

UNIVERSITÀ  
DEGLI STUDI  
DI PADOVA



Dipartimento  
di Fisica  
e Astronomia  
Galileo Galilei

**University of Padova**

Department of Physics and Astronomy "Galileo Galilei"

Ph.D. course in Physics  
XXXVII cycle

# 40MHz scouting at the CMS experiment

PH.D. CANDIDATE

**Matteo Migliorini**

Student ID 2039306

SUPERVISOR

**Prof. Jacopo Pazzini**

CO-SUPERVISOR

**Prof. Andrea Triossi**

COORDINATOR

**Prof. Giulio Monaco**

ACADEMIC YEAR  
2023/2024







*Viazar descanta...*  
*ma chi parte mona torna mona.*  
- Proverbio Veneto



## Abstract

Experiments at the Large Hadron Collider rely on sophisticated trigger systems to select and store a small fraction of “interesting” collision events. These systems are crucial for managing the enormous data rates produced by the experiments, ensuring that their physics program can be effectively carried out. However, events that are not selected are lost forever, which may limit the discovery potential by discarding signals that fall outside the trigger criteria and budget.

This thesis presents the development and first implementation of a 40 MHz Level-1 trigger scouting system for the CMS experiment at the Large Hadron Collider. Designed to capture and process data flowing through the L1 trigger boards at the bunch crossing rate before any trigger selection, this system enables the exploration of uncharted regions of phase space by bypassing traditional trigger constraints. A complete demonstrator was successfully deployed and operated during the 2024 LHC  $pp$  run. Early data from this demonstrator has been utilized in the initial steps toward a low-mass dijet search, an analysis typically constrained by stringent trigger thresholds due to the overwhelming background from QCD multi-jet production.

In addition, this work explores new triggerless data acquisition and processing techniques, implemented within the CMS Drift Tubes muon detector Phase-2 demonstrator. This allowed the collection of valuable data for detector monitoring and background studies during  $pp$  collisions. Furthermore, a novel data acquisition system based on Remote Direct Memory Access technology was developed, offering an efficient data transfer mechanism with the potential to enhance both triggerless systems and standard data acquisition methods.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>The LHC and CMS experiment and their Upgrades</b>	<b>5</b>
2.1	The Large Hadron Collider . . . . .	5
2.1.1	The LHC accelerator complex . . . . .	5
2.1.2	Design and specifications . . . . .	6
2.1.3	LHC operations and High-Luminosity upgrade . . . . .	9
2.2	The Compact Muon Solenoid experiment . . . . .	11
2.2.1	Coordinate system . . . . .	12
2.2.2	Sub-detectors . . . . .	15
2.2.3	Trigger and Data Acquisition . . . . .	23
2.3	Phase-2 upgrade . . . . .	27
2.3.1	Sub-detectors . . . . .	27
2.3.2	L1 Trigger . . . . .	31
<b>3</b>	<b>40 MHz Trigger-less Readout of Detector Front-ends</b>	<b>35</b>
3.1	40MHz readout of the CMS Drift Tubes detector . . . . .	37
3.1.1	The miniDT detector . . . . .	37
3.1.2	Trigger-less readout chain . . . . .	38
3.2	An online processing system for streaming data . . . . .	42
3.2.1	Distributed processing and messaging systems . . . . .	43
3.2.2	Multistage architecture . . . . .	45
3.3	DT Phase-2 upgrade demonstrator . . . . .	48
3.3.1	Data Scouting results during collisions . . . . .	51
3.3.2	Readout of the final OBDT prototypes . . . . .	61
3.4	FEROCE project . . . . .	67
3.4.1	A 100G RoCEv2 Implementation . . . . .	69
3.4.2	Towards a light-RoCE HDL Implementation . . . . .	73

## CONTENTS

<b>4</b>	<b>The CMS L1-Trigger scouting system</b>	<b>79</b>
4.1	Data taking strategies at the CMS experiment . . . . .	79
4.2	The L1-Trigger scouting system . . . . .	82
4.3	L1 Scouting demonstrator for LHC RUN-3 . . . . .	86
4.4	Data acquisition . . . . .	88
4.4.1	Hardware . . . . .	89
4.4.2	Software . . . . .	91
4.5	Online processing system . . . . .	94
4.5.1	Orbit Building . . . . .	94
4.5.2	Online data streams . . . . .	96
4.5.3	Online data handling . . . . .	101
4.5.4	Offline data processing and distribution . . . . .	103
4.6	Operation and performance . . . . .	104
4.7	First data from the L1 scouting demonstrator . . . . .	106
4.7.1	Muons . . . . .	107
4.7.2	Calorimeter objects . . . . .	112
4.7.3	BMTF stubs . . . . .	126
<b>5</b>	<b>Towards a dijet analysis with L1 Scouting data</b>	<b>131</b>
5.1	Low-mass dijet searches at CMS . . . . .	132
5.2	Data and Simulated samples . . . . .	135
5.2.1	Data Sample . . . . .	135
5.2.2	Signal Simulation . . . . .	135
5.2.3	Background Simulation . . . . .	140
5.3	Dijet Data and QCD Background . . . . .	141
5.4	Modelling of the QCD background . . . . .	146
5.4.1	Parametric model . . . . .	147
5.4.2	Ratio method . . . . .	148
5.5	Future steps . . . . .	150
<b>6</b>	<b>Conclusions</b>	<b>155</b>
<b>A</b>	<b>L1 Scouting data formats</b>	<b>157</b>
	<b>References</b>	<b>159</b>



# Introduction

The Standard Model (SM) of particle physics[1, 2] represents our current best understanding of the subatomic world. Developed and refined throughout the second half of the 20th century, the SM has been extensively tested by numerous experiments under various conditions, and its predictions have been measured to high precision. A fundamental milestone in its validation was the discovery of the Higgs boson by the ATLAS and CMS Collaborations at the Large Hadron Collider (LHC) in 2012[3, 4].

Despite its remarkable success, the SM is not a complete theory of fundamental interactions[5]. It fails to explain several key phenomena, such as the origin of the electroweak scale, the masses of neutrinos, the flavor structure of quarks and leptons, and it cannot account for the existence of dark matter or dark energy, pointing to the need for extensions beyond the SM, commonly referred to as physics beyond the Standard Model (BSM)[6, 7].

High-energy accelerator-based experiments have played a key role in the development and validation of the SM and remain essential in the search for new physics. In these experiments, particles such as protons or electrons are accelerated and collided either on fixed targets or with other particles at high center-of-mass energies. The particles produced in these collisions provide insight into the fundamental forces and interactions that govern their behavior. Detectors are employed to identify and reconstruct the properties of these particles, allowing researchers to infer the nature of the underlying interactions.

Modern high-energy physics experiments produce enormous volumes of data. Whether in terms of the instantaneous data rates produced or the overall




Figure 1.1: Production cross-sections for several representative standard model processes at hadron colliders versus center-of-mass energy. Lines discontinuity indicates the transition from  $p\bar{p}$  at Tevatron to  $pp$  at LHC. The vertical dotted lines mark the energies 1.96 TeV at Tevatron, and 7 TeV, 8 TeV, 13 TeV, and 13.6 TeV at LHC[8].

data volumes analyzed, handling the vast quantities of information generated by these experiments poses significant challenges, especially in the search for rare processes or subtle deviations from SM predictions.

The sheer number of events collected in these experiments arises from the fact that the “interesting” processes, such as rare SM interactions or potential BSM signals, have much lower cross-sections than more common background processes. This is especially true for hadron colliders. In Figure 1.1, the cross-sections of several representative SM processes are shown as a function of center-of-mass energy  $\sqrt{s}$  for both  $pp$  and  $p\bar{p}$  colliders. The total  $pp$  cross-section and even processes like  $b\bar{b}$  production are orders of magnitude larger than, for example, the cross-section of Higgs boson production.

Storing all the data produced from every collision is not feasible, both from data acquisition and analysis perspectives. The sheer volume of data would overwhelm storage and processing capacities. For this reason, experiments employ online selection mechanisms, called triggers, to identify and retain only a small fraction of events deemed interesting. At the LHC, these triggers typically select approximately  $O(0.01\%)$  of all events, where the selection of interesting events is based on theoretical expectations and predefined criteria. This selective



approach enabled the discovery of the Higgs boson and has been instrumental in probing the SM. Moreover, since the beginning of LHC operations, many efforts were dedicated to search for BSM signals.

Nevertheless, the key takeaway from the LHC so far is that the universe remains remarkably consistent with the Standard Model: despite the increasingly large amount of data collected at the LHC, no conclusive evidence of BSM physics has been observed[9, 10]. Before LHC operations began, it was anticipated that new physics would emerge at the TeV scale, accessible even during the first LHC run at 7 TeV[7, 11]. However, no significant deviations from SM predictions have been detected.

This raises the question: Could our expectations of new physics be wrong? Could the signatures of BSM physics be more elusive or experimentally challenging, thereby escaping detection? One possibility is that new particles, such as long-lived particles[12], might decay far from the interaction point or at a delayed time, making them difficult to observe with conventional search strategies.

Another possibility is that new physics signals could be missed by the online selection system, ending up in the vast majority of discarded collision events. One notable example is in the context of dijet searches. Resonant structures in the invariant mass distribution of dijet events are among the most generic signatures for “exotic” new physics at hadron colliders[13, 14]. However, access to the low-mass region of the spectrum is hindered by the overwhelming QCD multijet background, which necessitates tight trigger thresholds, thereby limiting sensitivity to low-mass resonances. While it would be ideal to collect and analyze all the data, this is simply not feasible.

The goal of this work is to develop a system capable of “scouting” the full set of collision events, free from the bias introduced by trigger selections. This system would allow for the exploration of events that currently do not pass standard trigger criteria and thus are forever lost. Such ambitious goal can be achieved through the collection of data at the trigger level, with coarser resolution than offline reconstructed data, but at more manageable volumes.

## Overview of the Thesis and Author’s Contribution

The work presented in this thesis is conducted within the context of the CMS experiment at the LHC, which is introduced in chapter 2.

In chapter 3, the feasibility of collecting data directly from the detectors

frontend in a triggerless mode is explored, using the CMS Drift Tubes muon detector as a demonstrator. Additionally, new data acquisition technologies that could facilitate the triggerless approach are investigated. I contributed to the development and validation of firmware components, while I designed and implemented the software components and the analysis of the first data collected in this mode.

Chapter 4 describes the 40 MHz Level-1 trigger scouting system, designed to capture and process data from the CMS Level-1 trigger at the full bunch crossing rate without any filtering. I contributed to the implementation and commissioning of a demonstrator of this novel data acquisition system during Run-3 of the LHC, with a focus on designing and implementing the online processing system responsible for ingesting raw data from the L1 trigger and creating analysis streams.

Finally, chapter 5 presents the first steps toward a low-mass dijet analysis using data collected with the L1 trigger scouting demonstrator during the 2024 LHC  $pp$  run. My contributions include performing the first exploratory analysis of the collected data, and signal and background simulations along with developing the analysis strategy.



# The LHC and CMS experiment and their Upgrades

## 2.1 The Large Hadron Collider

The Large Hadron Collider (LHC), located at CERN in Geneva, Switzerland, is the world's largest and most powerful particle accelerator. Housed in a 27-kilometer tunnel previously occupied by the Large Electron-Positron (LEP) collider, the LHC lies between 45 and 170 meters below ground level. This circular accelerator is designed to accelerate both protons and heavy ions. Currently, it achieves a center-of-mass energy of 13.6 TeV in proton mode, making it the most advanced machine of its kind.

The development of the LHC, from its initial conception to full operation, took over two decades. First proposed in 1984, it began operations in 2008. The LHC is now in its third operational run, with plans for a high-luminosity upgrade that will extend its functionality into the early 2040s.

### 2.1.1 The LHC accelerator complex

The LHC serves as the final component in a series of accelerators[15], illustrated in Figure 2.1, that were enhanced to meet its demanding specifications, enabling the study of extremely high-energy particle collisions. The proton injection process starts with the extraction of protons from a hydrogen gas reservoir, which are then accelerated and bunched in a Radio Frequency Quadrupole

## 2.1. THE LARGE HADRON COLLIDER

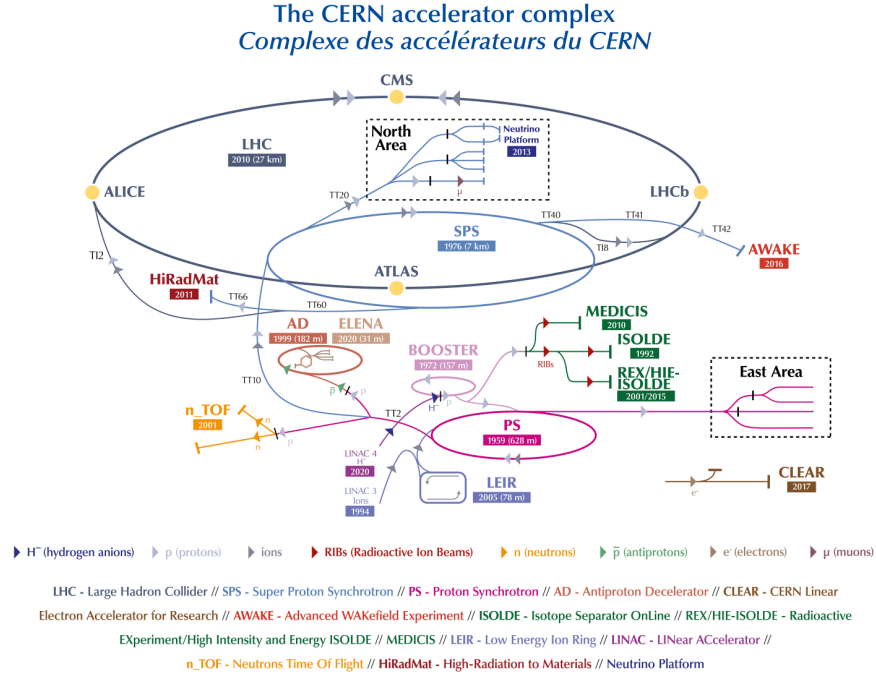


Figure 2.1: The accelerator complex at CERN.

(RFQ) and further boosted by the Linear Accelerator (LINAC 2). The protons then pass through a series of circular accelerators, namely the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS), where their energy is progressively increased up to 450 GeV. Once the protons reach this energy level, they are injected into the LHC, where their energy is further increased to the operational level of up to 7 TeV using high-frequency accelerating cavities. The LHC employs a sophisticated system of superconducting dipole and quadrupole magnets to maintain beam stability and focus, ensuring optimal conditions for collision at the four designated Interaction Points (IPs), where the particle detectors are located. To ensure optimal beam intensity and maximize the number of collisions, beam circulation and data acquisition typically last for approximately 10 hour, after which the beams are dumped. A new fill is then initiated, and the process is repeated.

