use coarsely segmented data from the calorimeter and muon modules. The local components are built around the front-end electronics, which are installed directly inside the experiment. The energy deposits in the calorimeters and hit patterns in the muon chambers are stored

as Trigger Primitive Generators (TPG) in the pipelined memories of the front-end electronics. Regional triggers combine TPG information in spatial regions to identify candidates such as muons or electrons. The global calorimeter and global muon triggers, respectively, combine information from the entire calorimeter and muon system. The global trigger, which is physically located 90 m underground from the CMS cavern, makes the final L1 decision. The L1 trigger system has a $3.2\ \mu\text{s}$ latency and reduces the data rate to around 100 kHz. To cope with the new LHC conditions, the system was upgraded for Run 2 [45]. If the trigger thresholds were not changed, the increased collision centre-of-mass energy, increased instantaneous luminosity, and higher pileup would increase the trigger rate up to six times the allowed rate. The new system was designed to be intrinsically flexible, with the ability to change architectures and algorithms via high-speed optical links and large FPGAs.

The HLT consists of a software-based trigger that runs on a computing farm. The data filtered by the L1 trigger is transferred to the HLT via a network connection capable of 100 GB/s data transfer volumes. The HLT employs algorithms that combine the L1 trigger objects with all detector module information. Physical objects such as muons, electrons, or jets are reconstructed with greater accuracy from the full detector information than objects reconstructed at the L1 trigger level. The HLT also provides trigger paths that combine data from multiple physical objects into a single event. These trigger paths include multi-lepton triggers and triggers that require the presence of a lepton as well as jets. All HLT filtering algorithms are based on a constant adaptation to collision conditions. The HLT reduces the overall rate of events stored to several hundred hertz. Figure 2.11 depicts a schematic representation of the trigger chain of the CMS experiment.




Figure 2.11: Data flow of the CMS Trigger (a) and the DAQ system (b) [46]. The CMS trigger system is divided into two layers: the Level 1 trigger, which reduces the data rate from 40 MHz to 100 kHz, and the HLT, which reduces the rate to 400 Hz.

The data collected by the CMS experiment and accepted by the trigger system must be saved and distributed. Despite the fact that the trigger has already significantly reduced the amount of data, the final data that must be stored corresponds to several tens of petabytes per year. The CMS computing model was created to distribute data as well as simulated events (so-called Monte Carlo samples) and to provide computing resources for data analysis to physicists worldwide. The Worldwide LHC Computing Grid (WLCG) [47], [48], a collaboration of the LHC experiments and several computing centers, organizes data distribution and computing resource allocation. The hierarchical structure of the WLCG is composed by Tier-0, Tier-1, and Tier-2 computing centers. The one-of-a-kind Tier-0 center is a computing facility located directly at CERN. The data is distributed from Tier-0 to Tier-1 centers (*e.g.* gridKA [49] at KIT). If improved reconstruction algorithms are available, the data can be reprocessed at Tier-1 computing centers. They also have massive storage capacities for data from LHC experiments as well as simulated data produced at either Tier-1 or Tier-2 facilities. Tier-2 centers are primarily intended to store copies of reduced data sets, which can then be accessed and analyzed directly by individual physicists or local analysis groups. Tier-3 sites are local computing devices that are connected to the grid for interactive analyses, despite the fact that they are not formally part of the WLCG. Each analyzer in the grid structure has access to any data set at any of the Tier-2 sites.

3 High-luminosity LHC

The High-luminosity Large Hadron Collider (HL-LHC) project aims to upgrade the LHC collider according to the long term schedule in Figure 3.1, the upgrades to the accelerator complex and the experiments are divided in two Long Shutdown (LS) periods from 2019-2021 (LS2) and 2025-2027 (LS3) after which the HL-LHC configuration is operating. The upgrades bring the luminosity of the LHC up to $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, or even $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in the ultimate performance scenario. To reach this goal, the HL-LHC is exploring new beam configurations and advanced technologies in the domain of superconductivity, cryogenics, radiation-hard materials, electronics and remote control, some of the upgrades to the accelerator complex are mentioned in Section 3.2. To prepare for the challenging conditions that are expected during the entire LHC operation period, the CMS experiment has planned two upgrade phases executed on the LS2 and LS3, which are described in Section 3.3. CMS will, in particular, design a completely new silicon tracking detector capable of discriminating high energetic particles and use that information to reconstruct tracks in real-time to contribute to the L1 trigger production. A detailed description of the new CMS tracking detector is provided in Section 3.4.




Figure 3.1: LHC operation and HL-LHC installation schedule, adapted from [50].

3.1 Physics Motivation

The upgrade for the HL-LHC opens exciting possibilities for beyond-SM discoveries, in spite of bringing substantial experimental challenges due to the increase in instantaneous luminosity. High radiation levels, particularly in the inner most detector layers, motivate a complete redesign of the silicon tracker with radiation tolerance and high spatial and time resolutions. Many other sub-detector systems would need to be significantly upgraded or be completely replaced to better exploit the discovery potential of the HL-LHC conditions.

One of the main challenges, particularly to CMS and ATLAS, is the substantial increase in the number of simultaneous proton-proton (pp) collisions per bunch crossing (pileup). The number of simultaneous collisions grows from the current average of 25 to 200 during the HL-LHC operations. The increment in pileup produces many more low energetic particles, increasing the background measurements in the calorimeter and generating an increased number of low-momentum muons, that could be miss-identified as high-momentum ones. Therefore, the trigger system needs to be improved with better handles to identify high-energy particles at the hardware L1 trigger, receiving tracker primitives at the collision rate.

These new trigger handles will help in the identification of extremely rare processes. A big part of the program for the HL-LHC will consist of more precise measurements of known phenomena, for example further precision of the W -mass, WW , WWZ , and $WW\gamma$ couplings. Other mechanisms of great interest are the Higgs couplings to the second-generation fermions via the process $H \rightarrow \mu\mu$, which using measurements from Runs 1-2 was observed with a 3σ statistical significance [51] and finally, the Higgs self-coupling via the rare di-Higgs process $H + H \rightarrow \mu\mu$. The goal of the LHC is therefore to extend the measurements of the Higgs properties, and to reach a precision in the range of few percent by the end of the HL operations [11].

The above studies rely heavily on the ability to efficiently identify at the hardware L1 trigger all those physics objects by correlating information from various detectors. To maintain the triggering capability with a manageable data rate, the trigger system needs to be upgraded and improved with the inclusion of charged-particle track information as a critical component. The large input data rate from the silicon tracker, the reduced latency at which tracks need to be identified, and the significant processing requirements for track reconstruction are some of the technological challenges which motivate the design of the tracker back-end electronics system described throughout this thesis.

3.2 The HL-LHC Machine

The high-luminosity design involves the upgrade of multiple systems. Existing systems may not be able to withstand the increasingly harsh conditions that the highest luminosity performance will generate. Increasing wear and radiation damage are some of the primary concerns. Many modifications will be required simply to keep the machine running in a regime of nominal or ultimate luminosity. Replacements for certain systems may be made with better-performing equipment rather than spares of the same specification. This ‘improvement’ in performance goes far beyond the basic reconfiguration that is already planned for the LHC.

For other systems, replacement, while triggered by technical reasons, provides the opportunity to carry out a complete change in layout or performance and may be considered a true upgrade. The most striking example, and the cornerstone of the upgrade, is the replacement of the inner triplet magnets with new magnets based on a Niobium-tin (Nb_3Sn) superconductor. Another example is the replacement of a significant portion of the current collimation system with a better design with lower impedance jaws.

In other cases, new equipment not currently included in the LHC layout will be installed to improve performance, either in terms of peak luminosity or availability. The most notable example is the superconducting RF crab cavities, which are constructed with a compact design as required for the HL-LHC, a completely novel development and the first for a proton collider. In this section, a brief description of some of the systems that will need to be upgraded or significantly improved in performance will be provided. This section is based on [10] and [11].

3.2.1 Insertion region magnets

The LHC is expected to reach an integrated luminosity of around 350 fb^{-1} by 2025, resulting in doses of up to 30 MGy to some components in the high-luminosity interaction regions. Inner triplet quadrupoles should be able to withstand radiation up to an integrated luminosity of 400 fb^{-1} to 700 fb^{-1} , but some nested-type corrector magnets may fail at around 300 fb^{-1} . The most likely failure mode is a sudden electric breakdown, which will require extensive and time-consuming repairs. Replacement of the triplet must be planned for before radiation damage grows to the point where a serious failure may occur.

The replacement can be combined with an increase in the quadrupole aperture.

However, larger aperture inner triplet quadrupoles and an increased luminosity, imply higher radiation levels, needing a redesign of the entire interaction region zone. This redesign includes larger D1 and D2 dipoles, a new electrical feedbox (DFBX), and a much improved maintenance access to various components. Furthermore, a redesign of the collimation system in the high-luminosity insertions, will be required.




Figure 3.2: The overall configuration of the insertion region for the HL-LHC between the IP and the orbit corrector Q4. The two beam envelopes are represented by the dark blue and red areas. The light regions correspond to a beam envelope value of 12σ [10].

The modification of the inner triplets in the high luminosity insertions, as seen in the Figure 3.2, is absolutely the central component of the LHC upgrade. The decision for the HL-LHC has been to rely on the success of advanced Nb₃Sn technology, which allows access to magnetic fields well beyond 9 T, allowing the aperture of the inner triplet quadrupoles to be maximized.

3.2.2 Collimation

The collimation system was developed for the first phase of the LHC operation. It is currently operating in accordance with the plan. However, if beam instabilities are triggered at intensities close to or just above nominal, the impedance of the collimation system may need to be reduced. The safe handling of a beam of 1 A or more with beta functions at collision exceeding the design value will be uncharted territory. The triplet must be protected throughout the large change in collision beam parameters (β^* transition from 6 m to 10-15 cm). This will be one of the most critical phases of HL-LHC operation: the beam halo alone could exceed the damage limit. As a result, the collimation system must be upgraded. The main additional requirements associated with the upgrade are improved alignment precision and materials capable of withstanding higher power.

3.2.3 Crab cavities

Superconducting RF crab cavities are required in the HL-LHC to compensate for the geometric reduction factor, shown in Figure 3.3, therefore increasing the effective cross section with a very low β^* angle. The unconventional, compact design of HL-LHC crab cavities cannot be achieved using current state-of-the-art technology based on elliptical cavities. It is also required the precise control the RF phase with a precision higher than 0.001° so that the beam rotation provided prior to the collision is precisely canceled on the other side of the Interaction Point (IP). As a result, machine protection will face new challenges with the installation of these new crab cavities. Compact crab cavities will be installed on both sides of IP1 (ATLAS) and IP5 (CMS). On each side of the IP, there are four crab cavities per beam. To rotate the beam in the crossing plane, all four cavities can be used; alternatively, a single cryomodule (two cavities) can be used, with the cavities in the second cryomodule providing a deflection in the orthogonal plane, enabling the so-called crab kissing scheme for reducing pileup density [52]. At the moment, the standard practice is to use all crab cavities for geometric compensation, *i.e.* rotation in the crossing plane as seen in Figure 3.3(b).




Figure 3.3: a) Geometrical luminosity reduction factor vs β^* for LHC and HL-LHC. b) Effect of the crab cavity on the beam. Small arrows indicate the torque on the bunch generated by the transverse radiofrequency field. [10].

3.3 The CMS Upgrade

The CMS experiment, as explained on Chapter 2.3, proposed an upgrade plan for the detector divided in two phases to cope with the LHC operation challenges. The Phase-1 upgrade was executed during the Extended Year-End Technical Stop (EYETS) between 2016 and 2017. For Run 2, a new pixel detector was installed, the HCAL forward (HF) readout was upgraded, and the upgraded L1 trigger was

included. Phase-1 upgrades continued in the LS2 period between 2019 and 2021. The upgrades included replacing the photo detectors on the HB and HE sections of the hadron calorimeter, exchanging the scintillating tiles in the HE, and installing new Front-End (FE) electronics in the Cathode Strip Chambers (CSC) of the endcap muon detector.

The Phase-2 upgrades motivated by the extreme new conditions with very high pileup are planned for the LS3 period between 2025 and 2027. As the beam width at the collision point is reduced, the luminosity increases to $L = 5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This leads to an increase in simultaneous inelastic proton-proton collisions up to 200. A negative side effect of this increased luminosity is a higher radiation exposure. To counteract this effectively, more radiation-resistant materials must be used in the construction of the detectors closer to the IP.

Another problem related to 200 pileup events is that individual tracks have to be distinguished from one another within 25 ns. This is countered with an increased granularity in the respective components. However, the increased granularity will also lead to a massive increase in data flow. Another factor that raises the amount of data is the increased coverage, which ensures that particles in the beam direction up to $|\eta| < 4$ can be detected. The current trigger system would not be able to process the amount of data that arises, the trigger electronics have to be revised to be able to withstand the requirements mentioned above.

Tracker Due to radiation damage, the current CMS tracking system must be replaced during LS3 [53]. The new detector is made up of silicon sensors, pixels, and strips with significantly higher granularity and a wider forward acceptance. Furthermore, a significant new feature of the Phase-2 tracker is the ability to provide tracking information at the L1 trigger level at the collision rate of 40 MHz. Currently, the silicon tracker information is only available at the HLT. Including tracks for the L1 decision allows the trigger rates to be kept at a manageable level while preserving the physics potential. Section 3.4 contains a more detailed description of the CMS Phase-2 tracker upgrade.

Timing detector An extra detector layer between the outer tracker and the ECAL is being considered [54]. This will provide extremely accurate time measurements in the range of 30 ps for each track. The detector in the barrel will be based on lutetium-yttrium orthosilicate crystals, activated with cerium (LYSO:Ce) and read out with SiPMs. The endcap region will use MIP-sensitive silicon technology arranged in a hermetic single layer. Several studies have demonstrated the advantages that the timing layer can bring to object reconstruction and physics analysis, such as improved

track and vertex reconstruction abilities, mitigate extreme pileup conditions, higher lepton efficiencies and diphoton vertex location, and improved jet identification.

Barrel calorimeters Data from LHC runs indicate that radiation damage to the EB and HB calorimeters will be acceptable for full HL-LHC operations [55]. The upgraded HCAL detector is planned to be harder to radiation damage. The lead tungsten crystals in the ECAL are cooled down to a lower temperature. To maintain the exact same physics performance, the FE electronics of both sub-detectors must be improved to meet the new L1 trigger requirements. Furthermore, in the case of the EB, the new FE boards will enable the use of information from single crystals in the L1 trigger, providing more precise timing resolution and aiding in the reduction of photodetector noise.

Endcap calorimeters Endcap electromagnetic and hadronic calorimeters will suffer significant radiation damage and therefore replaced during the LS3 [56]. The new detector will be a sampling calorimeter with electromagnetic and hadronic sections that have excellent transverse and longitudinal segmentation, known as a High Granularity Calorimeter (HGC). The high granularity will allow for the reconstruction of a highly detailed 3D image of the electromagnetic and hadronic particle showers. A large portion of the endcap calorimeters use hexagonal silicon sensors with cell sizes ranging from 0.5 to 1 cm² as active elements, towards the back of the endcap in the z direction the calorimeter uses highly segmented scintillators with Silicon-Photo-Multiplier (SiPM) readout electronics.

Muon detectors Several studies have been conducted since 2015 to determine whether the current muon chambers can handle the increased particle rates of the HL-LHC [57]. Because no significant deterioration of key chamber parameters has been observed, the chambers can be used until the end of the HL-LHC operation. Some upgrades in the front end electronics are foreseen to improve radiation tolerance, readout speed, and performance. Furthermore, the muon system will be extended in the pseudorapidity region $1.5 < |\eta| < 2.4$ with improved Resistive Plate Chambers (RPCs) and new chambers based on Gas Electron Multiplier (GEM) technology. The primary objectives of this forward upgrade are to add redundancy, improve triggering and reconstruction performance, and increase acceptance in the forward detector region.

Trigger and data acquisition Choosing interesting physics events at the L1 trigger stage becomes extremely difficult due to the increased trigger rate and greater complexity of the events in high pileup conditions. To meet the HL-LHC challenges,

the trigger and DAQ system will be upgraded [58], [59]. The L1 trigger latency will increase from approximately $4\ \mu\text{s}$, in the current system, to a maximum of $12.5\ \mu\text{s}$, allowing for track reconstruction in programmable hardware. Furthermore, the maximum readout rate will be increased to $750\ \text{kHz}$ from the current $100\ \text{kHz}$. As a result, the trigger and DAQ system will be upgraded to accommodate the increased bandwidth and computing power required by the larger event sizes and the new L1 trigger rate.




Figure 3.4: Functional diagram of the Phase-2 CMS L1 trigger system [60].

In Figure 3.4 a functional diagram of the entire L1 trigger system is presented, three main sub-detector systems contribute to the trigger production: the calorimeter trigger, the muon trigger, and the track trigger, which is explained in more detail in Section 3.5. From the Figure 3.4, the global calorimeter trigger (GCT) include inputs from the barrel calorimeter trigger (BCT) reading the barrel calorimeter (BC), the hadron forward calorimeter (HF), and the high-granularity calorimeter (HGCAL). The muon trigger receives input from various detectors, including drift tubes (DT), resistive plate chambers (RPC), cathode strip chambers (CSC), and gas electron multipliers (GEM). The muon trigger is composed of a barrel layer-1 processor and muon track finders processing data from three separate pseudorapidity regions and referred to as BMTF, OMTF and EMTF for barrel, overlap and endcap, respectively. The muon track finders transmit their muon candidates to the global muon trigger (GMT), where combination with tracking information is possible. The track finder (TF) provides tracks to various parts of the design including the global track trigger

(GTT). The correlator trigger (CT) in the center (yellow area) is composed of two layers dedicated to particle-flow reconstruction. All objects are sent to the global trigger (GT) issuing the final L1 trigger decision. External triggers feeding into the GT are also shown including potential upslope (mentioned as "others") such as inputs from the MTD. The dashed lines represent links that could be potentially exploited. The various levels of processing are indicated on the right: trigger primitives (TP), local and global trigger reconstruction, particle-flow trigger reconstruction (PF) and global trigger decision (GT).

3.4 The CMS Tracker Upgrade

The current silicon tracker is not suitable for operation under the HL-LHC conditions, it was originally designed to resolve between 20 to 30 simultaneous collisions (pileup), which is about a factor of ten less than the new conditions of up to 200. Its performance will degrade to an unacceptable level after Run 3 due to radiation damage exceeding its design value of 1000 fb^{-1} [61]. Another motivation to replace the tracker is the requirement to include track information in the L1 trigger production. The new silicon modules can select pairs of hits compatible with high-transverse-momentum tracks ($p_T > 2$ or 3 GeV), distinguishing them among a large number of background hits of little interest for use in the L1 trigger. The L1 trigger system is designed to have an output rate of 750 kHz , which appears to be impossible to achieve using only data from the calorimeters and muon detectors. Due to all these reasons, the CMS silicon tracker will be completely replaced with a novel design, based on double-sided detector modules arranged in a tilted geometry [53] shown in Figure 3.5. An alternative flat barrel layout was initially considered and used to start the development of the track reconstruction algorithms in Chapter 4 and Chapter 5, which later were updated to the tilted geometry. Similarly, performance results in Chapter 6 are often referred to either one or both configurations.

3.4.1 Requirements

The new Phase-2 tracking detector is composed by an Inner Tracker (IT) made of silicon pixel modules and an Outer Tracker (OT) made of silicon modules with two types of modules. The OT has dual layer detector modules containing strips and macro-pixel or double-strip sensors. The following are the main requirements for the tracker upgrade.




Figure 3.5: Sketch of a quarter of the CMS phase-2 silicon tracker in the r - z plane for the flat (upper) and tilted (lower) barrel configurations. The blue lines represent the Pixel-Strip (PS) modules, the red lines represent the Double Strip (2S) modules, the green lines represent pixel modules made of two readout chips, and the yellow lines represent pixel modules with four read out chips [53], [61].

High radiation tolerance The new tracker must be able to operate efficiently for the duration of the HL-LHC operation, up to an integrated luminosity of 3000 fb^{-1} . In the case of the OT, no maintenance interventions are planned for the duration of the HL-LHC operation, whereas the pixel modules of the Inner Tracker (IT) are expected to remain accessible and replaceable as they accumulate significant radiation damage. According to simulations, the integrated particle fluence in 1 MeV neutron equivalent for the inner most region of the IT is $2.3 \times 10^{16} \text{ n}_{eq}/\text{cm}^2$ and for the Outer Tracker (OT) is $9.6 \times 10^{14} \text{ n}_{eq}/\text{cm}^2$. the IT might have to be replaced after five years at full luminosity.

Increased granularity Under the harsher conditions at the HL-LHC, each beam crossing will produce approximately 6000 charged particles with transverse momentum greater than 300 MeV produced by approximately 200 collisions on average. To ensure efficient tracking performance and to limit the problem caused by combinatorics in both the seeding and track finding phases, channel occupancy should be kept at or below the percent level in the OT and below the one in a thousand level in

the IT. This requirement also allows for the avoidance of cluster merging and the maintenance of a good two-track separation in the IT, which is especially important for track finding performance in high-energy jets.

Extended tracker acceptance The HL-LHC physics program includes many analyses that will undoubtedly benefit from a larger tracker acceptance region, such as those involving vector-boson scattering measurements. As a result, the new tracker is intended to extend coverage in pseudorapidity up to $|\eta| = 4$.

Reduced material The amount of material in the tracking volume has a significant impact on tracking performance, secondary interactions of particles with matter causes multiple scattering, which overall decreases the physics performance. As a result, a lighter tracker option was chosen.

Participation in the L1 trigger As previously stated, the CMS trigger will operate at the HL-LHC with significantly increased latency and output rate. It is obvious that the tracker must be involved in the L1 trigger decision in order to maintain an acceptable output rate despite the increase in luminosity of about 10 times.

3.4.2 The Inner Tracker

Pixel modules will be used in the construction of the Inner Tracker (IT). The new IT will have to withstand the harsh conditions of the HL-LHC environment, where the high radiation dose and increased pileup will put a strain on the radiation hardness and front-end electronics. The current proposal uses 100-150 μm thin pixel modules segmented into pixel sizes of $25 \times 100 \mu\text{m}^2$ or $50 \times 50 \mu\text{m}^2$. They are expected to exhibit the required radiation tolerance as well as the desired performance in terms of detector resolution, occupancy, and two-track separation. ATLAS and CMS are working together under the framework of RD53 [62] to design a pixel chip in 65 nm CMOS technology.

The Inner Tracker is arranged in four cylindrical layers in the barrel region denominated as Tracker Barrel Pixel Detector (TBPX), eight small disks in the forward region of the pixel detector within the OT barrel denominated as Tracker Forward Pixel Detector (TPFX), and four larger disks in the endcaps with the name Tracker Endcap Pixel Detector (TEPX). The overall structure covers the acceptance region to $|\eta| = 4$ and can be seen in Figure 3.5. The design of the Inner Tracker detector will allow the replacement of degraded parts over an Extended Year-End Technical Stop (EYETS), therefore, the design has the ability to extract and insert the detector

without removing the CMS beam pipe. This is accomplished by placing the detector on inclined rails.

3.4.3 The Outer Tacker

The new layout of the outer track detector has three regions: the Tracker Barrel with Pixel-Strip (PS) modules (TBPS), the Tracker Barrel with Double Strip (2S) modules (TB2S), and the Tracker Endcap Double-Disks (TEDD) containing both PS and 2S modules. Three barrel layers are assigned for the TBPS region, three layers are also used for the TB2S region which are arranged at a constant r around the pipe. Five endcap layers are provided in the TEDD region. The PS modules and 2S modules, shown in more detail in Figure 3.7, are specially manufactured using planar silicon sensors. This layout ensures that particles emerging from the interaction region with $|z| < 70$ mm cross at least six layers in the rapidity range $|\eta| < 2.4$. The only exception is the transition region between barrel and endcap, located around $|\eta| \simeq 1.0$, where the number of crossed modules is reduced to five. Crossing several detector layers is critical for providing a reliable track finding for the L1 trigger. Additionally, to provide some initial vertex estimation at least three macro-pixel layers are required. In order to reduce the amount of material used in the barrel and increase the resolution, some of the PS modules are arranged in a progressive slope depending on their r and z coordinates according to the tilted barrel geometry in Figure 3.5.




Figure 3.6: The p_T module concept is illustrated. (a) Signal correlation in closely spaced sensors enables rejection of low- p_T particles; the green channels represent the selection window for defining an accepted stub (double correlated hit). (b) For a given sensor spacing, the same transverse momentum corresponds to a larger distance between the two signals at large radii. (c) A larger spacing between the sensors is required for the endcap disks to achieve the same accuracy as in the barrel at the same radius [61].

The p_T modules are the foundation of the entire OT and play a critical role in the ability to send tracks for the L1 trigger production. The main idea is to correlate two signals from two closely spaced sensors and test their compatibility with the hypothesis that both are from a track with a p_T above a certain threshold. To put this concept into practice, the p_T modules are made up of two silicon sensors with a gap between them in the range of 1.6 mm to 4 mm. Both sensors are read out by a single FE electronics unit that can combine the signal pairs. A programmable selection window is used at the FE chips to define the compatibility. If the signal pairs are within the selection window, they are combined to form stubs, as seen in Figure 3.6(a). At each bunch crossing, stubs are sent to the L1 trigger system. All other OT module hits are stored in the FE pipelines and read out when a trigger is received. Figure 3.8 depicts a simplified illustration of the concept.




Figure 3.7: Assembled 3D representation of the p_T modules including a side view showing the hybrid assembly and connectivity [53]. a) The 2S modules. b) The PS modules.

The 2S modules consist of two sensors, each with 2×1016 strips of 5 cm long with $90 \mu\text{m}$ pitch between them. The PS modules consist of a bottom sensor with 32×960 macro-pixels of $1.5 \text{ mm} \times 100 \mu\text{m}$ size and a top sensor with 2×960 strips each with a length of 2.4 cm and a $100 \mu\text{m}$ pitch. The sensitive elements of the modules in both cases are planar silicon sensors. The 2S sensors are mounted in such a way that their strips are parallel to each other. Figure 3.7 shows both types of modules with an indication of its structural construction. In the PS module type, the precise measurement of the z -coordinate coming from the macro-pixel is especially important in the L1 trigger for primary vertex discrimination and robust pattern recognition. Furthermore, to achieve efficient rejection of low- p_T particles

across the entire tracker volume, the selection window in the FE electronics must be programmable, and the modules must have different sensor gaps depending on their position in the tracker. The total power consumed by the 2S modules is about 6 W whereas for the PS modules it is about 9.4 W, power reported when using the sensors at its intended operational temperature of -20°C .

The front-end readout chips are connected to the sensor using wirebonding technology to a surrounding structure known as the service hybrid, as seen in Figure 3.7, the service hybrid contains various power supply modules and a radiation-hard optical transceiver [63] to transfer data to the back-end system. The 2S module has one readout Application-Specific Integrated Circuit (ASIC) known as the CMS Binary Chip (CBC). The PS modules has two types of FE ASICs, the Short Strip ASIC (SSA) and the Macro-Pixel ASIC (MPA), both implemented in a 65 nm CMOS process. The readout hybrid chips (CBC2 in the 2S module and SSA and MPA in the PS module) are buffered, aggregated, and formatted by the Concentrator Integrated Circuit (CIC) ASIC. Stub data are sent out after each collision between two proton packets, but all other data are retained and stored in the front-end pipelines for up to $12.5\text{ }\mu\text{s}$ until a positive trigger decision has been made to read the entire data in a selected event.

3.5 Track Finding at the Level-1 Trigger




Figure 3.8: Data-flow from OT detector modules through to the back-end electronics, including the different logical and processing elements and their interfaces [53].

The increased luminosity allows for a more precise search for clues to new physics. However, the performance of the L1 trigger and its ability to reconstruct the complex topology of events will depend heavily on the track finder. Track reconstruction in the OT back-end electronics system is done using a time-multiplexed architecture that requires the use of two layers of data processing, as shown in Figure 3.8. The

first layer is represented by the Data, Trigger, and Control (DTC) boards, which extract and pre-process the stub data, for example, by converting the position from module-local coordinates to experiment-wide 3-dimensional coordinates. After that, data is routed to the correct board in the Track Finding layer. The routing of stubs follows a set of rules aimed at splitting the data in the detector as a function of time and geometrical origin (time multiplexing is based on event number and azimuthal sector position). The DTC cards are also in charge of interpreting and repackaging raw hit data from triggered events before sending it to the DAQ and TTC Hub (DTH) cards. They also supply the tracker front-end modules with control and timing signals. The Track Finding (TF), the second layer, receives all data for a given event and reconstructs the trajectories of charged particles in the tracker. As shown in Figure 3.4, the TF sends track candidates at the collision rate to other L1 trigger boards such as the Global Track Trigger, Global Muon Trigger, and Layer-1 Correlator Trigger.




Figure 3.9: CMS tracker back-end system architecture using two neighboring detector nonants DTCs to time-multiplex and duplicate stub data across processing nonant boundaries for transmission to the Track Finding Processors (TFPs). With a time-multiplexing period of 18, a single TFP receives all data for one nonant for one event, every 18 bunch crossing [64].

The OT is physically divided in 9 distinct regions, covering roughly 40° in φ , each division is completely independent from the others by having dedicated services channels providing power and optical fibers routing. The DTC segmentation follows that scheme by connecting each ‘detector nonant’ to 24 DTC processing cards. The TF layer follows the segmentation into 9 regions, however the sector is rotated by half the width of the detector nonant, about 20° , to account for boundary conditions, therefore forming what is called as ‘processor nonant’, as shown in Figure 3.9. The light-green shaded area is data which belongs to a single processor nonant, the dark-green shaded area is formed by the curvature of two tracks with $p_T = \pm 2 \text{ GeV}$, which intersect at a given point determined by simulation. This shaded shape is

called the ‘hour-glass’, where data contained inside is duplicated and sent to both neighboring processor nonants. Two track reconstruction algorithms implemented on hardware following this configuration are introduced in Chapter 4 and detailed on Chapters 5 and 6. Both OT back-end electronics system layers, the DTC and TF, are build following the ATCA format and will be described in more detail in Chapters 7, 8, and 9.

4 The Track Finding Algorithms

Three distinct algorithms were initially considered as candidates for performing track reconstruction for the L1 trigger: The Associative Memory (AM) approach [65], the Tracklet approach [2], and the TMTT based on Hough Transform approach [1]. The AM approach used a combination of custom ASICs for performing pattern matching and FPGAs for doing the track fitting step. The other two approaches, Tracklet and Time-multiplexed Track Trigger (TMTT) used solely FPGAs for all steps of the data processing. All three approaches were implemented in Hardware Description Language (HDL), all split the detector into smaller geometrical regions, and all make use of time multiplexing to distribute the data into many different hardware units, where the algorithms take place. As a result of evaluating the proposals, CMS determined that all three were suited for the task of reconstructing tracks under the 4 μ s latency, however an all-FPGA approach was preferred due to risks associated with the fabrication of the associative memories in relatively new technologies such as the 28 nm node and 3D die integration [53]. In this Chapter an introductory description of the two all-FPGA approaches is given.

4.1 The Tracklet approach

The tracklet approach is an implementation of the "track following" pattern recognition algorithm [66] where a track candidate starts from a "seed" (i.e. a local short segment of track formed by two stubs in adjacent layers), here called 'tracklet'. The tracklets are then extrapolated to other detector layers in search for stubs compatible with those directions. Tracks are formed when four or more stubs are associated with the initial tracklet seed, then a linearized χ^2 fit is used to determine the final track parameters. Finally, the tracks are analyzed to identify those which share stubs and remove or merge tracks if found to be duplicates of one another.

4.1.1 Stub Organization: Layer Router - VM Router

A significant challenge in the tracklet algorithm comes from organizing and distributing the data in such a way that truncation and latency are kept within acceptable ranges. The detector is first divided along the azimuthal angle into narrow sections, called sectors, then each sector is further subdivided into finer regions, referred to as Virtual Module (VM). Incoming stub data is sorted by the Layer Router and stored in the input memories according to the layer to which it belongs, its φ position, and the specific z or r position, depending on whether it is found in a barrel or disk layer. Then the VM Router, receiving a start signal every 150 ns, reads the stubs from a three input layer memory, where the stub data is stored in full resolution in a 36-bit format, and forwards it to the corresponding output memory inside a particular VM, which corresponds to a particular small area of the detector. At the VM stub memory only coarse position information is stored in a 6-bit index pointing to the memory which contains the full hit resolution for further retrieving. This organization of the stubs allows the tracklet to initially be created by forming pairs of stubs which are consistent with $p_T > 2 \text{ GeV}$ and $|z_0| < 15 \text{ cm}$. The bend of the stubs should as well be compatible with the initial tracklet p_T [12].

4.1.2 Seeding: Tracklet Engine - Tracklet Calculator

In order to form a seed, or tracklet, pairs of stubs in selected adjacent layers or disks are combined. Each tracklet uses the beam interaction region as a constraint to build and calculate the initial trajectory, the projection to other layers and disks are computed using those requirements. Using the projected information, stubs within a small range of the trajectory are considered as matching stubs and therefore are added and used in the track fitting step. The seeding is performed by each VM independently and in parallel. Seeding from multiple combinations of layers is done to ensure a total coverage of the detector when finding potential tracks. Selected stub pairs are sent to the Tracklet Calculator where the precise track parameters are calculated.

4.1.3 Projections: Match Engine - Match Calculator

The projections to other layers are calculated with respect to the nominal detector module position, where the derivatives of the φ and r (z) positions are also evaluated. Stub matches are performed in two steps, with the Match Engine coarsely comparing the coordinates using the projection to the nominal radius and the Match Calculator

producing the exact position using precise residuals. Using the φ residual, this process selects only the best stub match per layer or disk.

4.1.4 Fitting: Linearized χ^2 fit

All stubs matched to the trajectory are fitted using a linearized χ^2 fit. To match stubs to the projection from the seed, the residuals between the projections and stubs are calculated in both φ and z (or r for disks). Using these residuals, δy_i , and linearizing the χ^2 fit, the final track parameters u can be expressed in terms of the inverse radius of curvature ρ^{-1} , the azimuthal angle φ_0 , $t = \sinh \eta$, and the longitudinal impact parameter z_0 [12].

$$u = \bar{u} + \sum_i M_i \delta y_i \quad (4.1)$$

where $\bar{u}$ are the track parameters from the seed and M is a weight matrix. The weight matrix is precomputed, for the most part it is independent of the parameters except for stubs in the disks where it depends on the parameter $t = \sinh \eta$. The weights are tabulated for different ranges of t and stored in Lookup Tables (LUTs). The linear form of this equation allows the updated track parameters to be implemented in hardware, where multiply and accumulate operations map directly to the Digital Signal Processor (DSP) blocks in the FPGA fabric.

A final track fit is considered performed when tracks have four or more stubs matched, further refining the track parameters from those initially provided by the seed using precise position of the stubs in full resolution.

4.1.5 Duplicate Removal: Purge Duplicate

The duplicate removal step aims at remove tracks which have been identified multiple times, all containing the same hits. This process occurs when the different seeding layers act in parallel, projecting their initial seed to the other layers of the detector and therefore finding the same stubs independently. Fitted tracks are sent to the Purge Duplicate algorithm where the tracks are compared pair-wise in search for stubs in common between them. When a track has less than 3 unique stubs it is eliminated. Other alternatives of the Purge Duplicate algorithm contemplated the idea of merging stubs from the compared tracks or use the χ^2 parameter from the fit to select which track to eliminate.

4.2 The Time-Multiplexed Track Trigger (TMTT) approach

The TMTT approach uses the Hough Transform (HT) [67], a well-known algorithm in image processing [68], to identify track candidates and a Kalman Filter [69] to fit the initial candidates into final track parameters. In firmware, however, the TMTT algorithm is composed by four distinct stages:

- The Geometric Processor which aims to increase the parallelization by splitting the detector nonant into finer subsectors in η and φ , then sorts the data depending on those parameters and forwards it to the next stage;
- The Hough Transform groups coarsely the stubs consistent with a high p_T track into candidates, it happens in parallel for each of the 36 subsectors defined in the previous stage;
- The Kalman filter processes the track candidate, branching out if multiple stubs are found per layer and removing stubs which do not fit well to the trajectory proposed by the inner most layers of the particular track;
- The Duplicate Removal using precise track fit information analyzes all found tracks in the Hough parameter space and remove any duplicate tracks created at the Hough transform stage.

4.2.1 Sector Assignment: Geometric Processor

The GP has two main purposes, to aid the Hough transformation by pre-processing the stubs from the DTC with a 48-bit format into a 64-bit extended format, and to sort each input stub into different subsectors. There are two sectors defined in φ , and 18 in η . Stubs can be associated with two or more subsectors, in which case they are propagated downstream to all of them. This is due to the curvature of tracks in the φ sectors or to the length of the luminous region along z axis for η sectors. In HDL firmware, the subsector calculation uses 23 DSPs and ~ 2 k LUTs per input link to generate the 36-bit address map for each stub, the subsequent routing stage takes this address map and sends the stub to all sectors marked in it. The routing block is implemented using a three-stage highly pipelined mesh where each stage is able to push stubs to different outputs according to their subsector address [53].

4.2.2 Track Finding: Hough Transform

In the CMS tracker, where a homogeneous magnetic field $B = 3.8 \text{ T}$ is present, the trajectory of a charged particle with $p_T > 3 \text{ GeV}$ produced at the interaction point moves in the r - φ projection according to the 2D helix linear equation [70]

$$\varphi_0 = \varphi - \frac{0.0015 qB}{p_T} \cdot r . \quad (4.2)$$

where r and φ are the trajectory coordinates, φ_0 the angle of the trajectory at the production point, q is the electric charge of the particle, and p_T is the transverse momentum. In track parameter space, the gradients of the straight lines are given by the radius of the stubs, and therefore are always positive. The stub radius is transformed to $r_T = r - T$ to optimize the distribution of lines in track parameter space. It is also necessary to use φ_T as a track parameter, which is the φ coordinate of the track at a radius T [cm]

$$\varphi_T = \varphi - \frac{0.0015 qB}{p_T} \cdot r_T , \quad (4.3)$$

For each stub, at a given r and φ , Equation 4.3 describes a straight line in the track parameter plane ($q/p_T, \varphi$), also known as Hough-space. For a given set of stubs any intersection of these lines in the track parameter plane would identify a circle or track in the q/p_T - φ_T plane, consistent with the origin and all participating stubs. Figure 4.1 visualizes the use of the Hough transform for six stubs produced by a primary track.

In an FPGA the HT track parameter space can be implemented by using an array of bins and find tracks when a bin has multiple entries. The array needs certain granularity which is constraint by the full range in φ_T of 2π and all possible $|q/p_T|$ values for a p_T above 3 GeV . The resolution of the bin array needs to be coarse enough to include the finite resolution of the silicon modules, multiple scattering effects, and approximations in the straight line equation, but it also needs to be sufficiently fine to deliver the most accurate track candidate possible. Each GP subsector delivers data to an independent HT array in firmware, where a total of 36 units are declared per TFP node. Each HT array is composed by 32 columns in q/p_T and 64 rows in φ . The pipelined nature of the system allows the processing of a stub per clock cycle where it traverse a set of daisy-chained columns, or bins in the HT space. The firmware then calculates if the stub is compatible with one or more of the HT rows using Equation 4.3 and the bend information precalculated by

the GP. Consequently, each stub is only added to those cells in the HT array, whose q/p_T column is compatible with this allowed range. This significantly reduces the probability of producing combinatorial fake candidates.

A track candidate is formed when a cell contains stubs from each detector layer. Nevertheless, to allow for detector or readout inefficiencies, and for geometric coverage, the threshold criteria used to identify a track candidate only requires stubs in at least five different tracker barrel layers or endcap disks, and this requirement is further reduced to four in the region $0.89 < |\eta| < 1.16$ to accommodate a small gap in detector coverage between the barrel and the endcaps [53].

4.2.3 Track Fitting: Kalman Filter

The Kalman Filter (KF) choice is motivated by simulation results which indicate that over half of the track candidates from the HT contain one or more stubs from other particles [53]. As a result, the fit stage would not only need to reduce the number of fake tracks but also to be capable of removing stubs from track candidates without affecting the track parameters. Moreover, the Kalman filter is suited for FPGA implementation compared to other global track fitting methods. The matrices are small and independent from the number of measurements, requiring only a single small matrix inversion. Additionally, to reduce the amount of logic required, the pipelined computational blocks can be contained within an iterative loop.

The Kalman filter begins by using the coarse track parameters provided by the HT as seed, it then calculates the covariance matrix of errors or the so called state of the candidate. The stubs previously sorted according to the layer are added iteratively starting from the most inner layers according to the Kalman formalism, with each iteration step, the state uncertainty improves. The update of the track parameters at each iteration is controlled by a limited Kalman gain product of using weighting factors derived from the relative uncertainties in the state and measurement. One simplification of this implementation of the Kalman filter allows fitting in the endcap and barrel regions to be treated algorithmically identically, provided the endcap measurement errors are correctly pre-calculated.

When multiple stubs per layer are found in a given candidate, the algorithm branches out creating multiple possible states. The algorithm allows for up to two non-consecutive missing layers on the state. Candidate tracks with stubs in all layers also use this condition to ensure that combinatorial background and stub finding inefficiencies are handled correctly. The χ^2 of the state can be calculated every time a new stub is added to the state, providing a handle for selection and reduction of the

number of combinations. Multiple KF instances can be used in parallel to maximize data throughput, which can be tuned according to the rate of track candidates from the HT.

4.2.4 Duplicate Removal

The Duplicate Removal (DR) is the final stage of the TMTT algorithm chain. It is designed to eliminate the tracks which have been identified multiple times due to the limited resolution of the Hough parameter space. Each of those duplicate candidate tracks will end up being fitted to tracks contained inside the same Hough parameter space [53]. The precise estimate of the helix parameters from the KF are then compared against the coarse parameters from the HT. In a first step, those which are not consistent are filtered. In a second step, the filtered tracks are then checked for uniqueness recovering efficiency in track reconstruction vs. Monte Carlo simulations. Ensuring that the track parameters calculated by the Kalman filter are consistent with the subsector boundaries that the candidate was initially found in helps reduce the duplicate rate to a few per cent. This algorithm is extremely lightweight, and its implementation has a negligible impact on the overall FPGA resource usage.

4.3 The Hybrid approach

Sections 4.1 and 4.2 showed that both approaches share comparable data flow structures, first the detector is divided and segmented in parallel data streams, then tracks are found containing rough parameter estimates which later get fitted in a subsequent stage, tracks which are reconstructed multiple times are finally filtered. Hybrid algorithms utilizing building blocks from each of the algorithms were tested and evaluated with the intention to standardize boundaries and data formats. The collaboration between the two groups helped in sharing ideas to improve the algorithms and fostered discussions towards a common definition of the reference algorithm.

One natural place to standardize is the pre-processing of the stubs at the first layer of time-multiplexing happening at the DTC level. The sector definitions along the r - φ plane are done in accordance with what is known as the "hour-glass" shape, see Figure 3.9. Stubs inside the duplication region in dark green, formed by tracks with $p_T = \pm 2 \text{ GeV}$ are sent to both processing nonants, therefore eliminating further inter-nonant communication. Processing nonants are rotated by half the width of




Figure 4.1: Track finding algorithms stages, left Tracklet approach, right TMTT approach with data representation at each step [64], [71].

a detector nonant so that they need information from maximum two neighboring detector sectors each. This configuration simplifies the cabling organization at the back-end electronics system [64].

CMS is now pursuing a hybrid algorithm approach as "baseline", where the tracklet method is used for finding seeds and matching those to other stubs in other detector layers or disks to form track candidates. The track candidates are then fitted using the Kalman filter adding iteratively stubs in other layers. The combination of the tracklet approach with precise seeds and the iterative Kalman filter provides an optimal performance compared to other combinations like Hough transform and linearized χ^2 fit. The hybrid system assumes a time-multiplexing factor of 18 and a division in φ into nine "nonants" [12].

Some of the potential additions or improvements to this approach are the possibility to find displaced tracks, which originate outside the current beam spot constraint of $|z_0| < 15$ cm by means of using "triplet" seeds and an improved five parameter Kalman Filter to fit them. This would be impossible to perform with the Hough transformation as it relies heavily on the interaction point as a constraint for the transformation.




Figure 4.2: a) Hybrid algorithm tracking efficiency as a function of η for tracks above 2 GeV (filled markers) or 8 GeV (open markers). b) Track z_0 resolution as a function of $|\eta|$, for tracks in $t\bar{t}$ events overlaid with an average of 200 pileup interactions. The z_0 resolutions correspond to intervals that encompass 68 % (filled markers) or 90 % (open markers) of all tracks with $p_T > 2$ GeV [72].

The hybrid algorithm performance is currently being studied in software using simulated $t\bar{t}$ events with an average pileup of 200. In Figure 4.2 the efficiency to reconstruct truth particles with stubs in at least four detector layers and the track z_0 resolution as a function of η is plotted. The shown performance in simulation of the

hybrid algorithm is similar to those of the Tracklet and TMTT algorithms.

4.4 Summary

Two FPGA-based track reconstruction algorithms using time-multiplexing and detector division into sectors to process the data, were developed in parallel by two independent working groups, which after a wide review of the algorithms are now exploring together the idea of a hybrid algorithm, implementing features from both original proposals. The Hybrid approach is, however, not yet completely implemented in hardware. Many of the processing blocks previously described in Verilog were subsequently re-written in High Level Synthesis (HLS) and are facing major implementation challenges.

The duplicate removal stage in the tracklet approach was never implemented in hardware and it is likely to require significant FPGA resources as explained in Section 5.2.1, since it relies on pair-wise comparisons between tracks and stubs contained in them. The duplicate removal proposed by the TMTT group relies heavily on the Hough transform stage and can not be used in conjunction with tracklet. Bringing the Hybrid approach to hardware is still work in progress. The TMTT chain however was fully validated in hardware and continued receiving gradual improvements to have the algorithms operating at higher clock speeds in newer target devices, two of its four processing stages were implemented by the author and will be described in more detail in the next chapter.

5 Track Finder Algorithm Contributions

The TMTT approach for the CMS L1 Track Finder is highly modular. Each processing stage can be developed independently from the others, guaranteeing code isolation and the possibility of designing different variations on certain particular stages. The data formats at each stage boundary were previously agreed and different topologies are possible. In this Chapter, a detailed look will be given on two of the four processing stages of the TMTT algorithm as explained in Section 4.2, the Geometrical Processor described in Section 5.1 and the Duplicate Removal presented in Section 5.2, which contains descriptions of two possible alternatives along the processing chain, before or after the Kalman Filter. Descriptions on this chapter will refer to the current baseline topology for the tracker, the nonant division with a tilted barrel configuration. Some selected results are shown for the earlier configuration of flat barrel and octant division.

5.1 The TMTT Geometric Processor (GP)

The Geometric Processor (GP) is the first stage in the TMTT algorithm, it receives all the required information corresponding to a ‘processing nonant’ as explained in Section 3.5. The input data, coming from the DTC, represents the stub in a 48-bit word (shown in Table 5.2) which is pre-processed and unpacked into an extended 57-bit word (in Table 5.3) to reduce the processing load on the subsequent processing stage, the Hough Transform (HT). The GP firmware consist of two blocks: a pre-processing kernel block that determines the appropriate subsector for each stub based on its global coordinate position, and a layered routing block. The stubs are assigned to each subsector and then transported to the following processing stage using dedicated outputs. In this way, the data from each subsector can be processed by a separate HT array.

As shown in Figure 5.1, the GP subdivides the processing nonant (or octant in the flat barrel configuration) into several subsectors, also known as (η, φ) subsectors, the GP consists of two divisions in the r - φ plane and 18 divisions in the r - z plane. The division of the nonant into subsectors simplifies the task of the downstream logic by

reducing the amount of data they need to process. Each track search can be carried out independently and in parallel within each of the subsectors. The use of relatively narrow subsectors in η has the additional advantage that every track found by the HT must be roughly consistent with a straight line in the r - z plane, despite the fact that the HT itself only does track finding in the r - φ plane.




Figure 5.1: The segmentation of the tracker volume into subsectors φ (above) and subsectors η (below). The white numbered areas represent the areas that are assigned to a single sector, while the colored areas (there is no difference in meaning between green or blue) represent the overlap area between adjacent sectors in which it is possible to have stubs being assigned to both sectors. The bold lines show the shape of a given sector. The two cylinders mentioned in the text with the radius $T = 67.2$ cm and $S = 50$ cm are indicated by dashed lines in the upper and lower figures.

The Track Finding Processor (TFP) uses each subsector to find tracks in different areas of φ_T and z_S , where φ_T (z_S) is defined as the φ (z) coordinate of a track trajectory with respect to the point where it crosses a cylinder of radius T (S) centered on the beam line. The values of these two parameters are chosen to be $T = 67.2$ cm and $S = 50$ cm, as this reduces the portion of stubs that are assigned to more than one

subsector. The ranges in φ_T or z_S covered by neighboring subsectors are contiguous and do not overlap. In the r - φ plane, the subsectors are the same size, while in the r - z plane their size varies in order to keep the number of stubs in each subsector approximately the same. The GP must assign each stub to a subsector based on whether the stub could have been created by a charged particle with a path within the φ_T or z_S range of that subsector while originating from the beam line. It is possible to calculate z_S using Equation 2.5 and the sector boundary in η with the formula

$$z_S = S \cot(2 \arctan(e^{-\eta})) \quad (5.1)$$

Table 5.1: η subsector boundary specification. The corresponding z_{50} value is given to the nearest cm.

η boundary	0	0.31	0.61	0.89	1.16	1.43	1.7	1.95	2.16	2.4
z_{50} cm	0	16	32	51	72	98	132	172	214	273

giving us the subsector boundaries described in Table 5.1. If the stub is consistent with more than one subsector, then the GP duplicates it. This can occur due to the curvature of tracks within the magnetic field (constrained by the configurable track finding p_T threshold, chosen to be $p_T^{\min} = 3 \text{ GeV}$) or due to the length of the interaction region along the beam axis (defined by the configurable parameter w , which is set to 15 cm, and defines half the width of the beam spot along z). Using the algorithm described below, each stub is assigned to an average of 1.8 subsectors.

A stub with coordinates (r, φ, z) is compatible in the r - z plane with a subsector that covers the range $z_S^{\min} < z_S < z_S^{\max}$ if

$$\frac{r \cdot z_S^{\min}}{S} - w \cdot \left| \frac{r}{S} - 1 \right| < z < \frac{r \cdot z_S^{\max}}{S} + w \cdot \left| \frac{r}{S} - 1 \right|. \quad (5.2)$$

In order to further improve the sector granularity in the TFP without using significant additional FPGA resources, each of these η subsectors can be further divided by an additional factor of two in the r - z plane, with this division being in the middle between the boundaries of the subsector ($z_S^{\min}, z_S^{\max}$). When a stub is assigned to a subsector, the GP checks the consistency of the stub with each of these subsector halves, taking into account some overlap, and stores this information as two bits within the stub data, for later use by the HT.

The corresponding equation for the compatibility in the r - φ plane of the stub with a subsector is

$$|\Delta\varphi| < 0.5 \cdot \frac{2\pi}{N_\varphi} + \varphi_{\text{res}}, \quad (5.3)$$

where $\Delta\varphi$ is the difference in the azimuth angle between the stub and the center of the subsector and N_φ is the number of φ subsectors and is always a multiple of the track-finder division (*e.g.* 18 for nonants or 16 for octants). The azimuthal angle of the center of subsector i is $\varphi_i = \frac{2\pi i}{N_\varphi}$, where $1 \leq i \leq N_\varphi$. The parameter φ_{res} includes the range of the curvature of the track which is permitted in φ by the threshold p_T^{min} , and is equal to

$$\varphi_{\text{res}} = \frac{0.0015 qB}{p_T^{\text{min}}} \cdot |r_T|, \quad (5.4)$$

where $r_T = r - T$, q is the charge of the particle in e units, and the variables p_T , B , and r_T are measured in units of GeV, Tesla and cm respectively. With a value of $N_\varphi = 18$ for nonants (or 16 for octants), no single stub can be compatible with two adjacent subsectors in φ , as long as p_T^{min} is not reduced below 2 GeV.

Here it is important to note that the stub can also be tested against a second condition in the r - φ plane, to reduce the number of stubs that need to be duplicated. By using the stub bend b , measured in units of the strip pitch, provided by the p_T -modules. The bend further restricts the allowable q/p_T range of the track to be inside $(q/p_T)_{\text{min}} < (q/p_T) < (q/p_T)_{\text{max}}$ where:

$$(q/p_T)_{\text{max/min}} = \frac{(b \pm k_b) \rho}{0.0015 r B}, \quad (5.5)$$

$\rho = (p/s)$ for flat barrel stubs, $\rho = (p/s) \cdot (z/r)$ for endcap stubs, and $\rho = (p/s) \cdot [(z/r) \cdot \cos(\alpha_m) + \sin(\alpha_m)]$ for tilted barrel stubs, where p and s are the pitch and separation of the two sensors in a module respectively and α_m is the angle which best represent all the tilted modules found in simulation. Only eight possible values of (p/s) are possible, this information is obtained from a look-up table in the firmware. This equation assumes that the resolution in the bend, when measured in units of the sensor pitch, is the same everywhere in the tracker. Simulations confirm that is the case giving an approximate value of $\sqrt{2/12}$ for the bend resolution, (which is expected due to the fact that each stub comprises two clusters, each of which should have a position resolution similar to $\sqrt{1/12}$ times the sensor pitch,

since the tracker has binary readout [73]). It is assumed that the true bend lies within k_b of the measured value, where k_b is a configurable cut parameter which is chosen to be 1.25 (approximately three standard deviations).

This constraint on q/p_T leads to the condition:

$$|\Delta\varphi'| < 0.5 \cdot \frac{2\pi}{N_\varphi} + \varphi'_{\text{res}} \quad (5.6)$$

where

$$\Delta\varphi' = \Delta\varphi + b\rho \frac{r_T}{r} \quad (5.7)$$

and φ'_{res} , which allows for the resolution in the stub bend, is given by

$$\varphi'_{\text{res}} = k_b \rho \left| \frac{r_T}{r} \right|. \quad (5.8)$$

From the Equation 4.3 used to describe the Hough Transformation, the Hough-space needs to be subdivided into a finite array of bins, which ultimately defines the digitization factors used in all the data formats. The cells are bounded horizontally by $|q/p_T| < q/p_T^{\text{min}}$ and vertically by the range in φ_T covered by the individual stub-sector. A good compromise between track reconstruction performance and FPGA resources have been found to be the array with 32×64 bins of granularity in the Hough-space ($q/p_T \times \varphi_T$).

If Equation 4.3 is written in its digital form we get

$$\frac{\varphi_{T\text{stubdigi}}}{\varphi_{TMult}} = \varphi_{digi} \cdot \Delta\varphi + q/p_{Tdigi} \cdot \Delta_{q/p_T} \cdot \frac{r_{T\text{stubdigi}}}{r_{TMult}} \quad (5.9)$$

after grouping the constants it looks like

$$\varphi_{T\text{stubdigi}} = \varphi_{digi} \cdot \underbrace{\Delta\varphi \cdot \varphi_{TMult}}_{2^A} + q/p_{Tdigi} \cdot \underbrace{\Delta_{q/p_T} \cdot \frac{\varphi_{TMult}}{r_{TMult}}}_{2^B} \cdot r_{T\text{stubdigi}}, \quad (5.10)$$

here the fixed width of each of the bins in the Hough-space are determined by the following equations

$$\Delta_\varphi = \frac{2\pi}{N_{\varphi\text{segments}} \cdot N_{\Delta_\varphi}} \quad \Delta_{q/p_T} = \frac{B \cdot c}{p_T^{\min} \cdot N_{\Delta_{q/p_T}}} \quad (5.11)$$

where $N_{\varphi\text{segments}}$ is the number of divisions in the r - φ plane chosen to be 18 for the nonant configuration, N_{Δ_φ} is the number of cells in the φ axis, $N_{\Delta_{q/p_T}}$ the number of bins in the q/p_T axis, B is the magnetic field inside the solenoid, and c the speed of light. Replacing φ_{TMult} in the 2^B part of the Equation 5.10 and solving for r_{TMult} one gets the firmware multiplier for φ_{TMult} and r_{TMult}

$$\varphi_{TMult} = \frac{2^A}{\Delta_\varphi} \quad r_{TMult} = \frac{\Delta_{q/p_T}}{\Delta_\varphi} \cdot 2^{A-B} \quad (5.12)$$

In this way both multipliers are related to each other with a power of 2. In order to have the biggest dynamic range possible $A - B$ has been chosen to be the number of bits minus one, used to represent digitally the value of r_T , similarly for A in the φ_{TMult} .

The multiplication factor for z can as well be derived from the existing r_{TMult} by a factor of 2. This is done with the aim of representing the whole z range in which detector modules are present, which is about ± 3200 mm. The input data format with the aforementioned firmware multipliers, resolution, and ranges are given in Table 5.2. It is important to note that φ_O is representing the angular space of two neighboring detector nonants and later φ_S in the output format, shown in Table 5.3, represents only the angular range of one processing nonant as defined in Section 3.5 keeping the same digital multipliers but needing one bit less.

Table 5.2: GP firmware input format, where the z multiplier is related to that of r_T by a power of 2 to simplify the downstream firmware. The φ_O and r_T multipliers depend on the number of HT divisions in the parameter space and the selected number of bits for the digitization.

Name	Unit	Multiplier	Resolution	Range	Bits
φ_O	rad	23468.35	0.000043	-0.6981317 to 0.6981317	15
r_T	mm	3.9752	0.25156	-515.191 to 515.191	12
z	mm	1.987613	0.50312	-4121.529 to 4121.529	14
$bend$	strip	4	0.25	-8 to 7.75	6
stub valid					1
total					48

Table 5.3: GP firmware output format. The bend information is transmitted to the HT in terms of the minimum and maximum possible mBin.

Name	Unit	Multiplier	Resolution	Range	Bits
φ_S	rad	23468.35	0.000043	-0.349066 to 0.349066	14
r_T	mm	3.9752	0.25156	-515.191 to 515.191	12
$layerID$		1	1	0 to 7	3
z	mm	1.987613	0.50312	-4121.529 to 4121.529	14
$\Delta_{q/p_T}min$	1/mm	1/mBin	mBin	-0.03388 to 0.03388	5
$\Delta_{q/p_T}max$	1/mm	1/mBin	mBin	-0.03388 to 0.03388	5
$sub\eta_{Low}$		1	1	0 to 1	1
$sub\eta_{High}$		1	1	0 to 1	1
stub valid					1
Total					57

The GP kernel block was initially written in the high-level synthesis language MaxJ [74] and as part of this thesis ported to VHDL with improved performance and better matching with the corresponding C++ emulation code. The tracking performance for CMSSW $t\bar{t} + 200$ PU events is 99.88 % with an improvement of 0.05 % representing the not negligible 1000 extra tracks found in only the first 100 events of simulation, therefore indicating its significance in tracking performance. The VHDL module in comparison with the MaxJ implementation utilizes ~ 25 % less DSPs and requires ~ 50 % fewer clock cycles to determine the corresponding subsector of a stub in the octant and flat barrel configuration. A sample implementation showing the difference in area usage between the MaxJ implementation and the one reported in this work is shown in Figure 5.2. Furthermore, the GP kernel was extended to support the division of the tracker into nonants and the inclusion of tilted barrel modules.

An example function which precalculates the angle from the interaction point for each of the eta regions is shown in Listing 5.1. The function takes the pseudorapidity boundaries listed in Table 5.1 and calculates the angles using the Equation 5.1. The resulting values are stored as constants and used together with Equation 5.2 to assign each stub to a given η sector. This method of precalculating values is used throughout the various calculations of the GP Kernel block implemented fully in VHDL. The code is also written with parametric values, in a generic way, so that later updates of the code, like the DSP number-of-bits, are easy to adjust accordingly.




Figure 5.2: Comparison implementation of the Geometrical Processor processing stage in a Xilinx Virtex-7 XC7VX690T FPGA [75] for the flat barrel configuration and octant division of the outer tracker. a) MaxJ implementation. b) VHDL implementation presented in this work.

Listing 5.1: VHDL code to generating precomputed values to asses the boundaries in z especified by the selected η regions and allow the asignment of a η sector for a given r and z coordinate.

```

1
2 function generate_tanlSec return signed_1d_array is
3   variable etaSec, tanlSec : real;
4   variable int_tanlSec : integer;
5   variable val : signed_1d_array
6       ((N_ETA_NODES / 2 - 1) downto 0)
7       (DSP_A'high downto 0);
8 begin
9   for i in 0 to (N_ETA_NODES / 2 - 1) loop
10     etaSec := etaSecBounds(i);
11     tanlSec := 1.0 / TAN(2.0*ARCTAN(EXP(-etaSec)));
12     int_tanlSec := integer(tanlSec * rTBase
13                           / zBase / tanlSecBase);
14     val(i) := to_signed(int_tanlSec, val(i)'length);
15   end loop;
16   return val;
17 end generate_tanlSec;

```

The GP routing block is implemented as a highly pipelined three-stage network. It can route stubs from any number of inputs, currently configured with 48 inputs as

shown in Figure 5.3, but implementations can have up to 72 inputs, one per DTC (with up to 36 DTCs assumed in each of the two detector nonants from which the GP receives data). It routes stubs to one of the 36 outputs, each corresponding to a subsector. The first layer organizes stubs into six groups of three subsectors in η , which later each is arranged according to their final η subsector in the second layer. The third layer routes the stubs according to their φ subsector. Each routing block is highly configurable, and can easily be adapted for alternative subsector boundaries.




Figure 5.3: Block diagram of the GP Router. Three stages perform the routing of the stubs to its corresponding output link depending on the pre-calculated address from the GP kernel, the first stage does a coarse sort in η , followed by the second stage which does a fine sort in η and finally the third stage performs a fine sort in φ . The output links on each of the blocks represent the η and φ subsector number. At the output there are 36 distinct subsectors, one per output link on which the HT is independently performed.

The GP firmware corresponding to an entire processing nonant has been implemented within different FPGA devices including a single Xilinx Virtex-7 XC7VX690T FPGA used in the hardware demonstrator shown in Section 6.1. The FPGA resource usage is shown in Table 5.4. An implementation running at 240 MHz has a latency of 58 ns for the kernel and 193 ns for the routing blocks. This processing latency is fixed and independent of pileup or occupancy. A version of the GP kernel and router running at 480 MHz uses extra pipelining registers to meet timing but has an overall latency reduction of 60 ns.

Table 5.4: Resource usage of each GP kernel block (with 48 required per TFP, depending on the number of inputs) and the entire GP routing block implemented on the Xilinx Virtex-7 XC7VX690T FPGA [75] running at 480 MHz. The usage as a percentage of the total resources in the device is shown in parenthesis.

	MaxJ Flat Barrel Kernel block	Flat Barrel Kernel block	Tilted Barrel Kernel block	Routing block
LUTs	1762 (0.41%)	1971 (0.45%)	2050 (0.47%)	27700 (6.4%)
LUTRAM	112 (0.06%)	81 (0.05%)	117 (0.07%)	0 (0.0%)
DSPs	31 (0.86%)	22 (0.61%)	23 (0.64%)	0 (0.0%)
FFs	2536 (0.29%)	2576 (0.30%)	2046 (0.24%)	89531 (10.3%)
BRAM	0.5 (0.03%)	1 (0.07%)	1 (0.07%)	174 (11.8%)

From the Table 5.4 can be observed how the update of the algorithm from the flat barrel and octant division to the tilted geometry with a nonant division takes a negligible impact on the amount of logic needed, it increases the DPS number just by one, dedicated to calculate the φ sector of the tilted modules specifically.

5.2 The Duplicate Removal Algorithms for the TMTT approach

Removing duplicated tracks can be performed in two different locations along the TMTT Track Finding Processor (TFP), before and after the fitting stage. In the hardware demonstrator described in Section 6.1, the Duplicate Removal (DR) algorithm is the last of four processing stages. Several duplicate removal architectures and designs were explored in terms of resource usage, processing latency, complexity, and tracking efficiency. In this section, two of those are described in more detail.

5.2.1 Pair-Wise Track comparison Duplicate Algorithm

Stubs are duplicated and assigned to multiple sectors or track candidates in several locations of the GP and HT firmware. This leads to the possibility to identify multiple tracks containing the same fundamental information but with different initial conditions. A very simple approach to find duplicated tracks is to compare them in a pair-wise fashion with each other, until all possible combinations have

been compared. In software, this is very easy to do with nested `for` loops, in HDL logic however, one must consider latency and particularly the amount of resources needed for such ‘brute force’ approach.