### 2.1.2 Design and specifications

The LHC was designed to collide protons, overcoming the LEP's limitation of synchrotron radiation, which significantly affects lighter particles like electrons. Because the power lost to synchrotron radiation scales with the inverse of the

fourth power of the particle's mass, i.e.  $\propto m^{-4}$ , protons experience much less energy loss, enabling the LHC to reach a nominal beam energy ( $E$ ) of 7 TeV and a  $\sqrt{s}$  of 14 TeV. This, on the other hand, leads to a much more complex collision dynamics, due to the composite structure of the colliding protons.

In the LHC, protons are grouped into discrete packets known as *bunches*, each containing about  $10^{11}$  protons. The spacing and number of filled bunches are carefully controlled to optimize collision rates and manage the accelerator's performance. The *filling scheme* refers to their arrangement and distribution within the LHC ring[16]. Filled bunches can either be *isolated*, meaning they have empty bunches around them, or organized into *trains* consisting of a sequence of closely spaced filled bunches, with gaps between different trains. The moment when two bunches from the two beams cross is referred to as a *bunch crossing* (BX). The bunches are designed to be separated by 25 ns, resulting in a BX rate of 40 MHz. The filling scheme repeats itself every 3564 bunches, known as an *orbit* of the LHC ring, with a frequency of 11 kHz. At the end of each orbit is the so-called *abort gap*[17], a 3  $\mu$ s long gap intentionally left free of particles. This gap provides a buffer time that allows the machine's beam dump system to safely extract the entire beam in case of an emergency or at the end of a fill. Figure 2.2 shows a representative example of a filling scheme for illustrative purposes. Filling schemes at the LHC are often subject to modification to adjust the operating conditions according to the needs of the experiments.

A central parameter of the LHC machine is the *instantaneous luminosity*  $\mathcal{L}$  at which collisions take place:

$$\mathcal{L} = F \cdot \frac{n_b \cdot f_b \cdot N_b^2 \cdot \gamma}{4\pi \cdot \epsilon_n \cdot \beta^*} \quad (2.1)$$

Such quantity is directly connected to the number of interactions at a given time  $t$ . It depends on the number of colliding bunches  $n_b$ , their frequency of revolution  $f_b$  and the number of protons in the bunches  $N_b$ . Here the transverse emittance and the focal length of the beam at the IP are quantified by  $\epsilon_n$  and the beta function  $\beta^*$ [18], respectively. The factor  $F$  accounts for the luminosity loss due to the geometry of the IP, depending on the beam crossing angle and the longitudinal and transverse r.m.s. bunch sizes at collision. The LHC is designed to function at a nominal instantaneous luminosity  $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ .

The figure of merit commonly used to quantify the number of collisions

## 2.1. THE LARGE HADRON COLLIDER




Figure 2.2: Representative example of filling scheme at the LHC. An orbit contains 3564 bunches, spaced 25 ns apart from each other. Only some specific bunches actually contain protons, forming the so-called trains of particles.

delivered or recorded by an experiment over a specific time interval is the *integrated luminosity*  $L$ , calculated by integrating the instantaneous luminosity over the desired time span.

$$L = \int_{t_0}^{t_1} \mathcal{L} dt \quad (2.2)$$

As collecting a large amount of events is crucial for probing rare processes, the LHC is designed to maximize the luminosity in order to provide the experiments with a large dataset for physics analysis. However, this design also introduces the phenomenon known as *pileup* (PU). Pileup refers to the occurrence of multiple proton-proton collisions happening simultaneously within a single bunch crossing. Given the high instantaneous luminosity of the LHC, each bunch crossing can result in several overlapping interactions in the same spatial region, surrounding the primary collision event. At a center-of-mass energy of 13 TeV, the average pileup at the LHC was designed to be around  $\langle PU \rangle \sim 22$ . However, over the years, this value has been pushed to significantly higher levels, reaching approximately three times the design value.

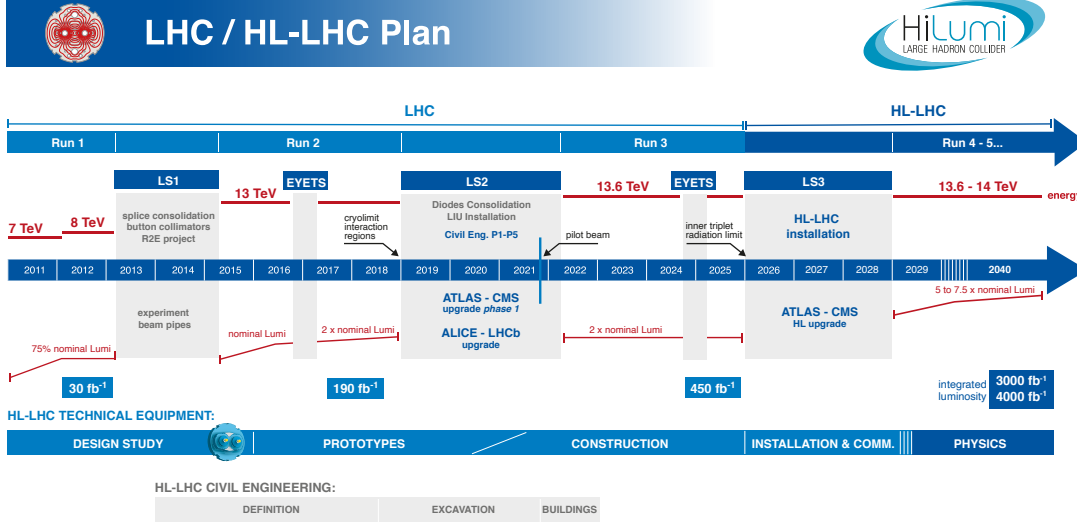


Figure 2.3: The LHC / High-Luminosity LHC plan[20].

### 2.1.3 LHC operations and High-Luminosity upgrade

Figure 2.3 summarizes the past and future timeline of operations at the LHC. At the time of writing, the LHC is approaching the end of its first period of operation, known as *Phase 1*, which began over a decade ago. Following this, there will be a long shutdown period dedicated to upgrading the machine in preparation for the high-luminosity upgrade of the LHC[19], known as *Phase 2*, which aims to deliver an unprecedented amount of data to the experiments.

#### LHC operations during Phase 1

The LHC began delivering luminosity to the experiments in 2010 and operated for two years, reaching center-of-mass energies up to  $\sqrt{s} = 8$  TeV. The dataset collected during this period, known as Run 1, led to a fundamental milestone in particle physics: the discovery of the Higgs Boson in 2012 by the ATLAS and CMS collaborations[3, 4].

Two additional data-taking campaigns followed, interleaved with shutdown periods for maintenance and development operations: namely, Run 2 (2016-2018) and Run 3, which began in 2022 and is ongoing at the time of writing (2024). During Run 2 and Run 3, the center-of-mass energy was increased to 13 TeV and 13.6 TeV, respectively. Additionally, the LHC's nominal instantaneous

## 2.1. THE LARGE HADRON COLLIDER

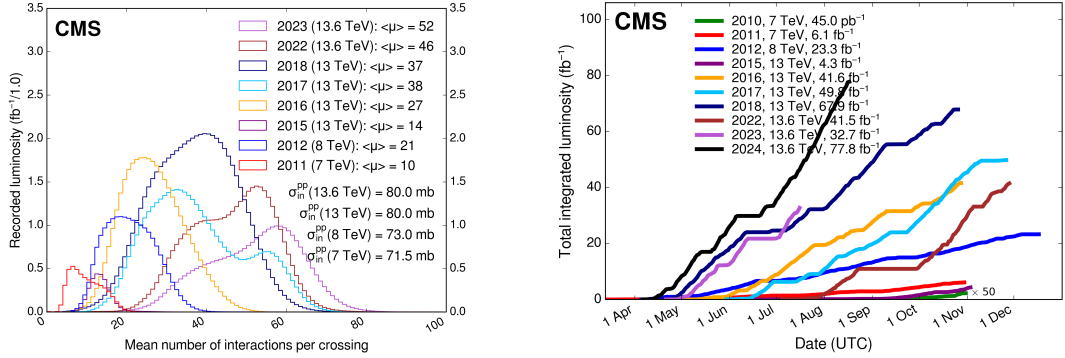


Figure 2.4: Pileup (left) and integrated luminosity (right) throughout the years of Runs 1, 2 and 3[21].

luminosity has been exceeded by a factor of 2, with the average PU now reaching up to 60, as shown in Figure 2.4. This has allowed for the collection of a vast dataset for physics analyses, but it has also posed significant challenges for both the accelerator and the experiments.

### LHC upgrade for the High-Luminosity phase

In order to deliver the substantially larger dataset needed to significantly improve the current physics reach of the LHC experiments, the accelerator will transition into its High-Luminosity phase, known as the HL-LHC, planned to start in 2029.

The target instantaneous luminosity at the beginning of HL-LHC is  $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ , and it is expected to reach up to  $7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$  within the first years of operation. This translates also in unprecedented PU levels, ranging from 140 to 200 collisions per BX. At the same time, the center-of-mass energy will be increased up to the original design value of 14 TeV, which was never reached during Phase 1.

Achieving such ambitious luminosity target requires significant upgrades to the LHC's accelerator components, including injector systems, kickers, and cryogenic mechanisms. The primary hardware for the HL-LHC will be installed during Long Shutdown 3 (LS3) from 2026 to 2028, with the goal of completing hardware commissioning by the machine's restart in 2029 and ensuring efficient operations through 2040. This upgrade will rely on cutting-edge technologies, such as 11-12 T superconducting magnets, ultra-precise superconducting



radio-frequency cavities for beam manipulation, and 100-meter-long superconducting links with zero energy loss. Enhanced luminosities will also demand improvements in vacuum systems, cryogenics, machine protection, collimation, diagnostics, and beam modeling[19].

Over the currently foreseen 12 years of operation, the HL-LHC is expected to deliver a total integrated luminosity within 3000 and 4000 fb<sup>-1</sup>, thus opening the way to unprecedented studies of exceedingly rare phenomena and possibly unveiling new physics never observed.

## 2.2 The Compact Muon Solenoid experiment

Strategically positioned within the LHC ring are four primary experiments at designated interaction points. Among those four, ATLAS[22] and CMS[23] are general-purpose detectors, while ALICE[24] focuses on heavy-ion physics and quark-gluon plasma investigations, and LHCb[25] targets mainly b-physics studies. This thesis focuses on the CMS experiment, whose broad physics programme ranges from the study of Standard Model (SM) particles to the search for new physics beyond the SM. The CMS experiment, as ATLAS, was designed to probe a large spectrum of physics phenomena, and played a fundamental role in the discovery of the Higgs Boson.

The name CMS (Compact Muon Solenoid) reflects its key design features. *Compact* refers to the detector's dense design, which is six times smaller in volume than ATLAS but twice as heavy. *Muon* highlights the large volume dedicated to muon tracking chambers, and finally *Solenoid* refers to the central feature of the detector, namely its superconducting solenoid magnet that generates a magnetic field of up to 3.8 Tesla.

The CMS detector is located in Cessy, France, and it is housed in a cavern approximately 100 meters below the surface. The underground area is divided into two main caverns: the Experimental Cavern (UXC) and the Service Cavern (USC), shown in Figure 2.5. The UXC is the larger of the two caverns. It houses the CMS detector itself and it was designed to allow for the assembly, operation, and maintenance of the detector. It is thus equipped with several support structures, such as platforms and cranes, to facilitate access to various parts of the detector. Located adjacent to the experimental cavern, the USC is slightly smaller but equally critical to the operation of the CMS experiment. This cavern houses a range of essential infrastructure, including power supplies,

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT

cooling systems, trigger and data acquisition systems. Finally, the control room is located on the surface building above the experiment, as well as several other service buildings, including the server farm used as the final step of the data acquisition chain.




Figure 2.5: Underground areas of the CMS experiment: the service cavern (USC) hosts infrastructures as cooling and power systems and the off-detector electronics, while the experimental cavern (UXC) hosts the detector itself.

### 2.2.1 Coordinate system

The CMS experiment describes detector geometry and collisions through a right-handed Cartesian coordinate system, centered on the nominal bunch crossing point. In this system, the  $x$  and  $y$  axes form a plane perpendicular to the proton beamline, known as the transverse plane; the  $x$ -axis points toward the geometric center of the LHC, while the  $y$ -axis points upward, perpendicular to the LHC plane. The  $z$ -axis serves as the longitudinal coordinate, aligned with the direction of the anticlockwise proton beam. Given the detector's cylindrical symmetry, a polar coordinate system is also defined based on this Cartesian framework. Here, the radial coordinate  $r$  is measured from the nominal interaction point, the polar angle  $\theta$  is defined as the angle between  $r$  and the  $z$ -axis, and the azimuthal angle  $\phi$  is the angle between  $r$  and the  $x$ -axis in the transverse

plane. Figure 2.6 provides a schematic illustration of the CMS Cartesian and polar coordinate systems.




Figure 2.6: Coordinate system adopted in the CMS experiment.

Each parton in the colliding protons carries a fraction of the proton's longitudinal momentum, while their momentum in the transverse plane is always negligible. As a result, the collision reference frame can only experience a boost along the  $z$ -axis. Therefore, one can define Lorentz-invariant variables by projecting the momentum onto the transverse plane, leading to the definitions of *transverse momentum* ( $p_T$ ) and *transverse mass* ( $m_T$ ):

$$p_T = \sqrt{p_x^2 + p_y^2} \quad (2.3)$$

$$m_T = \sqrt{m^2 + p_x^2 + p_y^2} = \sqrt{E^2 - p_z^2} \quad (2.4)$$

As the transverse momentum of the center-of-mass framework is equal to zero and the total momentum is conserved, the sum of the transverse momenta of particles is expected to be zero. However, due to their small interaction cross-section, neutrinos go undetected inside the CMS detector: one can therefore reconstruct their global momentum in the transverse plane as the so called *missing transverse momentum*, also known as *missing transverse energy* ( $MET$ ,  $\cancel{E}_T$ , or  $p_T^{miss}$ ):

$$p_T^{miss} = - \left| \sum_{i=1}^N \mathbf{p}_{i,t} \right| \quad (2.5)$$

Another set of fundamental quantities at hadron colliders are the *rapidity*  $y$

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT

and *pseudorapidity*  $\eta$  of a particle, which are defined as:

$$y = \frac{1}{2} \left( \frac{E + p_z}{E - p_z} \right) \quad (2.6)$$

and

$$\eta = -\log \left[ \tan \left( \frac{\theta}{2} \right) \right] \quad (2.7)$$

These two variables represent spatial coordinates used to describe the angle of a particle's trajectory relative to the beam axis. Even though  $y$  is not a Lorentz-invariant quantity, the rapidity difference between two particles  $\Delta y$  is. The importance of pseudorapidity, on the other hand, is due to the fact that on relativistic regimes particle production tends to be relatively uniform as a function of  $\eta$ . A visualization of  $\eta$  is shown in Figure 2.7.