The pair-wise duplicate removal algorithm can be described by comparing all tracks in a given event, two tracks at a time, for stubs in common between them. If the same stubs are present in more than 5 tracker layers, then those tracks are considered to be duplicates and one is eliminated. The track with stubs in more layers is preserved. If both tracks have the same number of layers, then the track with the highest p_T is kept. This algorithm however has several disadvantages. It implies that an event with N candidates would need to perform $N(N - 1)/2$ comparisons. Fortunately, N is limited by the length of the time-multiplex period and the maximum number of stubs possibly assigned to a given track. Several alternative comparison criteria were explored in software and found to have either lower efficiency or to be more resource-intensive when implemented in HDL logic.




Figure 5.4: Pair-Wise Duplicate Removal algorithm architecture as implemented in the HDL firmware description. The ‘Data Handler’ receives all stubs sequentially every clock cycle and stores them in the ‘Track/Stub Mem’ memory, selected data like the ‘stubID’ and track parameters are sent to 35 identical cells connected in a chain where each cell stores one track. Two memories per cell are required to handle consecutive events. After a fixed delay, the tracks are read starting from cell 1, by retrieving the full stub information from the memory and send it to the output if found to not be a duplicate.

The firmware implementation consists of an input module called ‘Data Handler’ which has the task to identify tracks from the incoming flow of stubs, one every clock cycle. It stores the full stub information in the ‘Track/Stub Mem’ memory, forwarding the track parameters and stubIDs to a chain of 35 identical ‘cells’. Each

cell stores in memory the stubs of a particular track per event. A maximum of 35 tracks per subsector have been found in simulation and are therefore used here as an upper limit. Each cell contains two 256-bit ‘stubMem’ memories. There is a maximum of 210 cycles per event, therefore stubIDs cannot be larger than that. Stubs are given a unique ‘stubID’ in the HT which is used as an address in the 256-bit memories. When a cell is empty, it uses the stubIDs to mark a ‘1’ in the particular location of the stubMem. If the cell is not empty it uses the stubID to see if the track passing by contains the same stubs as the track stored in the cell. When a given track j arrives, it is propagated through the chain of cells (1 to $j - 1$), representing previously stored tracks. It is then compared to all of them and finally if found to be unique, it is stored in the cell j .




Figure 5.5: Implementation of the Duplicate Removal algorithm based on pairwise comparisons between stubs inside tracks as implemented in a Xilinx Virtex-7 XC7VX690T FPGA [75]. No DR Router is implemented as the DR algorithm is implemented right after the HT and therefore each input link contains only information for an specific subsector.

This implementation suffers from the very high latency of a whole time-multiplex period, about 900 ns when running at 240 MHz. It requires that all data from a given subsector is present as input for the algorithm, therefore needing to sort the tracks beforehand. The biggest problem is the large number of resources that are needed to implement the algorithm. Despite implementation optimizations like instantiating specific FPGA primitives to avoid memories being synthesized in logic elements, the implementation still uses significant amount of resources. An example implementation of the described algorithm without routing stage is shown in Figure 5.5 and the resource usage is reported in Table 5.5.

Due to the non-deterministic nature of the GP router, stubs can be organized in any way. Similarly, higher p_T tracks can be found in any subsector. To properly

Table 5.5: Resource usage of a single Duplicate Removal block as implemented in the Xilinx Virtex-7 XC7VX690T FPGA [75]. The usage as a percentage of the available resources in the device are shown in parenthesis. These values do not include any kind of routing network or infrastructure firmware.

	One Pair-Wise Duplicate Removal block	Duplicate Removal block for one TFP
LUTs	4864 (1.12%)	175104 (40.42%)
LUTRAM	379 (0.22%)	13644 (7.92%)
DSPs	0 (0.0%)	0 (0.0%)
FFs	4224 (0.49%)	152064 (17.64%)
BRAM	4.5 (0.31%)	162 (11.16%)

assess the potential effectiveness of the duplicate removal algorithm, a modification was put in place to pre-sort the input tracks by p_T . Thus, software and hardware comparisons could be made with a higher degree of accuracy.




Figure 5.6: Firmware vs. software comparison for the Pair-Wise Duplicate Removal Algorithm run on a Xilinx Virtex-7 XC7VX690T FPGA and compared with its C++ CMSSW software implementation. Higher agreement is difficult to reach due to slight differences between firmware and software algorithm descriptions.

With the aim to improve the latency needed by the algorithm, restrictions regarding the range in which tracks are compared were put in place, with a pre-sorted input data, the first tracks are more valuable than the proceeding ones. The latency of the algorithm no longer needed to be an entire time-multiplex period, but instead a fixed selectable value which was found in simulation to be 58 clock cycles. After the fixed time, the tracks in the first cells are readout if they were not yet found to be duplicates. This firmware was implemented and compared to its software

implementation, results are shown in Figure 5.6.

5.2.2 A Duplicate Removal Algorithm based on the HT parameter space

The Hough-space based Duplicate Removal algorithm is aimed to run after the Kalman Filter stage where more than half the track candidates are duplicate tracks created at the Hough Transform processing stage. The goal of the DR algorithm is to eliminate them. During the formation of tracks, stubs are assigned to multiple HT cells due to granularity effects. Several tracks are formed in different HT cells containing the same stubs. In Figure 5.7 this process is shown, where three distinct HT cells, marked in yellow and green, contain information of four or more stubs depicted by the blue lines, therefore creating three track candidates, all having the same stubs. These three tracks after being fitted by the Kalman filter will contain track parameters matching the central green cell where all the stubs intersect, therefore yielding three identical tracks with exact same helix parameters.




Figure 5.7: r - φ Hough Transform showing formation of duplicates. The green cell represents the genuine track-candidate, whereas the yellow cells depict duplicate track candidates generated within the HT by the same set of stubs.

A simple duplicate removal algorithm has been designed with the knowledge of how tracks are being created multiple times by the HT and later fitted to a single HT cell by the KF. The algorithm can be described as follows: Track parameters in the Hough-space are compared before and after the Kalman Filter, any track which is not consistent with the cell it was initially found in is eliminated. In the example of Figure 5.7 the tracks in the yellow cells will be eliminated and the track in the green cell will be kept.

The main advantage of this algorithm is that it is able to identify duplicated tracks by only looking at information contained within the track, without needing to compare it to other tracks or the stubs contained within pairs of tracks. However,

this algorithm loses a few percent in efficiency due to resolution effects. A second phase is used to recover that efficiency, where tracks which were found by the HT in a particular cell and later fitted to another one, are recovered if they belong to a cell not yet occupied in the Hough-space. Recovered tracks must also be within one cell of distance from the original cell it was found. In Figure 5.8 is shown the most common mode of forming duplicate tracks on those sectors which are towards the high pseudorapidity regions, such as eta sector 0. Each long trace of candidate tracks is represented in the Hough-space as multiple cells aligned and activated, they most likely contain the same stubs inside each of the candidates. After the fitter, it is often the case that those will be located in a single cell containing several tracks. The DR makes sure that only one track per cell is allowed.




Figure 5.8: Example of the Hough-space Duplicate Removal algorithm for ϕ subsector 0 and η subsector 0 with all nonants aggregated in one plot. a) Data output of the HT shows several HT cells being activated but most of them contain the same stubs. b) The DR algorithm output after the fitting stage, several tracks remained in only few bins, only one track per parameter space is allowed.

In the TFP data is initially routed to multiple HT blocks which perform track finding in parallel and independently from other subsectors, each HT contains only information about a given subsector. After the HT, the KF is implemented in a way that it distributes data to any given free independent KF-worker unit, as these units need significant amount of resources, there are less units than subsectors. Each KF-worker also receive information from multiple events to better deal with high occupancy events containing jets (high energetic particles surrounded by several other particles with seemingly similar direction). At the end of the KF and before the Duplicate Removal block a DR Router was designed based on the GP Router to sort the data and allow the DR algorithm to be performed independently and in parallel for each subsector.

Table 5.6: Resource usage of the DR router block configured with 36 possible inputs and 18 or 3 possible outputs depending if the DR algorithm is performed on a single sector or six subsectors at a time. The usage as a percentage of the Xilinx Virtex-7 XC7VX690T available resources are shown in parenthesis.

	36 to 18 router block	36 to 3 router block
LUTs	25259 (5.83%)	16611 (3.83%)
DSPs	0 (0.0%)	0 (0.0%)
FFs	38436 (4.44%)	24655 (2.85%)
BRAM	144 (9.8%)	108 (7.35%)

Several configurations of the DR-Router were implemented depending on the number of KF-workers being used, in the Table 5.6 is shown a couple of those configurations, The final configuration used in the demonstrator system has 18 inputs and 6 outputs depicted in Figure 5.9. The specific η sectors contained in each of the connecting links are mentioned in the image, two layers are used to sort and merge the different data streams. Two sets of the depicted router are needed to fully cover all 36 subsectors, the router does not need to be wider. The KF-workers from 0 to 8 have information for all η subsectors but only for φ sector 0, and KF-workers from 9 to 17 contain similarly all η subsector but only φ sector 1.




Figure 5.9: Architecture of the DR Router where 9 KF-workers containing information of a single φ sector and 18 η sectors. Data is sorted to interleaved η sectors to reduce high occupancy in neighboring sectors due to high occupancy events containing high energetic particles. Two instances of the shown block need to be instantiated to cover both φ possibilities in the 36 subsectors.

The Figure 5.10 shows the implementation of the duplicate removal algorithm based on the Hough-space. Here the DR block processes the tracks found by the KF in six subsectors, six of these DR blocks must be instantiated in order to process the tracks from the 36 subsectors in a processing nonant. Designing the DR block to process six subsectors instead of one minimizes further the resource usage.




Figure 5.10: Architecture of the Duplicate Removal algorithm implementation based on the Hough-space. A single DR logic block is shown, which processes the KF tracks from six subsectors. Therefore, six such blocks are needed to process all 36 subsectors in the processing nonant.

Within the DR block, a ‘HT Matrix’ representing the HT arrays of the six subsectors is implemented in a 18 Kb memory, and is addressed using the subsector number and the $(q/p_T, \varphi)$ cell location within the HT array. Any fitted track that is flagged as ‘consistent’ (i.e. its fitted helix parameters correspond to the same HT cell as the HT originally found the track in) is forwarded to the output channel, and in addition, the corresponding matrix address is marked. In contrast, tracks which are ‘inconsistent’ are added to a FIFO (named ‘Rejected FIFO’). The address of each cell marked in the matrix is also stored in the corresponding ‘clearFIFO’ to reset the matrix A or B in a pin-pong fashion.

There are two methods to recover the tracks in a second phase as shown in Figure 5.11, as soon as all tracks from the KF have arrived, informed by a flag in the data format or by using a pre-configured delay based on the tail of the number of tracks per event found in simulation. The inconsistent tracks are read out from the ‘Rejected FIFO’ and compared to the already filled matrix. If one track has fitted parameters corresponding to a HT cell location not yet marked in the matrix, the track is restored by forwarding it to the output channel and marking the corresponding

address in the matrix. Another condition to rescue the tracks is that their HT bin needs to be only one distance away from the initial seed.




Figure 5.11: Different second phase implementations of the Duplicate Removal algorithm based on Hough-space. Rejected tracks are recovered only towards the end of the time-multiplexed period or right after receiving the last track for this event from the Kalman Filter flagged with a bit in the last track, in this case track 'C2'.

A full matrix reset is required before processing tracks from another LHC bunch crossing. Therefore, two matrices (labeled 'HT Matrix A' and 'HT Matrix B' in Figure 5.10) are instantiated, which take in turns to process alternate LHC events. Hence, there is always an active-matrix and a reset-matrix. Along with them, two 'clearFIFOs' are used, one for each matrix, to store the addresses that have been marked and therefore need to be cleared in preparation for a new event. Each matrix plus its corresponding clear FIFO occupy one 36 Kb memory block as shown in Figure 5.12.

The FIFO in which the 'inconsistent' tracks are temporarily stored uses two 36 Kb block RAMs. Therefore, a total of four 36 Kb block RAMs are used for the entire DR block design handling six subsectors. In addition to be a very lightweight design, it also has a low latency of just four clock cycles. Table 5.7 shows the total resource usage, including other types of resources.

Comparing the total resources used in the description of the Duplicate Removal Algorithm based on the Hough Transform Space for an entire TFP to that of a traditional approach comparing stubs in pairs of tracks, like in Section 5.2.1, there is an improvement by using only about 1% of the resources declared in the traditional approach.

The Duplicate removal Firmware implementations have been used to run several thousands of simulated physics events and compared the results with its C++ CMSSW counterpart. The algorithm is so simple and deterministic, that it can be described accurately in software, reaching a perfect matching with simulated events. Performance plots comparing the firmware and software ability to remove duplicates are




Figure 5.12: a) Implementation of the Duplicate Removal algorithm based on the Hough-space with multi-sector and the DR Router with 36 inputs and 3 outputs layouted in a Xilinx Virtex-7 XC7VX690T FPGA [75]. b) Zoomed in implementation highlighting one DR algorithm based on Hough-space with labeled BRAM resources according to Figure 5.10.

shown in Figure 5.13.

For this comparisons, the input data was extracted from the software framework after the Kalman Filter, then ran through both the firmware implementation in the FPGA and the software algorithm for later comparison. The algorithm efficiency can be seen in Figure 5.14, it is calculated only considering the Duplicate Removal stage, not the overall reconstruction efficiency as will be described in the Chapter 6.

5.3 Summary

The TMTT reconstruction algorithm is composed by four stages, where the Geometric Processor (GP) is located in the first position and the Duplicate Removal (DR) in the last. The GP divides the data into many different and independent regions that can be processed in parallel, it performs various calculations with the goal of offloading work from the HT and aid in its implementation within the chosen device constraints. The GP divides the processing nonant into 18 divisions in pseudorapidity (η) and 2 in the azimuthal angle (φ). The GP is fundamental to increase the parallelization and therefore reduce truncation effects of the downstream stages.

Two alternatives were presented for implementation of the Duplicate Removal




Figure 5.13: Firmware vs. Software comparison for the Hough-Space Duplicate Removal Algorithm run in a Xilinx Virtex-7 XC7VX690T FPGA and compared with its C++ CMSSW software implementation. Agreement can be seen to be 100 % in all three cases; Total number of tracks found vs. η region (a), ϕ region (b), and q/p_T (c) of the track.




Figure 5.14: Hough-Space Duplicate Removal Algorithm efficiency as a function of η . Here only the DR stage is considered for the efficiency calculation. Firmware and Software values overlap in 99.5 % tracking efficiency.

Table 5.7: Resource usage of a single Duplicate Removal block based on Hough-space for six subsectors, as implemented in the Xilinx Virtex-7 XC7VX690T FPGA [75]. The usage as a percentage of the available resources in the devices are shown in parenthesis. The entire TFP needs six of these DR blocks.

	One Hough-Space Duplicate Removal block	Duplicate Removal block for one TFP
LUTs	291 (0.1%)	1746 (0.69%)
DSPs	0 (0.0%)	0 (0.0%)
FFs	496 (0.1%)	2976 (0.34%)
BRAM	4 (0.3%)	24 (1.63%)

stage. First, the DR can be implemented after the HT and before the KF, where its goal is to reduce the number of tracks which contain the same stubs to lower the processing load on the KF. This implementation required substantial resource usage and long latency in the order of a time-multiplex period as presented earlier. Lastly, considering that the KF is properly load balanced and can handle the increased number of tracks without a pre-filter, the DR can be implemented at the output of the KF to lower the processing requirements of the downstream elements in the trigger system. This implementation relied on the HT-space to identify duplicates by only looking at individual tracks, its implementation uses only 1 % of the resources compared to the ‘brute-force’ option of comparing tracks and stubs. The DR firmware implementation has a latency of only four clock cycles. Both the GP and DR based on HT-space are part of the ultimate configuration of the Demonstrator system presented in the next chapter.

6 The TMTT Hardware Demonstrator

This Chapter contains results for the TMTT chain explained in Section 4.2 and presented previously in collaboration [1], the firmware implementation of the GP and DR processing stages were developed by the author according to the descriptions in Chapter 5. In this chapter, a description to the hardware demonstrator and software setup will be given, different track reconstruction results will be presented, including efficiency, resolution, data rates, resource usage, and robustness against failure modes.

6.1 The Hardware Demonstrator Slice

A demonstration system using a μ TCA [76] crate was built to implement a section of the L1 track finder on existing FPGA-based hardware. The main objectives of the demonstration system are to test and validate the firmware implementation of the TMTT algorithm, to ensure its latency is within the allocated budget, to gauge whether the resource usage is within reasonable limits, and to assess the performance and resolution of fixed point logic.

Figure 6.1 depicts the firmware components and their connections, which correspond to the components mentioned in Section 4.2. The demonstrator chain employ eight Master Processor 7 (MP7) [77] boards. Each of the two boards labeled ‘sources’, represents data from a collection of up to 36 DTCs. Each source board is designed as a large buffer for storing stub data from a detector octant, with data loaded directly from the simulation framework via IPBus [78]. Each output stream from the source boards represents a separate DTC capable of playing up to 30 consecutive events via the demonstrator by injecting pre-formatted 48-bit stubs into the Geometric Processor. To simulate how data from two adjacent detector slice can feed a single TFP for tracks that cross the detector boundary, two sources are needed. The TFP is implemented on five different boards: one for the GP, two for the HT, and two more for the KF and DR. The sink is an additional board that captures the track finder output for later read out via IPBus. Furthermore, three additional boards are placed in the demonstrator crate for standalone testing of firmware blocks or parallel data




Figure 6.1: The demonstrator system is constructed by interconnecting eight MP7 boards as shown in the diagram. Boards are placed in a single μ TCA crate and marked according to the diagram. The direction of data flow is from left to right [1].

taking alongside the entire chain. Figure 6.1 in the bottom shows a picture of the demonstrator crate.

The track finder slice corresponds to one Track Finding Processor (TFP) as defined in Section 4.2. Initially, the tracker was divided in eight physical divisions, following the flat barrel configuration. Later, this changed to nine physical divisions and a tilted configuration of the barrel modules. The demonstrator slice was originally built using the first configuration, where each TFP processes data from $1/8$ of the tracker in φ , and all of the tracker in η , since each TFP is independent of the others, data for each φ -octant can be processed sequentially, allowing the entire event to be reconstructed in hardware. Updates to the algorithms were later done to run using the nonant configuration.

The demonstrator system is located at the Tracker Integration Facility (TIF) at CERN, it is composed by one dual-star μ TCA crate. Inside the crate it contains a commercial network adapter as the μ TCA Carrier Hub (MCH) for Gigabit Ethernet communication through the back-plane, and a CMS special auxiliary card known as the AMC13 [79] for synchronization, timing, and power. TFP algorithms were implemented on five MP7 boards, each with a Virtex-7 FPGA and several Avago Technologies MiniPOD optical transmitters/receivers [80]. The links in the demon-

strator are set to run at 10 Gb/s with 8b/10b encoding for an effective 8 Gb/s data transfer rate, half the required throughput to model a time-multiplex of 18, but can be used to simulate a time-multiplex of 36 for the demonstrator system.

The MP7 includes infrastructure tools such as the core firmware for managing transceiver serialization/de-serialization, data buffering, I/O formatting, board and clock setup, and external communication through the Gigabit Ethernet interface. The firmware in charge of these tasks is separate from the track finding firmware. This makes it possible to quickly construct a device like the demonstrator. Individual track finding blocks in the demonstrator are run on single MP7 boards that are daisy-chained together with high-speed optical fibers. The partition of the demonstrator in this fashion allows for the easy division of firmware responsibilities across individuals, provided I/O formats between the processing blocks are specified. Total FPGA resources can be extracted from the aggregated values from all the boards utilized in the demonstrator system.

6.2 Software Setup

CMS simulation software (CMSSW) is used to generate input data for the demonstrator, using Monte Carlo physics events generated under the HL-LHC conditions, typically top quark pair production with a Pileup (PU) of up to 200 proton-proton (pp) interactions per bunch crossing, the simulations include the modeling of particle interactions with the detector and the mechanism for stub generation within the detector modules. Stubs were initially generated following the flat tracker geometry with the division of the tracker in octants, as it was the default configuration at the time of the implementation. Selected results are presented as well for the scenario where the tracker follows the tilted barrel configuration and it is divided in nonants.

To insert stubs from these samples into the demonstrator chain, software built to produce and analyze the output of the hardware slice is used, converting the simulated events to text files before transmission over IPBus. Tracks that the demonstrator has reconstructed using these stubs are retrieved via IPBus at the end of the chain and saved for later study and comparison.

A software emulation of the hardware chain has been created, which in conjunction, can process the same integer formatted stub data used as input to the demonstrator to generate tracks for offline validation with hardware output. In order to model time-dependent results, the simulator employs fixed-point mathematical operations

where necessary and closely simulates the logic implemented within the FPGAs. However, since it is not a clock-cycle accurate emulation, minor variations between hardware and emulation are to be anticipated on occasion. The emulator code can be modified to use absolute floating-point precision for comparison if desired. The comparison software examines both emulated and hardware tracks for accuracy on an event-by-event basis.

6.3 Track Reconstruction Efficiency

The track reconstruction efficiency is calculated relative to all charged particles produced by the primary interaction that generated stubs in at least four layers of the tracker and are within the kinematic acceptance region ($p_T > 3 \text{ GeV}$, $|\eta| < 2.4$, $|z_0| < 30 \text{ cm}$ and $L_{xy} < 1 \text{ cm}$, where L_{xy} is the transverse distance from the beam line to the particle vertex). A charged particle is described as successfully reconstructed and contributing to efficiency if and only if the following conditions are met:

- the reconstructed track contains stubs in at least four different tracker layers associated with the truth particle;
- the reconstructed track contains no incorrect stubs (*i.e.* all its stubs were produced by the same particle).

The second condition only applies when the results for the full processing chain are mentioned. It is natural to find incorrect stubs when only a sub-section of the chain is used. Tracks that are successfully reconstructed are also referred to as ‘matched tracks’. Tracks not correctly associated are also called ‘fake tracks’. And tracks matching to more than one reconstructed track are defined as ‘duplicates’.

Table 6.1 depicts how tracking performance changes as data moves through the reconstruction chain, these results are obtained without applying the second of the two matching conditions listed above. The HT finds tracks with a high efficiency, but many of them are either fake or duplicate tracks. The KF removes the vast majority of fake tracks, while the duplicate removal algorithm removes nearly all duplicates. For the entire chain, all the tracks that satisfy the first matching requirement above also satisfy the second, particularly due to the ability of removing incorrect stubs in the KF.

The HT initially required stubs in a minimum of five tracker layers. However, this resulted in a loss of efficiency in the detector region between the barrel and the

Table 6.1: Track finding performance on simulated $t\bar{t}$ events with 200 PU for each stage of the demonstrator chain. The track finding efficiencies are according to the efficiency definitions given in the text. The average number of reconstructed tracks per event in the whole tracker, the number of fake and duplicates are also mentioned for the flat and tilted barrel configurations.

Flat barrel	Efficiency [%]	# of tracks	# of fakes	# of duplicates
HT	97.1	331	139	126
KF	95.1	190	27	103
DR	94.4	79	16	3
Full chain	94.4	79	16	3

Tilted barrel	Efficiency [%]	# of tracks	# of fakes	# of duplicates
HT	97.1	295	104	124
KF	96.3	159	16	84
DR	95.1	73	10	4
Full chain	95.1	73	10	4

endcap. As a result, in the corresponding sectors, this criterion was reduced from five to four layers, avoiding this efficiency loss. This updated description is used in all of the findings discussed in this section.

The mean tracking efficiency over all applicable $|\eta|$, in $t\bar{t}$ events with 200 PU, as measured in hardware, is 94.4 % for the flat barrel configuration and 95.1 % for the tilted barrel configuration. This number agrees with 99.5 % of the results generated by emulation. Figure 6.2 depicts the agreement as a feature of the particle kinematic properties. The efficiency of tracking from emulation is the same if floating-point or integer precision stub data are used.

As seen in Figure 6.3, the efficiency to reconstruct leptons in $t\bar{t}$ events exceeds 97 % for muons over the entire acceptance region but it is slightly lower for electrons. The loss of electron efficiency is to be anticipated, and it is caused primarily by bremsstrahlung effects, which cause the particle trajectory to deviate from the helix trajectory assumed by the tracking algorithm. Some of this efficiency loss should be recoverable. The KF algorithm, for example, can be modified to allow for multiple scattering. The agreement between hardware and emulation shows similar results for leptons.




Figure 6.2: Track reconstruction efficiency, as a function of p_T (a) and η (b), for tracks originating from the primary interaction in $t\bar{t}$ events with 200 PU, calculated in both hardware and emulation [1].




Figure 6.3: Track reconstruction efficiency, for electrons and muons as a function of p_T (a) and η (b). These results are obtained from emulation from $t\bar{t}$ events with 200 PU [1].




Figure 6.4: Track reconstruction efficiency as a function of η , for $t\bar{t}$ events at 200 PU. Tracks originating from the primary interaction (black dots), and tracks contained within a primary jet that has a total transverse momentum exceeding 100 GeV (red open circles), these results are obtained from emulation. a) No incorrect stubs are allowed on the track. b) Only one incorrect stub allowed [1].

The consistency of tracks in the center of dense jets is marginally degraded due to the increased probability of incorrect stubs being included in the track candidate. Figure 6.4 shows that there is a small efficiency loss when selecting on charged particles in jets with total transverse momentum greater than 100 GeV, especially in the region $|\eta| > 1$. This effect is reduced when the second matching condition is adjusted to enable reconstructed tracks with at most one incorrect stub to contribute to the efficiency, as shown in Figure 6.4. This suggests that the reduced track purity in these high-energy jets accounts for a large portion of the loss in the region $|\eta| > 1$. Better rejection of these incorrect stubs by the KF should help boost the overall track finder performance.

6.4 Track Parameter Resolution

The resolution of each of the four track parameters ($p_T, \varphi, \cot \theta, z_0$) is shown in Figure 6.5 for tracks reconstructed in hardware and emulation for $t\bar{t}$ events with 200 PU. It is possible to have a fifth track parameter (d_0) which is currently not generated by the KF, but could be a future improvement. The shown resolutions are similar to those obtained with the offline track reconstruction software [53], and show to have sufficient performance which is useful to the L1 trigger production. In general, there is reasonable agreement between hardware and emulation, with some variations due to the use of floating-point arithmetic in some areas of the emulator code. The reduction in resolution with increasing pseudorapidity is assumed as a result of

a combination of the decreased effective accuracy for hits in the endcap, and the influence of crossing more material by the particles.




Figure 6.5: Relative p_T resolution (a), φ resolution (b), z_0 resolution (c), and $\cot \theta$ resolution (d) determined by both hardware and emulation for tracks resulting from the primary interaction in $t\bar{t}$ events with 200 PU [1].

It is also useful to quantify the parameter resolutions for particles of varying transverse momenta. Figure 6.6 shows the resolution of the track parameters for single isolated muons determined only with the simulator. Multiple scattering effects predominate at low transverse momenta, as demonstrated by the approximation of the φ parameter, where the resolution is greater than 0.4 mrad for muons with $15 < p_T < 100$ GeV and between 0.7 mrad and 1.5 mrad for muons with $3 < p_T < 5$ GeV. At high pseudorapidities, similar results are observed for the $\cot \theta$ resolution. Multiple scattering limits the relative accuracy of the transverse momentum for muons with $p_T < 15$ GeV, but it degrades with increasing p_T due to the decreasing radius of track curvature and the limited detector granularity.

The precision of the four parameters in Figure 6.6 also contrasts favorably with the offline simulation [53], which can use all available data from the tracking system and more advanced reconstruction algorithms. In offline simulations, the φ resolution




Figure 6.6: Relative p_T resolution (a), ϕ resolution (b), z_0 resolution (c), and $\cot\theta$ resolution (d) for single isolated muons with $3 < p_T^\mu < 5$ GeV, $5 < p_T^\mu < 15$ GeV, and $15 < p_T^\mu < 100$ GeV. These outcomes are the product of emulation [1].

for 10 GeV muons going through the center of the tracker barrel is approximately 0.2 mrad, while the p_T resolution is $\sim 0.5\%$. The accuracy of the remaining two parameters, on the other hand, is more than an order of magnitude higher in offline simulations than in the demonstrator, primarily due to the addition of hit information from the pixel detector.

When the z_0 resolution of single isolated muons in Figure 6.6 is compared to earlier simulation studies of the track finder [61], the demonstrator tends to display roughly half the predicted precision in the barrel. The demonstrator system resolution degrades as a result of deciding to format the r and z stub coordinates too coarsely (using 10 and 12 bits, respectively), resulting in a reduction in the ultimate accuracy of the track parameters. Figure 6.7 demonstrates that enhancing the encoding of the stub coordinates by adding two extra bits to both r and z restores the missing accuracy, indicating that the z_0 resolution achieves ~ 1 mm for muons with $5 < p_T < 15$ GeV and $|\eta| < 2$. A slight change can also be seen in the $\cot\theta$ resolution. This modification can be made without affecting the efficiency of the rest

of the demonstrator system. Figure 6.7 also demonstrates that with this enhanced encoding scheme, the accuracy of all four parameters equals to that obtained by the demonstrator floating-point simulation.




Figure 6.7: Relative p_T resolution (a), φ resolution (b), z_0 resolution (c), and $\cot \theta$ resolution (d) measured for single isolated muons with $5 < p_T^\mu < 15$ GeV obtained from emulation, using different levels of precision in simulation: default encoding (10-bit r , 12-bit z , 15-bit φ stub coordinates); improved encoding (12-bit r , 14-bit z , 15-bit φ stub coordinates); and full floating-point simulation [1].

6.5 Data Rates

As shown in Figure 6.8, the number of tracks reconstructed per event increases with increasing pileup. On average, 79 tracks per event are reconstructed in $t\bar{t}$ events with 200 PU. The demonstrator system can easily accommodate the high data rates present in these events. The distribution of the number of stubs per event transmitted from the GP to the HT in each sub-sector is seen in Figure 6.9(a). If these stubs could not all be transmitted within the 900 ns time frame specified by the time-multiplexed factor in the demonstrator, data would be truncated. Since stubs from each sector




Figure 6.8: Total number of reconstructed tracks per event when processing $t\bar{t}$ events superimposed with 0, 140, and 200 PU. These results are obtained from emulation, The effects of truncation, caused by excess data flow through the system, are both included and excluded [1].

are transmitted at 240 MHz, there is a potential limit of 216 stubs per sub-sector, but in the present system the limit is closer to 175, due to gaps in the output data. This maximum exceeds the average data rate by nearly a factor of two, so truncation effects are minimal in this section of the system: At 200 PU, 0.3 % of stubs are lost in $t\bar{t}$ events, resulting in a 0.5 % loss of tracking efficiency.




Figure 6.9: Data rates in two crucial sectors in the system for $t\bar{t}$ with 200 PU events. a) The amount of stubs transmitted from the GP to the HT per event and subsector. If this value is higher than 175, truncation effects arise. b) The number of reconstructed tracks from the HT for each subsector and event [1].

The number of reconstructed tracks per event in each sector generated by the HT is seen in Figure 6.9(b). It is important to note that 70 % of the subsectors have no reconstructed tracks, thanks to the fact that the 3 GeV p_T threshold used for track

reconstruction is very good at removing tracks from PU associations, and 97.5 % percent of sectors have less than ten reconstructed tracks. Since the HT typically assigns about seven stubs to each track, outputting tracks from the HT under the time-multiplexed limit is generally not a problem. The only difficult case is caused by collimated, high energetic jets produced by the $t\bar{t}$ mechanism itself, which produce several particles and stubs in a narrow angular region, accounting for the tails shown in Figure 6.9. This problem is addressed by the load-balancing implemented after the HT. When processing $t\bar{t}$ events with 200 PU, the reduction in tracking efficiency due to HT output truncation is less than 0.1 %.




Figure 6.10: Track reconstruction efficiency in $t\bar{t}$ events with 200 PU originating from the primary interaction, as a function of p_T (a) and η (b) showing the effect due to truncation. These results are obtained by emulation with and without truncation effects [1].

In the KF, the latency is set to a value that allows it to allocate stubs from four tracker layers in almost all tracks. When processing $t\bar{t}$ events at PU of 200, less than 0.1 % of efficiency is lost in the KF, mostly due to events containing high-energy jets. The loss in tracking efficiency due to truncation effects for the whole tracking chain is less than 0.6 % when processing $t\bar{t}$ events with 200 PU, as shown in Figure 6.10.

As seen in Table 6.2, the tilted barrel geometry reduces significantly the output data rate of the HT, the number of fake tracks is also reduced considerably. This is mostly due to the lower amount of stubs generated in the innermost tilted barrel layers. The total number of tracks in both configurations are rather similar, showing that the following reconstruction stages, the KF and the DR are performing as expected reducing the number of fake stubs, and the number of duplicated tracks. The table also shows that the efficiency of both geometries at the end of the entire chain is identical. This indicates that substantial savings in terms of resources could potentially be made as there are less data to be processed on each subsequent stage.

Table 6.2: Comparison of the performance of flat and tilted barrel tracker geometries reconstructing single muons ($p_T(\mu) = 10 \text{ GeV}$) superimposed with 200 PU. The average number of tracks identified and the tracking efficiency are given.

	Flat geometry	Tilted geometry
# of tracks after HT	229	161
# of fakes after HT	92	35
# of tracks after full chain	55	48
# of fakes after full chain	9	4
Efficiency after full chain	97.3 %	97.3 %

6.6 Flexibility and Robustness of the System

The extremely low truncation rates reported in Section 6.5 at each point of the demonstrator chain under challenging conditions are an indication of the ample margin of the system. To further evaluate the effect on increased data flow in the system, particularly on the HT and KF stages, one can reduce the criteria to form tracks from five detector layers to four over the whole tracking acceptance simulating for example a dead detector layer. By doing this, there is an increase number of track candidates being generated in the HT by a factor of 3.6, equal to 1 190 track candidates for a $t\bar{t}$ 200 PU event.

As seen in the Figure 6.11(a), the demonstrator chain will process this much higher data rate while keeping the tracking efficiency loss due to truncation to around 1.7 %. This minor additional loss of $\sim 1 \%$ in comparison to the nominal configuration is not caused by the HT, but is instead due to the KF inability to assign stubs from four layers to all tracks in time. If necessary, increasing the accumulation time in the KF will recover this loss at the cost of latency, but other architecture optimizations adding extra KF-worker units or sorting tracks more efficiently would mean this is not necessary.

To understand the level of performance in a situation where a fraction of modules in the tracker do not generate stubs, the demonstrator was evaluated on samples that simulated this scenario. Figure 6.11(b) depicts the localized efficiency loss assumed when all modules on barrel layer four, between $-1 < \eta < 0$ and $0 < \varphi < \pi$, are prevented from producing stubs in simulation. As seen in Figure 6.11(b), this efficiency loss can be restored by lowering the threshold criterion on the amount of hit layers in the HT from five to four in the affected (η, φ) subsectors only. As seen in




Figure 6.11: a) Track reconstruction efficiency as a function of η when processing $t\bar{t}$ events with 200 PU and a global threshold criterion of four hit layers in the HT, measured in both hardware and emulation. To demonstrate the total efficiency loss due to hardware truncation, the emulation result excludes truncation effects. b) Track reconstruction efficiency as a function of η , measured in emulation, when processing $t\bar{t}$ events with 200 PU, where the tracker is affected by a failure of all modules in the region $-1 < \eta < 0$ and $0 < \varphi < \pi$ of barrel layer 4. The results are compared before (black dots) and after (red open circles), as described in the text, relaxing the threshold criterion on the number of hit layers in the affected region [1].

Table 6.3, changing the threshold results in just a minor rise in data rate. The spike is due to an increase in fake tracks as a result of the lower threshold.

Table 6.3: Mean number of tracks from the HT when processing $t\bar{t}$ events with 200 PU considering a module failure.

No module loss	With module losses	
	Before recovery	After recovery
330	304	347

The detector modules may be configured to generate stubs for particles with 2 GeV, track reconstruction under those conditions may be of interest for the L1 trigger despite the increase data rate generated. The demonstrator system has been used to study the effects of that scenario. It is possible to extend the reconstruction range in the demonstrator from 3 GeV down to 2 GeV by changing the GP parameters which guarantee that adequate duplications happen in the $r - \varphi$ plane, modifications in the HT granularity are also required increasing the q/p_T columns by 50 % to maintain the precision of the track estimates. These modifications in the HT relate to an increase in FPGA resource usage by 50 % and an increase in output data rate by a factor of 2.2. Comparing the results with the ones previously obtained for $p_T > 3$ GeV, there is a

loss of tracking efficiency in the $2 < p_T < 2.7$ GeV range, due to multiple scattering where stubs do not always converge within a single HT cell and thus fail to meet the threshold requirement to generate track candidates. To tackle this limitation, it is possible to reduce the resolution of the HT along q/p_T and φ_T by a factor of two only for the range $2 < p_T < 3$ GeV. This variable precision HT, which was introduced in a separately firmware, is able to recover some of the loss (increasing efficiency from 65 % to 75 % in the range of $2 < p_T < 2.7$ GeV). If required, further changes to the HT, such as p_T based thresholds, should be able to recover further efficiency. Optimization of the KF to the new minimum p_T threshold can also result in higher reconstruction efficiency.

The efficiency and application of the bend filter cut used inside the HT and calculated in the GP has also been investigated. A single parameter in the firmware allows you to easily change the assumed bend resolution. In the unlikely event that the stub bend information becomes unreliable, the filter can be switched off completely. As shown in Table 6.4, disabling the bend filter increases the rate of misreconstructed and duplicate track candidates provided by the HT but does not reduce efficiency or create truncation in the HT processing. The overall performance of the device is comparable to that defined by Figure 6.11.

Table 6.4: The mean number of tracks found and the tracking efficiency after the HT, with and without application of the bend filter, for tracks originating from the primary interaction in $t\bar{t}$ events with 200 PU

	Bend filter	No bend filter
Tracking efficiency after Hough Transform	97.1	97.9
Track candidates after Hough Transform	331	1285

6.7 Latency

In Table 6.5 the latency measurements for each of the processing blocks in the demonstrator chain are shown, similarly the value for the entire chain is reported. The sum of all the individual latencies is equal to the latency observed for the whole chain. Latency numbers include the delay caused by the optical transmission channel and the serialization/de-serialization (SERDES) block in each of the links. The overall system latency is constant, regardless of PU or the number of tracks in each subsector. The table indicates not just the time difference between the first stub entering the system and the first track leaving it, but also the time difference

between the first stub entering and the last track leaving. The L1 trigger, which will sit downstream of the track-finder device, is interested in each of these latency specifications.

Table 6.5: Latency of the each of the firmware components of the track reconstruction chain, including serialization/de-serialization (SERDES) and optical transmission delays between each board.

System latency	Latency [ns]
SERDES + optical length 1	143
Geometric Processor	251
SERDES + optical length 2	144
Hough Transform	1025
SERDES + optical length 3	129
Kalman Filter	1641
Duplicate Removal	17
SERDES + optical length 4	129
Total: First out - First in	3479
Last out - First out	225
Total: Last out - First in	3704

6.8 FPGA Resource Usage

By building the hardware demonstrator out of several MP7 boards, the logic constraints imposed by a single device is avoided. On the other hand, it is critical to keep the overall resource consumption realistic so that a final system can be installed at a reasonable cost using FPGAs that are anticipated to be available on the production timescale. The total FPGA resource use for each demonstrator component is shown in Table 6.6 (where the numbers given for the HT and KF implementations are summed across the two boards used for each component). The sum of the four components provides the resources required to demonstrate the functionality of one entire TFP with a time-multiplexed factor of 36. The MP7 core infrastructure firmware, which is necessary for board configuration, connection buffering, and error checking, is also run on each FPGA in the demonstrator. This firmware was created for the CMS calorimeter trigger, and although it does not make up a large portion of the logic in the TFP (as seen in Table 6.6), it is assumed that with some

optimization, it could be reduced in size while still providing the functionality required by the track-finder. To meet timing and routing constraints in the Virtex-7, designs often prioritized block RAM over LUT-based distributed memory. This balance will need to be re-adjusted in the future as the design is updated for newer FPGAs with different resources ratio. In the lower part of Table 6.6 are various potential devices which can be used to implement the TFP algorithms in newer FPGA architectures (*i.e.* Xilinx Ultrascale+ [81]), in the following chapters several hardware developments will be presented where these devices are used to build prototype hardware to continue developing the rest of the necessary firmware and software components of the Phase-2 CMS tracker readout electronics system.

Table 6.6: Total resource usage for the demonstrator TFP (with time-multiplexed factor of 36), as implemented in the Xilinx Virtex-7 XC7VX690T FPGA [75]. The resources required to build a full TFP are the sum of the numbers in the four rows labeled GP, HT, KF, and DR. In the bottom are listed the total resources available in some particular FPGA devices. Important to note, some devices contain 'UltraRAM' memory, where the total memory value including it is shown between brackets.

	LUTs [10^3]	DSPs	FFs [10^3]	BRAM (36 Kb)
GP	121	1056	205	222
HT	244	2304	299	1188
KF	430	5112	363	1984
DR	2	0	3	24
Infra. per MP7	90	0	91	291
TFP Total (excl. infra.)	795	8472	870	3392
TFP Total (incl. infra.)	1245	8472	1325	4857
V7-690	433	3600	866	1470
KU-15P	523	1968	1045	984 (2008)
VU-9P	1182	6840	2364	2160 (9840)
VU-13P	1728	12288	3456	2688 (12928)

6.9 Summary

The results presented in this chapter provide clear confirmation that the TMTT reconstruction algorithm is feasible to be implemented on hardware, it utilizes a reasonable amount of FPGA resources and meets the target latency of 4 μ s with ex-

cellent performance. In general, the resolution performance for all track parameters are in accordance between the hardware and emulated values. An increased digital resolution in r and z by 2 bits allowed the recovery of z_0 resolution, which is critical for vertex reconstruction.

The performance of both the processing stages developed in this dissertation, the GP and DR, are fundamental for the even distribution of data, avoiding efficiency loss due to truncation effects, and reducing track candidates by two methods: preventing excessive data rates, a factor of 3 higher, in the HT by using a bend filter in the GP; and by reducing the KF output tracks by a factor of 2 eliminating the duplicate tracks. This significantly reduces the processing requirements of the downstream components, such as the global track trigger and the correlator trigger. The demonstrator system was built using previous FPGA technology that was readily available at the time of implementation. The algorithms are being updated to work with newer FPGA architectures found in the hardware prototypes described in the following chapters.

7 The CMS Tracker Back-End System

In the tracker back-end system, track reconstruction is performed using a time-multiplexed architecture requiring two layers of data processing as explained in Section 3.5. This Chapter introduces the overall back-end system by showing the proposed CMS tracker rack configuration in the underground service cavern in Section 7.1. Then, the ATCA crate system chosen by CMS for its Phase-2 upgrade is described in Section 7.2. Moreover, Section 7.3 presents a prototype for the Hub slot of the crate, which was developed centrally by CMS to provide the distribution of the timing and control signals to the front-end modules. Later, Section 7.4 describes two initial ATCA development boards for the tracker back-end system. Finally, a brief description of a common and generic firmware framework will be presented in Section 7.5, which describes the necessary infrastructure for any FPGA device using a single source.

7.1 CMS Tracker Rack Configuration

For the CMS Phase-2 upgrade, all the electronics located at the CMS Underground Service Cavern (USC), primarily reading out the many sub-detector systems at the CMS experiment, are designed following the ATCA form factor [82]. A preliminary baseline configuration for the racks located at the CMS USC is shown in Figure 7.1 where the air and water flow is depicted together with the preliminary configuration for the DTC and TFP ATCA crates. The rack infrastructure provides power delivery and cooling capability rated for a maximum heat load of 10 kW per rack. Furthermore, the USC HVAC system is configured to handle a cumulative heat leakage to the cavern of about 50 W per rack [60]. Based on this, node boards should be configured for a maximum power dissipation of 250 W for the OT-DTC, 292 W for the IT-DTC, and 350 W for the TFP [83].

The 13 200 OT detector modules need a total of 216 DTCs to be readout. Each ATCA shelf has twelve node slots and two hub slots. Each rack can allocate 2 shelves. As a result, the OT-DTC system will take up nine racks at the USC of CMS. Each shelf will correspond to 2S or PS modules from the same OT detector nonant as seen




Figure 7.1: Proposed rack configuration for the tracker back-end electronics at the CMS Underground Service Cavern (USC). The red stars denote locations at which the air or water temperature are measured, the blue arrows show the air or water flow inside the rack. An additional measurement of the output water flow is as well taken for each rack. The two images in the right shown the proposed rack configuration for the DTC and the TFP where 12 and 9 processing blades are used respectively.

in Figure 3.9, different cabling mapping studies based on data rates are currently being investigated [84] to know how each shelf will connect to specific detectors using separate trunk fibers. Each DTC shelf will also contain Timing Control and Distribution System (TCDS) and Data Acquisition (DAQ) Hub boards known as DTH-400 and DTH-800 respectively, a prototyping board containing those functionalities is described in Section 7.3.

The Track Finding (TF) system needs 18 boards per nonant, for a total of 162 boards. Each nonant will be hosted in one USC rack, with the 18 TFPs distributed equally between two ATCA shelves. One slot in each TF shelf will be used to accommodate a DTH-400, leaving four spare slots per shelf for redundant TFP nodes. The optical patch panel is used to connect each TF rack assigned to a processor-nonant to the output of two OT-DTC reading each a detector-nonant. The DTC and TF racks would also include the -48V power supplies, heat exchangers, air deflectors, and turbines for cooling the boards.

7.2 The ATCA Shelf

An ATCA shelf is a chassis with a fixed form factor that accepts Field Replaceable Units (FRUs). This may include a variety of intelligent components such as cooling fan trays, power supplies, or user-designed electronic boards. To ensure proper interoperability, FRUs must comply with a set of mechanical, electrical, and interface requirements specified in the PCI Industrial Computer Manufacturing Group (PICMG) standard [82]. The services made available by the shelf from the perspective of an ATCA electronic board are:

Power: Each board receives a dual redundant -48 V power rail with a maximum power rate of 400 W per board. This is a limit from the crate system, not by the CMS cavern as mentioned before. Each of the power rails are fed from independent sources, maximizing the system availability.

Cooling: To eliminate the heat produced by the electronics, redundant fan trays in the air inlet at the bottom of the crate and in the air outlet at the top of the crate drive a forced air stream through it. These fans are integrated in the crate and are independent from the turbine units shown in Figure 7.1.

Data and clock bus: ATCA defines two backplane topologies, the *dual-star*, in which two boards connect to all other boards, and the *full-mesh*, in which all boards connect to all other boards. The standard defines the "base interface" for 1000BASE-T Ethernet and the user defined "fabric interface" for any kind of point-to-point 100 ohm differential protocol, where 10-Gigabit Ethernet, InfiniBand, PCI Express are some of the protocols that can be implemented. In CMS, the dual-star topology is chosen, allowing for the propagation of low-jitter clock signals and dedicated high-speed connections from two central hubs to the rest of the node boards according to the Figure 7.3.

IPMB: The dual redundant Intelligent Platform Management Bus A and B (IPMB-A and IPMB-B), also referenced collectively as IPMB-0 [85], operate in open collector mode with a signaling level of 3.3 V and compatible with the I²C protocol operated in multi-master mode. The two buses are used by the FRUs and the Shelf Manager Controller (ShMC) to relay messages, listening as slaves in the I²C bus for messages addressed to them and taking control of the bus in a I²C multi-master topology to reply.

Other management signals: The address of the actual slot into which a board is inserted is electrically encoded by eight pins in the backplane with odd parity.

7.2.1 Hardware Platform Management

The primary function of the Hardware Platform Management (HPM) system is to track the health of the hardware by gathering sensor data (voltages, power consumption, temperatures, fan speeds, etc.) and taking corrective measures (increasing fan speed, turning off power, triggering alarms, etc.) if the measurements are outside the nominal range. The system is also responsible for managing the power draw of the FRUs in a shelf so that the total power budget does not exceed the permitted operating envelope.




Figure 7.2: ATCA HPM Architecture, adapted from [82].

The HPM system is made up of Intelligent Platform Management Controllers (IPMCs) [85] that handle one or more FRUs, a Shelf Manager Controller (ShMC) that operates on each shelf, and an optional System Manager, as seen in Figure 7.2. The System Manager is a global high-level controller that manages ShMCs in several shelves connected by a network; the ShMC is a device that orchestrates the behavior of all the IPMCs running within a crate; and the IPMC is a controller local to each FRU that is responsible for controlling all aspects specific to the FRU operational state and providing real-time hardware status and sensor information to the ShMC. Messages based on the Intelligent Platform Management Interface (IPMI) [86] protocol are used for HPM communication between FRUs and ShMCs. The IPMI protocol is expanded in the ATCA implementation with the inclusion of remote board monitoring, fault detection, and fault management features.

7.2.2 Intelligent Platform Management Controller

The IPMC is responsible for managing the *hot-swap* process, which involves inserting or removing the FRU from a powered and operating shelf [82]. During a

hot-swap, the IPMC aided by the ShMC, must gracefully activate or deactivate the board. When a board is inserted into a shelf, the IPMC powers up instantly, alerting the ShMC of its presence. The user can request the board activation by locking the mechanical handle on the front plate of the board, which activates a switch connected to the IPMC. The IPMC subsequently communicates all information regarding its sensors, identification, and power needs to the ShMC. The ShMC then assesses the power budget in the shelf and decides whether or not to power on the board. If power-on authorization is given, the IPMC performs the particular procedures required to activate the electronics on the board, such as booting an operating system on one of its CPUs. Similarly, when the front handles are unlocked, the IPMC initiates the board-specific operation to gently shut down the electronics in the board, coordinating with the ShMC the process.

7.3 The CMS DTH Hub Prototype

The hub boards serve as the bridge between a Phase-2 CMS ATCA crate and the surrounding DAQ infrastructure. The baseline architecture specifies a DAQ and TCDS Hub (DTH) [87] with integrated 1 Gb/s Ethernet switch in the first hub slot, and a 10 Gb/s high speed Ethernet switch in the second hub slot as an alternative. If necessary, additional DTH blades may be added to a shelf to increase the DAQ throughput, those additional boards can be inserted either in the second hub slot or in any node slot.

In the baseline configuration of the CMS Phase-2 DAQ system, the data from the detectors and the level-1 trigger systems are aggregated in a per-crate bases using the DTH board. Data is sent using mid-board optics, creating point-to-point links using short optical patch fiber cables connected in the front panel of the boards. Custom protocols [88], [89] are used for the data transmission from the node boards to the DTH. The DTH then aggregates data from several backend boards into a single backend crate for transfer to the data-to-surface network for processing and storage. The DTH decouples the detector and networking domains: variations in event size or trigger rate are buffered in the backend boards, while variations in network throughput are buffered in the DTH. Dedicated point-to-point backplane connections are used by the DTH to distribute the LHC bunch clock, the Level-1 accepts, and the fast commands, as well as the trigger control instructions for calibration, synchronization, and other tasks to the backend boards. The backend boards communicate their data-taking readiness to the DTH, which then aggregates these signals for trigger throttling and data-taking recovery procedures.




Figure 7.3: CMS dual star backplane signal conventions [87]. Each connection shown represents a ‘star’ connection from each of the two hub slots to each of the node slots.

One of the primary tasks of the DTH is to distribute the CMS Trigger and Timing Control and Distribution System (TCDS). It is composed by two independent data streams moving data from the TCDS master and the sub-system electronics. The LHC bunch clock, beam-synchronous timing instructions, and L1 triggers are distributed via a Trigger, Timing and Control (TTC) stream from the master to the sub-systems. In the reverse direction, a Trigger Throttling Stream (TTS) notifies the master of the state in data-taking readiness of the particular sub-system. Both the timing (TTC2) and throttling (TTS2) streams for Phase-2 are based on high-speed serial back-plane links operating at a line rate of about 10 Gb/s and synchronous to the LHC bunch-crossing clock. The TTC2 stream for Phase-2 is an improved version of the TCDS TTC Phase-1 stream. The following are the most notable new features:

- The distribution of Level-1 event/trigger types, including a multi-bit physics trigger type from the Global Trigger.
- Simultaneous delivery of several synchronization instructions within the same bunch crossing, therefore eliminating the majority of the scheduling restrictions present in the existing TTC system of CMS [59].
- Beam Radiation, Instrumentation, and Luminosity (BRIL) type DAQ functionality [90] is being extended to all receivers. This enables its usage by any sub-system in luminosity runs and/or beam-induced background measurements. This feature would also enable sub-systems to (re)synchronize to ongoing data-taking runs without having to wait for a start-of-run synchronization moment, making commissioning and troubleshooting easier.

- The distribution of dedicated ‘BRIL’ triggers for luminosity, and/or background measurements [91].

The throttling stream is also an improved version of the TCDS TTS. The primary distinction is the expanded bandwidth and large number of backend links aggregated per TTS node.

The CMS central DAQ group is working with the CERN electronics group to design the DAQ and TCDS Hub. A picture of a prototype DTH in a test frame is seen in Figure 7.4.




Figure 7.4: The DTH v1.0 prototype in a test fixture. The front panel, towards the left part of the picture contains the SFPs cages used for TCDS2 connectivity and QSFPs cages used for the data-to-surface network. The main FPGAs are underneath the black heatsinks, the top one is used for DAQ and the bottom for TCDS2 distribution in the back-plane. The board contains a COM-Express computer on module under the heatsink towards the top-right side. The links to the node blades are established using Firefly mid-board optics, two of which are plugged in the front panel with blue optical fiber cable [60].

7.4 Tracker Hardware Development Platforms

The back-end processing system for the Phase-2 Tracker will be handled by three types of ATCA node blades (OT-DTC, IT-DTC, and OT-TFP). Most of the design and operating requirements for these blades are common. To ease future operations and long-term maintenance, common solutions have been sought for the various technical challenges involved in the design of these blades. Initially, two hardware prototypes were developed as part of the R&D program within the tracker data

processing community [58]. A further prototype, presented in Chapter 9, was designed with the aim to aggregate and simplify the board architecture looking towards the final system specification.

7.4.1 Apollo

The Apollo [92] board, designed by Boston and Cornell Universities, separates physically the ATCA blade into generic infrastructure and application-specific parts, therefore allowing the use of different PCB materials and stackups in their fabrication. Using co-planar back-plane connectors allows the use of the full height of the ATCA slot for custom made heatsinks as seen in Figure 7.6.

In Figure 7.5 a high-level block diagram of the Apollo board is shown. The Apollo Service Module (SM) provides all the necessary infrastructure from the point of view of the ATCA system. It can either mount the CERN-IPMC [93], OpenIPMC-HW [94] (explained in more detail in Section 8.4), or the UW-IPMC [95] by changing some resistor configurations in the board. The IPMC mezzanine controls the different power modules present in the blade and controls the booting of a commercial off-the-shelf (COTS) System-on-Module (SoM) [96] which communicates with the main FPGAs via two lanes of AXI Chip to Chip (C2C). The SM also contains an ethernet switch mounted on a mezzanine module.




Figure 7.5: a) Apollo block diagram showing both the SM and CM sections. b) Assembled Apollo Service Module (SM) module [97].

The Apollo Command Module (CM) could be tailored to each application and therefore can in principle host any number of processing elements, in Figure 7.5

the current configuration containing two FPGAs is shown. the CM contains as well several hundred optical links with speeds up to 16 or 25 Gb/s, where Samtec Firefly modules are used for the optical connections. The CM contains a microcontroller performing temperature monitoring and low-level configuration for different devices. Both, the SM and CM can operate independently on the bench without requiring an ATCA shelf, therefore helping in the initial commissioning phase of each board.




Figure 7.6: *Apollo Service Module (SM) and Command Module (CM) assembled together [92].*

Related to the power distribution, the SM board isolates the dual-redundant -48 V from the backplane using a GE PIM400 module, which also produces the standby power at 3.3 V . Then the isolated -48 V is regulated down to 12 V by the GE ‘Barracuda’ DC-DC converter. The 12 V feeds two LTM4622 dual DC-DC converters with 4.5 V and 3 V outputs, which are later regulated down to 3.3 V and 1.8 V respectively using Point of Load (POL) linear regulators. The 12 V feeds the CM connectors with a maximum current of 30 A . In the CM, LGA80D and IND072 power modules regulate the various voltages in use on the different power domains of the main FPGAs and optical engines.

7.4.2 Serenity-Z

Serenity is a family of ATCA boards which include the Serenity-A board, designed as part of this dissertation and explained in detail in Chapter 9, and the Serenity-Z board [98] designed primarily by Imperial College London. The Serenity-Z showcases alternative card design approaches by not soldering directly the main processing elements to the board (typically very large FPGAs). Instead, the processing elements are mounted in small Printed Circuit Boards (PCBs) in a $64\times 64\text{ mm}$ form factor and connected to the main board via a double sided spring-array inter-

poser manufactured by Samtec [99]. The interposer has been designed to operate at line rates up to 28 Gb/s, it has 1 mm thickness and about 2000 pins which should be used to also provide power to the daughtercard.

The architecture of the main board follows a traditional division of tasks among clearly defined boundaries between software and hardware/firmware. The board contains a COTS computer-on-module (CoM) in the industrial form factor "COM-Express type 10" which using an Intel Atom CPU is capable of running x86 CentOS Linux, the same operating system as most of the IT infrastructure at CERN. The CPU is assisted by a small Artix7 FPGA which is central to all the slow control tasks and provides the appropriate level shifting and fan out required to all the interfaces (e.g. I^2C , JTAG). The CPU has a PCIe root complex to which three end points are connected via single lane PCIe Gen2 links: the small Artix7 FPGA and both of the main FPGAs mounted on the two daughtercard sites. In addition, all endpoints are linked through an Ethernet network provided by the on-board ethernet switch. The board has a DDR3 Very Low Profile (VLP) DIMM connector on to which either the CERN-IPMC [93] or the OpenIPMC-HW [94] can be plugged in. The main board hosts a total of 32 connectors where Samtec Firefly [100] optical transceivers can be plugged in. The general high-speed connectivity is shown in the left of Figure 7.7, two sites denominated north and south are symmetrical, 64 high-speed links are currently used for the inter-site bridge but could also be assigned to the spare Firefly connectors in a slight revision of the board.

The baseboard remains generic and only through specific daughtercard design, the resources present in the board are plugged to the FPGA logic or high-speed transceivers. Various daughtercards have been designed targeting different particular applications. For the OT-DTC case, a Xilinx Kintex KU15P FPGA was used to designed a daughtercard targeting the specific requirements of the OT-DTC as explained in Section 3.5 and detailed in Section 8.1. In the Figure 7.8 can be observed how a fully populated board with these daughtercards looks like. The fiber optic cables coming out of the Firefly transceivers need to be routed through the middle of the board holding them in place with plastic brackets. Multi-Fiber Push-On (MTP) connectors located outside the front panel provide a location for coupling optically other boards or the OT detector modules in the case of the OT-DTC.

7.5 The EMP framework infrastructure firmware

Several CMS upgrade projects will use FPGAs as the main processing device in the blades, designing a common firmware and software framework will significantly




Figure 7.7: a) The top side layout of the Serenity-Z1.0 ATCA carrier highlighting the two daughter-card sites (labeled North and South), the 64 inter-site channels, the positions of the Firefly optical connectors, and the QSFP for DAQ applications [60]. b) The Serenity-Z1.0 ATCA carrier hosting two daughtercards utilizing the Xilinx Kintex KU115 FPGA [72].




Figure 7.8: Serenity-Z1.1 populated with two Xilinx Kintex KU15P daughtercards and firefly optics. Two large heatsinks able to dissipate up to 100 W cover the processing sites. Heatsinks for the optical transceivers are assembled as well. In the top-right a heatsink for the Com-Express module is observed. The IPMC DIMM connector is in the center-right. In the left, the front panel mounts outside the regular ATCA limits 12 MTP12 or MTP24 optical connectors [72].

reduce the development and maintenance effort required across all those projects. The Extensible, Modular data Processor (EMP) framework is an evolution of the MP7 framework, which was used in all firmware implementations shown in Chapter 5 and Chapter 6. The EMP framework allows users to focus only in the application-specific development extracting away the details of the infrastructure used to synchronously send and receive data to and from other FPGAs. The details about a specific board are communicated via build-time constants which specify settings like clock frequencies or the number of high-speed transceivers or I/O buffers used in the application. The framework is designed in a generic way so that it can be ported to multiple devices, these include several Serenity-Z daughtercards like the one presented in Section 8.1 and the ATCA board Serenity-A described in Chapter 9.




Figure 7.9: Diagram showing the main functional components in the EMP firmware framework [72].

The main components of the EMP framework are shown in Figure 7.9, the top-level firmware instantiating all the firmware blocks is a generic design containing no application information. The Payload block in the center of the diagram contains all the application-specific sources which together with the build-time constants configures the framework for a specific task. The Datapath block contains the buffer memories and high-speed transceiver implementations, its internal architecture allows at run-time the configuration to synchronously send either external data or a particular pattern generated locally, the high-speed transceivers are configured to also operate at a fixed latency, therefore having a deterministic data transfer latency. The Trigger, Timing and Control (TTC) firmware block receive external clocks and TTC signals, it is capable to record the command history and inject command for debugging purposes. The Readout module takes care of sending data to the aggregation module DTH for further forwarding to the DAQ system. Finally the Control block is instructed via an object-oriented software library which includes a command-line interface to control and monitor all block in the EMP framework including the application-specific payload.

7.6 Summary

The ATCA form factor has been selected as the crate system for the Phase-2 upgrade of CMS. This crate system has a number of features aiming at providing a reliable and dependable system with high uptime, features fitting very well with the intended use at CMS. The CMS underground service cavern contains several racks which will host the many ATCA crates needed for the readout of the various sub-detector systems. The outer tracker, in particular, will use two crates per rack, where one rack is sufficient to read one-ninth of the detector, data is later send to the track finder layer utilizing also one rack per one-ninth of the detector. As a result, the OT back-end electronics system is made up of 18 racks in total. This chapter introduced three ATCA development boards that were used for early firmware and software development. The Serenity-Z ATCA board in particular was used by the author as the host system to design various of the hardware developments which are presented in the next chapter like the daughtercard and the unified management architecture based on a ZynqUS+ device.

8 Hardware R&D Contributions

In this chapter, several developments contributing to the overall hardware Research and Development (R&D) program of the tracker back-end electronics system will be described. In Section 8.1 is described the development of an FPGA daughtercard to implement the OT-DTC functionality on the Serenity-Z board, which was presented in Section 7.4. Next, in Section 8.2, the concept of an unified management architecture for controlling ATCA boards is presented by the development of a PCB adapter using as host the Serenity-Z board and a commercially available Zynq Ultrascale+ module. In addition, Section 8.3 implements the concepts previously presented about the unified architecture in a full-sized ATCA board in preparation for the development of a full-featured ATCA board outlined in Chapter 9. Furthermore, Section 8.4 shows the development of a common mezzanine for implementing an Intelligent Platform Management Controller (IPMC) which can be used by several HEP boards. Finally, Section 8.5 contains initial but important work related to qualification efforts to other high-speed optical engines, which many sub-detector systems can benefit from.

8.1 FPGA Daughtercard for Serenity-Z

As previously stated in Section 7.4.2, the Serenity-Z board has two processing sites, each of which employs the Samtec Z-Ray interposer technology. The board can be configured in a variety of topologies, including daisy chained, parallel, and parallel with shared bus, all of which are determined solely by the daughtercard design. For the tracker-specific application, two different configurations can be built, each of which corresponds to one of the two processing layers required by the time-multiplexed track reconstruction algorithm demonstrated in Chapter 6.

The interconnection needed for the OT-DTC was implemented in a daughtercard featuring the biggest Xilinx Kintex Ultrascale+ device (KU15P). Each daughtercard is capable of reading the data from up to 36 outer tracker front-end modules via optical links at 5 or 10 Gb/s according to the module type. Each daughtercard is also connected to the Track Finder layer via 24 fiber optical links running at 25 Gb/s. The




Figure 8.1: High-Speed link interconnection architecture for realizing the OT-DTC requirements on a Serenity-Z with dual KU15P daughtercards.

design of the daughtercard was done in a symmetric manner so that the same board could be use in both processing sites, placed in parallel in one motherboard.

Due to the fact that charged particles bend when they travel through a homogeneous magnetic field, like it is the case inside the CMS detector, both the daughtercards, connected to different detector modules, need to also share data between each other making use of the inter-site bridge of the motherboard for that. According to the Figure 8.1, after all high-speed links have been distributed, the bridge links are assigned with 7 of them capable of running up to 16 Gb/s and 6 links capable of running up to 25 Gb/s, this limitation is according to the resources inside the device. The 4 bidirectional links communicating with the DTH are split between the two daughtercards, where each contributes with 2 lanes. Lastly, there is one lane dedicated for PCIe communication with the board controller.

It is possible to mount components in the bottom layer of the daughtercard as seen in Figure 8.2, those components need to have a maximum thickness of 1 mm as that is the total thickness of the interposer under compression. The footprint of the interposer is based on a 1 mm pitch diamond pad array. On the other side of the daughtercard PCB, a regular circular array with 1 mm pitch is used as the footprint for the FPGA. Appropriate silkscreen graphics help in the identification of the correct orientation for the daughtercard and the type of optical module that should be connected to the Firefly shoes in the motherboard.