Figure 2.7: Pseudorapidity visualization.

Another variable widely used to describe collision events is the angular distance between two particles  $\Delta R$ , which is defined as:

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \quad (2.8)$$

Finally, when considering the kinematics of a multi-body interaction, an extremely important quantity *invariant mass*  $m$ , defined as follows:

$$m = \sqrt{\left( \sum_{i=1}^n E_i \right)^2 - \left\| \sum_{i=1}^n \vec{p}_i \right\|^2} \quad (2.9)$$




Figure 2.8: 3D view of the CMS detector.

### 2.2.2 Sub-detectors

The CMS detector is designed to be highly granular, quick in response, and resistant to radiation. It comprises several concentric subdetectors, each designed to detect different particle signatures stemming from the collisions. Figure 2.8 show the structure of CMS through a 3D model and an octant of its transverse section, respectively. As charged particles fly outward from the collision point, their trajectories are bent by the magnet, which allows for the identification of their charge and measurement of their momentum; positive and negative particles bend in opposite directions, with high-momentum particles bending less than low-momentum ones. The precise identification of these particle tracks is accomplished by the silicon pixel and tracker, along with muon detectors. The energy of the particles is measured by two different calorimeters: the Electromagnetic Calorimeter (ECAL), which fully absorbs electrons and photons, and the Hadron Calorimeter (HCAL), which stops hadrons after they pass through the ECAL.

Muons, being minimum ionizing particles (MIPs), are not stopped by the calorimeters. Consequently, their momentum is measured both inside the superconducting coil by the tracking detectors and outside it by the muon chambers. Neutrinos, on the other hand, do not interact with the CMS detector and therefore can't be detected directly. Their presence is inferred from the

missing transverse energy, which is the imbalance in the total transverse momentum of all detected particles, as defined in Equation 2.5. Wide rapidity coverage, hermeticity and high-precision energy measurement are thus fundamental characteristics of the CMS detector. Further details on each subdetector will be outlined in the following paragraphs.

### Tracker

The Silicon Tracker represents the innermost layer of the CMS detector[26], designed to precisely track the paths of charge particles emerging from collisions, even in the high-density environment near the interaction point. It is designed to reconstruct muons, electrons, and hadrons with high momentum resolution and high efficiency. It is a fully silicon-based detector, sensitive to the passage of charged particles, which deposit energy through ionization inside the semiconductor material.

The tracker consists of two primary substructures: the pixel detector and the silicon strip detector, as illustrated in Figure 2.9.

It consists of two primary substructures: the pixel detector and the silicon strip detector, sketched in Figure 2.9.




Figure 2.9: Sketch of one quarter of the Phase-1 CMS tracking system in the  $r$ - $z$  view. The pixel detector is shown in green, while single-sided and double-sided strip modules are depicted as red and blue segments, respectively.

The pixel detector[27] forms the innermost part of the tracker, with each pixel measuring  $100 \times 150 \mu\text{m}^2$  in the  $\phi \times \eta$  direction. The modules are arranged in four barrel layers and three endcap disks. This configuration provides a spatial resolution of  $10 \mu\text{m}$  in the transverse plane and  $20 \mu\text{m}$  in the longitudinal direction.

Surrounding the pixel detector are the silicon micro-strip detectors[28], which are arranged in four inner barrel layers (TIB), two inner disks (TID), six concentric outer barrel layers (TOB), and two endcaps (TEC) that close off the tracker. The dimensions of the strips vary depending on their position: those closest to the interaction point have a pitch of  $60\ \mu\text{m}$  and a length of 7 cm, while those farther from the IP have a pitch of  $270\ \mu\text{m}$  and a length of 12.5 cm. This ensures a single-point resolution between 20 and  $50\ \mu\text{m}$  in the transverse plane and between 200 and  $500\ \mu\text{m}$  in the longitudinal direction.

### Calorimetry: ECAL and HCAL

Electrons, photons, and hadrons deposit their energies in the calorimeters. The working principle of the calorimeters is based on the fact that showers produced by particles interacting with the detector can be converted into light, with an intensity proportional to the energy of the particle. The first layer after the tracker, known as the Electromagnetic Calorimeter (ECAL)[29], is specifically designed for precise energy measurements of electrons and photons. The CMS ECAL is a compact, hermetic, fine-grained, homogeneous calorimeter consisting of approximately 76,000 scintillating lead tungstate ( $\text{PbWO}_4$ ) crystals. In the barrel section, avalanche photodiodes detect scintillation light, while vacuum phototriodes are used in the endcap regions due to their superior radiation tolerance. A preshower system (ES) is positioned in front of the endcaps to enhance photon detection and reject background from  $\pi_0$  decays to two photons.

The geometry of the ECAL is depicted in Figure 2.10.

The ECAL achieves exceptional energy resolution since nearly all the energy from electrons and photons is absorbed within the crystal matrix. The intrinsic energy resolution of the ECAL ( $E$ ) for electrons, as measured in a test beam, is given by[30]:

$$\frac{\sigma_E}{E} = \frac{2.8\%}{\sqrt{E(\text{GeV})}} \oplus \frac{12\%}{E(\text{GeV})} \oplus 0.3\%, \quad (2.10)$$

where the three terms correspond to the stochastic, noise, and constant components, respectively. The energy resolution for electrons from Z-boson decays is better than 2% in the central region of the ECAL barrel and ranges from 2% to 5% elsewhere. Similarly, the energy resolution for photons from 125 GeV Higgs boson decays varies across the barrel, ranging from 1.1% to 2.6%, and from 2.2% to 5% in the endcap regions[31].



## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT




Figure 2.10: Layout of the CMS Electromagnetic Calorimeter (ECAL).

Beyond the ECAL lies the Hadron Calorimeter (HCAL), which serves as the outermost calorimetric layer in CMS. The HCAL is primarily dedicated to measuring final-state hadrons and hadronic showers that, although depositing some energy in the ECAL, are not fully contained by it. The structure of the HCAL is shown in Figure 2.11.

The HCAL Barrel (HB) and Endcap (HE) detectors have a similar design, covering the pseudorapidity region of  $|\eta| < 3$ . Both are sampling calorimeters constructed with alternating layers of brass as an absorber and plastic scintillator tiles for energy detection. Scintillation light is collected by wavelength-shifting fibers and read out by Silicon Photomultipliers (SiPMs). These are complemented by the HCAL Forward (HF), which extends the pseudorapidity coverage to  $3 < |\eta| < 5$ . The HF, located at  $|z| = 11$  m, is subject to a much higher particle flux and is designed as a Cherenkov sampling calorimeter. It consists of steel absorbers alternating with quartz fibers as the active material, which are read out by Photomultiplier Tubes (PMTs).

Test beam results with pions show that the combined energy resolution of the HCAL and ECAL, when corrected for non-linearity effects, is [32]:

$$\frac{\sigma_E}{E} = \frac{84.7\%}{\sqrt{E(\text{GeV})}} \oplus 7.4\%. \quad (2.11)$$






Figure 2.11: Layout of the CMS Hadron Calorimeter (HCAL), consisting of the barrel and endcap sections (HB and HE, respectively), the forward component (HF), and an outer detector (HO) located outside the solenoid.

When combining information from all subdetectors, the typical jet energy resolutions at central rapidities are 15-20% at 30 GeV, about 10% at 100 GeV, and 5% at 1 TeV[33].

### Muon system

The outermost part of the CMS detector is the muon-tracking system, which is dedicated to the detection of muons[34]. With typical energies ranging from a few GeV to hundreds of GeV, muons at the LHC are close to minimum ionizing particles, losing minimal energy as they traverse the CMS detector and thus easily escaping it. The muon system is located outside the solenoid magnet volume, using the flux return yokes as mechanical support, which also serve as a barrier against hadrons. The presence of a  $\sim 2$  T return magnetic field allows for the measurement of the charge and momentum of muons, providing information complementary to that of the tracker. This capability is particularly useful for triggering purposes.

A schematic representation of the muon system is shown in Figure 2.12.

The muon system consists of different sets of gas detectors, each specifically designed according to their position within the CMS detector. In the barrel region ( $|\eta| < 1.2$ ), where muon rates are lower, Drift Tubes (DTs) with rectangular cells are arranged in four concentric stations within the flux return plates. In the

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT

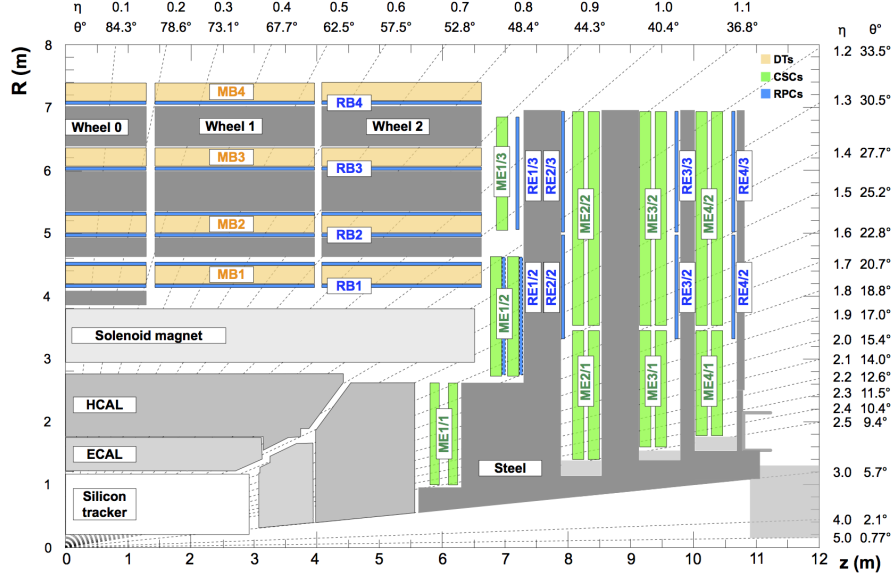


Figure 2.12: Schematic longitudinal view of the layout of a quarter of the muon detection system. It comprises the Drift Tubes (DT, yellow), the Cathode Strip Chambers (CSC, green), and Resistive Plate Chambers (RPCs, blue).

endcap regions ( $0.9 < |\eta| < 2.4$ ), which face higher muon rates and background, Cathode Strip Chambers (CSCs) are used due to their fast response, fine granularity, and radiation resistance. Four stations of CSCs are positioned in each endcap.

To further improve triggering, Resistive Plate Chambers (RPCs) are distributed across both the barrel and endcap regions ( $|\eta| < 1.8$ ). Although RPCs have slightly lower spatial resolution, their fast response and excellent time resolution are crucial for unambiguous bunch crossing assignment, enhancing both muon triggering and offline reconstruction.

The efficiency for muon identification and reconstruction is greater than 96%. Matching muons to tracks measured in the silicon tracker results in a relative transverse momentum resolution, for muons with  $p_T$  up to 100 GeV, of 1% in the barrel and 3% in the endcaps. The  $p_T$  resolution in the barrel remains better than 7% for muons with  $p_T$  up to 1 TeV[35].

### Drift Tubes

The drift tube muon detector will be described in more detail, introducing the components and terminology that will be used in chapter 3.

The DT system[36] is organized into five wheels along the  $z$ -direction, each

spanning approximately 2.5 m in width. As shown in Figure 2.12, the central wheel is Wheel 0, while Wheels  $\pm 2$  indicate the outermost ones on the positive and negative sides, respectively.

Each wheel is divided into 12 azimuthal sectors, each covering approximately  $30^\circ$ . Each sector hosts four DT chambers, interspaced by the iron yoke. In this arrangement, MB1 represents the chamber closest to the interaction point, while MB4 is the outermost chamber. This architecture is depicted in Figure 2.13.




Figure 2.13: Transverse view of the CMS detector. The 12 DT sectors, each containing four chambers (MB), are visible outside the magnet.

Each chamber, illustrated in Figure 3.11, consists of three super-layers (SL), each made up of four layers. The SL1 and SL3 super-layers measure the  $r - \varphi$  coordinate in the bending plane of CMS, while the SL2 superlayer measures the  $z$  coordinate, along the direction parallel to the beam (perpendicular to the plane of the figure). SL1 is the superlayer closest to the interaction point, separated from SL2 and SL3 by a spacer. In MB4, however, only two super-layers (SL1 and SL3) are used.

Each layer consists of contiguous drift tube cells, schematized in Figure 2.15. The number of tubes in the layers of each station varies, increasing for SL1/3 as one moves from MB1 to MB4, as they need to cover a larger surface in the  $r - \varphi$  plane.

Each tube is approximately  $42 \times 13 \text{ mm}^2$  in cross-section and 2.5 m in length. Its volume is filled with a gas mixture such that, when a charged particle crosses

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT




Figure 2.14: Schematic view of a DT chamber.

it, the gas is ionized. Due to the electric field, the electrons drift toward an anode wire installed in the center of the cell, where they are collected and amplified, producing a measurable charge pulse.




Figure 2.15: Layout of a drift cell showing the electric field lines in the gas volume.

Each DT chamber is equipped with a “minicrate”, mounted directly on the chamber, which contains the local trigger and readout electronics. These electronics follow two distinct paths[37]. The readout electronics are designed to perform precise time measurements of the chamber signals with a resolution of around 0.7 ns. These measurements are transmitted to the backend only when a global trigger signal is received, at a rate of around 100 kHz. On the other hand, the local trigger electronics process signals with a coarser time resolution of 25 ns but at the bunch crossing rate of 40 MHz. The local trigger is responsible for bunch crossing assignment and for determining the muon track parameters within a single chamber. This information is then transmitted to subsequent trigger stages that operate off-detector.