Figure 8.2: a) Bottom layer of the Serenity KU15P daughtercard showing the Samtec Z-Ray footprint. b) Top layer of the Serenity KU15P daughtercard with the Xilinx Kintex Ultrascale+ KU15P FPGA mounted.

8.1.1 PCB Layout High-Speed Differential Lines Tuning

The high-speed differential lines connecting the gigabit transceivers running at 25 Gb/s were tuned and de-skewed down to less than 100 μm of difference between the positive and negative traces. The corners were rounded and compensation jogs were added to every differential pair to balance the corners in the routing path as shown in Figure 8.3. As a consequence, the phase of the signal along the whole differential transmission line is maintained. The gigabit transceivers running at 16 Gb/s, or lower speeds, are less strict and a total maximum skew of 400 μm is tolerated, or the equivalent of one 90 degree corner.




Figure 8.3: Serenity KU15P daughtercard PCB layout showing the 25 Gb/s differential pairs with rounded corners and compensation jogs to reach less than 100 μm intra-pair skew.

8.1.2 High-Speed Optical Transceiver Qualification

Two Serenity-Z motherboards were assembled with two KU15P daughtercards each, along with six x4 Firefly optics at 25 Gb/s each, with the goal of measuring and qualifying the different high-speed channels present in the board, particularly those at higher data rates. Because the assembled multi-PCB structure with an interposer in between is not a common approach for high-speed designs, validating the quality of the links was of particular interest.

First, the inter-site bridge containing 14 links at 16 Gb/s and 12 links at 25 Gb/s was validated. This interconnection resides inside the motherboard, it is a DC coupled link using differential pairs routed in the inner layers of the motherboard. A pseudo random bit sequence (PRBS31) was used to generate known data on each end and validate the receiving stream. The links were tested with more than 1×10^{15} bits being transmitted with a total of zero errors. Eye diagram plots can be observed in Figure 8.4.




Figure 8.4: Dual KU15P daughtercard copper bridge links at 16 Gb/s (a) and 25 Gb/s (b), using PRBS31, LPM mode, and no pre/post-cursor tuning [101].

Later, two fully populated motherboards were used to perform a board-to-board data transfer resembling a more challenging scenario. Each board is equipped its own independent clock component, power modules and control devices. Both cards were inserted in a ATCA shelf. A total of 48 optical channels running simultaneously from four independent KU15P daughtercards were used. The Firefly optics were running with the Clock Data Recovery (CDR) feature enabled, a pseudo random bit sequence (PRBS31) was used for pattern checking and 2 meters of fiber optical patch cables were used. The total data transferred exceeded 200 Pb, the test lasted two days and no errors were observed in the received data. From the eye diagrams shown in Figure 8.5 it can be observed that the openings are quite large, and the eyes remained open in the range between 50-70 % of the Unit Interval (UI).




Figure 8.5: Two Serenity-Z boards (a) and (b), mounted with dual KU15P daughter-cards with 24 Firefly optical links each, all running at 25 Gb/s, using PRBS31, LPM mode, and no pre/post-cursor tuning [101].

8.2 Trenz to Serenity Adapter

The CMS level-1 track trigger community developed common hardware infrastructure to accommodate future reference algorithms for track reconstruction. Several novel technologies are to be evaluated for their use in an all-FPGA fully-multiplexed trigger architecture. Slow control and shelf management paradigms are widely different between early prototypes as seen in Section 7.4, they contain small FPGA devices, micro-controllers, and SoCs. Integration and unification of all these devices and tasks into a single Zynq Ultrascale+ (US+) System-on-Chip (SoC) propose benefits to the back-end system administrators as all information about the board status and configuration is in a single location.

8.2.1 The Unified slow control architecture

The unified slow control architecture, as shown in Figure 8.6, is composed primarily by a ZynqUS+ SoC, which integrates in a single package FPGA logic, high-performance ARM-A53 multi-core processors, and two ARM-R5 real-time capable processors. The ZynqUS+ is divided into four power domains, each of which can be independently enabled or disabled. The ARM-R5 cores are utilized to implement time-critical and deterministic tasks either in bare-metal applications or based on FreeRTOS like the OpenIPMC software explained in Section 8.2.3. The Intelligent Platform Management Controller (IPMC) functionality is implemented in the R5 cores; it interacts with the shelf manager via the backplane of the ATCA crate to negotiate card power-up and its subsequent stable operation. The ARM-R5 are also connected to the power supplies (via PMBus), to the voltage and current monitors, and to clock generators and the jitter cleaners (via I^2C , SPI).

When the shelf manager allows enabling the full power to the blade after negotiations with the IPMC, a CentOS Linux operating system is launched on the ARM-A53 cores. Some low-level interfaces, such as IPBus or glue-logic, are implemented using the FPGA, which are enabled at this point. The SoC serves as the primary interface to the main FPGAs in the motherboard through IPMB and TCP/IP-based network interfaces. The AXI chip-to-chip protocol is used for communication between the ZynqUS+ SoC and the main FPGAs.




Figure 8.6: Unified slow control architecture.

8.2.2 Hardware components of the Trenz-Serenity adapter

The adapter board shown in Figure 8.7, first described in [102], was developed to map the COM-Express and IPMC-DIMM footprints found on the Serenity-Z Board to a commercially available ZynqUS+ module produced by Trenz Electronic [103] containing the XCZU4EG device. The adapter employs two voltages, 3.3 V, and 1.8 V, which are used to power all components. The primary 3.3 V is generated by either regulating down the "Payload" 12 V voltage, or by connecting the "3.3 Standby" power available in the IPMC-DIMM connector. The adapter may effortlessly switch between them by using the TPS2121 priority power multiplexer.




Figure 8.7: Trenz-Serenity adapter v1.3 [102].

Several peripherals were incorporated within the adapter to provide various functions:

- **USB:** The USB3340 physical layer (PHY) is utilized to offer USB 2.0 host capabilities; it connects using the low pin interface (ULPI) to the MIO Pins on the ZynqUS+ device and a USB type A female connector on the motherboard.
- **Ethernet:** The 88E1512 Ethernet PHY provides Ethernet connectivity to the module, the RGMII interface is used to connect to the MIO pins of the ZU+ and the Media Dependent Interface (MDI) interface connects to the COM-Express connector. In the motherboard a broadcom BCM53134M Ethernet switch connects the two back-plane base interfaces as well as ETH connections to other devices on the board.
- **microSD:** The NVT4857UK chip was used in a 7×7 mm 24 pin QFN adapter also designed in-house. It supports the SD 3.0-SDR104 protocol with a data throughput of up to 104 MB/s. It substantially enhances the speed of memory reads and writes with respect to SD 2.0 speeds. The microSD adapter shown in Figure 8.8 was created to isolate the very small 0.4 mm pitch WLCSP package in its own project with suitable production quality measures (e.g. x-ray imaging for a complete manufacturing batch).
- **SATA:** One processor-high-speed-transceiver is linked to the M.2 footprint present in the motherboard for use with commercial M.2 SATA SSDs in the 2080 form factor. The drive may be recognized, formatted, and mounted as a persistent storage unit by the Linux operating system.
- **On-board Programmer:** A JTAG and two UART interfaces are provided via a micro USB port, which is powered by the USB connection and uses the FTDI




Figure 8.8: MicroSD controller module. a) panel fabrication. b) X-Ray image [102].

FT4232 chip. The 5 V on the micro USB connector is regulated down to 3.3 V using a dedicated regulator, therefore lowering the load from the 3.3 V standby power when the circuit is not in use.

- **AXI-C2C:** A high-speed Multi-Gigabit Transceivers (MGTs) in the programmable logic of the ZynqUS+ is connected to each of the main processing FPGAs and the auxiliary FPGA present on the motherboard, it is possible to communicate with them via AXI-C2C or PCIe.
- **IPMC:** The low power domain peripherals have two separate I^2C buses for communication with the shelf management through the hot-swappable controller LTC4300-1.

8.2.3 IPMC Software Implementations

IPMC is a crucial component for electronic boards that adhere to the Advanced Telecommunications Computing Architecture (ATCA) standard. It is in charge of monitoring the health parameters of the board, controlling its power states, and providing board control, debug, and recovery functionality to remote clients. The IPMC operates on the 3.3 V standby rail and has a maximum permitted power of 11 W. The Trenz-Serenity Adapter has been successfully utilized with two alternative IPMC software implementations, both targeting the ARM-R5 cores in lock-step mode and mapping the First Stage Boot Loader (FSBL) to the Low-Power Domain (LPD), as illustrated in Figure 8.9.

Using the Vivado system block design tool, a unique hardware project was developed, with an isolation configuration in place to map peripherals solely to the LPD. The Power Management Unit (PMU) was also modified to handle the Full-Power




Figure 8.9: Boot sequence of the Zynq Ultrascale+ device [102].

Domain (FPD) reset release after the payload power was activated. When a Field Replaceable Unit (FRU) is inserted, the 3.3 V standby power switches ON the LPD of the ZynqUS+, where the boot process loads the FSBL in the R5 cores, which then directs the PMU to reset and load the IPMC software. Communication with the Shelf Manager (ShM) begins at this stage.

The IPMC begins by signaling the transition from the virtual state M0 to the FRU inactive state M1. The ShM asks the Device ID as well as details about the installed FRU such as the name, firmware version, and available features. When the front panel handle is closed, the IPMC notifies the transition M1-M2 state, and the ShM replies by requesting the Sensor Data Record (SDR) of the blade. The ShM directs the M2-M3 transition. Power negotiation occurs in M3 to distribute available power in the crate based on the board needs. The ShM additionally requires FRU information in order to determine which E-Key interfaces to activate. At this stage, the board turns on the payload power, configures the programmable logic, and loads the uBoot loader, which is followed by CentOS7 on the A53 cores. The IPMC then tells the ShM of the active condition of the FRU M4 state. The deactivation procedure begins by opening the handle or by requesting that the ShM return the FRU state to M1.

Pigeon Point IPMC Software

Pigeon Point IPMC [104] is a commercially available software stack that has been adapted and customized for use on the ZynqUS+ architecture. It is based on an implementation that is intended for use with other processor architectures. As a result, the implementation on ZynqUS+ is not yet a ready-made option, but rather a work-in-progress [102]. There is capability inside the program to select where peripherals are attached using a Hardware Abstraction Layer (HAL), reducing the amount of files that require changes by the user. This however, does not apply to the

more fundamental features that, for example, need particular architectural functions to instruct the PMU of the ZynqUS+ device to turn ON or OFF certain parts of the multi-processor SoC.

The Pigeon Point implementation is vast and comprehensive. However, not all the features in the code apply to the setup and have not been validated using the ZynqUS+ architecture. Some features require particular hardware components that are not available in the current configuration of the adapter and motherboard. For instance, the local IPMC-L bus is not connected to any Advanced Mezzanine Card (AMC) or Rear Transition Module (RTM) as the motherboard does not provide those features. The remote update functionality is implemented in the code solely to a very specific reference of SPI flash memory which the version of the adapter used does not have. Features related to the E-Keying system of ATCA are as well not implemented in the adapter and therefore can not be tested beyond configuring them in software. The basic operation however was verified in two different crate systems from Schroff and Comtel.

OpenIPMC Software

OpenIPMC [105]–[107] is a free and open-source firmware designed to function as an Intelligent Platform Management Controller (IPMC) with the goal of serving as an alternative to commercial software, like the previously described Pigeon Point IPMC [104], operating on other IPMC solutions. It is built on the FreeRTOS real-time operating system, it is designed to be architecture-independent, allowing it to be compiled for a wide range of microprocessor architectures. The open-source nature of the code allows the user to completely adapt it to their own needs.

The OpenIPMC uses signal call-backs and a Hardware Abstraction Layer (HAL). The HAL allows to mask the details of the particular hardware platform from the primary IPMC functionality. The code was tested on three distinct hardware platforms (Tensilica Xtensa LX6, ARM Cortex-R5, and ARM Cortex-M7) [106] and yielded equivalent results for all the implemented functions on each of them. The code was developed to prioritize those fundamental features of the ATCA standard that are absolutely necessary and eventually proceed towards others that are required but not crucial for the CMS OT application.

At the time of writing, the OpenIPMC software implementation in the ZU+ platform enabled the dual multi-master IPMB_A/B interfaces, bringing the board to the active state. The power-ON and -OFF sequences were tested as well, using the blue ATCA LED on the front panel to signal the changeover states. The FRU data is

saved and read from an EEPROM memory. Sensor Data Record (SDR) are reported to the Shelf Manager (ShM), the temperature and voltage sensors in the PIM400 module are read and reported too. A Further adaptation of the OpenIPMC software was made for implementation on the ATCA ZynqUS+ IPMC Test board presented in Section 8.3 and the management mezzanines presented in Section 9.2.

8.3 ATCA ZynqUS+ IPMC Test Board

The ATCA ZynqUS+ IPMC Test Board was designed primarily with the aim to provide a prototyping platform which could be used for Firmware and Software development related to the integrated management solution for controlling ATCA-based systems, as explained earlier in Section 8.2.1, without having the added complexity of various adapter layers. This board also provided valuable experiences in the fabrication and assembly of multi-layer ATCA-sized PCBs which proved to have their own challenges as explained later in this Section.

The ATCA ZynqUS+ IPMC Test Board is designed around a Zynq UltraScale+ module from the company Trenz Electronic [103], which is the central processing element used to control all the components in the blade. This Board uses the unified management architecture defined by the Serenity-Trenz adapter in Section 8.2 and implements various other necessary components to interface with the ATCA backplane signals. Following the CMS ATCA specification, the board integrates a seven port Ethernet switch, using a dual port magnetic coupler, the dual-redundant Ethernet links in the backplane are connected to the switch. Other CMS specified interfaces like the LHC clocks are routed to Silicon Labs high-performance jitter attenuators [108] and high-performance fan-out buffers [109]. These devices provide the necessary clock distribution network to implement the TCDS2 firmware and prototype modifications to its firmware. Modifications are needed to be able to also include inside the same MGT fabric block the AXI-C2C interface. This feature is of great importance as the number of FPGA MGTs present in this ZynqUS+ module is limited to only four, but it is also important as the much larger device used in the board design of Chapter 9 has all the links assigned and these two interfaces need to also share the same MGT fabric block.

Another purpose of designing this board was to determine the total maximum power required by the 3V3_Standby rail if the entire infrastructure providing the LHC clocks and Ethernet were supplied by it. For this purpose several differential voltage measurement devices [110] were implemented in the board to determine the current consumption of each rail. A total current of 3 A was registered at the

3V3_Standby rail when the ZynqUS+ module, the clock jitter cleaners, both ETH PHYs, the ETH switch, and the USB PHY were active and running, this is equivalent to 83.3 % of the total maximum allowed current of 3.6 A. It is important to note that the ZynqUS+ module in this case was fully powered and configured. If the power is only supplied to the LPD performing the IPMC functionality including the current measurements, then the total consumed current in the 3V3_Standby rail is 2 A, a 55.5 % of the maximum allowed. With these results, it can be concluded that any ATCA board is able to be designed to operate the management infrastructure including the IPMC controller solely from the 3V3_Standby power and still have an ample margin.

Designing PCB boards of this size is not a trivial endeavor, the assembly of the press-fitted backplane connectors poses an additional challenge in the case of the ATCA form factor. These connectors require an additional fabrication step to plate the inner side of the through holes with a layer of hard gold, allowing the connector pins to fit under pressure inside the opening and therefore have a stronger connection with less resistance than soldered connections. It is also absolutely necessary to design and use a mechanical adapter to press the connectors inside the PCB. Such mechanical tool needs to be aware of all other components in the board to avoid damages. During the assembly of the board, there are four fabrication steps; first, the Surface Mount Devices (SMD) are assembled in the top and bottom faces of the board, then the press-fitted connectors are assembled, next the Through-Hole Technology (THT) components are mounted, and finally the board is cleaned. Something important to note here, is that the Power Input Module (PIM400) from General Electric (GE) [111] has an aluminum heatsink integrated, that was observed to react with the chemicals and deionized water used in the cleaning process. As a result, the PIM400 module was only soldered after the cleaning process in subsequent builds.

The ATCA ZynqUS+ IPMC Test board shown in Figure 8.10 was designed placing all of its components, mostly related to infrastructure services, towards the back side of the PCB, leaving the front-end for the target application circuit, in this case, a Xilinx VCU118 development board was mechanically mounted to the base board PCB. In the front panel the UART and JTAG interfaces are available via a FTDI component and a micro-USB connector. The VCU118 board receives power from the ATX connector and the included jumper cable shown in the picture. A custom length copper Firefly cable connects both the ZynqUS+ and the Virtex VU9P FPGA on the development board. This platform, as mentioned before, was used to prototype and develop of advanced boot methods for the ZynqUS+ device in accordance to the split use of its multi-core processor units. This developments lead to the patent application in [112].




Figure 8.10: The ATCA ZynqUS+ IPMC Test Board on a bench top mode operation with the simple ATCA backplane providing ETH connectivity and power. The board is also mounted with a Xilinx VCU118 development board that receives power from the base board. Both boards are connected via a x4 bidirectional copper Firefly cable providing an AXI-C2C management link.

8.4 OpenIPMC-HW DIMM

The OpenIPMC-HW DIMM was designed in accordance with the pinout and footprint previously established in the HEP community by the LAPP-IPMC [113] and followed by the Pulsar-IPMC [114], and the CERN-IPMC [93]. This connector is also present in both ATCA hardware development platforms showed in Section 7.4. The protocols and interfaces present in the DIMM are depicted in Figure 8.11, where the bottom ones represent those mandated by the PICMG 3.0 standard and those in the top are additional interfaces present in the module to help perform the various management tasks, such as communicating via Ethernet or configuring other devices via JTAG.

The OpenIPMC-HW DIMM is designed around a STM32 microcontroller with two high-performance integrated ARM cores, a Cortex M7 running at 480 MHz and a Cortex M4 running at 240 MHz. The OpenIPMC-HW DIMM has six 16 bit Serial Peripheral Interface (SPI) I/O expanders to control the nine AMC ports available. The board contains a 100 Mb/s ETH PHY, a 1 Gb SPI flash memory and a power OR-ing switch to select from two different power sources, therefore allowing it to be fully powered from the on-board micro USB connector as well.




Figure 8.11: *OpenIPMC-HW block diagram.*

The initial production of the DIMM version v1.0 was done in three different countries: Brazil, USA and Germany. This small production of 40 units was used to further develop the software and to perform qualification test with different motherboards and the ATCA Tester from Polaris Networks [115]. The results from the tester are overall positive with 56 % of the tests being passed, 17 % failed, and 26 % were skipped [94]. All of the tests performed were generated automatically by the test system, and the tests requiring operator intervention were not executed. Most of those tests which were either failed or skipped are related to features not yet developed or peripherals not available at the baseboard used for the tests, like no AMC ports available or RTM connectors.

In Figure 8.12 a minor revision of the DIMM is presented, where the only additions are the dedicated boot button, the Single Wire Debug (SWD) header, and a small bug fix for the power OR-ing circuit. This fix now allows the in-hand powering and programming of the DIMM via the micro Universal Serial Bus (USB) connector using the *dfu-util* Linux tool. This feature is an enormous time saving when programming many modules in the production line.

Due to the popularity of the Dual In-line Memory Module (DIMM) form factor in many HEP baseboards, the OpenIPMC-HW has the potential to be used in many different applications, currently the three boards planned for used in the tracker upgrade; Serenity-A, Serenity-Z and Apollo, are considering using it as the baseline system regarding the IPMC of the blades. These platforms are not exclusive to the CMS tracker upgrade, they will potentially be used in other CMS trigger boards like the high granularity calorimeter and the level-1 correlator [116], and by ATLAS in the muon drift tube trigger processor [92].

A slightly larger production of OpenIPMC-HW v1.1 DIMMs was made with the aim to supply the future revision of the development platforms. A total of 70 units



Figure 8.12: OpenIPMC-HW revision v1.1 Top and Bottom views. The main changes with respect to v1.0 are the addition of the BOOT button in the top left and the SWD connector.

were fabricated. In this way future revisions of the boards can be mounted with one OpenIPMC-HW DIMM from the start. The project has evolved and now the software support for the DIMM platform has increased significantly with new end-users helping on finding and fixing bugs or adding functionality tailored to each board.




Figure 8.13: a) Manufacturer Tester Board (MTB) as viewed from the top. b) MTB mounted with two OpenIPMC-HW v1.1 DIMMs where one acts as the Tester and the other as the DUT.

In the v1.1 production run, a focus on testability and traceability of each DIMM module has been made, a custom designed 10×10 cm PCB board with two DIMM connectors is used at the assembly step to test and validate with the help of a software script all connections and functionality of the module. After being manufactured, every single board is labeled with a unique ID number, which can be used throughout the whole tests process and lifetime of the module. The Manufacturer Tester Board

(MTB) seen in Figure 8.13 is a four layer PCB that routes every single I/O from the Tester to the Device Under Test (DUT), in this way, the Tester can request a specific pattern in the I/O ports from the DUT and then verify it. Other interfaces like I2C or UART are connected with each other on each module, the ETH is implemented connecting both ETH PHY together via capacitive coupling. The ETH PHYs allow this mode of connection as internal bias resistors are present within the device, no magnetics are needed as both devices are in the same board.

8.5 FMC+ Board for 25 Gb/s Optical Evaluation

As explained previously in Chapter 7, in the CMS phase-2 upgrade several sub-detector systems will use high-speed serial-links with a line rate of about 25 Gb/s to transfer data between multiple layers of electronic boards. Some of the early prototypes made different choices in terms of the exact optical engine to use [116], which may result in incompatibilities with one another if not tested properly. Furthermore, some of these modules are also potential alternatives to future detector readout systems in various other high-energy physics (HEP), neutrino physics or astro-particle physics experiments. As a consequence, 25 Gb/s communication links are very likely to become a commonality in many future DAQ systems. A characterization of the available alternatives is required and will hopefully favor a particular component for implementation on future designs.




Figure 8.14: Six different FMC+ mezzanines hosting 25 Gb/s optical transceivers in different configurations. Some of these are commercially available modules, others need specific development. Ultimately, for comparison purposes, it is best if all are developed with a common selection of materials and geometry on the differential pairs.

A deeper evaluation of the optical transceivers is underway with the goal of quali-

fying quantitatively the capabilities of different optical transceivers from different vendors. There are currently six options for implementing a 25 Gb/s serial link as seen in Figure 8.14; The Samtec Firefly [100] in x4 or x12 configuration (where the x12 is currently in the alpha phase of development), the Amphenol LEAP [117] in x12 configuration, the Finisar Board-Mount Optical Assembly (BOA) [118] in x12 configuration, the QSFP28 [119] form-factor in x4 configuration, available from multiple manufacturers, and the QSFP-DD [120] with a x8 configuration. Some of those configurations are commercially available [121]–[123]. Alternatively, FPGA Mezzanine Cards Plus (FMC+) can be fabricated to reduce the number of variables for comparison purposes. At such high data rates, several factors may play a significant role in the signal integrity, such as power supply selection (induced switched noise), board layout, differential pair geometry, and even Printed Circuit Board (PCB) material selection.

Aside from differences induced by the PCB design, each optical transceiver also has particular characteristics which may or may not be compatible between each other. Looking at the specifications for each part and doing an initial simplistic interoperability analysis, the minimum (P_0) and maximum (P_1) optical power levels are different between them, as illustration see Figure 8.15. These values determine the Optical Modulation Amplitude (OMA) and the Extinction Ratio (R_e) in Equation 8.1, which are some of the metrics to quantify the power, margin, and quality of the transmitted signal. Similarly, the receiver sensitivity, also sometimes specified as the minimum OMA, on the RX side is an interesting value to determine a possible interoperation between parts.




Figure 8.15: Relationship between the average power, the optical modulation amplitude, and the extinction ratio for two different transceivers A and B, adapted from [124].

Transceivers from different manufacturers propose incompatibilities at the specification level which sometimes even cross the maximum allowable values. From the technical specifications, the Firefly, LEAP and BOA are the three parts which seem to be designed to interoperate more easily; they all have a default operation mode which guarantees a communication link at 10^{-12} BER without the need to use

Forward Error Correction (FEC) algorithms, therefore being sufficient for use at the CMS Tracker back-end processing system, where the repetition of failed messages is often not possible. Other transceivers, mostly in the QSFP form factor that are use primarily with Ethernet protocols, have typically a 10^{-5} BER as they heavily rely on being able to re-transmit lost packets.

$$OMA(dBm) = 10 \log_{10} \frac{P1 - P0}{1 \text{ mW}} \quad R_e(dB) = 10 \log_{10} \frac{P1}{P0} \quad (8.1)$$




Figure 8.16: The Finisar BOA [118] FMC+ optical evaluation mezzanine. a) Bottom side showing the DC-DC regulator supplying the 2.5 V, the level translators, and the decoupling capacitors as well as the FMC+ connector. b) The Top side showing just the BOA footprint where the optical transceiver is mounted via 4 screws which compresses the springs in the interposer between the PCB and the transceiver.

To evaluate one of the aforementioned parts, an FMC+ board was designed containing one Finisar BOA transceiver with 12 channels TX-RX at 25 Gb/s. The board is a 12 layer PCB fabricated with Panasonic MEGTRON6 [125] dielectric. Several Bit Error Rate (BER) measurements were performed using currently available development FPGA boards. A test firmware was implemented in the ZCU111 evaluation board from Xilinx [126] to measure and qualify the signal integrity of the Finisar BOA transceivers and compare them to those obtained with DC coupled copper links or with Firefly optical modules at the same line rate using the daughtercard design of Section 8.1. In Figure 8.18 the optical performance is shown with Decision Feedback Equalizer (DFE) and without it. It is expected that the opening of the eye with DFE disabled is smaller and at the same time the depth of the center is shallower. If the link is tuned with the correct parameters, somewhere in the center of the blue region, it should operate at a lower BER. The MGTs used in the test were scanned in various parameters and set with those values which achieved the




Figure 8.17: a) Overall routing of the BOA FMC+ mezzanine. b) The routing detail under the BOA footprint and compensation jogs for the 25 Gb/s lines of the Finisar BOA FMC+ mezzanine.

best results for one of the links (TX-POST = 1.16 dB and TX-SWING = 828 mV), that configuration was then extended to all 12 links. From the Figure 8.18, it is rather hard to establish if there is in fact any difference between DFE enabled or disabled. Therefore, showing that the optical module behaves rather well in both scenarios, in Table 8.1 are listed the quantitative numbers of the percentage of the opening of the eye with respect to a Unit Interval (UI).




Figure 8.18: Finisar BOA x12 Optical loopback eye measurements with PRBS7. All MGTs were configured with TX-POST = 1.16 dB and TX-SWING = 828 mV a) DFE enabled. b) DFE disable.

A second test was done using a Samtec FMC+ loopback card [127] and the same ZCU111 evaluation board. The MGT transceivers in the FPGA remained tuned with the same settings as before but now the MGT transceivers needed to be configured to operate with a DC coupled link, as opposed to the Optical Transceivers which integrate capacitors inside the device on each differential pair. From the opening of the eyes, the effect of DFE disabled is more noticeable this time, link #10 is specially

susceptible to the change, but it still remained with more than 51 % of open area and therefore more than enough for a working link.




Figure 8.19: FMC+ mezzanine with x12 copper differential pairs capacitively coupled in lookback mode. Eye measurements were made with PRBS7, and all MGTs were configured with TX-POST = 1.16 dB and TX-SWING = 828 mV a) DFE enabled. b) DFE disable.

Table 8.1: Quantitative results from the opening of eye diagram measurements in percentage of a Unit Interval (UI) as presented in Figure 8.18 and Figure 8.19 from the Vivado Hardware Manager. Results were obtained from the Finisar BOA x12 FMC+ mezzanine card and the Samtec FMC+ loopback card using the Xilinx ZCU111 evaluation board.

Mode / Channel	1	2	3	4	5	6
BOA DFE	54.55 %	57.58 %	54.55 %	54.55 %	51.52 %	60.61 %
BOA LPM	51.52 %	60.61 %	60.61 %	54.55 %	48.48 %	54.55 %
LoopBack DFE	72.73 %	72.73 %	75.76 %	66.67 %	72.73 %	72.73 %
LoopBack LPM	66.67 %	69.70 %	72.73 %	63.64 %	63.64 %	63.64 %
Mode / Channel	7	8	9	10	11	12
BOA DFE	66.67 %	57.58 %	60.61 %	60.61 %	63.64 %	48.48 %
BOA LPM	66.67 %	51.52 %	66.67 %	63.64 %	60.61 %	51.52 %
LoopBack DFE	72.73 %	72.73 %	72.73 %	72.73 %	72.73 %	69.70 %
LoopBack LPM	69.70 %	72.73 %	66.67 %	51.52 %	72.73 %	63.64 %

The preliminary results on the performance of the Finisar BOA x12 transceiver presented here are the starting point of the larger investigation mentioned earlier. The collected findings demonstrated with a high level of confidence, that this is an important choice to consider in addition to the Samtec Firefly parts. Before the final designs for all CMS phase-2 DAQ systems are completed, all other optical solutions

will be studied and compared against and with each other in order to clarify any potential interoperability concerns.

8.6 Summary

Several hardware developments were presented in this chapter as part of this dissertation to contribute to the many R&D topics within the CMS tracker back-end community. The work covers a wide range of topics, including slow control and management infrastructure for ATCA boards, high-speed material selection and design, characterization of FPGA high-speed transceivers with optical transmission engines, fabrication techniques, quality management, and traceability of large scale PCB manufacturing. All of these topics were developed independently with each design, but they are reflected as part of a larger system described in Chapter 9 about the design of a fully-featured ATCA board for implementing the HL CMS outer tracker DTC functionality. The unified management architecture described in this chapter is the central element regarding the management mezzanines shown next. The various experiences designing and measuring high-speed transceivers is fundamental for the qualification of the many high-speed links present in the ATCA board of the next chapter.

9 Serenity-A2577 ATCA Board

The Serenity-A2577 ATCA board is one of the hardware prototypes envisioned to be used for the CMS Phase-2 tracker back-end electronics system. It is designed from the collected experiences on the different electronics sub-systems presented on Chapter 8. The board is thought to have a simplified architecture that aimed to be a consolidation of the services infrastructure required for any ATCA board. It combines elements of the two previous development platforms described in Section 7.4. In this Chapter, a general architecture of the board will be presented in Section 9.1. After that, in Section 9.2 the design of two mezzanines with a common layout for both Serenity motherboards will be described, it aggregates all required slow-control tasks in a single System-on-Module (SoM). Then, a thermal analysis and simulation of the Serenity-A board is shown in Section 9.3, followed by several performance measurements of high-speed Multi-Gigabit Transceivers (MGTs) using even alpha versions of commercial optical transceivers. Finally, Section 9.5 describes the synchronous timing and clock distribution architecture of the Serenity-A board.

9.1 Serenity-A2577 Architecture

The Serenity-A2577 board has a large selection of MGTs to provide the necessary connectivity to all other interconnected boards, as seen in Figure 9.1. There are up to 120 MGTs available for general use routed in a x12 fashion to Samtec Firefly mid-board optical transceivers [100], which are split for the OT-DTC case in 72 to the front-end modules and 48 to the TF layer. Here, four MGTs are dedicated for the DAQ path with a x4 bidirectional Firefly, one MGT is used for the detector Trigger and Timing Control and Distribution System (TCDS) distributed centrally by the CMS experiment [128], two more are used for the Advanced eXtensible Interface (AXI) chip-to-chip (C2C) communication with the slow control mezzanine described later in Section 9.2. Finally, one more transceiver is connected to the back-plane for a potential 10 Gb/s Ethernet link to the second hub slot. Overall this board is capable of a total data throughput of more than 3.1 Tb/s. According to the current reference topology for the CMS Outer Tracker, assembling the 216 OT-DTC boards together constitutes an impressive ~ 670 Tb/s capable system. If we suppose

that all 72 links from the outer tracker use 10 Gb/s the data rate is ~ 156 Tb/s, which can be handled even by running the 48 links to the TFP at 16 Gb/s.




Figure 9.1: High-speed link interconnection architecture for realizing the OT-DTC requirements on a Serenity-A with a single VU9P or VU13P FPGA in the A2577 package.

9.2 Integrated Slow Control Management Modules

Following the Unified Management Architecture presented in Section 8.2 and Section 8.3, two mezzanine boards were designed and fabricated to act as slow control for the Serenity-A and Serenity-Z ATCA motherboards. Both mezzanines share the same schematics and components, but have different board-to-board connectors to match the specific requirements of each of them. The mezzanines have been designed to share the same core layout; keeping the DDR-memory routing and differential pair tuning common between both designs and perform a re-arrangement of the other components to match the specific board footprint.

In Figure 9.3 the overall architecture of both mezzanines is depicted. Both types of connectors interfacing to the motherboards are represented by the top box, where the different interfaces are routed. There are two gigabit-capable interfaces using SGMII and RGMII protocols with an included ethernet PHY on the mezzanine. An USB 2.0 PHY and one MGT from the hard processor block are used to provide USB 3.0 connectivity, that in conjunction to a USB 3.0 multiplexer fully populate an USB-C connector. An on-board clock synthesizer and jitter cleaner with 10 outputs




Figure 9.2: Top face of the Serenity-A2577 ATCA board. The power conditioning components, the Ethernet switch, and the FMC+ connector for the Management Control Mezzanine are located on the right side of the board. Several Samtec Fire-fly [100] connectors are placed in the center region of the board. Clock distribution components and auxiliary power regulators are found on the left. Finally, the Xilinx Virtex Ultrascale+ [129] VU9P or VU13P in the A2577 package is located in the center-left.




Figure 9.3: Architecture of the Integrated Slow Control Management Modules based on a ZynqUS+ device with 2 GB of RAM and several interfaces connected to two types of board-to-board connectors, the FMC+ [130] and the Com-Express [131] plus Samtec ADF6 [132].

is used to clock each of the 16 MGT clock inputs in the PL region. 2 GB of DDR4 RAM are routed to the PS side of the ZynqUS+ device providing the main system memory. Finally, programmable power supplies provide all the required voltages on the mezzanines.

Both mezzanines are designed including the IPMC functionality inside the R5 cores of the ZynqUS+ platform. The A53 cores are used to launch a CentOS based Linux operating system. The mezzanines have a power distribution network capable of powering independently the Processing System (PS) and the Programmable Logic (PL). The PS is composed by the Low-Power Domain (LPD) and the Full Power Domain (FPD). The PL region also includes the Multi-Gigabit Transceivers (MGTS). By doing this, the LPD and FPD can be turned ON from the 3.3 V standby power of the PIM400 DC-DC converter, and the PL and MGTS at a later stage from the 12 V payload power. The total power consumption for the standby rail is in the range of 2 W, where 11 W is the maximum available.

9.2.1 FMC+ version

The FMC+ version of the mezzanine as the name implies has an FMC+ connector and form factor. However, the VITA 57.4 [130] specifications are not followed entirely; the VIO_B (3.3 V) is higher than V_ADJ (1.8 V) and VIO_B is provided from the base board as well. Other requirements were followed as much as possible, for example the HB bank is only referenced to the 3.3 V of VIO_B, while HA and LA are only referenced to 1.8 V of V_ADJ. Another non-complaint feature with the standard, is that the 3.3V_AUX, 3.3 V and VIO_B are all connected to the 3.3V_Standby power from the PIM400 module and therefore always present.

In Figure 9.4 it is possible to see how the density of the mezzanine board is quite high, only leaving significant space at the other side of the connector, this is due to the fact that the board has been designed with only standard through hole vias, and therefore the BGA breakout of the connector does not permit any other component to be placed opposite to it. Despite the presence of through hole vias, decoupling capacitors have been placed behind the BGA footprint of the Xilinx ZynqUS+ component, arranged between the vias in a 1 mm grid.

The mezzanine contains a couple of debugging headers from which the configuration for the integrated power supply controller IRPS5401 can be loaded using the provided vendor tool. The IRPS5401 has 16 user-definable memory slots where specific settings can be stored and recalled after boot using a specific resistor value as a selector. This feature is quite convenient as the regulator can be programmed




Figure 9.4: ZynqUS+ Mezzanine in the FMC+ formfactor. a) Top layer including the Xilinx XZCU4EG device in the B900 package in the center position and all other components around it. B) The bottom layer showing the position of the FMC+ connector and the many decoupling capacitors.

once more after performance measurements have determined the more demanding configuration for the board, therefore guaranteeing that the desired voltage reaches the destination with minimal deviation. For this purpose, remote sensing is also used in all outputs.

9.2.2 CMX-EXT version

The CMX-EXT version of the mezzanine has two different sets of connectors, on one side, the standard com-express [131] connector with a final board-to-board height of 5 mm is used. In the other side, a new high-density connector (ADM6/ADF6) from Samtec [132] with 4 columns and 60 rows is used for routing of the high-speed transceivers. Next to the ADM6 connector is located a power connector UMPS [133] with 4 blades from Samtec which allows up to 21 A per blade for power delivery.

As can be seen from Figure 9.5 the layout and component placement of the CMX-EXT is largely identical to that of the FMC+. However, selected modifications have been made, like the inclusion of a micro SD-card slot as this interface is not present in the Serenity-Z motherboard and other PL signals routing to the CMX-EXT connector, signals which are otherwise not present in the FMC+ version. The design and routing of the CMX-EXT variant is complete, however its fabrication has been delayed after evaluation and optimization of the FMC+ version is conducted.




Figure 9.5: 3D model of the ZynqUS+ Mezzanine in the CMX-EXT form factor showing roughly the same component position of the FMC+ version. a) Top layer. b) Bottom layer with both connectors, the Com-Express [131] and the Samtec ADF6 [132].

9.3 Thermal Analysis

The position of the different components on the Serenity-A2577 board were arranged to provide an optimal cooling capability for the many optical transceivers. These devices are expected to run for over 10 years if their temperature is kept below 50 °C, otherwise, their expected life time is reduced considerably. The optical engines were placed far away from the FPGA as it is the main heat source in the board. The optical transceivers could be cooled off by a singular very large heatsink spanning from top to bottom of the ATCA board, where the intake of cool air is located. No other components doing shadowing effects on the optics were placed in their vicinity. Another option, as the firefly shoes are grouped in two locations, is to have two large heatsinks, one for each group as seen in the simulation results on the right side of the Figure 9.6.

In the left side of Figure 9.6 it is depicted a wire frame diagram representing the location of the many components placed on the board and marked in red are the different locations where the heat is being generated or the temperature is being measured. The simulation results shown in the right side of Figure 9.6 indicate the different heat zones of the board. The FPGA was simulated to generate 120 W of heat and the optics each was modeled with 6 W, the heatsinks of the FPGA and optics are shown in green representing values around 50 °C. The DC-DC regulator for the core voltage of the FPGA and the clock distribution components were represented with a red color showing values above 80 °C. Finally, in blue is the rest of the board

that for simulation purposes did not include all the smaller detailed heat sources. This simulation was performed [134] using the Ansys Icepak cooling simulation software [135], the geometry of a single ATCA slot was modeled with different vertical air speeds. In simulation an improvement of 5 °C was observed when a thermal compound sheet with a thickness of 0.5 mm and a thermal conductivity of 2.1 W/mK was added as heat transfer element between the heatsinks and the active components, in reality these are the same properties of the thermal sheet assembled between the optical transceivers and the heatsinks.




Figure 9.6: Serenity-A2577 Icepack thermal simulations [134].

Simulation results from the Icepack topology model were confirmed using a test-stand where measurements were made at various fan speeds and thermal loads [60]. These measurements showed that a fan speed of 10 is perhaps an operational sweet spot between cooling capability vs. noise generated and power needed by the fan units. If the shelf is configured to run at the maximum fan speed of 15, just the fan units require 2 kW of power, that compared to the shelf electronics which are about 4 kW, it is a significant amount. At level 15, the noise generated is also significantly higher at more than 90 dBA providing only a marginal cooling improvement of 15 % temperature reduction compared to ambient temperature. At level 10, the fans consume only 0.5 kW and the noise is about 75 dBA. All the measured values from the test-stand are depicted in Figure 9.7, where it can be seen the logarithmic behavior of both the noise and the power consumed by the fan units vs. the fan setting.

Further thermal, power consumption, and power distribution measurements were carried out with the Serenity-A ATCA board [136]. A custom designed firmware was used together with all possible MGTs instantiated. The firmware artificially generates an excess of power by declaring a recursive LUT oscillator [137], which without the need for a very high clock frequency, it is able to generate a significant




Figure 9.7: a) The fan noise as a function of shelf fan setting for a closed shelf, but not completely closed rack. The increase is ~ 3.1 dB per fan setting. b) Fan power consumption as a function of fan setting [60].).

amount of local heat in the FPGA fabric. Several LUTs in the device were configured in this manner and controlled using Virtual Input and Output (VIO) registers. In Figure 9.8 are depicted the results from the measurements done when activating them in an increased fashion, each of the heater units that is programmed in the firmware constitutes a 65 k LUT block. When seven of these blocks are active, the FPGA reaches the chosen maximum temperature for this test of 100°C , while having a local fan unit on a benchtop setup. The total power consumed by the whole board was about 192 W. During this test, it was found that some of the power was lost in the power distribution network for the VCCINT rail, it was losing about 20 W in heating the PCB itself. The VCCINT rail used only one PCB layer, with a thickness of $17\text{ }\mu\text{m}$ plus a plating of $25\text{ }\mu\text{m}$, to distribute the total 115 A consumed during the test at its maximum load point. An improved placement of the DC-DC regulator and a better power distribution network is foreseen for the revision of the board.




Figure 9.8: FPGA temperature, VCCINT power, and total board power vs. number of active heater units in the firmware. Each heater unit is a 65k oscillating LUT implementation.

9.4 MGT performance

Locating the optical transceivers far away from the FPGA is the most efficient location in terms of cooling performance, however this has the potential to degrade the performance of the high-speed signals running at up to 25 Gb/s. Consequently, characterizing MGT performance is a top priority. Appropriate stackup and dielectric material selection is very important, it was decided to use the Isola I-Tera MT40 [138], which has a relative permittivity of 3.45 and a tangent loss of 0.0031 at 10 GHz, as the dielectric material for the fabrication of the PCB. In addition, the design was rotated 22.5° with respect to the orientation of the weave in the material.

Initial MGT performance was established using Firefly copper loopback cables fabricated to match the specific length between the TX and RX connectors within each FPGA bank. All 120 MGTs were connected with a x12 configuration loopback cable. The DAQ dedicated slot with a x4 configuration was also connected with a copper loopback cable. All transceivers were configured to run at either 16 or 25 Gb/s. Using the IBERT test firmware [139], two different eye diagrams can be observed in Figure 9.9 corresponding to different pseudo-random generated data streams and verification patterns. Here, it is possible to observe how wide the eye diagram looks with either one of the shown PRBS streams. These results constitute the best case scenario, where the link is established through high-quality copper pairs in a short run. All links performed similarly and therefore only one MGT is shown here.




Figure 9.9: Serenity-A2577 MGT IBERT eye diagram test with Samtec Firefly ECUE copper loopback cable [140]. Different patterns were used PRBS7 (a) and PRBS31 (b).

Consequent tests were performed with the help of several optical patch cables of different lengths from 1 m to 25 m, as can be seen from the Figure 9.10. Each of the MTP24 connectors in the front plate is looped to another set of TX-RX pairs and therefore the signal is received in opposite FPGA banks. 120 individual links

were validated at 16 Gb/s. Standard 16 Gb/s Firefly transceivers were utilized to populate all possible x12 slots in the board. No significant difference in performance was observed when using different fiber cable lengths. The performance is also quite similar to that reported with copper cables. Fully loading the board with all possible transceivers and an appropriate firmware to run them in parallel brought an additional validation of the capabilities of the power supplies.




Figure 9.10: *Serenity-A2577 MGT optical test setup.*

9.4.1 25 Gb/s Optical Evaluation with Samtec x12 alpha-v2 parts

Further tests were carried out using Samtec Firefly x12 alpha-v2 optical transceivers, parts which are not yet commercially available and were kindly supplied by Samtec for evaluation under their alpha distribution program. The parts are designed to run at about 25 Gb/s. the performance was evaluated using a single pair of TX and RX in different slots on the board. Early findings indicated a good performance of the links, therefore validating the dielectric material, the board stack-up and the routing of the high-speed transceivers for up to 25 Gb/s operation, even with optical engines. When testing these optical transceivers, several different settings are available for configuration. Figure 9.11 shows a couple of those settings that were investigated. The MGTs inside the FPGA can be configured to operate in two distinct modes: Low Power Mode (LPM) and Decision Feedback Equalization (DFE). There are also several options on the optical transceiver where the Clock Data Recovery (CDR) setting was explored.

In Figure 9.12 and Figure 9.13 are collected some of the various measurements performed with the alpha-v2 parts. In the left side of both Figures, the optical engine was operated with the CDR circuit ON, as it is intended in normal operation at this line rate. In the right side of both Figures, the CDR circuit was turned OFF in both TX and RX devices, therefore showing potentially the performance of the whole optical path, this assumption is not entirely true, and optical sampling oscilloscopes




Figure 9.11: Eye diagram of 25 Gb/s Firefly alpha parts.

have been used in other systems to characterize the optical path as well [141]. The results collected in both, Figure 9.12 and Figure 9.13, were performed running the eye scans to a BER of 1×10^{-7} , the quantitative values are collected in Table 9.1.




Figure 9.12: Eye diagram performance measurements of 12 channels with the FPGA MGT transceiver configured to operate on DFE mode and Samtec Firefly 25 Gb/s x12 alpha-v2 part running with its internal CDR circuit ON (a) or OFF (b).




Figure 9.13: Eye diagram performance measurements of 12 channels with the FPGA MGT transceiver configured to operate on LPM mode and Samtec Firefly x12 at 25 Gb/s alphav2 part running with its internal CDR circuit ON (a) or OFF (b).

Table 9.1: Quantitative results from the opening of eye diagram measurements in percentage of a Unit Interval (UI) as presented in Figure 9.12 and Figure 9.13 from the Vivado Hardware Manager. Results were obtained from the Firefly alpha-v2 x12 optical transceiver and the Serenity-A2577 ATCA board.

Mode / Channel	1	2	3	4	5	6
DFE - CDR ON	72.73 %	66.67 %	66.67 %	66.67 %	69.70 %	66.67 %
DFE - CDR OFF	30.30 %	36.36 %	48.48 %	48.48 %	36.36 %	30.30 %
LPM - CDR ON	66.67 %	69.70 %	63.64 %	66.67 %	66.67 %	66.67 %
LPM - CDR OFF	36.36 %	45.45 %	42.42 %	42.42 %	39.39 %	30.30 %

Mode / Channel	7	8	9	10	11	12
DFE - CDR ON	69.70 %	72.73 %	66.67 %	72.73 %	66.67 %	69.70 %
DFE - CDR OFF	42.42 %	45.45 %	36.36 %	24.24 %	24.24 %	27.27 %
LPM - CDR ON	66.67 %	69.70 %	66.67 %	63.64 %	63.64 %	66.67 %
LPM - CDR OFF	45.45 %	45.45 %	39.39 %	33.33 %	30.30 %	30.30 %

9.4.2 Bathtub analysis with 25 Gb/s Samtec x12 alpha-v2 parts

The bathtub plot is another method for analyzing the performance of a high-speed transceiver using the IBERT test firmware for MGTs [139]. As seen in Figure 9.14, the eye diagram is a 2D plot which scans the parameters of the transceiver in amplitude (voltage) and skew (sampling point across the unit interval). The bathtub plot is a cross section at the mid-point voltage level, a 1D curve which is much faster to obtain when several hundreds of transceivers need to be analyzed. Furthermore, the bathtub curve provides an important benefit in terms of its extrapolation abilities; data with relatively low statistics at BER 1×10^{-8} can be analyzed and extrapolated to a substantially deeper bathtub at BER 1×10^{-12} . This is a significant improvement since the first takes 4 seconds to capture, while the second takes around 2 hours.

The dual-Dirac model [143] is used in the extrapolation procedure to estimate the total jitter (TJ) specified at a low bit error ratio $TJ(BER)$. This model is based on a few assumptions, which are mostly illustrated in the Figure 9.15. There are two types of jitter: random jitter (RJ) and deterministic jitter (DJ). The RJ is a Gaussian distribution that is defined by the σ parameter, representing the width of the Gaussian distribution. The DJ distribution is made up of two Dirac-delta functions separated by roughly the unit interval, located at μ_L and μ_R . Finally, convolution is used to combine both components to obtain a model for the total jitter distribution.




Figure 9.14: Example 2D eye diagrams vs. 1D bathtub curves for two different types of eyes (not performed on any hardware described in this thesis). The specific values of the plots are not important, but rather the relationship between them to understand their correlation. a) Objectively good eye and its corresponding bathtub curve. b) Objectively bad eye and its corresponding bathtub curve at the midpoint voltage [142].

$$\begin{aligned}
 & \mu_L \quad \mu_R \quad \sigma \quad DJ \\
 & [\delta(x - \mu_L) + \delta(x - \mu_R)] * \exp\left[-\frac{x^2}{2\sigma^2}\right] = \exp\left[-\frac{(x - \mu_L)^2}{2\sigma^2}\right] + \exp\left[-\frac{(x - \mu_R)^2}{2\sigma^2}\right]
 \end{aligned}$$

Figure 9.15: Jitter distribution model using the dual-Dirac approximation. Jitter can be modeled by the convolution of the sum of two delta functions separated by Deterministic Jitter (DJ) and a Gaussian Random Jitter (RJ) distribution of width σ [143].

Continuing the model description, the BER scan on the left of Figure 9.16 can be modeled by two sections where $BER(x) = BER_L(x) + BER_R(x)$, the left side is expanded in Equation 9.1 where ρ_T is the transition density representing the ratio of the number of logic transitions to the total number of bits. From the Figure 9.16, the opening of the eye at a given BER is simply the separation between both curves, in this case the opening at $BER 1 \times 10^{-12}$ is given by $x_R - x_L$. A linearized version introducing the variable $Q = (\mu_L - x)/\sigma$ is used to plot the curve in the right side of Figure 9.16 where it is evident that now the Gaussian jitter distribution is a straight line in $Q(x)$, therefore it is easier to manipulate and extrapolate.

$$BER_L(x) = \rho_T \frac{1}{\sqrt{2\pi}\sigma} \int_x^\infty \exp\left[-\frac{(\mu_L - x')^2}{2\sigma^2}\right] dx' \quad (9.1)$$




Figure 9.16: a) A bathtub plot, the bit error ratio as a function of sampling point delay x . b) The Q -scale version of a bathtub plot ($Q(x)$ rather than $BER(x)$) where Gaussian effects are straight lines of slope $1/\sigma$. The dashed line gives the dual-Dirac approximation to $Q(x)$ [143].

A software tool was developed [142] to use the industry standard dual-Dirac model previously explained to extrapolate the measurements done at $BER\ 1 \times 10^{-8}$ to values equivalent to a $BER\ 1 \times 10^{-12}$. In Figure 9.17 results from the extrapolation process are presented. In the left side, Samtec Firefly optical transceivers running at about 25 Gb/s in the x4 bidirectional configuration are shown, in the right side the copper loopback cable is used. These measurements were done as a reference when comparing to the alpha x12 devices.

The software was written in a way that projections with less than 30 % opening at $BER\ 1 \times 10^{-12}$ will be displayed in a red background for easier identification when looking at a large collection of plots, none of the results shown here are in that range, therefore the yellow characteristic background. In the x4 configuration, the average opening for the optical transceivers is 68.9 % while the copper cable is 53.6 %. The difference could be attributed to an improved link achieved by the use of the retiming CDR circuit inside the optical transceiver. In Figure 9.18 the average of the eye openings for the extrapolation values using the Samtec Firefly in x12 configuration running at 25 Gb/s is 62 % while the equivalent measurements with a x12 copper loopback cable is 52.14 %, similar behaviors to the values of the x4 configuration. These values are also smaller but not significantly lower as those found in the eye diagram scans at 1×10^{-7} . A much larger analysis using this method of extrapolating the bathtub curve will be employed in the future in a much larger evaluation where all transceivers are used in the board with 16 and 25 Gb/s Firefly optical transceivers, similar to what was done in Figure 9.10.




Figure 9.17: Extrapolations of the bathtub plots at $BER\ 1 \times 10^{-12}$ for: a) The 25 Gb/s Firefly in x4 configuration running in DFE mode. b) The Firefly copper x4 loopback cable in DFE mode [142].




Figure 9.18: Extrapolations of the bathtub plots at $BER\ 1 \times 10^{-12}$ for the 25 Gb/s Firefly alpha-v2 part in x12 configuration with CDR enabled and running in DFE mode.

9.5 Timing Control and Distribution System (TCDS2)

The Serenity-A board is intended to be used as the OT-DTC board, reading out the OT Tracker detectors and sending them the timing and synchronization signals from the overall CMS experiment, this important role requires a high-performance clock




Figure 9.20: CMS Clock and Timing distribution for the Serenity-A2577 ATCA Board.

In the Figure 9.20 is also depicted the MGT configuration for the TCDS2 link from the backplane to the main VU9P FPGA and a potential forwarding of that information from the main FPGA to the ZynqUS+ mezzanine illustrated by the thicker dashed pink link. This feature is so far not supported by the IP-core provided by the DTH & TCDS2 working group, a collaborative effort, still work-in-progress, was established as several back-end boards have multiple end-points which need to receive TCDS2 information. It can also be seen that not only synchronous protocols are connected to the B219 MGT bank of the VU9P FPGA but the AXI-C2C communication with the mezzanine is also there. This brings an extra level of complication as the MGT-Common block located inside the IP-core needs to be shared with any other external protocol, also as seen from Figure 9.19 the IP-core is locking both the MGT-REFCLK inputs, therefore a modification to allow this use needs to be put in place as well.

9.6 Summary

The first revision of the Serenity-A2577 ATCA board showed potential as a possible candidate for implementing the OT-DTC functionality for the Phase-2 CMS upgrade. The board features one of the largest FPGA footprints in the Xilinx portfolio, with up to 128 MGTs available for the user. The links are distributed efficiently according to the target application. The selection of appropriate dielectric materials and board stackup construction allows the Serenity-A board to perform well even at the demanding high-data rates of about 25 Gb/s of the most advanced optical transceivers in the market. The board features an optimal component position to allow the better cooling performance for those temperature-sensitive devices, like the optical engines. The board presents an innovative management architecture, which centralizes all interfaces in a single System-on-Module (SoM). The management module is designed around an heterogeneous Zynq Ultrascale+ device utilizing its different types of processing cores for various fundamental applications related to ATCA operations and user applications for controlling the on-board hardware. In general, the Serenity-A ATCA board is a ground breaking design providing a total maximum data throughput of ~ 670 Tb/s, more than enough to handle the future HL OT requirements.

10 Conclusion

The CMS experiment must be significantly upgraded during the long shutdown period between 2025 and 2027 in order to fully meet the HL-LHC conditions. At CMS, the increased simultaneous proton-proton collisions of up to 200 per bunch crossing will produce a massive amount of data, proposing enormous challenges to the Level-1 trigger system to remain efficient at selecting those events with interesting physics signatures. The updated hardware-based trigger will have a longer latency of up to $12.5\ \mu\text{s}$ and an increased output trigger rate of up to 750 kHz. The improved performance of the trigger is largely due to the inclusion of tracker information in its production, *i.e.* track finding at the collision rate.

To realize track finding at the collision rate, the CMS silicon tracker will be completely replaced with a novel design, based on double-sided detector modules arranged in a tilted geometry. The detector modules are capable of discriminating high energetic double hits known as ‘stubs’. The outer tracker readout system receives the stub data at the collision rate and performs track reconstruction algorithms to form track candidates. Three different methodologies were initially considered for the task, out of which the Time-multiplexed Track Trigger (TMTT) algorithm was fully demonstrated from end-to-end on digital logic, implemented on FPGAs and meeting the processing latency requirements of less than $4\ \mu\text{s}$. The contributions by the author explained in Chapter 5 were fundamental in the successful implementation of two out of four processing steps of the TMTT algorithm. The Geometrical Processor (GP), the first stage in the TMTT algorithm, subdivides the tracker into several independent sectors to improve the parallelization, thus avoiding efficiency loss due to truncation. The GP reduces by a factor of 3 the output data rate in the subsequent processing stage by calculating the bend filter on each input stub. The Duplicate Removal (DR) is the last stage on the TMTT algorithm, it reduces by 50 % the number of output tracks without losing reconstruction efficiency. It identifies those duplicated tracks by only looking at the information contained inside a single track. As a result, it uses only 1 % of the resources compared to other methods comparing pairs of tracks and the stubs contained in them. Both the GP and DR were initially described for the flat barrel configuration and the division of the tracker in octants. Later, the entire TMTT chain, including the GP and DR were updated for the tilted barrel configuration and the physical division of the tracker in nonants.

Furthermore, various performance optimizations and more efficient logic resources usage were developed for the GP and DR stages as part of this dissertation. The performance results of the entire chain showed a reconstruction efficiency of 95.1 % when considering $t\bar{t}$ events with 200 PU and 97.3 % for single muon events with 200 PU. Also, the analysis about reconstructing tracks under different scenarios proved the reliability and ample margin of the system.

For the phase-2 upgrade of CMS, the readout electronics system at the back-end of the tracker detector must be designed in accordance with the ATCA standard. Three development platforms were presented as overall contributions from the community where common developments are sought. In Chapter 8 are listed some of the many contributions developed under this dissertation to the general hardware research and development program within the community. For instance, the OpenIPMC DIMM presented in Section 8.4 has the potential to be used in several CMS sub-detector systems beyond the CMS tracker. Its open-source nature permits complete ownership and allows effortless modifications to the code by the end-user. The evaluation program for different optical transceivers proposes an interesting endeavor, 25 Gb/s communication links are very likely to become a commonality in many future DAQ systems. The integrated management architecture motivated the development of advanced boot methods for the ZynqUS+ architecture. The proposed boot scheme lead to a patent application related to the use of a partially configured device and a subsequent full configuration from a remote storage media.

Finally, the Serenity-A2577 ATCA board, which was developed as part of this thesis and is described in Chapter 9, provides an enormous processing capability which greatly exceeds the requirements of the high-luminosity outer tracker. The board has an impressive data throughput of 3.1 Tb/s. The ATCA system allows the aggregation of several boards inside a crate to build a more capable system. In the case of the OT-DTC, the 216 boards in the reference architecture construct an impressive ~ 670 Tb/s system using 18 racks. The Serenity-A board has a pin compatible footprint with two high-end Virtex Ultrascale+ devices: the VU9P and the VU13P. These devices currently offer the most cost effective alternative for implementing the outer tracker DTC (OT-DTC) and the Track Finding Processor (TFP) with an overall cost savings of potentially several hundreds of thousands of Swiss Francs (CHF) from the current pledged budget baseline [83]. Detailed descriptions and performance measurements of the key components of the Serenity-A board were presented. The excellent optical performance, the optimal cooling capabilities and the lean yet fully featured slow-control infrastructure, based on a ZynqUS+ mezzanine, propose this board as a key candidate for use in the high luminosity CMS tracker back-end electronics system.

Acronyms

2S	Double Strip
ALICE	A Large Ion Collider Experiment
AMC	Advanced Mezzanine Card
ARM	Advanced RISC Machine
ASIC	Application-Specific Integrated Circuit
AM	Associative Memory
ATX	Advanced Technology eXtended
ATLAS	A Toroidal LHC Apparatus
ATCA	Advanced Telecommunications Computing Architecture
AXI	Advanced eXtensible Interface
BER	Bit Error Rate
BSM	Beyond Standard Model
CBC	CMS Binary Chip
CERN	European Organization for Nuclear Research
C2C	chip-to-chip
CIC	Concentrator Integrated Circuit
CMS	Compact Muon Solenoid
CSC	Cathode Strip Chambers
DAQ	Data Acquisition

DFE	Decision Feedback Equalizer
DIMM	Dual In-line Memory Module
DSP	Digital Signal Processor
DTC	Data, Trigger, and Control
DTH	DAQ and TCDS Hub
DR	Duplicate Removal
ECAL	Electromagnetic Calorimeter
EEPROM	Electrically Erasable Programmable Read Only Memory
EMP	Extensible, Modular data Processor
ETH	Ethernet
FSBL	First Stage Boot Loader
FE	Front-End
FMC+	FPGA Mezzanine Cards Plus
FPD	Full Power Domain
FPGA	Field Programmable Gate Array
FRU	Field Replaceable Unit
GEM	Gas Electron Multiplier
GP	Geometric Processor
HCAL	Hadronic Calorimeter
HEP	High-energy Physics
HDL	Hardware Description Language
HGC	High Granularity Calorimeter
HL	High Luminosity
HLT	High-Level Trigger
HL-LHC	High-luminosity Large Hadron Collider
HLS	High Level Synthesis

HPM	Hardware Platform Management
HT	Hough Transform
HVAC	Heating, Ventilation and Air Conditioning
IT	Inner Tracker
IP	Interaction Point
IPMB	Intelligent Platform Management Bus
IPMC	Intelligent Platform Management Controller
IPMI	Intelligent Platform Management Interface
JTAG	Joint Test Action Group
KF	Kalman Filter
L1	Level-1
L2	Level-2
LED	Light Emitting Diode
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty
LPD	Low-Power Domain
LS	Long Shutdown
LUT	Lookup Table
MGT	Multi-Gigabit Transceiver
MPA	Macro-Pixel ASIC
MP7	Master Processor 7
MTP	Multi-Fiber Push-On
OT	Outer Tracker
PCB	Printed Circuit Board
PMU	Power Management Unit

PHY	Physical Layer Device
pp	proton-proton
PU	Pileup
PS	Pixel-Strip
QCD	Quantum Chromodynamics
QGP	Quark-Gluon-Plasma
RF	Radio Frequency
RISC	Reduced Instruction Set Computing
RPC	Resistive Plate Chamber
RTM	Rear Transition Module
SDR	Sensor Data Record
ShM	Shelf Manager
ShMC	Shelf Manager Controller
SiPM	Silicon-Photo-Multiplier
SM	Standard Model
SoC	System-on-Chip
SoM	System-on-Module
SPS	Super Proton Synchrotron
SSA	Short Strip ASIC
SWD	Single Wire Debug
SPI	Serial Peripheral Interface
TCDS	Trigger and Timing Control and Distribution System
TFP	Track Finding Processor
TF	Track Finding
TIB	Tracker Inner Barrel
TID	Tracker Inner Disks

TIF	Tracker Integration Facility
TEC	Tracker EndCaps
TMTT	Time-multiplexed Track Trigger
TOB	Tracker Outer Barrel
TPC	Time Projection Chamber
TPG	Trigger Primitive Generators
TTC	Trigger, Timing and Control
TTS	Trigger Throttling Stream
UART	Universal Asynchronous Receiver-Transmitter
UI	Unit Interval
USB	Universal Serial Bus
USC	Underground Service Cavern
US+	Ultrascale+
VM	Virtual Module
WLCG	Worldwide LHC Computing Grid

List of Figures

2.1	Schematic representation of the LHC [18].	4
2.2	The superconducting quadrupole magnet [20].	5
2.3	Cumulative delivered luminosity versus time [23].	6
2.4	Modeled cutaway view of the CMS detector [34].	8
2.5	CMS coordinate system [35].	9
2.6	Layout of the CMS Tracker [36].	11
2.7	The electromagnetic calorimeter of the CMS experiment [41].	13
2.8	Cross-section view in the r - z plane of the CMS HCAL [42].	15
2.9	Cross-section through a slice of the CMS detector [43].	17
2.10	Architecture of the L1 trigger [15].	18
2.11	Data flow of the CMS Trigger and DAQ system [46].	19
3.1	LHC operation and HL-LHC installation schedule, adapted from [50].	21
3.2	The overall configuration of the insertion region for the HL-LHC [10].	24
3.3	Geometrical luminosity reduction factor vs β^* for LHC and HL-LHC [10].	25
3.4	Functional diagram of the Phase-2 CMS L1 trigger system [60]. . . .	28
3.5	Sketch of a quarter of the CMS phase-2 silicon tracker in the r - z plane for the flat and tilded barrel configurations [53], [61].	30