### 2.2.3 Trigger and Data Acquisition

Considering that the aggregated size of data from all subsystems amounts to approximately 1 MB per event, and assuming a collision rate of 30 MHz<sup>1</sup>, this would lead to a data throughput of  $\sim 30$  TB/s. Even if the technical difficulties of reading and managing this amount of data online were overcome, storing such a massive volume would be unfeasible. Furthermore, most of this data comes from low-energy proton-proton interactions and other processes outside the CMS physics program's interests, which primarily targets processes occurring at much lower rates. For example, as shown in Figure 1.1, at  $\sqrt{s} = 13$  TeV, the inclusive proton-proton cross-section is approximately  $\sigma(pp) \sim 10^{11}$  pb, whereas W boson production is around  $\sigma(W) \sim 10^5$  pb, a difference of six orders of magnitude. Similarly, for lower cross-section processes like Higgs boson production via gluon-gluon fusion, the gap is about eight orders of magnitude.

Thus, to implement the CMS physics program, an efficient online selection system is needed. This is the role of the trigger system, which reduces the event rate from the input of 40 MHz to approximately 1 kHz by selecting only the most interesting collision events for offline analysis[38].

In CMS, this is achieved using a two-tiered architecture consisting of the Level-1 Trigger (L1T) and the High-Level Trigger (HLT), as illustrated in Figure 2.16.

The data produced by the detectors for every bunch crossing, at 40 MHz, is digitized and stored in front-end pipelines. A subset of this data is processed locally by subsystems to generate Trigger Primitives (TPs), which provide a coarse-granularity view of the event for that bunch crossing. These TPs are processed by the L1T, which is implemented on custom electronics. Within a few  $\mu$ s after the bunch crossing, the L1T must decide whether to keep or discard the event data. The available latency for the decision depends on the depth of the front-end buffers. The L1T can select events up to a rate of 100 kHz. If an event is accepted, the full event data is sent to the event builder network, which collects and combine raw data from all subdetectors and delivers it to the HLT. The HLT then performs a more refined analysis, using the full detector resolution and

---

<sup>1</sup>The bunch spacing is 25 ns, but not all bunches are filled. In a typical filling scheme, resulting in a “full-machine”, the rate of filled bunch crossings at the interaction point is around 30 MHz.

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT




Figure 2.16: Diagram of the CMS experiment Trigger and Data Acquisition system[39].

running on commercial servers equipped with GPUs[40], reducing the selected event rate further down to about 1 kHz. The time needed to fully process on event by an HLT process is approximately  $\sim 500$  ms. Accepted events are sent to the CERN computing center where they are finally stored.

The focus of this section will be on the L1T, which plays a central role in this thesis. Specifically, we will examine the Phase-1 upgrade of the trigger system[41], deployed since Run 2. It consists of about 100 custom boards, processing approximately 5 Tb/s of data. The L1T must complete processing within  $3.8 \mu\text{s}$ , while accepting new inputs every 25 ns.

A schematic representation of the L1T architecture is visible in Figure 2.17. The system is split in two separate sections, one processing inputs from the muon systems and one from the calorimeters.

### Muon Trigger

The Muon trigger[42] processes local TPs from the CMS muon systems to combine them into a single muon track. It is divided into three regions, each covering a specific  $\eta$  range. The Barrel Muon Track Finder (BMTF) handles  $|\eta| < 0.8$  in the barrel, starting from the DTs and RPCs trigger primitives. The Overlap Muon Track Finder (OMTF) processes TPs from DTs, RPCs, and CSCs in the transition region between barrel and endcap,  $0.8 < |\eta| < 1.24$ . Finally, the Endcap Muon Track Finder (EMTF) processes TPs from the endcap RPCs and

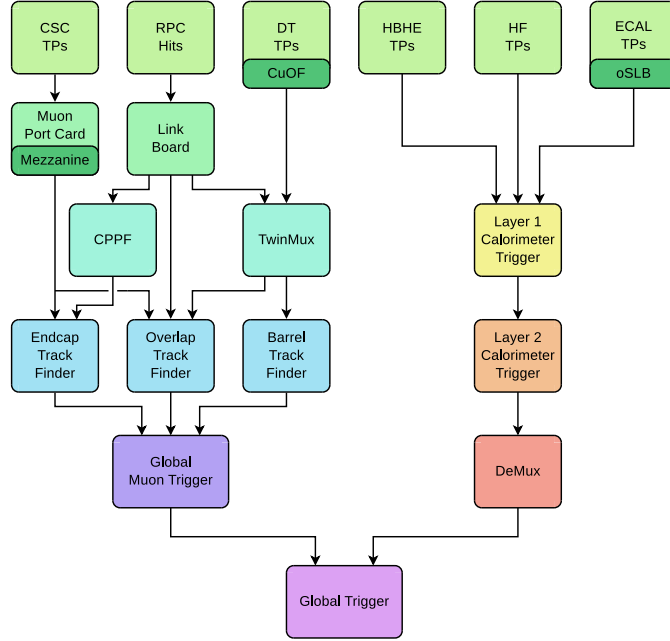


Figure 2.17: Schematic representation of the Phase-1 upgrade Level-1 trigger system architecture in Run-2 and Run-3.

CSCs, covering  $1.24 < |\eta| < 2.4$ . Candidates from these systems are sent to the Global Muon Trigger (GMT), which removes duplicates, sorts candidates based on  $p_T$  and track quality, and sends the best candidates to the Global Trigger.

### Calorimeter Trigger

The calorimeter trigger[43] receives inputs from the ECAL, HCAL, and HF calorimeters and uses them to reconstruct  $e/\gamma$ ,  $\tau$ , and jet candidates, as well as calculate global energy sums, such as  $p_T^{\text{miss}}$ . In the current L1 trigger system, it is not possible to distinguish electron and photons, as no tracker information is available. The input data is organized into Trigger Towers (TTs), each covering a region of  $(\eta, \phi) \approx (0.087 \times 0.087)$  in the barrel, with more complex geometry in the endcaps. Each TT encodes the energy deposited in a group of ECAL crystals and the HCAL geometrically behind them. The system is structured into two layers, called *Layer-1* and *Layer-2*, and utilizes a time-multiplexed architecture[44]. In this architecture, with a period of  $N$ , a multiplexing layer sends data from the same event to a single board every  $N$  events. The advantage of this approach is that each board processes the full event. In the calorimeter trigger, a factor of  $N = 9$  is used. Layer-1, which receives the local TPs from the calorimeters, serves as the multiplexing stage, while in Layer-2, algorithms for identifying

## 2.2. THE COMPACT MUON SOLENOID EXPERIMENT




Figure 2.18: Fractions of the 100 kHz rate allocation for single- and multi-object triggers and cross triggers in the baseline Run 3 menu[46].

calorimeter objects and energy sums are implemented. The outputs are then sent to a demultiplexing layer, *DeMux*, which sorts the objects based on their energy before transmitting them to the global trigger.

### Global Trigger

The Global Trigger (GT)[45] collects the best candidates from the muon and calorimeter triggers to decide whether an event should be accepted. The decision is based on algorithms defined by the *Trigger Menu*, which includes criteria ranging from  $p_T$  thresholds on single objects to multi-object relations like invariant mass or  $\Delta R$ . Up to 512 distinct algorithms can be evaluated in parallel. After evaluating all the algorithms, the GT performs a logical OR on their outputs and, if the event is accepted, issues the *Level-1 Accept* (L1A) signal, triggering the readout of all subsystems and sending the raw data to the event builder and HLT.

The rate allocation for single- and multi-object triggers in the baseline Run 3 menu is shown in Figure 2.18 [46]. The system has been extended to include new algorithms for long-lived particles [47], anomaly detection [48], along with novel data-taking strategies such as data parking and HLT scouting [49], which are described in chapter 4. These developments significantly broaden the physics reach of the CMS experiment.

Examples of trigger algorithms include a single-muon trigger, which requires at least one muon with  $p_T > 22 \text{ GeV}$  and results in a rate of approximately 10 kHz. The transverse momentum requirement can be relaxed by considering, for instance, a pair of oppositely charged muons with  $p_T > 3 \text{ GeV}$ , thereby lowering the trigger rate. Similarly, for isolated calorimeter objects such as  $e/\gamma$  candidates with  $p_T > 30 \text{ GeV}$ , the rate is around 20 kHz. Single-jet triggers also exist, with a typical threshold starting at  $p_T > 150 \text{ GeV}$ . Multi-jet events are



often selected based on the energy sum  $H_T$ , defined as the sum of the transverse momenta of jets with  $p_T > 30 \text{ GeV}$  and  $|\eta| < 2.5$ . A common requirement is  $H_T > 280 \text{ GeV}$ . These values are continuously adjusted and optimized during data taking to maintain the trigger rate within the system's capacity, though this could result in thresholds that are too stringent for certain analyses.

## 2.3 Phase-2 upgrade

The ambitious goal of HL-LHC in terms of the foreseen instantaneous luminosity delivered to the experiments poses several challenges for the various detectors on the LHC ring. The detectors will be required to withstand much higher level of radiation, and the unprecedented amount of simultaneous collision will demand a much better noise discrimination. The detectors will need to have faster readout and longer pipelines, and precise timing will become an essential feature to ensure a correct and precise reconstruction of collision events. To face such challenges, CMS will undergo deep changes[50]: some of its sub-detectors will be dismantled and replaced, new layers will be added, and every subsystem will undergo major upgrades. Trigger and data acquisition logics will be changed and improved, leveraging all sorts of cutting-edge technologies both on the hardware and software sides. A major upgrade is the introduction of tracking information into the trigger logics.

### 2.3.1 Sub-detectors

#### Tracker upgrade

The Phase-2 tracker upgrade will involve a complete replacement of the existing strip and pixel detectors with a more advanced system[51]. This new system will follow the phase-1 implementation, consisting of an Inner Tracker, equipped with small-size pixel sensors, and an Outer Tracker, featuring strip and macro-pixel sensors, extending coverage up to  $|\eta| = 4.0$ .

The Inner Tracker will have a highly granular structure, incorporating thinner sensors, smaller pixels, and faster, more radiation-tolerant electronics compared to the existing Phase-1 detector. Specifically, the pixel size will be reduced by a factor of 6, down to  $100 \times 25 \mu\text{m}^2$ , to maintain low mean occupancy and improve spatial resolution.

### 2.3. PHASE-2 UPGRADE

In the Outer Tracker, the introduction of stacked strip modules, with the size of each strip increasing with the distance from the beam pipe, will help reduce hit multiplicity and allow for the reconstruction of track candidates (L1 tracks) at the L1 trigger level, an operation previously accessible only to the High-Level Trigger, covering the region up to  $|\eta| = 2.4$ . In this setup, stubs with a  $p_T$  larger than a specific threshold (e.g.,  $p_T > 2 \text{ GeV}$ ) will be transmitted to the trigger at the bunch crossing rate. This selection is performed by the so-called “ $p_T$  modules”, with the threshold chosen to fit within the system’s bandwidth limits.

#### **Precise timing measurements: MTD**

To perform extremely precise timing measurement, the CMS collaboration has developed a whole new layer of particle detection: the MIP Timing Detector (MTD)[52]. Such system will comprise a barrel timing layer (BTL), made of scintillating LYSO crystals coupled to Silicon Photomultipliers, and an endcap timing layer (ETL), featuring planar silicon devices with internal gain, the Ultra Fast Silicon Detectors, coupled to low gain avalanche diodes. MTD will be placed between tracker and calorimeters, covering up to  $|\eta| < 3.0$ . It aims at providing time resolution of 30 ps for minimum ionising particles, which is expected to degrade to up to 60 ps by the end of Phase-2.

The addition of precision timing information for tracks will provide an important tool for mitigating the effects of pile-up, which is expected to reach values of up to 200 at the HL-LHC. This is based on the fact that individual collisions within a bunch crossing do not occur simultaneously but are spread over a time interval of 180-200 ps. The MTD is designed to exploit this time spread to more accurately associate tracks from charged particles with their corresponding vertices.

#### **Calorimetry upgrades**

To address the increased radiation levels and meet the stricter demands for trigger latencies and rates, the barrel calorimeters are set for a major electronics upgrade. Both the on-detector and off-detector electronics of the ECAL and HCAL will be completely replaced[53]. This refurbishment will also enable finer granularity, improved noise suppression, and in the case of the ECAL, timing information with a resolution of up to 30 ps for electrons and photons

above 50 GeV.

The existing endcap calorimeters will be dismantled and replaced with the new High-Granularity Calorimeter (HGCAL)[54], covering the pseudorapidity region  $1.5 < |\eta| < 3$ . The proposed design uses silicon sensors as active material in the front sections and plastic scintillator tiles, with the scintillation light read out by SiPMs, towards the rear. In addition to providing very precise 3D positions, HGCAL also offers precise timing information for the energy deposits, allowing discrimination of clusters from pile-up within a single bunch crossing. The five-dimensional information (energy,  $x$ ,  $y$ ,  $z$ , and time) is ideally suited for particle-flow reconstruction. A schematic representation of the detector layout is shown in Figure 2.19.




Figure 2.19: Schematic view of the longitudinal cross section of the HGCAL detector, divided into the electromagnetic compartment (CE-E) followed by the hadronic compartment (CE-H). The green area employs silicon detectors, while the blue area makes use of plastic scintillators. The absorber material is represented in gray.

The detector is divided into two sections: the electromagnetic compartment (CE-E), which consists of layers of silicon active material interleaved with CuW, Cu, and Pb absorbers, and the hadronic compartment (CE-H), which contains layers of silicon or scintillator active material alternated with steel absorbers. HGCAL contributes to the L1 trigger decision by producing highly granular trigger primitives (TPs), which, thanks to the longitudinal and transverse structure of the energy deposits, can be used to efficiently discriminate between electromagnetic showers, taus, jets, and pile-up already at the L1 trigger level.

### 2.3. PHASE-2 UPGRADE

#### Upgrade of the muon systems

To cope with the challenging conditions of increasing luminosity expected at HL-LHC, several upgrades of the muon detectors and trigger system are planned[55] in order to maintain the current performance in the harsh conditions. A scheme of the upgraded muon system is represented in Figure 2.20, highlighting the position of both the present detectors and the ones planned as part of the Phase-2 upgrade.