3.6	The p_T module concept [61].	32
3.7	Assembled 3D representation of the p_T modules [53].	33
3.8	Data-flow from OT detector modules through to the back-end electronics [53].	34
3.9	CMS tracker back-end system architecture [64].	35
4.1	Track finding algorithms stages for Tracklet and TMTT.	44
4.2	Hybrid algorithm tracking efficiency and track z_0 resolution as a function of η [72].	45
5.1	The segmentation of the tracker volume into φ and η subsectors. . . .	48
5.2	Comparison implementation of the Geometrical Processor between MaxJ and VHDL	54
5.3	Block diagram of the GP Router.	55
5.4	Pair-Wise Duplicate Removal algorithm architecture.	57
5.5	Implementation of the Duplicate Removal algorithm based on pair-wise comparisons between stubs inside tracks.	58
5.6	Firmware vs. software comparison for the Pair-Wise Duplicate Removal Algorithm.	59
5.7	r - φ Hough Transform showing formation of duplicates.	60
5.8	Example of the Hough-space Duplicate Removal algorithm for φ subsector 0 and η subsector 0.	61
5.9	Architecture of the DR Router.	62
5.10	Architecture of the Duplicate Removal algorithm implementation based on the Hough-space.	63
5.11	Different second phase implementations of the Duplicate Removal algorithm based on Hough-space.	64

5.12	Implementation of the Duplicate Removal algorithm based on the Hough-space with multi-sector and the DR Router.	65
5.13	Firmware vs. Software comparison for the Hough-Space Duplicate Removal Algorithm.	66
5.14	Hough-Space Duplicate Removal Algorithm efficiency as a function of η	66
6.1	The demonstrator system using eight MP7 boards in a μ TCA crate [1].	70
6.2	Track reconstruction efficiency as a function of p_T and η for $t\bar{t}$ events with 200 PU [1].	74
6.3	Track reconstruction efficiency for electrons and muons as a function of p_T and η [1].	74
6.4	Track reconstruction efficiency as a function of η highlighting tracks with p_T higher than 100 GeV [1].	75
6.5	Relative p_T , φ , z_0 , and $\cot \theta$ resolution for $t\bar{t}$ events with 200 PU [1]. .	76
6.6	Relative p_T , φ , z_0 , and $\cot \theta$ resolution for for single isolated muons [1].	77
6.7	Relative p_T , φ , z_0 , and $\cot \theta$ resolution for single isolated muons with $5 < p_T^\mu < 15$ GeV with improved r and z encoding [1].	78
6.8	Total number of reconstructed tracks per event when processing $t\bar{t}$ events superimposed with 0, 140, and 200 PU [1].	79
6.9	Data rates for $t\bar{t}$ with 200 PU events after the GP and HT [1].	79
6.10	Track reconstruction efficiency in $t\bar{t}$ events with 200 PU as a function of p_T and η showing the effect due to truncation [1].	80
6.11	Track reconstruction efficiency as a function of η considering truncation effects and regional failure modes [1].	82
7.1	Proposed rack configuration for the tracker back-end electronics at the CMS Underground Service Cavern (USC).	88
7.2	ATCA HPM Architecture, adapted from [82].	90

7.3	CMS dual star backplane signal conventions [87].	92
7.4	The DTH v1.0 prototype in a test fixture [60].	93
7.5	Apollo block diagram and assembled Apollo Service Module [97]. . .	94
7.6	Apollo Service Module (SM) and Command Module (CM) assembled together [92].	95
7.7	The Serenity-Z block diagram and assembled board hosting two daughtercards [72].	97
7.8	Serenity-Z1.1 populated with two Xilinx Kintex KU15P daughtercards and firefly optics [72].	97
7.9	Diagram showing the main functional components in the EMP firmware framework [72].	98
8.1	High-Speed link interconnection architecture for realizing the OT-DTC requirements on a Serenity-Z with dual KU15P daughtercards. .	102
8.2	Top and Bottom images of the Serenity KU15P daughtercard.	103
8.3	Layout showing the 25 Gb/s differential pairs of the Serenity KU15P daughtercard.	103
8.4	Eye diagram measurements of the copper bridge using dual KU15P daughtercards [101].	104
8.5	Eye diagram measurements of Firefly optical links at 25 Gb/s of dual serenity boards each with two KU15P daughtercards [101].	105
8.6	Unified slow control architecture.	106
8.7	Trenz-Serenity adapter v1.3 [102].	107
8.8	MicroSD controller module [102].	108
8.9	Boot sequence of the Zynq Ultrascale+ device [102].	109
8.10	The ATCA ZynqUS+ IPMC Test Board mounted with a Xilinx VCU118 development board.	113
8.11	OpenIPMC-HW block diagram.	114

8.12	Top and Bottom views of the OpenIPMC-HW v1.1.	115
8.13	The Manufacturer Tester Board (MTB).	115
8.14	Six different FMC+ mezzanines hosting 25 Gb/s optical transceivers in different configurations.	116
8.15	Optical modulation amplitude for two different transceivers [124]. .	117
8.16	The Finisar BOA FMC+ optical evaluation mezzanine.	118
8.17	Routing of the BOA FMC+ mezzanine.	119
8.18	Finisar BOA x12 Optical loopback eye measurements with PRBS7. . .	119
8.19	FMC+ loopback mezzanine eye diagram measurements.	120
9.1	High-speed link interconnection architecture for realizing the OT- DTC requirements on a Serenity-A.	124
9.2	Top face of the Serenity-A2577 ATCA board.	125
9.3	Architecture of the Integrated Slow Control Management Modules based on a ZynqUS+ device.	125
9.4	Top and Bottom views of the ZynqUS+ Mezzanine in the FMC+ form- factor.	127
9.5	Top and Bottom 3D views of the ZynqUS+ Mezzanine in the CMX- EXT form factor.	128
9.6	Serenity-A2577 Icepack thermal simulations [134].	129
9.7	Fan noise and power consumption as a function of fan setting. . . .	130
9.8	FPGA temperature, VCCINT power, and total board power vs. num- ber of active heater units in the firmware.	130
9.9	Serenity-A2577 MGT IBERT eye diagram test with Samtec Firefly ECUE copper loopback cable.	131
9.10	Serenity-A2577 MGT optical test setup.	132
9.11	Eye diagram of 25 Gb/s Firefly alpha parts.	133

9.12 Eye diagram performance measurements in DFE mode with Firefly 25 Gb/s x12.	133
9.13 Eye diagram performance measurements in LPM mode with Firefly x12 at 25 Gb/s.	133
9.14 Example 2D eye diagrams vs. 1D bathtub curves for two different types of eyes.	135
9.15 Jitter distribution model using the dual-Dirac approximation.	135
9.16 A bathtub plot model and the Q-scale version as a function of sam- pling point delay x	136
9.17 Extrapolations of the bathtub plots for Firefly x4 at 25 Gb/s, optical and copper loopback.	137
9.18 Extrapolations of the bathtub plots for Firefly x12 at 25 Gb/s.	137
9.19 CMS TCDS2 high-precision clock recovery architecture.	138
9.20 CMS Clock and Timing distribution for the Serenity-A2577 ATCA Board.	139

List of Tables

5.1	η subsector boundary specification.	49
5.2	GP firmware input format.	52
5.3	GP firmware output format.	53
5.4	Resource usage of each GP kernel and the GP router.	56
5.5	Resource usage of the Pair-Wise Duplicate Removal.	59
5.6	Resource usage of the DR router in two configurations.	62
5.7	Resource usage of the Duplicate Removal based on Hough-space. . .	67
6.1	Track finding performance for each stage of the demonstrator chain.	73
6.2	Comparison of the performance of flat and tilted barrel tracker geometries	81
6.3	Mean number of tracks from the HT considering a module failure. .	82
6.4	The mean number of tracks and the tracking efficiency after the HT when using or not the bend filter.	83
6.5	Latency of the each of the firmware components of the track reconstruction chain.	84
6.6	Total resource usage for the demonstrator TFP.	85
8.1	Finisar BOA x12 and loopback FMC+ card opening of eye diagram measurements in percentage of a Unit Interval (UI).	120

9.1	Firefly alpha-v2 x12 opening of eye diagram measurements in percentage of a Unit Interval (UI).	134
-----	---	-----

Bibliography

- [1] R. Aggleton, L. Ardila-Perez, F. Ball, *et al.*, “An FPGA based track finder for the L1 trigger of the CMS experiment at the High Luminosity LHC”, *Journal of Instrumentation*, vol. 12, no. 12, P12019–P12019, Dec. 2017. DOI: [10.1088/1748-0221/12/12/p12019](https://doi.org/10.1088/1748-0221/12/12/p12019). [Online]. Available: <https://doi.org/10.1088/1748-0221/12/12/p12019>.
- [2] E. Bartz, G. Boudoul, R. Bucci, *et al.*, “FPGA-based tracking for the CMS Level-1 trigger using the tracklet algorithm”, *Journal of Instrumentation*, vol. 15, no. 06, P06024–P06024, Jun. 2020. DOI: [10.1088/1748-0221/15/06/p06024](https://doi.org/10.1088/1748-0221/15/06/p06024). [Online]. Available: <https://doi.org/10.1088/1748-0221/15/06/p06024>.
- [3] G. Aad, T. Abajyan, B. Abbott, *et al.*, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Physics Letters B*, vol. 716, no. 1, pp. 1–29, Sep. 2012, ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.08.020](http://dx.doi.org/10.1016/j.physletb.2012.08.020). [Online]. Available: <http://dx.doi.org/10.1016/j.physletb.2012.08.020>.
- [4] S. Chatrchyan, V. Khachatryan, A. Sirunyan, *et al.*, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”, *Physics Letters B*, vol. 716, no. 1, pp. 30–61, Sep. 2012, ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.08.021](http://dx.doi.org/10.1016/j.physletb.2012.08.021). [Online]. Available: <http://dx.doi.org/10.1016/j.physletb.2012.08.021>.
- [5] S. Weinberg, “A model of leptons”, *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 21 Nov. 1967. DOI: [10.1103/PhysRevLett.19.1264](https://link.aps.org/doi/10.1103/PhysRevLett.19.1264). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.19.1264>.
- [6] S. L. Glashow, “Partial-symmetries of weak interactions”, *Nuclear Physics*, vol. 22, no. 4, pp. 579–588, 1961, ISSN: 0029-5582. DOI: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0029558261904692>.

- [7] G. 't Hooft and M. Veltman, "Regularization and renormalization of gauge fields", *Nuclear Physics B*, vol. 44, no. 1, pp. 189–213, 1972, ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(72\)90279-9](https://doi.org/10.1016/0550-3213(72)90279-9). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0550321372902799>.
- [8] P. Higgs, "Broken symmetries, massless particles and gauge fields", *Physics Letters*, vol. 12, no. 2, pp. 132–133, 1964, ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)91136-9](https://doi.org/10.1016/0031-9163(64)91136-9). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0031916364911369>.
- [9] P. W. Higgs, "Broken symmetries and the masses of gauge bosons", *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 16 Oct. 1964. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [10] G. Apollinari, I. Béjar Alonso, O. Brüning, *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. DOI: [10.5170/CERN-2015-005](https://doi.org/10.5170/CERN-2015-005). [Online]. Available: <http://cds.cern.ch/record/2116337>.
- [11] L. Rossi and O. Brüning, *The High Luminosity Large Hadron Collider: the new machine for illuminating the mysteries of Universe*, ser. Advanced series on directions in high energy physics. Hackensack, NJ: World Scientific, 2015. DOI: [10.1142/9581](https://doi.org/10.1142/9581). [Online]. Available: <https://cds.cern.ch/record/1995532>.
- [12] A. Ryd and L. Skinnari, "Tracking triggers for the hl-lhc", *Annual Review of Nuclear and Particle Science*, vol. 70, no. 1, pp. 171–195, 2020. DOI: [10.1146/annurev-nucl-020420-093547](https://doi.org/10.1146/annurev-nucl-020420-093547). eprint: <https://doi.org/10.1146/annurev-nucl-020420-093547>. [Online]. Available: <https://doi.org/10.1146/annurev-nucl-020420-093547>.
- [13] J. Hegeman, J.-M. André, U. Behrens, *et al.*, "Design and development of the DAQ and Timing Hub for CMS Phase-2", *PoS*, vol. TWEPP2018, p. 129, 2019. DOI: [10.22323/1.343.0129](https://doi.org/10.22323/1.343.0129).
- [14] M. E. Pozo Astigarraga *et al.*, "Evolution of the ATLAS Trigger and Data Acquisition System", *Journal of Physics: Conference Series*, vol. 608, p. 012006, May 2015. DOI: [10.1088/1742-6596/608/1/012006](https://doi.org/10.1088/1742-6596/608/1/012006). [Online]. Available: <https://doi.org/10.1088/1742-6596/608/1/012006>.

- [15] CMS Collaboration, “The CMS experiment at the CERN LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08004–S08004, Aug. 2008. DOI: [10.1088/1748-0221/3/08/s08004](https://doi.org/10.1088/1748-0221/3/08/s08004). [Online]. Available: <https://doi.org/10.1088/1748-0221/3/08/s08004>.
- [16] The ATLAS Collaboration *et al.*, “The ATLAS Experiment at the CERN Large Hadron Collider”, *Journal of Instrumentation*, vol. 3, no. 08, S08003–S08003, Aug. 2008. DOI: [10.1088/1748-0221/3/08/s08003](https://doi.org/10.1088/1748-0221/3/08/s08003). [Online]. Available: <https://doi.org/10.1088/1748-0221/3/08/s08003>.
- [17] *LEP design report*. Geneva: CERN, 1984. [Online]. Available: <https://cds.cern.ch/record/102083>.
- [18] E. Mobs, “The CERN accelerator complex - 2019. Complexe des accélérateurs du CERN - 2019”, Jul. 2019. [Online]. Available: <https://cds.cern.ch/record/2684277>.
- [19] L. Evans and P. Bryant, “LHC Machine”, *Journal of Instrumentation*, vol. 3, no. 08, S08001–S08001, Aug. 2008. DOI: [10.1088/1748-0221/3/08/s08001](https://doi.org/10.1088/1748-0221/3/08/s08001). [Online]. Available: <https://doi.org/10.1088/1748-0221/3/08/s08001>.
- [20] G. Datzmann, “Aufbau und Charakterisierung des Hochenergie Rasterionenmikroskops SNAKE”, Dissertation, Technische Universität München, München, 2002.
- [21] L. Guiraud, “Model of an LHC superconducting quadrupole magnet. Aimant quadripôle supraconducteur”, Jan. 2000, [Online]. Available: <http://cds.cern.ch/record/40918>.
- [22] B. Povh, K. Rith, C. Scholz, *et al.*, *Particles and Nuclei*. Springer, Berlin, Heidelberg, Jan. 2015, ISBN: 978-3-662-46320-8. DOI: [10.1007/978-3-662-46321-5](https://doi.org/10.1007/978-3-662-46321-5).
- [23] CMS Collaboration. (2021). Public CMS Luminosity Information, [Online]. Available: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults#Run_2_annual_charts_of_luminosit (visited on 09/07/2021).
- [24] The LHCb Collaboration *et al.*, “The LHCb Detector at the LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08005–S08005, Aug. 2008. DOI: [10.1088/1748-0221/3/08/s08005](https://doi.org/10.1088/1748-0221/3/08/s08005). [Online]. Available: <https://doi.org/10.1088/1748-0221/3/08/s08005>.
- [25] The LHCb Collaboration *et al.*, *The Large Hadron Collider beauty (LHCb) experiment*. [Online]. Available: <https://home.cern/science/experiments/lhcb> (visited on 09/07/2021).

- [26] C. Fitzpatrick, J. M. Williams, S. Meloni, *et al.*, “Upgrade trigger: Bandwidth strategy proposal”, CERN, Geneva, Tech. Rep., Feb. 2017. [Online]. Available: <http://cds.cern.ch/record/2244313>.
- [27] The ALICE Collaboration *et al.*, “The ALICE experiment at the CERN LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08002–S08002, Aug. 2008. DOI: [10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002). [Online]. Available: <https://doi.org/10.1088/1748-0221/3/08/S08002>.
- [28] The ALICE Collaboration *et al.*, *A large ion collider experiment*. [Online]. Available: <https://alice-collaboration.web.cern.ch/> (visited on 09/07/2021).
- [29] CMS Collaboration, “Search for standard model production of four top quarks with same-sign and multilepton final states in proton–proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C*, vol. 78, no. 2, p. 140, 2018. DOI: [10.1140/epjc/s10052-018-5607-5](https://doi.org/10.1140/epjc/s10052-018-5607-5). arXiv: [1710.10614](https://arxiv.org/abs/1710.10614) [hep-ex].
- [30] CMS Collaboration, “Search for supersymmetry in proton-proton collisions at 13 TeV in final states with jets and missing transverse momentum”, *JHEP*, vol. 10, p. 244, 2019. DOI: [10.1007/JHEP10\(2019\)244](https://doi.org/10.1007/JHEP10(2019)244). arXiv: [1908.04722](https://arxiv.org/abs/1908.04722) [hep-ex].
- [31] CMS Collaboration, “Search for direct pair production of supersymmetric partners to the τ lepton in proton-proton collisions at $\sqrt{s} = 13$ TeV. Search for direct pair production of supersymmetric partners to the τ lepton in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C*, vol. 80, 189. 52 p, Jul. 2019. DOI: [10.1140/epjc/s10052-020-7739-7](https://doi.org/10.1140/epjc/s10052-020-7739-7). arXiv: [1907.13179](https://arxiv.org/abs/1907.13179). [Online]. Available: <http://cds.cern.ch/record/2684459>.
- [32] CMS Collaboration, “Search for dark photons in decays of Higgs bosons produced in association with Z bosons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP*, vol. 1910, 139. 35 p, Aug. 2019. DOI: [10.1007/JHEP10\(2019\)139](https://doi.org/10.1007/JHEP10(2019)139). arXiv: [1908.02699](https://arxiv.org/abs/1908.02699). [Online]. Available: <http://cds.cern.ch/record/2685273>.
- [33] A. M. Sirunyan *et al.*, “Principal-component analysis of two-particle azimuthal correlations in PbPb and pPb collisions at CMS”, *Phys. Rev. C*, vol. 96, p. 064902, 6 Dec. 2017. DOI: [10.1103/PhysRevC.96.064902](https://doi.org/10.1103/PhysRevC.96.064902). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevC.96.064902>.
- [34] T. Sakuma and T. McCauley, “Detector and Event Visualization with SketchUp at the CMS Experiment”, *J. Phys. Conf. Ser.*, vol. 513, D. L. Groep and D. Bona-

- corsi, Eds., p. 022 032, 2014. DOI: [10.1088/1742-6596/513/2/022032](https://doi.org/10.1088/1742-6596/513/2/022032). arXiv: [1311.4942](https://arxiv.org/abs/1311.4942) [physics.ins-det].
- [35] University of Zurich. (2021). CMS coordinate system, [Online]. Available: https://wiki.physik.uzh.ch/cms/_detail/latex:cms_coordinate_system.png?id=latex%3Atikz (visited on 08/18/2021).
- [36] G. Sguazzoni, "Performance of the CMS silicon tracker", *PoS*, vol. Vertex 2011, p. 013, 2012. DOI: [10.22323/1.137.0013](https://doi.org/10.22323/1.137.0013).
- [37] A. Dominguez, D. Abbaneo, K. Arndt, *et al.*, "CMS Technical Design Report for the Pixel Detector Upgrade", Tech. Rep., Sep. 2012. [Online]. Available: <http://cds.cern.ch/record/1481838>.
- [38] V. Tavolaro, "The phase1 CMS pixel detector upgrade", *Journal of Instrumentation*, vol. 11, no. 12, pp. C12010–C12010, Dec. 2016. DOI: [10.1088/1748-0221/11/12/c12010](https://doi.org/10.1088/221/11/12/c12010). [Online]. Available: <https://doi.org/10.1088/1748-0221/11/12/c12010>.
- [39] *The CMS electromagnetic calorimeter project: Technical Design Report*, ser. Technical design report. CMS. Geneva: CERN, 1997. [Online]. Available: <https://cds.cern.ch/record/349375>.
- [40] *The CMS hadron calorimeter project: Technical Design Report*, ser. Technical design report. CMS. Geneva: CERN, 1997. [Online]. Available: <https://cds.cern.ch/record/357153>.
- [41] G. L. Bayatian *et al.*, "CMS Physics: Technical Design Report Volume 1: Detector Performance and Software", 2006.
- [42] CMS Collaboration, "Performance of the CMS hadron calorimeter with cosmic ray muons and LHC beam data", *Journal of Instrumentation*, vol. 5, no. 03, T03012–T03012, Mar. 2010. DOI: [10.1088/1748-0221/5/03/t03012](https://doi.org/10.1088/1748-0221/5/03/t03012). [Online]. Available: <https://doi.org/10.1088/1748-0221/5/03/t03012>.
- [43] CMS Collaboration, "CMS slice raw illustrator files", Aug. 2016. [Online]. Available: <https://cds.cern.ch/record/2204899>.
- [44] CMS Collaboration, "The CMS trigger system", *JINST*, vol. 12, no. 01, P01020, 2017. DOI: [10.1088/1748-0221/12/01/P01020](https://doi.org/10.1088/1748-0221/12/01/P01020). arXiv: [1609.02366](https://arxiv.org/abs/1609.02366) [physics.ins-det].
- [45] A. Tapper and D. Acosta, "CMS Technical Design Report for the Level-1 Trigger Upgrade", Tech. Rep., Jun. 2013. [Online]. Available: <https://cds.cern.ch/record/1556311>.

- [46] CMS Collaboration, “CMS: The TriDAS project. Technical design report, Vol. 2: Data acquisition and high-level trigger”, P. Sphicas, Ed., Dec. 2002.
- [47] K. Bos, N. Brook, D. Duellmann, *et al.*, *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*, ser. Technical design report. LCG. Geneva: CERN, 2005. [Online]. Available: <https://cds.cern.ch/record/840543>.
- [48] I. Bird, P. Buncic, F. Carminati, *et al.*, “Update of the Computing Models of the WLCG and the LHC Experiments”, Tech. Rep., Apr. 2014. [Online]. Available: <http://cds.cern.ch/record/1695401>.
- [49] *Grid Computing Centre Karlsruhe (GridKa)*. [Online]. Available: <http://www.gridka.de/cgi-bin/frame.pl?seite=/welcome.html> (visited on 08/17/2021).
- [50] HL-LHC Industry. (2021). Project Schedule, [Online]. Available: <https://project-hl-lhc-industry.web.cern.ch/content/project-schedule> (visited on 09/07/2021).
- [51] CMS Collaboration, “Measurement of Higgs boson decay to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, CERN, Geneva, Tech. Rep. CMS-PAS-HIG-19-006, 2020. [Online]. Available: <https://cds.cern.ch/record/2725423>.
- [52] S. Pattalwar, A. J. May, P. A. McIntosh, *et al.*, “Key Design Features of Crab-Cavity Cryomodule for HiLumi LHC”, in *5th International Particle Accelerator Conference*, Jul. 2014, WEPRI045. DOI: [10.18429/JACoW-IPAC2014-WEPRI045](https://doi.org/10.18429/JACoW-IPAC2014-WEPRI045).
- [53] The CMS Collaboration, D. Abbaneo, J. Alexander, *et al.*, “The Phase-2 Upgrade of the CMS Tracker”, CERN, Geneva, Tech. Rep. CERN-LHCC-2017-009. CMS-TDR-014, Jun. 2017. [Online]. Available: <https://cds.cern.ch/record/2272264>.
- [54] CMS Collaboration, “Technical proposal for a MIP timing detector in the CMS experiment Phase 2 upgrade”, CERN, Geneva, Tech. Rep., Dec. 2017. [Online]. Available: <https://cds.cern.ch/record/2296612>.
- [55] CMS Collaboration, “The Phase-2 Upgrade of the CMS Barrel Calorimeters”, CERN, Geneva, Tech. Rep., Sep. 2017. [Online]. Available: <https://cds.cern.ch/record/2283187>.
- [56] CMS Collaboration, “The Phase-2 Upgrade of the CMS Endcap Calorimeter”, CERN, Geneva, Tech. Rep., Nov. 2017. [Online]. Available: <https://cds.cern.ch/record/2293646>.

- [57] CMS Collaboration, “The Phase-2 Upgrade of the CMS Muon Detectors”, CERN, Geneva, Tech. Rep., Sep. 2017. [Online]. Available: <https://cds.cern.ch/record/2283189>.
- [58] CMS Collaboration, “The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report”, CERN, Geneva, Tech. Rep. CERN-LHCC-2017-013. CMS-TDR-017, Sep. 2017. [Online]. Available: <https://cds.cern.ch/record/2283192>.
- [59] CMS Collaboration, “The Phase-2 Upgrade of the CMS DAQ Interim Technical Design Report”, CERN, Geneva, Tech. Rep., Sep. 2017. [Online]. Available: <https://cds.cern.ch/record/2283193>.
- [60] CMS Collaboration, “The Phase-2 Upgrade of the CMS Level-1 Trigger”, CERN, Geneva, Tech. Rep. CERN-LHCC-2020-004. CMS-TDR-021, Apr. 2020. [Online]. Available: <https://cds.cern.ch/record/2714892>.
- [61] D. Contardo, M. Klute, J. Mans, *et al.*, “Technical Proposal for the Phase-II Upgrade of the CMS Detector”, Geneva, Tech. Rep. CERN-LHCC-2015-010. LHCC-P-008. CMS-TDR-15-02, Jun. 2015. [Online]. Available: <https://cds.cern.ch/record/2020886>.
- [62] J. Chistiansen, M. Garcia-Sciveres, *et al.*, “RD Collaboration Proposal: Development of pixel readout integrated circuits for extreme rate and radiation”, CERN, Geneva, Tech. Rep., Jun. 2013. [Online]. Available: <http://cds.cern.ch/record/1553467>.
- [63] J. Troska, A. Brandon-Bravo, S. Detraz, *et al.*, “The VTRx+, an Optical Link Module for Data Transmission at HL-LHC”, *PoS*, vol. TWEPP-17, p. 048, 2018. DOI: [10.22323/1.313.0048](https://doi.org/10.22323/1.313.0048).
- [64] L. E. Ardila-Perez, “Level-1 track finding with an all-FPGA system at CMS for the HL-LHC”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 936, pp. 329–330, 2019, Frontier Detectors for Frontier Physics: 14th Pisa Meeting on Advanced Detectors, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.10.174>. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0168900218314888>.
- [65] G. Fedi, “Associative memory pattern matching for the l1 track trigger of cms at the hl-lhc”, *EPJ Web Conf.*, vol. 127, p. 00 008, 2016. DOI: [10.1051/epjconf/201612700008](https://doi.org/10.1051/epjconf/201612700008). [Online]. Available: <https://doi.org/10.1051/epjconf/201612700008>.

- [66] R. Frühwirth and A. Strandlie, *Pattern Recognition, Tracking and Vertex Reconstruction in Particle Detectors*. Springer, Cham, Switzerland, 2021, pp. XVII, 203, ISBN: 978-3-030-65771-0. DOI: [10.1007/978-3-030-65771-0](https://doi.org/10.1007/978-3-030-65771-0). [Online]. Available: <https://link.springer.com/book/10.1007%2F978-3-030-65771-0>.
- [67] P. V. C. Hough, “Method and means for recognizing complex patterns”, US3069654A, Dec. 1962. [Online]. Available: <https://patents.google.com/patent/US3069654>.
- [68] N. Dahnoun, *Hough transform*. John Wiley & Sons, Ltd, 2018, ch. 19, pp. 591–603, ISBN: 9781119125587. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119125587.ch19>.
- [69] R. E. Kalman, “A New Approach to Linear Filtering and Prediction Problems”, *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960, ISSN: 0021-9223. DOI: [10.1115/1.3662552](https://doi.org/10.1115/1.3662552). eprint: https://asmedigitalcollection.asme.org/fluidsengineering/article-pdf/82/1/35/5518977/35_1.pdf. [Online]. Available: <https://doi.org/10.1115/1.3662552>.
- [70] C. Amstutz, F. A. Ball, M. N. Balzer, *et al.*, “An FPGA-based track finder for the L1 trigger of the CMS experiment at the high luminosity LHC”, in *2016 IEEE-NPSS Real Time Conference (RT)*, 2016, pp. 1–9. DOI: [10.1109/RTC.2016.7543102](https://doi.org/10.1109/RTC.2016.7543102).
- [71] CMS Tracker Data Processing Group. (2018). CMS Phase II Tracker Backend SysDev Workshop, [Online]. Available: <https://indico.cern.ch/event/689620/> (visited on 07/29/2021).
- [72] L. Ardila, D. Gastler, K. Hahn, *et al.*, “Specification of the Phase-2 Tracker Backend Electronics”, CERN, Geneva, Tech. Rep. CMS DN-18-011, 2020. [Online]. Available: https://espace.cern.ch/Tracker-Upgrade/Data-Processing/Shared%20Documents/DN-18-011_23_03_2020.pdf.
- [73] F. Wang, B. Nachman, and M. Garcia-Sciveres, “Ultimate position resolution of pixel clusters with binary readout for particle tracking”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 899, pp. 10–15, 2018, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.04.053>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218305709>.

- [74] S. P. Summers, “Application of FPGAs to Triggering in High Energy Physics”, PhD thesis, 2018. [Online]. Available: <http://cds.cern.ch/record/2647951>.
- [75] Xilinx, *7 Series FPGAs Overview, Product Specification*, DS180 (v1.17), May 2015. [Online]. Available: http://www.xilinx.com/support/documentation/data_sheets/ds180_7Series_Overview.pdf (visited on 09/07/2021).
- [76] PICMG, *Micro telecommunications computing architecture base specification*, 2020. [Online]. Available: <https://www.picmg.org/product/microtca-base-specification-r2-0/> (visited on 09/07/2021).
- [77] K. Compton, S. Dasu, A. Farmahini-Farahani, *et al.*, “The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS”, *JINST*, vol. 7, p. C12024, 2012. DOI: [10.1088/1748-0221/7/12/C12024](https://doi.org/10.1088/1748-0221/7/12/C12024).
- [78] C. G. Larrea, K. Harder, D. Newbold, *et al.*, “Ipbu: A flexible ethernet-based control system for xtca hardware”, *JINST*, vol. 10, p. C02019, 2015. DOI: [10.1088/1748-0221/10/02/C02019](https://doi.org/10.1088/1748-0221/10/02/C02019).
- [79] E. Hazen, A. Heister, C. Hill, *et al.*, “The AMC13XG: a new generation clock-/timing/DAQ module for CMS MicroTCA”, *JINST*, vol. 8, p. C12036, 2013. DOI: [10.1088/1748-0221/8/12/C12036](https://doi.org/10.1088/1748-0221/8/12/C12036).
- [80] Broadcom, *MiniPOD 12x10G Receiver Module (300m OM3)*. [Online]. Available: <https://www.broadcom.com/products/fiber-optic-modules-components/networking/embedded-optical-modules/minipod/a-fbr-821vx3z> (visited on 08/17/2021).
- [81] Xilinx, *UltraScale Architecture and Product Data Sheet: Overview*, DS890 (v4.0), Mar. 2021. [Online]. Available: https://www.xilinx.com/support/documentation/data_sheets/ds890-ultrascale-overview.pdf (visited on 09/07/2021).
- [82] PICMG, *Advanced telecommunications computing architecture base specification*, 2008. [Online]. Available: <https://www.picmg.org/product/advancedtca-base-specification/> (visited on 09/07/2021).
- [83] M. Pesaresi and P. Wittich. (2020). DPS meeting: update on power, cooling & FPGAs, [Online]. Available: https://indico.cern.ch/event/957762/contributions/4025902/attachments/2106104/3541977/DPS_9_2020_News.pdf (visited on 07/29/2021).