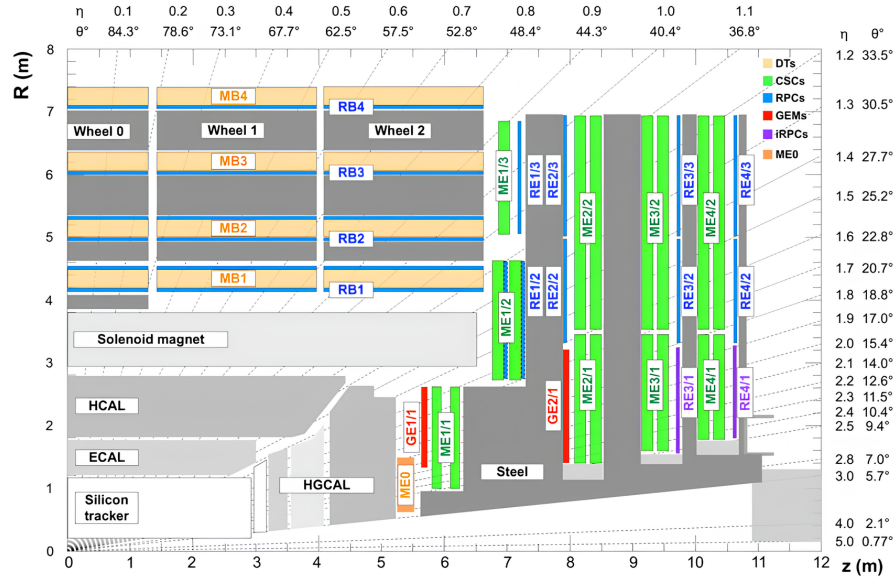


Figure 2.20: Quadrant of the CMS muon system, highlighting the position of both the present detectors and the ones planned as part of the Phase-2 upgrade.

The Phase-2 upgrade for the Drift Tube system primarily focuses on modernizing the on-detector electronics. While the DT chambers themselves will remain unchanged, the electronics needs to be changed in order to handle the increased L1T rate and the higher radiation levels expected under the HL-LHC conditions. In this upgraded setup, the new electronics will transmit time digitization (TDC) data directly to a new backend electronics system located in the service cavern for both trigger and readout.

Similarly, the Resistive Plate Chambers (RPC) off-detector electronics will be replaced. This upgrade will allow for increasing the readout frequency from 40 MHz to 640 MHz, allowing each RPC hit sent to the muon track finders to include more detailed time information. Additionally, to improve the Muon System's performance in the forward region, new RPC chambers (iRPCs) will be added to stations 3 and 4 (RE3/1 and RE4/1), extending the RPC pseudorapidity

coverage up to  $|\eta| < 2.4$ . These chambers are optimized to handle the extremely high rates expected in these areas.

In the Cathode Strip Chambers (CSCs), the upgrade will involve replacing the on-chamber cathode boards in the inner rings within  $1.6 < |\eta| < 2.4$ . The new boards will be able to manage the increased trigger and data output rates. Additionally, most on-chamber anode boards will receive new FPGA mezzanine boards to adjust to the higher L1T latency. Boards responsible for trigger and data processing will also be updated to handle the higher rates.

As part of the CMS upgrade strategy for the muon system towards the High-Luminosity LHC, a new station of triple-Gas Electron Multiplier (GEM) detectors, known as GE1/1, has been installed in both endcaps of the CMS detector. This station, installed in 2020, consists of two layers of triple-GEM detectors and has been operational during Run 3, serving as a demonstrator for the future upgrade. The GE1/1 station enhances the muon system by providing two additional muon hit measurements, which improve muon tracking and triggering performance while helping to control the muon trigger rate. Further GEM chambers (GE2/1 and ME0) will be introduced in the forward region as part of the Phase-2 upgrade, widening the system coverage to  $1.6 < |\eta| < 2.8$ . The overall GEM system is designed to provide precise measurements of the muon's bending angle in the first two stations, crucial for managing the increased muon trigger rate expected under HL-LHC conditions.

### 2.3.2 L1 Trigger

If the Phase-1 L1 Trigger were used after the Phase-2 upgrade with the same trigger menu and algorithms, the event acceptance rate at an instantaneous luminosity of  $\mathcal{L}_{\text{inst}} = 5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$  would be 1.5 MHz, increasing to 4 MHz at  $\mathcal{L}_{\text{inst}} = 7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [50]. Operating under these conditions would be impossible for the CMS detector.

The Phase-2 upgrade of the L1 trigger[56] is designed to achieve, and potentially extend, the same physics acceptance as the original CMS detector, despite the harsher pile-up environment of the HL-LHC. This is made possible by incorporating track reconstruction, which, combined with high-granularity calorimetry and muon detector information, allows for physics object reconstruction with significantly improved resolution. The enhanced processing capabilities are provided by modern FPGAs and fast optical links, along with

### 2.3. PHASE-2 UPGRADE

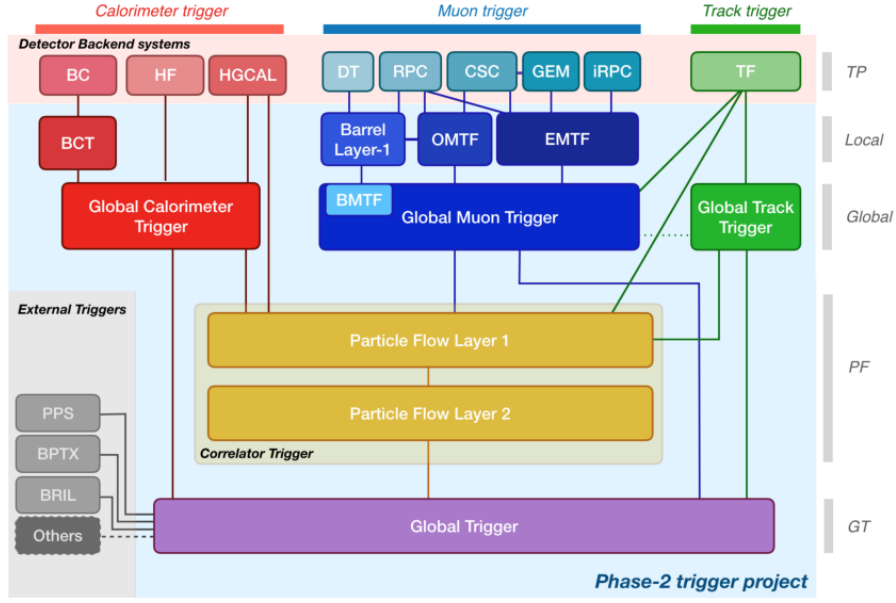


Figure 2.21: Functional diagram of the CMS L1 Phase-2 upgraded trigger design.

an increased maximum fixed latency from  $4\ \mu\text{s}$  to  $12.5\ \mu\text{s}$ . Moreover, the total maximum L1 trigger rate is increased from 100 kHz to 750 kHz.

With these improvements, the Phase-2 L1T can maintain the same physics acceptance as the Phase-1 trigger, achieving trigger rates of 260 kHz and 500 kHz at  $\mathcal{L}_{\text{inst}} = 5 \times 10^{34}\ \text{cm}^{-2}\text{s}^{-1}$  and  $\mathcal{L}_{\text{inst}} = 7 \times 10^{34}\ \text{cm}^{-2}\text{s}^{-1}$ , respectively. This is due to the higher quality of the online reconstructed objects, with algorithms that closely resemble those used in offline analysis. Alongside standalone objects such as muons and jets, which are reconstructed using only the muon and calorimeter systems, particle flow (PF) objects[57] will be reconstructed by combining information from all subsystems. Moreover, PF objects are resilient to the challenging pile-up conditions of the HL-LHC, thanks to the pile-up per particle identification (PUPPI) algorithm[58].

The functional diagram of the CMS L1 Phase-2 upgraded trigger design is shown in Figure 2.21. Similar to the Phase-1 trigger, three branches receive and process local trigger primitives from the subsystems. From the left of the image, the calorimeter trigger receives inputs from the barrel calorimeter (BC), the high-granularity calorimeter (HGCAL), and the hadron forward calorimeter (HF). The barrel calorimeter trigger (BCT) combines TPs from the barrel calorimeters to form clusters, which are sent along with HGCAL clusters to the global calorimeter trigger (GCT).

The muon trigger maintains the organizational structure of the Phase-1 sys-

tem, divided into barrel, overlap, and endcap regions. Local standalone muon candidates are sent to the global muon trigger (GMT), which performs duplicate cleaning of the candidates. Additionally, track stubs from the track trigger are matched to the muon tracks in the muon system.

The muon trigger keeps the organization of the phase-1 system, being divided in barrel, overlap and endcap. Local standalone muons are sent to the global muon trigger, which perform duplicate cleaning of the candidates. Additionally, it receives track stubs from the track trigger which are matched to the muon tracks in the muon system.

The track finder (TF), operating in the tracker backend, sends track information to the Global Track Trigger (GTT), which identifies the primary vertex of the event and computes hadronic jets and energy sums using only tracker information.

Finally, the three global systems transmit their standalone objects to the correlator trigger (CT). The CT, composed of two layers, reconstructs objects using the PF algorithm and applies the PUPPI algorithm to subtract pile-up contributions.

The global trigger then receives objects from the GCT, GMT, GTT, and CT to run the algorithms defined in the trigger menu, ultimately making the final event selection decision. The trigger algorithms are designed to use tracking information to achieve near-offline resolution. The performance of the reconstructed trigger objects is close to that of offline physics objects, providing optimized response and resilience to high pile-up conditions. These advanced capabilities, which were previously confined to offline reconstruction, make the objects reconstructed by the trigger more suitable for analysis, rather than serving merely as trigger objects. This is the idea behind L1 Trigger scouting, which will be introduced in the next chapters.





# 3

## 40 MHz Trigger-less Readout of Detector Front-ends

The Phase-2 upgrade of the CMS detector, as discussed in chapter 2, involves significant upgrades to the front-end and trigger electronics of many subsystems, often relocating a large portion of the logic off-detector. One of the key components enabling this transition is the introduction of high-speed data links, which allow the entire set of raw data to be streamed to backend electronics located in the service cavern. Custom FPGA-based boards are responsible for collecting and processing the data using advanced algorithms, providing performance levels that were previously only accessible to the High-Level Trigger and offline processing.

This setup opens up the possibility of intercepting and analyzing the data stream before it undergoes the filtering imposed by the Level-1 trigger, thereby granting access to the full dataset at the bunch crossing rate.

This approach can be pursued in two ways, both of which will be covered in this thesis. The first approach leverages the high quality of the objects reconstructed by the Level-1 Trigger. For example, it involves collecting the complete set of Level-1 PUPPI candidates produced by the correlator trigger, at the full bunch crossing rate, independently of the trigger decision. This method will be explored in chapter 4.

The second approach involves intercepting the data flow before it reaches the trigger, effectively processing data directly from the detector's front-end electronics. This method is equivalent to a triggerless readout, and it will be

discussed in this chapter.

The focus will be on the Drift Tubes muon detector, which offers an ideal test bed for developing and validating new ideas that can potentially be applied across all CMS subsystems. The new FE boards, called OBDT, designed for the Phase-2 upgrade, will stream data from the experimental to the service cavern, with duplicates sent to a scouting system using spare links. Prototypes of these boards have already been installed in four DT chambers during Run-3, making them a perfect candidate for initial phase of these tests.

Multiple scenarios could benefit from trigger-less readout:

- **Extended Physics Potential:** Continuous data acquisition at 40 MHz enables the capture of rare or unexpected events which could evade the trigger chain. For instance, raw data can complement the CMS L1-trigger scouting system, providing a richer dataset for analysis in contexts such as showers produced by long-lived particles or out-of-time products.
- **Enhanced Data Quality Monitoring (DQM):** The ability to continuously monitor detector data at 40 MHz offers an unparalleled tool for ensuring data quality. By collecting raw hits at this rate, the system can serve as a powerful DQM tool, enabling real-time analysis and immediate response to any detector anomalies.
- **Detector Background Studies:** The high-rate readout, including where no trigger was produced, allows for detailed studies of detector backgrounds. This is essential for optimizing detector performance and understanding environmental impacts, especially in the high-radiation environment of the HL-LHC, where background rates are expected to increase.

This chapter will describe the progression from the initial testbed setup at the INFN Legnaro National Laboratories (LNL) to its implementation in CMS as part of the DT Phase-2 demonstrator slice test[59, 60].

The implementation of the triggerless readout system on a large scale requires a new DAQ architecture, more performant and efficient in terms of CPU utilization dedicated to moving the data. For these reasons, a new data acquisition scheme based on Remote Direct Memory Access (RDMA) from FPGA, known as the FEROCES project[61] (Front End Rdma Over Converged Ethernet), has been introduced. The FEROCES project aims to improve the performance of data acquisition architecture for modern experiments requiring large bandwidth. First, an initial implementation based on existing libraries will be presented, followed by a focus on a custom implementation of the protocol, targeting the implementation on both backend and front-end FPGA.

### 3.1 40MHz readout of the CMS Drift Tubes detector

In this section, we describe the detector and electronics chain that enables triggerless data collection at 40 MHz. This setup lays the groundwork for the implementation of the online processing stage, which will be discussed in the following section of this chapter.

#### 3.1.1 The miniDT detector

A set of gaseous detectors, located in the LNL, is used in the testbed experimental setup. These detectors are based on the design of the Drift Tubes of the CMS but with a reduced area. They were employed for tracking and particle identification of muons during the 2019 test-beam campaigns of the LEMMA project [62].

Each detector, referred to as a *miniDT* chamber<sup>1</sup>, is composed of 4 layers of 16 tubes (*cells*), each with a transverse size of  $42 \times 13 \text{ mm}^2$ . The cells in adjacent layers are mounted in a staggered configuration, with a relative shift of half the cell width, to provide tracking capabilities (Figure 3.1 (right)[36]). As a charged particle passes through a cell, it induces ionization in the 85%-15% Ar – CO<sub>2</sub> gas mixture, causing the drift of primary electrons towards the anodic wire, which is kept at a positive potential.

Close to the wire, secondary ionization occurs with a typical gain of  $10^5$ , and the electron avalanche is collected by the sensing wire. The electric field inside the cells is shaped using a multi-electrode design to ensure a constant drift velocity of the electron cloud throughout the entire cell ( $v_{\text{drift}} \approx 54 \mu\text{m/ns}$ ).

Up to four time-coherent signals (*hits*) can be produced by a muon crossing a chamber. The time coherence among the hits is determined by the maximum time allowed for the signal in each cell to form and be collected by the wire, as well as the angle at which the muon crosses the tubes. The charge collected on each wire is amplified, shaped, and discriminated against a fixed threshold by an ASIC [63] chip housed within the gas volume of the chamber. Hits passing the discriminator threshold are finally converted to LVDS-compatible signals and sent to the time-digitization electronics.

---

<sup>1</sup>A miniDT chamber is equivalent to a single super-layer of a CMS DT chamber, with a reduced number of drift tubes

### 3.1. 40MHZ READOUT OF THE CMS DRIFT TUBES DETECTOR




Figure 3.1: Left: Array of 4 miniDT chambers, where the bottom 3 are oriented in the same direction while the top-most one provides tracking in the orthogonal plane. Right: Schematic view of the layers' staggered structure (top) and a DT cell (bottom). The anodic wire is visible in the center of the cell, along with the strips providing the uniform electric field inside the gas volume.

For these tests, an array of 4 miniDTs stacked on top of each other was used, as shown in Figure 3.1 (left), each with a roughly square shape of  $70 \times 70 \text{ cm}^2$  and made of layers of 16 cells, totaling 64 cells per chamber. The top-most chamber was oriented orthogonally to the other three, providing additional tracking information by offering access to the position along the wire that would be inaccessible if all the chambers had the same orientation.

#### 3.1.2 Trigger-less readout chain

The front-end readout electronics are responsible of performing the time-to-digital conversion of the LVDS signals, tagging the time of arrival of the hits relative to a reference clock. The components used in the testbed are compliant with the electronics standards defined for the HL-LHC DT Phase-2 Upgrade, specifically the so-called *OB DT* (On Board electronics for Drift Tubes) [64]. These components employ the same communication protocol on the optical link and adopt the same strategy for the time digitization of the front-end signals, performed on FPGA. In the testbed two Xilinx evaluation boards (VC707), each hosting a Virtex-7 FPGA, are used to receive the discriminated LVDS signals from the FE and implement the TDC for each firing channel. An external oscillator provides a 120 MHz clock, distributed across both VC707 boards. This frequency is also upsampled and downsampled to generate all the synchronous clocks required by the TDC cores. Even though the clock is distributed in phase to the two boards, a synchronization signal has been implemented to synchro-

nize the downscaled clock. This is necessary because, when generating a lower frequency from the same clock, a phase indeterminacy is present. To emulate the CERN Timing, Trigger and Control (TTC) Systems [65], which distribute timing information for all LHC experiments, a 40 MHz clock is generated, compatible with the LHC BX frequency. The TDC provides a time measurement of the rising edge of the input signal relative to the rising edge of the 40 MHz clock. A coarse counter then assigns the BX to the TDC measurement, giving the hit timestamp a granularity of  $1/30^{\text{th}}$  of a BX. For each VC707, a total of 138 TDCs are implemented using serial-to-parallel converters<sup>2</sup> running at 1.2 Gsps. A single VC707 can perform the digitization of all the channels from two miniDTs. Additional TDCs are available for digitizing other external signals. For example, a pair of scintillator tiles equipped with photomultipliers can provide an unbiased external timing reference. The coincidence NIM signal from the two photomultipliers is reshaped into an LVDS signal and injected into one VC707 for digitization.

No filter or selection is applied to the signals collected by the front-end boards, and the entire data stream of TDC hits is serialized using the GBTx-FPGA protocol [66] to SFP+ transceivers for transmission over optical links to a back-end board, at 4.8 Gbps.

The data stream is continuous and asynchronous, as all signals received by the front-end boards are digitized and transmitted to the optical links as soon as they are available. The two optical fibers, each carrying TDC hits from two miniDTs, are collected into a back-end (BE) board, a Xilinx KCU1500 evaluation board equipped with a Kintex UltraScale FPGA (XCKU115) and hosted in a Dell PowerEdge R730 server.

Data from each link is deserialized using the GBTx-FPGA protocol. The clock is recovered directly from the received data by the CDR (Clock Data Recovery) of the fast transceivers on the FPGA. The stream from each link is then processed independently by an algorithm implemented on the FPGA, which performs the reconstruction of the trajectory compatible with the passage of muons through each chamber.

The reconstruction algorithm [60] combines logic from Neural Networks and analytical relations to identify signals compatible with local muon tracks

---

<sup>2</sup>ISERDESE2 is a serial-to-parallel converter available in AMD 7-series FPGAs. For more details, see [https://docs.amd.com/v/u/en-US/ug471\\_7Series\\_SelectIO](https://docs.amd.com/v/u/en-US/ug471_7Series_SelectIO).

### 3.1. 40MHZ READOUT OF THE CMS DRIFT TUBES DETECTOR

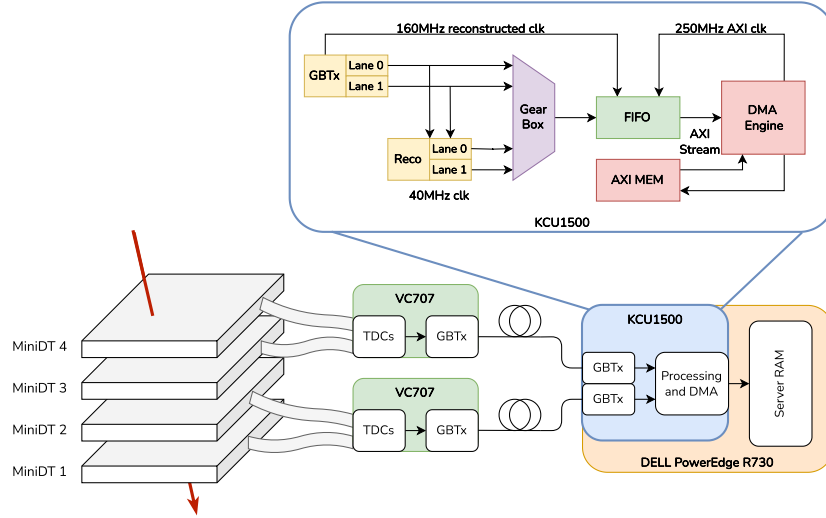


Figure 3.2: Simplified diagram representing the trigger-less readout chain installed at the INFN Legnaro laboratory. Left: a cosmic muon (in red) crossing the four detectors (grey). Center: two VC707 front-end boards (green) implementing the TDC and GBTx serialization. Right: the readout server (yellow), hosting the back-end board KCU1500 (blue), where online segment reconstruction and DMA readout is performed. On the top, a schematic representation of the firmware implemented in the back-end board is visible. From the left, the stream of hits and products of the reconstruction algorithm are merged into a single stream and inserted into a FIFO, where they are streamed to a DMA engine. A set of registers is used for control and monitoring of the board.

(referred to as *track segments*) and extract relevant information: crossing position, angle, and time of passage of the muon inside the chamber. All resulting parameters are merged with the original stream of hits, enriching the data stream with higher-level information. However, they are not directly used to filter or select data at this stage. Data from both lanes are subsequently gear-boxed into a single data stream and fed into a FIFO acting as a buffer.

The data is finally streamed to a Direct Memory Access (DMA) engine via the Advanced eXtensible Interface (AXI) stream protocol over a PCIe Gen 3 bus [67]. The Xilinx AXI-DMA engine [68] is used for transmitting the data stream to the server memory. An overview of the readout chain used for the testbed experimental setup, with a focus on the KCU1500 firmware, is presented in Figure 3.2.

The VC707 boards used in this setup are functionally equivalent to the preliminary version of the OBDT boards installed in the DT slice-test Phase-2 demon-

strator [69]. Tests were conducted by replacing the VC707 with a prototype of the OBDT once it became available, and from a readout perspective, no differences were observed. However, it is important to note that this is not the final version of the board that will be installed during the Phase-2 upgrade. While the functionalities remain the same, the board design will differ. For instance, the communication with the back-end, currently performed using GBTx, will be replaced by the new lpGBT [70] for both readout and slow-control [71].

The main bottleneck of the trigger-less readout system could potentially be the available bandwidth for the DMA transfer. However, when performing the DMA transfer over a PCIe Gen3 bus and using a sufficiently large transfer size, the throughput can easily reach 4 GB/s, with peak throughput reaching around 6 GB/s. Measurements were performed using both “dummy” transfers and data obtained by generating fake signals in the detector FE, confirming the expected performance provided by Xilinx<sup>3</sup>, which is far greater than what the detector could realistically produce. Given that 64 bits<sup>4</sup> are used to represent each hit, and assuming that a muon crossing the four mini drift tubes leaves four hits in each, the measured DMA transfer throughput would correspond to a rate of 47 million muons crossing the telescope every second. This is orders of magnitude larger than the expected rate from cosmic muons, which at sea level is approximately 1 muon per cm<sup>2</sup> per minute [72]. At the same time, this rate is also larger than the peak hit rate expected for the chamber with the highest rate (MB1 of Wheels±2) during the High-Luminosity LHC, where up to 50 kHz are anticipated per drift cell, leading to approximately 30 million hits per second for that chamber.

After the DMA transfer, the data stream can be either stored locally on the server disk or delivered to an online processing system, which is described in the next section.

The data collected using this system served as the foundation for subsequent research focused on applying machine learning techniques to perform anomaly detection in data streams [73, 74].

---

<sup>3</sup>[https://support.xilinx.com/s/article/68049?language=en\\_US](https://support.xilinx.com/s/article/68049?language=en_US)

<sup>4</sup>64 bits allows for the addition of redundant information for every hit, enabling each hit to be treated independently. In fact, 32 bits would be sufficient to carry all the necessary information, but this would require a more complex firmware design. For the purposes of this work, a design simplifying the implementation was adopted.



## 3.2 An online processing system for streaming data

In the previous section, the electronics acquisition chain was described, which generates an asynchronous data stream that is made available in server memory. Multiple instances of this chain may be present, for example, reading data from different detectors, producing data streams without any inherent synchronization. To effectively manage this data, a system must be designed that can receive all incoming data, combine outputs from multiple sources, and process them in a unified manner.

The system should be horizontally scalable, meaning it should be capable of accommodating additional data sources by scaling the number of components used for data reception and processing. Additionally, the design must prioritize ease of deployment and management, ensuring that no specialized resources are required for effective operation.

To meet these requirements, two alternative architectures for data transfer and online processing have been developed and implemented. The first approach is based on a widely adopted distributed event streaming system from outside the HEP domain, while the second draws inspiration from the standard architecture used in HEP experiments [75].

High-Energy Physics experiments are based on the concept of an “event”, which includes everything happening in a bunch crossing and is thus time-tagged. However, in this setup, a continuous stream of hits is produced without the concept of an event. Events will be identified in later stages of processing, after data from all sources have been aggregated. It is important to note that in this test setup, any reconstruction algorithm running in hardware is used solely to insert their products into the data streams, acting as “taggers” rather than filtering the data. For example, in the case of the Drift Tubes discussed earlier, the hits related to the same muon are produced across multiple bunch crossings due to the detector intrinsic slow signal collection time, which can take up to 16 BXs. Consequently, merging data based on the BX number alone is not feasible in this setup, as it would be in a system where the hits are tagged by a trigger.

For these reasons, data from all sources are merged based on coherent “time slices”. Following the example of the LHC TTC system, a common counter shared by all boards is the `Orbit Counter`, which resets its value when a synchronization signal is sent to all boards and increments by 1 every 3564 BXs. Data collected with the same `Orbit Counter` value thus correspond to the same



time fragment, lasting approximately  $90\ \mu\text{s}$ . In the case of mini-DTs collecting cosmic data, this is an arbitrary choice, as the distribution of cosmic rays is flat in time. However, when working with collisions delivered by the LHC, the abort gap clearly marks a region of the orbit where no collision will occur, thereby providing a natural division between consecutive time slices.

### 3.2.1 Distributed processing and messaging systems

The idea behind this implementation, illustrated in Figure 3.3, is to adopt horizontally scalable frameworks for the brokerage and processing of front-end TDC signals. Apache Kafka [76], a distributed event streaming platform based on a “pub-sub” messaging model, is used as the data brokerage system to serve data from the backend servers (the *publisher*) to the workers (*subscribed* to consume the data). Kafka represents a widely adopted, scalable, and fast distributed system optimized to build real-time streaming data pipelines, acting as a middle layer decoupling the entities producing data and those consuming it.

This model fits well with the described system: the sources producing data are the backend systems, which could potentially be from different subdetectors, where the messages are the raw data from that source for a given orbit number, and the consumers are the workers that need to process the data from a particular orbit. Adopting Kafka as a middle layer has the advantage of relieving the user from designing the data distribution/buffering model, as this is already handled by the brokers, which take care of receiving the data and placing them in a queue from which the consumers can read.

In the backend node, an application acting as a Kafka Producer continuously polls data from the FPGA, performs zero suppression, encapsulates the data in a message using the Kafka API, and publishes it to a remote *topic*, distributed in several partitions across multiple broker nodes. The partition where the data will land is selected based on the orbit number via a modulo operation. In this way, data from the same orbit will land in the same partition, acting as a “FIFO” structure for all the messages. This can be seen as a temporary buffer for the raw data fragments.

At this point, a set of workers can subscribe to the topic and retrieve data from either a single or multiple partitions if the number of workers is smaller than the number of partitions. In Kafka, multiple consumers cannot subscribe

### 3.2. AN ONLINE PROCESSING SYSTEM FOR STREAMING DATA

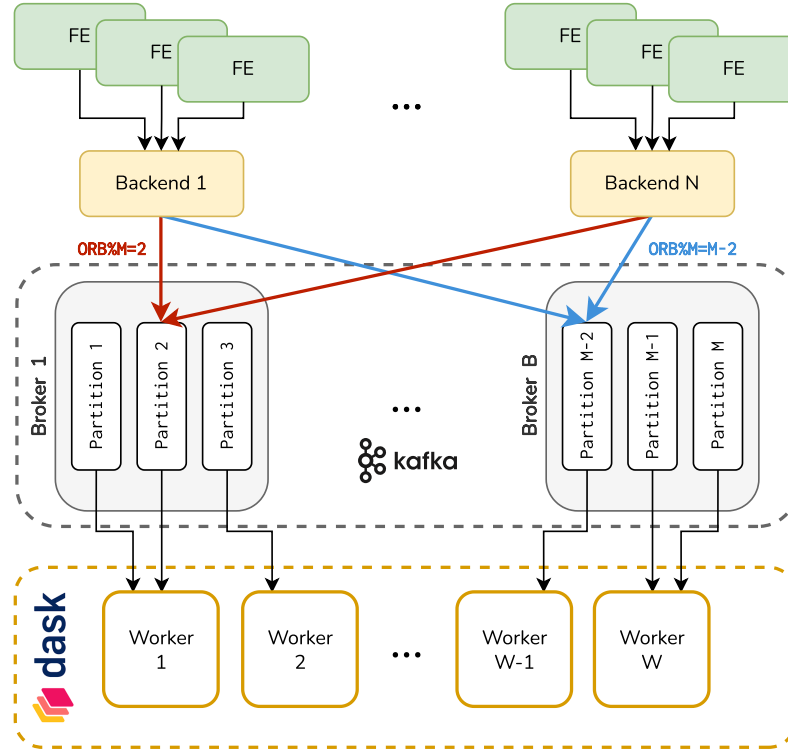


Figure 3.3: Schematic representation of the architecture based on a messaging system. At the top,  $N$  instances of the trigger-less acquisition chain, as presented in section 3.1, are depicted, each collecting data from a set of front-end (FE) boards and transmitting it to a backend server. The raw data from these backend servers are then sent to an Apache Kafka cluster, composed of  $B$  Kafka brokers, where the data is partitioned based on the orbit number. This ensure that data with the same orbit number is consistently directed to the same partition. The data from these partitions are subsequently consumed by a set of  $W$  workers, which can subscribe to one or more partitions. In this example, the workers are components of a Dask cluster, enabling distributed processing of the data streams.

to the same partition, which suits this use case as it prevents scattering orbit data across multiple workers.