- [84] L. Calligaris. (2020). Generating new cabling map payloads from tkLayout files, [Online]. Available: https://indico.cern.ch/event/957966/contributions/4026820/attachments/2115859/3560093/Generating_cabling_map_payloads_from_tkLayout.pdf (visited on 08/21/2021).
- [85] Intel, Hewlett-Packard, NEC, *et al.* (Nov. 1999). Intelligent Platform Management Bus Communications Protocol Specification, [Online]. Available: <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/ipmp-spec-v1.0.pdf> (visited on 05/30/2021).
- [86] Intel, Hewlett-Packard, NEC, *et al.* (Feb. 2002). Intelligent Platform Management Specification, [Online]. Available: <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/second-gen-interface-spec-v1.5-rev1.1-2.pdf> (visited on 05/30/2021).
- [87] J. Hegeman, R. Blažek, U. Behrens, *et al.*, “Phase-2 CMS DAQ and TCDS Hub – Specifications and design outline for Prototype 1”, CERN, Tech. Rep. CMS-IN20018/001, Mar. 2018.
- [88] D. Gigi, J. Hegeman, and C. Schwick, “The Phase-2 CMS SLinkRocket DAQ link - Protocol and format description”, CERN, Tech. Rep. [Online]. Available: <https://edms.cern.ch/document/2502737/>.
- [89] E. Brandao De Souza Mendes, S. Baron, and M. Taylor, “TCLink: A Timing Compensated High-Speed Optical Link for the HL-LHC experiments”, *PoS*, vol. TWEPP2019, p. 057, 2020. DOI: [10.22323/1.370.0057](https://doi.org/10.22323/1.370.0057).
- [90] A. Dabrowski, “Upgrade of the cms instrumentation for luminosity and machine induced background measurements”, *Nuclear and Particle Physics Proceedings*, vol. 273-275, pp. 1147–1154, Apr. 2016. DOI: [10.1016/j.nuclphysbps.2015.09.180](https://doi.org/10.1016/j.nuclphysbps.2015.09.180).
- [91] CMS Collaboration, “The Phase-2 Upgrade of the CMS Beam Radiation, Instrumentation, and Luminosity Detectors: Conceptual Design”, CERN, Geneva, Tech. Rep., Jan. 2020. [Online]. Available: <https://cds.cern.ch/record/2706512>.
- [92] A. Albert *et al.*, “The Apollo ATCA Platform”, *PoS*, vol. TWEPP2019, p. 120, 2020. DOI: [10.22323/1.370.0120](https://doi.org/10.22323/1.370.0120). arXiv: [1911.06452](https://arxiv.org/abs/1911.06452).
- [93] J. M. Mendez, V. Bobillier, S. L. Haas, *et al.*, “CERN-IPMC solution for AdvancedTCA blades”, in *Topical Workshop on Electronics for Particle Physics (TWEPP17): Santa Cruz, CA, USA, September 11-15, 2017*, vol. TWEPP-17, 2018, p. 053. DOI: [10.22323/1.313.0053](https://doi.org/10.22323/1.313.0053).

- [94] L. Ardila-Perez, L. Calligaris, A. Cascadan, *et al.*, “The OpenIPMC project”, in *16th CERN xTCA Interest Group Meeting*, 2021. [Online]. Available: https://indico.cern.ch/event/1021679/contributions/4333600/attachments/2242916/3803473/20210511__xTCA_IG__OpenIPMC_Project.pdf.
- [95] M. Vicente, T. Gorski, A. Svetek, *et al.*, “Next generation ATCA control infrastructure for the CMS Phase-2 upgrades”, *PoS*, vol. TWEPP-17, p. 102, 2017. DOI: [10.22323/1.313.0102](https://doi.org/10.22323/1.313.0102).
- [96] Enclustra GmbH. (2020). Mercury ZX1 Xilinx Zynq 7030/7035/7045 SoC Module, [Online]. Available: <https://www.enclustra.com/en/products/system-on-chip-modules/mercury-zx1/> (visited on 02/10/2021).
- [97] Dan Gastler. (2021). Apollo Platform Update, [Online]. Available: <https://indico.cern.ch/event/996093/contributions/4376632/attachments/2260716/3837066/Dgastler-2021-06-09-ApolloUpdate.pdf> (visited on 07/26/2021).
- [98] A. Rose, D. Parker, G. Iles, *et al.*, “Serenity: An ATCA prototyping platform for CMS Phase-2”, *PoS*, vol. TWEPP2018, p. 115, 2019. DOI: [10.22323/1.343.0115](https://doi.org/10.22323/1.343.0115). [Online]. Available: <https://doi.org/10.22323/1.343.0115>.
- [99] Samtec Inc., *Z-RAY™ ultra-low profile arrays*. [Online]. Available: <https://www.samtec.com/connectors/high-speed-board-to-board/compression-interposers/zray> (visited on 08/13/2021).
- [100] Samtec Inc., *Samtec Firefly™ micro flyover system™*. [Online]. Available: <https://www.samtec.com/optics/optical-cable/mid-board/firefly> (visited on 08/13/2021).
- [101] A. Rose. (2020). Serenity - From ATCA prototyping platform towards a final product, [Online]. Available: <https://indico.cern.ch/event/875862/> (visited on 07/29/2021).
- [102] L. Ardila-Perez, A. Cascadan, L. Calligaris, *et al.*, “A novel centralized slow control and board management solution for ATCA blades based on the Zynq Ultrascale+ System-on-Chip”, *EPJ Web Conf.*, vol. 245, p. 01 015, 2020. DOI: [10.1051/epjconf/202024501015](https://doi.org/10.1051/epjconf/202024501015). [Online]. Available: <https://doi.org/10.1051/epjconf/202024501015>.
- [103] Trenz Electronic GmbH. (2018). TE0803 - Zynq UltraScale+, [Online]. Available: <https://shop.trenz-electronic.de/en/Products/Trenz-Electronic/TE08XX-Zynq-UltraScale/TE0803-Zynq-UltraScale/> (visited on 08/18/2021).

- [104] Pigeon Point. (2020). Pigeon Point's PICMG products, [Online]. Available: http://www.pigeonpoint.com/products_picmg.html (visited on 02/10/2021).
- [105] A. Cascadan, L. Calligaris, L. Ardila, *et al.* (2020). OpenIPMC: a free open source Intelligent Platform Management Controller for AdvancedTCA, [Online]. Available: <https://gitlab.com/openipmc/openipmc> (visited on 03/10/2021).
- [106] L. Calligaris, A. Cascadan, L. E. Ardila-Perez, *et al.*, "OpenIPMC: A Free and Open-Source Intelligent Platform Management Controller Software", *IEEE Transactions on Nuclear Science*, vol. 68, no. 8, pp. 2105–2112, 2021. DOI: [10.1109/TNS.2021.3092689](https://doi.org/10.1109/TNS.2021.3092689). [Online]. Available: <https://doi.org/10.1109/TNS.2021.3092689>.
- [107] L. Ardila-Perez, L. Calligaris, A. Cascadan, *et al.*, "OpenIPMC: a free and open source Intelligent Platform Management Controller", in *ACES 2020 - Seventh Common ATLAS CMS Electronics Workshop for LHC Upgrades*, 2020. [Online]. Available: https://indico.cern.ch/event/863071/contributions/3856106/attachments/2046289/3428449/POSTER__OpenIPMC_an_Open_Source_Intelligent_Platform_Management_Controller.pdf.
- [108] Skyworks High-Performance Jitter Attenuators. [Online]. Available: <https://www.skyworksinc.com/en/Products/Timing/High-Performance-Jitter-Attenuators> (visited on 08/17/2021).
- [109] Skyworks, Any Format Clock Buffers. [Online]. Available: <https://www.skyworksinc.com/en/Products/Timing/Any-Format-Clock-Buffers> (visited on 08/17/2021).
- [110] Texas Instruments, INA3221 Triple-Channel, High-Side Measurement, Shunt and Bus Voltage Monitor with I2C- and SMBUS-Compatible Interface. [Online]. Available: <https://www.ti.com/product/INA3221> (visited on 08/17/2021).
- [111] General Electric (GE), PIM400 Series; ATCA Board Power Input Modules. [Online]. Available: <https://library.industrialsolutions.abb.com/publibrary/checkout/PIM400?TNR=Data%20Sheets%7CPIM400%7Cpdf> (visited on 08/17/2021).
- [112] L. Ardila Perez, M. Fuchs, T. Mehner, *et al.*, "Verfahren zum Konfigurieren einer integrierten Schaltung, Verfahren zur Bereitstellung von Lade-Software für eine integrierte Schaltung, und integrierte Schaltung", K 6878 - ro / ksc, Jun. 2021.

- [113] S. Lafrasse, “LAPP IPMC Overview”, in *IPMC Workshop - ATLAS upgrade*, 2018. [Online]. Available: https://indico.cern.ch/event/737733/contributions/3077000/attachments/1730606/2796822/LAPP_IPMC_Overview.pdf.
- [114] L. A. Ramalho, T. C. Paiva, R. L. Iope, *et al.*, “Development of an intelligent platform management controller for the pulsar iib”, in *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2015, pp. 1–2. DOI: [10.1109/NSSMIC.2015.7581788](https://doi.org/10.1109/NSSMIC.2015.7581788).
- [115] Polaris Networks Inc. (2020). AtcaTest3, [Online]. Available: https://members.picmg.org/kshowcase/view/view_item/f66bd9bbd767677d1e7ece11ca07c31a5ffca72d (visited on 02/24/2021).
- [116] CMS Collaboration, “The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger”, CERN, Geneva, Tech. Rep., Mar. 2021. [Online]. Available: <https://cds.cern.ch/record/2759072>.
- [117] *Amphenol Leap™ On-Board Transceiver*. [Online]. Available: <https://www.amphenol-icc.com/product-series/leap-on-board-transceiver.html> (visited on 08/13/2021).
- [118] Finisar, *25G BOA (Board-Mount Optical Assembly)*. [Online]. Available: <https://ii-vi.com/product/25g-boas-board-mount-optical-assembly/> (visited on 08/04/2021).
- [119] *The QSFP28 (Quad Small Form Pluggable) transceiver*. [Online]. Available: <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/qsfp28> (visited on 08/13/2021).
- [120] *QSFP-DD (Quad Small Form Pluggable Double Density) transceiver*. [Online]. Available: <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/qsfp-dd> (visited on 08/13/2021).
- [121] HiTech Global, LLC, *600Gig Leap™ On-Board Transceiver FMC+ Module*. [Online]. Available: http://www.hitechglobal.com/FMCModules/FMC_Leap-600Gig.htm (visited on 08/17/2021).
- [122] HiTech Global, LLC, *6-Port Samtec FireFly (6x100G) FMC+ Module (Vita57.4 compliant)*. [Online]. Available: <http://www.hitechglobal.com/FMCModules/FMC+FireFly.htm> (visited on 08/17/2021).
- [123] HiTech Global, LLC, *6-Port QSFP28 (6x100G) / QSFP+ (6x40G or 6x56G) FMC+ Module (Vita57.4)*. [Online]. Available: <http://www.hitechglobal.com/FMCModules/x6QSFP28.htm> (visited on 08/17/2021).

- [124] *Measuring Extinction Ratio of Optical Transmitters*, Agilent Technologies, Inc. [Online]. Available: <http://literature.cdn.keysight.com/litweb/pdf/5966-4316E.pdf> (visited on 08/17/2021).
- [125] Matrix Circuit Board Materials. (2021). Panasonic MEGTRON6 High Speed, Low Loss Multi-layer Materials, [Online]. Available: <https://www.matrixelectronics.com/products/panasonic/megtron-6/> (visited on 09/10/2021).
- [126] Xilinx. (2021). Zynq UltraScale+ RFSoc ZCU111 Evaluation Kit, [Online]. Available: <https://www.xilinx.com/products/boards-and-kits/zcu111.html> (visited on 09/14/2021).
- [127] Samtec Inc., *VITA 57.4 FMC+ HSPC Loopback Card*. [Online]. Available: <https://www.samtec.com/kits/optics-fpga/hspc-fmcp> (visited on 08/04/2021).
- [128] J. Hegeman. (2021). TCDS2 features and implementation progress, [Online]. Available: https://indico.cern.ch/event/1012587/contributions/4263552/attachments/2213890/3747565/20210323_ec_phase2_backend_workshop.pdf (visited on 08/21/2021).
- [129] Xilinx, *Virtex UltraScale+™ devices*, Aug. 2021. [Online]. Available: <https://www.xilinx.com/products/silicon-devices/fpga/virtex-ultrascale-plus.html> (visited on 09/07/2021).
- [130] Samtec Inc., *VITA 57.4 FMC+ Standard Products and Support*. [Online]. Available: <https://www.samtec.com/standards/vita/fmc-plus> (visited on 08/13/2021).
- [131] *ComExpress*, PICMG. [Online]. Available: <https://www.picmg.org/openstandards/com-express> (visited on 07/13/2021).
- [132] Samtec Inc., *AcceleRate HD ultra-dense, multi-row mezzanine strips*. [Online]. Available: <https://www.samtec.com/connectors/high-speed-board-to-board/ultra-micro/accelerate-hd> (visited on 08/13/2021).
- [133] Samtec Inc., *mPOWER Ultra Micro Power Socket*. [Online]. Available: <https://www.samtec.com/connectors/micro-pitch-board-to-board/rugged/ultra-micro-power> (visited on 08/13/2021).
- [134] G. Fedi. (2020). Preliminary thermal and power tests on Serenity, [Online]. Available: https://indico.cern.ch/event/964426/contributions/4148431/attachments/2161465/3646997/Thermal_studies_Dec11.pdf (visited on 08/20/2021).

- [135] *Ansys Icepak Cooling Simulation Software for Electronic Components*. [Online]. Available: <https://www.ansys.com/products/electronics/ansys-icepak> (visited on 08/13/2021).
- [136] L. Ardila. (2021). KIT-IPE Hardware updates, [Online]. Available: https://indico.cern.ch/event/1003095/contributions/4212135/attachments/2189382/3700048/2021-02-11_DPS-HW.pdf (visited on 08/20/2021).
- [137] A. Agne, H. Hangmann, M. Happe, *et al.*, "Seven recipes for setting your FPGA on fire - A cookbook on heat generators", *Microprocessors and Microsystems*, vol. 38, no. 8, Part B, pp. 911–919, 2014, ISSN: 0141-9331. DOI: <https://doi.org/10.1016/j.micpro.2013.12.001>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0141933113002007>.
- [138] *Isola I-Tera MT40: Very Low-Loss Laminate and Prepreg*. [Online]. Available: <https://www.isola-group.com/pcb-laminates-prepreg/i-tera-mt40/> (visited on 08/17/2021).
- [139] *IBERT for UltraScale/UltraScale+ GTH Transceivers*. [Online]. Available: https://www.xilinx.com/products/intellectual-property/ibert_ultrascale_gth.html (visited on 08/21/2021).
- [140] *Samtec FireFly™ Low Profile Micro Flyover System™ Cable Assembly*. [Online]. Available: <https://www.samtec.com/products/ecue> (visited on 08/13/2021).
- [141] Greg Iles. (2021). Overview of optical engines for back-end board interconnects, [Online]. Available: <https://indico.cern.ch/event/1012587/contributions/4262932/attachments/2214229/3748230/2021-03-23%20CMS%20Phase%20%20Backend%20Workshop%20v3.pdf> (visited on 08/21/2021).
- [142] L. Ardila, D. Tcherniakhovski, A. Howard, *et al.* (2021). 12x25 G Firefly alpha-1 sample tests, [Online]. Available: <https://indico.cern.ch/event/993270/contributions/4223652/attachments/2190333/3701798/12x25%20G%20Firefly%20Alpha-1%20sample%20test.pdf> (visited on 08/21/2021).
- [143] *Jitter analysis: The dual-Dirac model, RJ/DJ, and Q-scale*, Agilent Technologies, Inc. [Online]. Available: https://www.keysight.com/upload/cmc_upload/All/dualdirac1.pdf (visited on 08/17/2021).

List of Publications

Patents

- [1] Luis Ardila Perez, Marvin Fuchs, Torben Mehner, and Oliver Sander. “Verfahren zum Konfigurieren einer integrierten Schaltung, Verfahren zur Bereitstellung von Lade-Software für eine integrierte Schaltung, und integrierte Schaltung”. K 6878 - ro / ksc. submitted, pending approval. June 2021.

Peer-Reviewed Articles

- [2] T. Aarrestad, D. Abbaneo, M. Abbas, W. Adam, J.-L. Agram, I. Ahmed, B. Akgun, S. Albergo, E. Albert, Y. Allard, G. Anagnostou, J. Andrea, K. Androsov, S. Arab, L. Ardila, et al. “The CMS Phase-1 pixel detector upgrade”. In: *Journal of Instrumentation* 16.02 (Feb. 2021), P02027–P02027. DOI: [10.1088/1748-0221/16/02/p02027](https://doi.org/10.1088/1748-0221/16/02/p02027). URL: <https://doi.org/10.1088/1748-0221/16/02/p02027>.
- [3] Luigi Calligaris, André Cascadan, Luis E. Ardila-Perez, Bruno Casu, Alison França da Costa, Ailton Akira Shinoda, Lucas Arruda Ramalho, and Oliver Sander. “OpenIPMC: A Free and Open-Source Intelligent Platform Management Controller Software”. In: *IEEE Transactions on Nuclear Science* 68.8 (2021), pp. 2105–2112. DOI: [10.1109/TNS.2021.3092689](https://doi.org/10.1109/TNS.2021.3092689). URL: <https://doi.org/10.1109/TNS.2021.3092689>.
- [4] W. Adam, J.-L. Agram, I. Ahmed, B. Akgun, S. Albergo, E. Albert, M. Aldaya, J. Alexander, M. Alhusseini, J. Alimena, Y. Allard, G. Altopp, C. Amsler, G. Anagnostou, J. Andrea, K. Androsov, A. Apresyan, L. Ardila, et al. “Experimental study of different silicon sensor options for the upgrade of the CMS Outer Tracker”. In: *Journal of Instrumentation* 15.04 (Apr. 2020), P04017–P04017. DOI: [10.1088/1748-0221/15/04/p04017](https://doi.org/10.1088/1748-0221/15/04/p04017). URL: <https://doi.org/10.1088/1748-0221/15/04/p04017>.
- [5] Luis Ardila-Perez, André Cascadan, Luigi Calligaris, Denis Tcherniakhovski, Matthias Balzer, Marc Weber, and Oliver Sander. “A novel centralized slow control and board management solution for ATCA blades based on the Zynq Ultrascale+ System-on-Chip”. In: *EPJ Web Conf.* 245 (2020), p. 01015. DOI:

- 10.1051/epjconf/202024501015. URL: <https://doi.org/10.1051/epjconf/202024501015>.
- [6] T. Aarrestad, D. Abbaneo, M. Abbas, J.G. Acosta, W. Adam, J.-L. Agram, I. Ahmed, B. Akgun, S. Albergo, E. Albert, M. Aldaya, J. Alexander, M. Alhusseini, J. Alimena, Y. Allard, G. Altopp, M. Alyari, C. Amsler, G. Anagnostou, J. Andrea, K. Androsov, L. Ardila, et al. “The DAQ and control system for the CMS Phase-1 pixel detector upgrade”. In: 14.10 (Oct. 2019), P10017–P10017. DOI: 10.1088/1748-0221/14/10/p10017. URL: <https://doi.org/10.1088/1748-0221/14/10/p10017>.
- [7] L. E. Ardila-Perez. “Level-1 track finding with an all-FPGA system at CMS for the HL-LHC”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 936 (2019), pp. 329–330. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.10.174>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900218314888>.
- [8] Andrew Rose, Duncan Parker, Gregory Iles, Ozgur Sahin, Pierre-Anne Bausson, Andromachi Tsirou, Giacomo Fedi, Pierro-Giorgio Verдини, Luis Ardila, Matthias Balzer, Thomas Schuh, Tom Williams, Alessandro Thea, Kristian Harder, Shashi Dugad, Raghunandan Shukla, and Irfan Mirza. “Serenity: An ATCA prototyping platform for CMS Phase-2”. In: *PoS TWEPP2018* (2019), p. 115. DOI: 10.22323/1.343.0115.
- [9] A. Álvarez Fernández, T.K. Aarrestad, D. Abbaneo, S.M. Abbas, G. Abbiendi, M. Abbrescia, S. Abdullin, A. Abdulsalam, S. Abu Zeid, W. Adam, P. Adzic, S. Afanasiev, M.N. Agaras, R. Aggleton, J.-L. Agram, M. Ahmad, A. Ahmad, M. Ahmad, I. Ahmed, I. Ahmed, S. Ahuja, M. Akbiyik, B. Akgun, C. Albajar, S. Albergo, A. Albert, E. Albert, M. Albrow, J. Alcaraz Maestre, M. Aldaya Martin, A. Aleksandrov, T. Alexander, J. Alexander, Y. Allard, J. Almond, G. Altopp, L. Alunni Solestizi, F.L. Alves, G.A. Alves, M. Alyari, N. Amapane, F. Ambrogio, C. Amendola, N. Amin, C. Amsler, G. Anagnostou, D. Anderson, J. Andrea, Yu. Andreev, V. Andreev, M.B. Andrews, K. Androsov, I. Antropov, Z. Antunovic, G. Apollinari, A. Apresyan, A. Apyan, D. Arcaro, R. Arcidiacono, L. Ardila, et al. “Precision measurement of the structure of the CMS inner tracking system using nuclear interactions”. In: *Journal of Instrumentation* 13.10 (Oct. 2018), P10034–P10034. DOI: 10.1088/1748-0221/13/10/p10034. URL: <https://doi.org/10.1088/1748-0221/13/10/p10034>.
- [10] Davide Cieri, Luigi Calligaris, Kristian Harder, Konstantinos Manolopoulos, Claire Shepherd-Themistocleous, Ian Tomalin, Robin Aggleton, Fionn Ball, Jim Brooke, Emyr Clement, Dave Newbold, Sudarshan Paramesvaran, Peter Hobson, Alexander Davide Morton, Ivan Reid, Geoff Hall, Gregory Iles, Thomas Owen James, Takashi Matsushita, Mark Pesaresi, Andrew William Rose, Antoni Shtipliyski, Sioni Summers, Alex Tapper, Kirika Uchida, Paschalis

- Vichoudis, Luis Ardila-Perez, Matthias Balzer, Michele Caselle, Oliver Sander, Thomas Schuh, and Marc Weber. "An FPGA-based Track Finder for the L1 Trigger of the CMS Experiment at the HL-LHC". In: *PoS TWEPP-17* (2018), p. 131. DOI: [10.22323/1.313.0131](https://doi.org/10.22323/1.313.0131).
- [11] R. Aggleton, L.E. Ardila-Perez, F.A. Ball, M.N. Balzer, G. Boudoul, J. Brooke, M. Caselle, L. Calligaris, D. Cieri, E. Clement, S. Dutta, G. Hall, K. Harder, P.R. Hobson, G.M. Iles, T.O. James, K. Manolopoulos, T. Matsushita, A.D. Morton, D. Newbold, S. Paramesvaran, M. Pesaresi, N. Pozzobon, I.D. Reid, A.W. Rose, O. Sander, C. Shepherd-Themistocleous, A. Shtipliyski, T. Schuh, L. Skinnari, S.P. Summers, A. Tapper, A. Thea, I. Tomalin, K. Uchida, P. Vichoudis, S. Viret, and M. Weber. "An FPGA based track finder for the L1 trigger of the CMS experiment at the High Luminosity LHC". In: *Journal of Instrumentation* 12.12 (Dec. 2017), P12019–P12019. DOI: [10.1088/1748-0221/12/12/p12019](https://doi.org/10.1088/1748-0221/12/12/p12019). URL: <https://doi.org/10.1088/1748-0221/12/12/p12019>.
- [12] H. Mohr, T. Dritschler, L. E. Ardila, M. Balzer, M. Caselle, S. Chilingaryan, A. Kopmann, L. Rota, T. Schuh, M. Vogelgesang, and M. Weber. "Evaluation of GPUs as a level-1 track trigger for the High-Luminosity LHC". In: *Journal of Instrumentation* 12.04 (Apr. 2017), pp. C04019–C04019. DOI: [10.1088/1748-0221/12/04/c04019](https://doi.org/10.1088/1748-0221/12/04/c04019). URL: <https://doi.org/10.1088/1748-0221/12/04/c04019>.
- [13] M. Caselle, L.E. Ardila Perez, M. Balzer, T. Dritschler, A. Kopmann, H. Mohr, L. Rota, M. Vogelgesang, and M. Weber. "A high-speed DAQ framework for future high-level trigger and event building clusters". In: *Journal of Instrumentation* 12.03 (Mar. 2017), pp. C03015–C03015. DOI: [10.1088/1748-0221/12/03/c03015](https://doi.org/10.1088/1748-0221/12/03/c03015). URL: <https://doi.org/10.1088/1748-0221/12/03/c03015>.
- [14] M. Caselle, L.E. Ardila-Perez, M. Balzer, A. Kopmann, L. Rota, M. Weber, M. Brosi, J. Steinmann, E. Bründermann, and A.-S. Müller. "KAPTURE-2. A picosecond sampling system for individual THz pulses with high repetition rate". In: *Journal of Instrumentation* 12.01 (Jan. 2017), pp. C01040–C01040. DOI: [10.1088/1748-0221/12/01/c01040](https://doi.org/10.1088/1748-0221/12/01/c01040). URL: <https://doi.org/10.1088/1748-0221/12/01/c01040>.
- [15] L. Rota, M. Vogelgesang, L.E. Ardila Perez, M. Caselle, S. Chilingaryan, T. Dritschler, N. Zilio, A. Kopmann, M. Balzer, and M. Weber. "A high-throughput readout architecture based on PCI-Express Gen3 and DirectGMA technology". In: *Journal of Instrumentation* 11.02 (Feb. 2016), P02007–P02007. DOI: [10.1088/1748-0221/11/02/p02007](https://doi.org/10.1088/1748-0221/11/02/p02007). URL: <https://doi.org/10.1088/1748-0221/11/02/p02007>.

Proceedings

- [16] Luigi Calligaris, André Cascadan, Bruno Casu, Alison França da Costa, Ailton Akira Shinoda, Lucas Arruda Ramalho, Luis E. Ardila-Perez, and Oliver Sander. “OpenIPMC: a free and open source Intelligent Platform Management Controller”. In: *22nd IEEE Real Time Conference*. Nov. 2020. arXiv: [2011.01088](https://arxiv.org/abs/2011.01088) [[physics.ins-det](https://arxiv.org/abs/2011.01088)].
- [17] R. Aggleton, L. Ardila-Perez, F. A. Ball, M. N. Balzer, J. Brooke, L. Calligaris, M. Caselle, D. Cieri, E. J. Clement, G. Hall, K. Harder, P. R. Hobson, G. M. Iles, T. James, K. Manolopoulos, T. Matsushita, A. D. Morton, D. Newbold, S. Paramesvaran, M. Pesaresi, I. D. Reid, A. W. Rose, O. Sander, T. Schuh, C. Shepherd-Themistocleous, A. Shtipliyski, S. P. Summers, A. Tapper, I. Tomalin, K. Uchida, P. Vichoudis, and M. Weber. “A novel FPGA-based track reconstruction approach for the level-1 trigger of the CMS experiment at CERN”. In: *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*. Sept. 2017, pp. 1–4. DOI: [10.23919/FPL.2017.8056825](https://doi.org/10.23919/FPL.2017.8056825).
- [18] M. Vogelgesang, L. Rota, L.E. Ardila-Perez, M. Caselle, S. Chilingaryan, and A. Kopmann. “High-throughput data acquisition and processing for real-time x-ray imaging”. In: vol. 9967. 2016. DOI: [10.1117/12.2237611](https://doi.org/10.1117/12.2237611). URL: <http://dx.doi.org/10.1117/12.2237611>.

Conferences

- [19] Luis Ardila-Perez, Luigi Calligaris, André Cascadan, Bruno Casu, Alison da França Costa, Lucas Ramalho, Oliver Sander, and Ailton Shinoda. *The OpenIPMC project*. 2021. URL: https://indico.cern.ch/event/1021679/contributions/4333600/attachments/2242916/3803473/20210511__xTCA_IG__OpenIPMC_Project.pdf.
- [20] Luis Ardila. *KIT-IPE Hardware updates*. 2021. URL: https://indico.cern.ch/event/1003095/contributions/4212135/attachments/2189382/3700048/2021-02-11_DPS-HW.pdf (visited on 08/20/2021).
- [21] Luis Ardila, Denis Tcherniakhovski, Alex Howard, Greg Iles, Rui Zou, and Charlie Strohmman. *12x25 G Firefly alpha-1 sample tests*. 2021. URL: <https://indico.cern.ch/event/993270/contributions/4223652/attachments/2190333/3701798/12x25%20G%20Firefly%20Alpha-1%20sample%20test.pdf> (visited on 08/21/2021).
- [22] Luis E. Ardila-Perez, Luigi Calligaris, André Cascadan, Bruno Casu, Alison da França Costa, Lucas Ramalho, Oliver Sander, and Ailton Shinoda. *OpenIPMC: a free and open source Intelligent Platform Management Controller*. 2020. URL: https://indico.cern.ch/event/863071/contributions/3856106/attachments/2046289/3428449/POSTER__OpenIPMC_an_Open_Source_Intelligent_Platform_Management_Controller.pdf.

-
- [23] L.E. Ardila-Perez and O. Sander. *CMS R&D for Phase-2 Tracker Back-end Electronics*. 2019. URL: https://indico.cern.ch/event/799275/contributions/3413734/attachments/1860874/3058571/SoC_RD_for_CM_S_Phase-2_Tracker_Back-end_Electronics.pdf.
- [24] L.E. Ardila-Perez, O. Sander, T. Schuh, D. Tcherniakhovski, D. Bormann, M. Balzer, M. Weber. *EureKA-Maru: an ATCA board for the CMS Phase 2 Tracker Upgrade with centralized slow control and board management solution based on a Zynq Ultrascale+ System-on-Chip*. 2019. URL: <https://indico.cern.ch/event/799025/contributions/3486510/>.
- [25] O. Sander, L. Ardila-Perez, D. Tcherniakhovski and M. Balzer, and M. Weber. *A novel centralized slow control and board management solution for ATCA blades based on the Zynq Ultrascale+ System-on-Chip*. 2019. URL: https://indico.cern.ch/event/773049/contributions/3474311/attachments/1933805/3211730/CHEP19_ZUSP_IPMC.pdf.
- [26] A. Akira Shinoda and L.E. Ardila-Perez and L. Arruda Ramalho and M. Balzer and D. Bormann and L. Calligaris and A. Cascadan and V. Finotti and A. França Queiroz da Costa and O. Sander and M. Schleicher and T. Schuh and D. Tcherniakhovski and S. de Souza and M. Weber. *Ultraflex: An ATCA prototype board for the CMS Phase 2 Tracker Upgrade*. 2018. URL: https://indico.cern.ch/event/697988/contributions/3056081/attachments/1718840/2773929/poster_TWEPP_ZYNQ.pdf.
- [27] L. E. Ardila-Perez. *The HL-LHC CMS Level-1 Track Trigger*. 2018. URL: <https://indico.desy.de/indico/event/19924/session/4/contribution/8/material/slides/1.pdf>.
- [28] L.E. Ardila-Perez. *Duplicate Removal Algorithm of the Time-Multiplied Track Trigger of CMS*. 2018. URL: https://indico.cern.ch/event/689620/contributions/2879964/attachments/1599315/2534971/DR_2008_02_13.pdf.
- [29] Luis Ardila-Perez. *Level-1 track fining with an all-FPGA system at CMS for the HL-LHC*. 2018. URL: <https://indico.desy.de/indico/event/19924/session/4/contribution/8/material/slides/0.pdf>.
- [30] L. Ardila-Perez, F. Ball, M. Balzer, M. Caselle, L. Calligaris, D. Cieri, E. Clement, K. Harder, G. Iles, T. James, K. Manolopoulos, A. Morton, D. Newbold, M. Pesaresi, I. Reid, A. Rose, O. Sander, C. Shepherd-Themistocleous, T. Schuh, S. Summers, I. Tomalin, and M. Weber. *Scalability of the Time Multiplexed CMS Level1 Track Trigger System*. Jan. 2017. URL: https://indico.cern.ch/event/566138/contributions/2466071/attachments/1407084/2153161/2017_INFIERI_SP_LA.compressed.pdf.
- [31] L. Ardila-Perez. *TMTT Duplicate Tracks Removal*. Oct. 2016. URL: https://indico.cern.ch/event/557734/contributions/2322836/attachments/1358139/2053961/TMTT_Duplicate_Track_Removal.pdf.

- [32] M. Caselle, L. Ardila-Perez, S. Chilingaryan, T. Dritschler, A. Kopmann, H. Mohr, L. Rota, M. Vogelgesang, M. Balzer, and M. Weber. *High-speed low-latency readout system with realtime trigger based on GPUs*. June 2016. URL: https://indico.cern.ch/event/390748/contributions/1825218/attachments/1281568/1917838/Trigger2_99_Caselle.pdf.
- [33] H. Mohr, L. Ardila-Perez, M. Balzer, M. Caselle, L. Rota, A. Kopmann, S. Chilingaryan, T. Dritschler, M. Vogelgesang, and M. Weber. *Evaluation of GPUs for High-Level Triggers in High Energy Physics*. 2016. URL: <http://indico.cern.ch/event/489996/contributions/2211076/>.

Collaboration Technical Reports

- [34] Luis Ardila, Dan Gastler, Kristian Hahn, Eric Hazen, Greg Iles, Kevin Lannon, MarkPesaresi, Oliver Sander, Thomas Schuh, Sarah Seif El Nasr-Storey, Ian Tomalin, Tom Williams, and Peter Wittich. *Specification of the Phase-2 Tracker Backend Electronics*. CMS DN-18-011. Geneva, 2020. URL: https://espace.cern.ch/Tracker-Upgrade/Data-Processing/Shared%20Documents/DN-18-011_23_03_2020.pdf.