In principle, any application implementing a Kafka consumer can be used to process the data. In this case, a set of workers managed by Dask [77] has been adopted. Dask is used as the distributed processing engine to schedule the data processing workload across a set of working nodes. This includes retrieving the data from the Kafka partitions, forming the orbit, identifying genuine hits produced by the passage of a muon, and writing the results for offline processing.

All system components were containerized and orchestrated using Kubernetes [78], enabling easy management and deployment on both bare-metal machines and cloud resources. A test aimed at measuring the system's processing capacity was conducted using resources provided by the CloudVeneto [79] computing infrastructure. This test demonstrated that a set of eight producers, each generating around 230 MB/s, corresponding to a conservative estimated throughput of a DT sector on the external wheels during Run 3, could inject an aggregate throughput of approximately 1.7 GB/s into a Kafka cluster composed of three brokers. The generated data from each backend server would correspond to nearly 30 MHz of hits, equating to almost 1.8 million muons crossing the test bed every second, assuming that no noise is present.

A Dask cluster made up of 10 workers was capable of consuming around 70 million hits per second. While more complex and time-demanding processing, such as the full reconstruction of a large number of particles in complex collision environments, would reduce the processing rate, the horizontal scalability provided by the Dask distributed processing framework can be exploited, allowing more workers to be added to share the processing load and maintain the necessary throughput.

This processing system is described in detail, including performance measurements, deployment, and control, in a dedicated paper [59].

### **3.2.2** Multistage architecture

As an alternative to the architecture described in the previous section, a multistage data aggregation architecture has been developed. This design is inspired by the event-building applications used in experiments such as CMS. A schematic representation of this architecture is shown in Figure 3.4.

### 3.2. AN ONLINE PROCESSING SYSTEM FOR STREAMING DATA

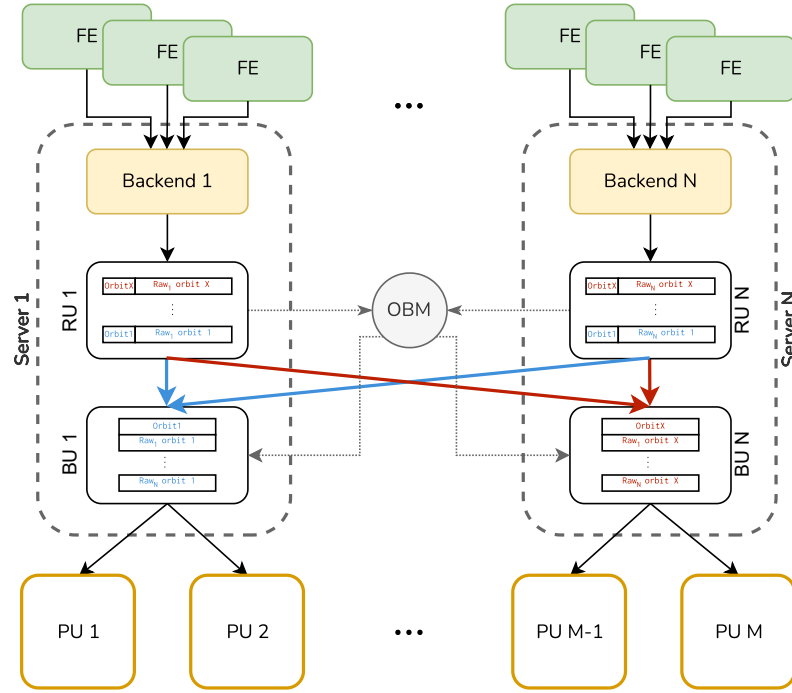


Figure 3.4: Schematic representation of the multistage data aggregation architecture. At the top,  $N$  instances of the trigger-less acquisition chain, as described in section 3.1, are depicted, each collecting data from a set of front-end (FE) in a backend board. The data from the backend boards are processed by a Readout Unit (RU) and stored in memory in a key-value store, where the orbit number serves as the key and the raw data as the value. An Orbit Manager (OBM) is responsible for tracking the availability of orbits across all RUs and assigning a set of orbit numbers to a Builder Unit (BU). The BU then retrieves all the fragments for those orbits from the RUs and assembles them into complete orbits. These complete orbits are subsequently delivered to a set of Processing Units (PU), which perform the final data processing. In this example, the RU and BU entities are located on the same server, but they could easily be distributed across independent machines.

Unlike the messaging-system-based architecture, this approach focuses on merging the data fragments collected by the backend (BE) boards before serving complete events to the processing units (PUs). In the servers where the backend boards are installed, a readout unit (RU) application continuously polls data from the FPGA. After preprocessing, including zero suppression and reformatting, the raw data is cached in an in-memory key-value store, where the key is the Orbit Number, and the value is the set of raw data associated with that orbit. Once this process is complete and the data has been inserted, the orbit number is sent to an entity called the Orbit Manager (OBM), which is responsible for the bookkeeping of orbit fragments across all the RUs.

Another entity, referred to as the Builder Unit (BU), is tasked with merging the data fragments corresponding to the same orbit, stored in all the RUs. When a BU is idle, it requests a set of orbits that have not yet been built from the OBM. After receiving the assignment, the BU becomes responsible for processing that specific set of orbits. To accomplish this, the BU contacts all the relevant RUs, retrieves the corresponding raw data from their caches, and performs the orbit-building process. Upon completion, the BUs make the aggregated sets of raw data available to a number of Processing Units (PUs) registered to it. These PUs are then responsible for analyzing the coherent set of hits and reconstructing all the features associated with the passage of muons through the detector and finally writing the results to a file.

Finally, a master entity orchestrates the control and monitoring all the components, including starting and stopping a run, configuring units for a particular run, and monitoring the performance of each component, such as average occupancy, the number of processed orbits, and overall throughput.

All the components have been implemented using standard C++ primitives, with the exception of the ZeroMQ [80] library, a high-performance asynchronous messaging library designed for use in distributed or concurrent applications. This library abstracts in-process, inter-process, and inter-node transport layers, simplifying development and increasing system flexibility. Indeed, units can be reconfigured to work with multiple transport-classes, such as TCP and IPC (Inter-process communication). It also provides standard communication primitives, such as the “pub-sub” model described in the previous section.

The system was implemented in CMS as part of the DT Phase-2 upgrade slice test. More details on this implementation will be provided in the next section. Although the RUs and BUs could, in principle, run on different machines, they

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR

can be deployed on the same machine for less demanding applications, such as the DT slice test. However, tests were also performed with components distributed across multiple machines to verify the system’s functionality.

The server used for the acquisition was a Dell PowerEdge R730, similar to the one used in the LNL testbed system, equipped with dual Intel Xeon E5-2620 octa-core CPUs and 64 GB of memory. During a typical collision run of 2022, with peak instantaneous luminosity of  $\mathcal{L}_{\text{inst}} \approx 1.9 \times 10^{34} \text{ cm}^{-2}$ , the monitoring system indicated that the orbit-building and processing rate was around 11.2 kHz, matching the LHC orbit rate. This suggests that, as expected, at least one hit was present in the DT chambers during each orbit. Moreover, the input throughput for each readout unit, each handling data from half of a DT sector, ranged between 150 MB/s and 170 MB/s. After zero-suppression and reformatting, this was reduced to around 40 MB/s. During this process, the CPU utilization was below 30%, indicating that the workload could have been increased. Using simulated data, it was observed that the system can handle a similar throughput as the one observed for the messaging system based architecture, hence around 5 times what is needed for the full slice test sector.

A “file-mover” service was responsible for periodically transferring data from the PUs’ local disks to permanent storage, ensuring that no backlog of files would gradually fill up the storage space.

## **3.3** DT Phase-2 upgrade demonstrator

During LS2, the second LHC long shutdown, which took place between the end of Run-2 and the beginning of Run-3 (2018-2022), one sector of the DT detector (specifically, the MB1 to MB4 chambers of Wheel+2, sector 12) was instrumented with OBDT prototypes to establish a demonstrator for the Phase-2 upgrade system. This setup is known as the DT slice test [56, 81, 82], highlighted in Figure 3.5 (left).

To avoid interfering with the ongoing Run-3 data-taking process, the prototypes of the OBDTs were installed in parallel with the current (referred to as Legacy) minicrate system, which hosts the electronics responsible for computing the TDC and local trigger primitives. This was achieved by installing passive splitters to divide the chamber’s LVDS signals: one end was connected to the legacy system, while the other was routed to the OBDT prototype (Figure 3.5

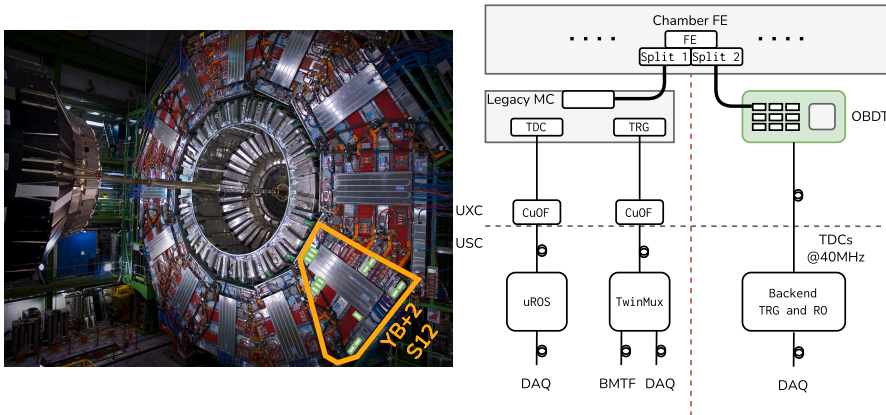


Figure 3.5: Left: Transverse view of the CMS detector highlighting the muon barrel Wheel+2, sector 12, where the DT slice test is installed. The approximate location of the 13 OBDT prototypes is marked in green. Right: Diagram illustrating the data flow from the chamber's frontend electronics. The raw signals are split into two copies: one copy is fed into the legacy minicrate (MC), where TDC and local trigger primitives are computed in the experimental cavern (UXC) and then transmitted to the trigger backend (TwinMux) and readout of TDC hits (uROS) systems located in the service cavern (USC), with the latter being triggered only when an L1A is received. The second copy of the signals is sent to an OBDT, which computes the TDC in FPGA and transmits all the hits to a backend processor demonstrator responsible for both triggering and DAQ, replicating the Phase-2 architecture.

right).

The Phase-2 upgrade system simplifies the Legacy architecture significantly: all TDC hits produced by the OBDTs will be streamed at 40 MHz to a backend located in the service cavern, which is responsible for both trigger and readout. In the slice test, a link to the DAQ is available, enabling the integration of the demonstrator data into the CMS event for offline studies of both Phase-2 hits and trigger prototype performance [83].

Figure 3.6 provides a schematic view of the slice test electronics, including the scouting component. To cover all the channels in the sector, 13 OBDTs have been installed: 3 in MB1, MB2, and MB3, and 4 in the MB4. In each OBDT, two links transmit the same set of TDC hits with different data formats: one for the slice-test demonstrator and the other for the scouting system. Optical fibers from the experimental cavern are routed to a custom optical patch panel, where the two links are split into two independent sets of fibers, one per each data format. One set is used in the slice-test demonstrator backend, received by

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR

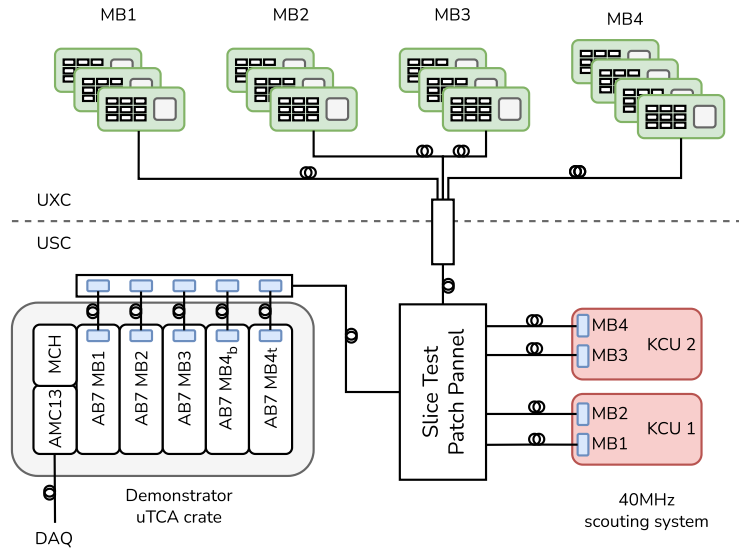


Figure 3.6: Schematic representation of the slice test electronics components. The stream of hits from the 13 OBDTs<sup>3</sup> installed in MB1, MB2, and MB3, and 4 in MB4 are transmitted from the experimental cavern (UXC) to the service cavern (USC), where the backend systems are located. An optical patch panel is used to split the links into two sets of fibers: one set connects to the scouting system, and the other to the upgrade demonstrator backend. The upgrade demonstrator receives the hits on a set of boards called AB7, which are spare TwinMux boards where the Phase-2 trigger algorithm prototype is implemented. A link to the DAQ allows the slice test data to be injected into the CMS event stream when an L1 Accept is received. The second set of fibers is received by the scouting system, which is implemented using two KCU1500 boards and is equivalent to the system described in subsection 3.1.2.

5 boards called AB7, running the Phase-2 reconstruction firmware prototype. The results of the reconstruction, i.e., the phase-2 DT local trigger primitives and hits are inserted into the CMS event as standard data when a L1 Accept from the L1-trigger is received. The same boards are used for timing distribution to the OBDT and slow control.

The other set of fibers is sent to two KCU1500 boards of the DT scouting system. This system is identical to the one used in the LNL testbed, as described in subsection 3.1.2. The only change consisted in moving the CDR frequency reference from 40 MHz used in LNL to the  $\sim 40.079$  MHz used by LHC. From the backend implementation perspective, no differences exist compared to the



setup with the VC707. The main difference between the two is the larger number of hits due to the increased number of chambers and channels compared to the miniDTs. Additionally, no firmware changes were required due to the different TTC system, as the KCUs are not connected to one and the clock is reconstructed from the data. One of the KCUs receives the links from the OBDTs in MB1 and MB2 (6 links), while the second receives the links from MB3 and MB4 (7 links). Raw data were received and merged by the “multistage” processing system, as described in subsection 3.2.2. The scouting readout system operates in a completely “opportunistic” manner, meaning that data from the boards can be collected independently of the central DAQ and trigger chain of the experiment. This design ensures that the system does not interfere with the standard CMS run for example, by generating back-pressure, i.e., pausing the data producer when the incoming rate exceeds the processing capacity. The full set of hits is received, and if the system cannot keep up with the rate, additional hits will be dropped. The system only depends on the OBDTs to be properly configured.

### 3.3.1 Data Scouting results during collisions

The DT scouting system has been tested under various data-taking conditions, beginning with the first fill of LHC Run-3, fill number 7920 [75]. During this fill, the LHC was filled with three bunches, two of which were colliding at Point 5 (IP5).

In Figure 3.7, the rate of hits as a function of the bunch crossing assigned to each hit by the OBDT is shown. The bottom panel marks the BXs that collided at P5 with a black line. Over the flat, constant-rate background caused by cosmic rays and detector noise, two higher-rate regions are visible, corresponding to the two colliding bunches where a larger rate of muons is produced, consequently inducing a greater number of hits.

Additionally, a third region with an increased rate is visible around BX 900, albeit with a much lower rate compared to the other two. This corresponds to the third bunch circulating in the LHC but not colliding in CMS. The hits produced by this bunch are likely due to muons from the beam halo or machine-induced background.

In Figure 3.8 (left), a zoomed-in view around the colliding bunch 1786 is shown. The standard distribution of hits produced by the passage of a muon in the CMS DTs, referred to as the *Time Box*, is clearly visible [84]. This shape arises

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR




Figure 3.7: Rate of TDC hits as a function of the bunch crossing tagged by the OBDTs of MB2 during the first fill of LHC Run-3 (Fill 7920). In this fill, 3 bunches were circulating in the LHC, with two of them colliding at CMS (IP5). The bottom panel marks the bunches colliding at P5 (Bunch ID 1 and 1786) with a black line. An additional increase in the rate is visible around BX 900, corresponding to the third bunch that did not collide in CMS but still left hits in the chambers as machine-induced background.

from the relatively slow signal collection time inside a DT cell, where primary electrons can take up to approximately 400 ns (16 bunch crossings) to traverse the full cell length. In other words, the signal from a muon crossing the chamber at, for instance, BX 100 can leave a signal that is collected and digitized up to BX 116. This is reflected in the time box visible in the figure, where hits are tagged with BXs occurring after the collision, marked with a black line in the bottom panel. Non-linear effects of the electric field inside the cell, combined with geometrical effects due to the chamber's position relative to the interaction point, cause the characteristic peak visible at the start of the time box. In an ideal scenario, this would produce a flat distribution, as the name “Time Box” suggests.

This shape is clearly visible thanks to the low background in the chambers, due to the small number of bunches and their intensity, combined with the presence of only one colliding bunch. Typically, a trigger combined with muon track reconstruction would be required to isolate genuine hits and produce a similar time box.




Figure 3.8: Zoomed view of Figure 3.7 showing the regions of bunch crossings around the colliding bunch (left) and non-colliding bunch (right) at CMS. In the colliding bunch region, the time box distribution of the TDC hits produced by the passage of a muon is clearly distinguishable over the flat background. This shape can also be recognized in the non-colliding bunch region.

Similarly, in Figure 3.8 (right), the time box shape can be recognized over the background. However, this does not correspond to any bunch colliding in CMS, as indicated by the bottom panel. This shape is instead caused by the beam halo induced by the third bunch. The time box shape suggests that genuine muons were crossing the chambers within that range of bunch IDs.

As the LHC gradually increased the number of bunches during intensity ramp-up steps, the previously analyzed clean scenario began to change. In Figure 3.9, data from LHC fill 8016 is shown, where 603 bunches were circulating, 590 of which were colliding in CMS. While two isolated bunches are present, visible as small increases in the rate around BX 800 and 2600, the others are organized in trains. The first train appears at the beginning of the orbit, consisting of 12 bunches, followed by 4 groups of 3 trains, each with 48 bunches. Within a train, consecutive bunches are filled, meaning they are separated by 25 ns. A large rate is visible when looking at the colliding bunches, but a similar rate is also seen when considering the non-colliding bunches, indicating the presence of out-of-time hits that are not directly associated with the interaction point but still contribute to the hit rate [85, 86, 87].

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR






Figure 3.9: Rate of TDC hits as a function of the bunch crossing tagged by the OBDTs during LHC fill 8016. In this fill, the machine was filled with 603 bunches, 590 of which were colliding at CMS. The train structure is clearly visible in the hit rate, while the two isolated bunches are marked by small rate increases around bunch crossings 800 and 2600. The rate of each muon chamber is shown in different colors.

Additionally, the behavior of the 4 stations differs. MB1, being closer to the interaction point, shows a higher rate of hits from both collisions and background due to activation effects compared to the other stations. MB2, benefiting from the additional shielding provided by the iron of the return yoke, has a much lower background but also a lower rate of hits in correspondence to colliding bunches. MB3, representing the most internal chamber, is shielded by MB1 and MB2 on one side and by MB4 on the other. This results in extremely low background levels, as well as fewer hits from genuine muons due to the geometric acceptance effects of this chamber. Interestingly, the outermost chamber, MB4, exhibits a background rate higher than that of MB2.


These effects are reflected in the chamber occupancy, as shown in Figure 3.10. The occupancy of each channel is displayed in the context of the CMS DT chamber geometry, visible in Figure 3.11. The SL1 of the MB1 chamber has the highest rate, slightly lower in SL2 and SL3. The larger occupancy in the SL2 of MB2 was previously observed [88], and points to the gap between the calorimeter barrel and endcap systems. One possible source of this occupancy is hadronic showers that cannot be fully contained by the calorimeter, allowing




(a) WH+2 S12 MB4



(b) WH+2 S12 MB3



(c) WH+2 S12 MB2



(d) WH+2 S12 MB1

Figure 3.10: Occupancy of DT slice test chambers during LHC fill 8016. The chamber's geometry is shown in Figure 3.11.

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR




Figure 3.11: Rate of TDC hits for MB1 and MB4 chambers as a function of time. On the left side of the plots, the last minutes of collisions from fill 8016 are shown. A first drop in the rate is visible around lumisection 287, where the beams were dumped. Following this, the rates of both chambers gradually decrease towards a similar value after approximately 14 minutes.

charged particles to penetrate the muon detector.

On the other hand, no significant difference is visible between the superlayers of MB2 and MB3. In the latter, one hot-spot is visible in SL2, gradually decreasing with increasing wire number, i.e., towards the inside of the detector.

In MB4, where the gap in the middle is due to the fact that SL2 is not installed, it is evident that the external superlayer has a higher rate compared to the innermost one. This suggests that the background is mainly coming from the experimental cavern rather than the inside of the detector. Indeed, hadronic showers also give rise to an intense neutron flux, which, due to the low neutron interaction cross-section, can easily traverse meters of steel, thermalize, and fill the experimental cavern with a so-called “neutron gas”. Due to the slow nature of this process, the background induced by this neutron gas is out of time with respect to the collision, resulting in contributions outside the colliding bunches [89].

The beams of this LHC fill were dumped due to an Unidentified Falling Object (UFO) event [90], where micrometer-sized dust particles lead to fast beam losses and a consequent protection dump. However, the same CMS run continued for several minutes, allowing for the observation of the transition between collision




Figure 3.12: Rate of TDC hits for MB1 and MB2 chambers as a function of time during fill 8484, corresponding to a “full-machine” with 2450 bunches colliding in CMS. All phases of the collision sequence are captured: from the start of collisions, the emittance scan, and the gradual increase in instantaneous luminosity. All these steps are reflected in the hit rate, from both genuine muons and noise.

and no-beam periods. In Figure 3.11, the rate of hits in chambers MB1 and MB4 is shown as a function of time. The time is represented in “Lumisections” (LS), a unit used in CMS corresponding to a block of  $2^{18}$  orbits, lasting around 23.3 s. On the left side of the plot, the last minutes of collisions are visible. A sudden drop in the rate is observed around lumisection 287, corresponding to the beam dump. Following this, a gradual decrease in the hit rate is seen in both chambers, with the rate slowly converging to a similar asymptotic value after approximately 14 minutes. This behavior is due to the detector’s activation, which gradually decreases after the collisions end.

The rate of hits collected from the chambers closely follows the LHC status. In Figure 3.12, the rate of hits collected from MB1 and MB2 is shown as a function of the lumisection number. Following the timeline, the rate starts to be visible only around lumisection 15. This is because the DT High Voltage is turned on only during the LHC ramp, where the 6.8 TeV energy level is reached. This occurs during the *squeeze* phase, where the betatron function at the collision points ( $\beta^*$ ) is reduced to increase luminosity. After this, the *adjust* phase begins, during which beams are adjusted and prepared for collisions. This phase is visible

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR




Figure 3.13: Rate of TDC hits per bunch crossing from fill 8484, corresponding to a “full-machine” with 2450 bunches colliding in CMS. The bunches colliding in CMS are marked with a black line in the bottom panel.

as the first increase in the rate around lumisection 25, before reaching *stable beams* around lumisection 35, representing stable conditions with collisions in the experiments, with beam backgrounds and lifetime under control. Following this, an *emittance scan* [91] is visible. During an emittance scan, the two beams are moved relative to each other in steps along the  $x$  and  $y$  planes, resulting in the two bell-shaped distributions visible between lumisections 40 and 60. The peak rate occurs when the beams collide head-on, i.e., with zero displacement. After this, visible around lumisection 100, the  $\beta^*$  is further reduced in steps to increase the instantaneous luminosity up to a predefined value, which is then maintained during the fill through luminosity leveling[92].

For the same fill, the rate of TDC hits per bunch crossing is shown in Figure 3.13. In this filling scheme, 2450 bunches were colliding in CMS. In this scenario, the rate of hits per bunch crossing is much higher than what was observed in previous filling schemes, including the non-colliding ones, representing out-of-time background activity. A “fast” component of the background is visible, for instance, in the orbit gap, where the hit rate decreases after the last collision of the orbit, around BX 3400.

A zoomed-in view of the bunch crossing range around 2600 and 2800, where the isolated bunch is located, is shown in Figure 3.14. The time box shape






Figure 3.14: Zoom of Figure 3.13 showing the BX range including an isolated bunch between two sets of trains.


described earlier is barely visible for the isolated bunch, sitting atop the large noise present between trains. However, this shape is not visible when looking at contiguous bunches. This is because in these cases, a convolution of time boxes one for each bunch is present.

Indeed, when observing the first bunch of the train starting around BX 2780, it is evident that this corresponds to the first increase in the hit rate compared to the non-colliding bunches. However, the second bunch of the train produces its own time box on top of the one generated by the first bunch. This process continues, generating the “stair” pattern visible in the initial part of the trains.


A saturation point is reached: each time box is approximately 16 BXs long, so after the 16th bunch of the train, the time box produced by the first bunch ends, while the one from the new bunch begins, leading to an overall constant rate. The descending stair visible after the last bunch of a train around BX 2685 is explained by the same process in reverse order.

The background activity, in addition to contributing to detector aging, also impacts trigger efficiency, as reconstructing segments when a large number of hits is present can be challenging. In Figure 3.15, the channels firing in time with the passage of a muon are shown for a typical event in MB3 (a), one in MB2 (b), and one in MB1 (c). In MB3, the “cleanest” chamber as it is shielded from the beam products by the innermost chambers and from the neutron gas by MB4,


### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR



(a) WH+2 S12 MB3



(b) WH+2 S12 MB2



(c) WH+2 S12 MB1

Figure 3.15: Event displays of channels firing after the passage of muons in MB3 (a), MB2 (b), and MB1 (c). In all cases, the muon traverses all the superlayers, leaving one hit per layer. In the MB1 chamber, where the noise level is significantly higher, a larger number of firing channels are observed compared to the typically clean events found in MB3. A window of 30 BXs from the passage of the muon is used to filter hits. The event display for MB2 highlights another common noise pattern: in SL1, four clean hits are present, while a large number of contiguous firing channels are found in SL2 and SL3 (two per layer, with some channels containing multiple hits).

only hits from the muon itself are visible, with four hits per superlayer. Two hits were registered in the same channel in SL3 Layer 4, possibly due to the emission of a  $\delta$ -ray, which produces an additional hit close in time to the original one.

On the other hand, the larger occupancy in the chamber for a similar event in MB1 is evident. In this case, the muon crosses all the superlayers, leaving four hits in each. However, additional hits are present in the event, particularly those close to the real track, which are the most problematic as they can bias the reconstruction of the muon segment. This is also apparent in the event display for MB2: while four clean hits were left in SL1, a larger number of hits were produced in SL2 and SL3. Multiple contiguous channels fired in time with the passage of the muon, with all hits occurring within 15 BXs relative to the time the muon crossed the chamber.

The data collected with this trigger-less DAQ system offers a unique opportunity to develop data quality monitoring (DQM) studies based on 100% of the hits produced by the detectors' front-end electronics, completely unbiased thanks to the absence of dedicated trigger selection logic. Indeed, many of the plots shown in this section can be generated with less than 1 minute of data taking. For example, the 40 MHz data scouting readout could provide CMS with an alternative unbiased monitor of the instantaneous luminosity delivered by the accelerator, based on the raw counts of all the detectors' signals.

Additionally, it can enable in-situ measurement of the detector occupancy, using data collected both during and outside of LHC bunch crossings. The off-collision data can also be used to perform an unbiased estimation of the detectors' background noise rate and its evolution with integrated luminosity.

### 3.3.2 Readout of the final OBDT prototypes

The DT slice test was instrumental as a demonstrator for the Phase-2 upgrade, allowing the exercise and validation of a prototype of the Phase-2 electronics chain. The successful operation of this demonstrator confirmed the feasibility of the new architecture, including preliminary results on the new trigger algorithm. Meanwhile, the development and finalization of the production versions of the OBDT boards were underway.

Although the final OBDT prototypes are functionally equivalent to the initial versions tested in the slice, they differ substantially in design and implementation. The final OBDT system consists of two board flavors: OBDT- $\phi$ , responsible

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR

for the digitization of the two  $\phi$  superlayers (SL1 and SL3), and OBDT- $\theta$ , which handles the  $\theta$  view (SL2). Both variants host the same Microchip Polarfire FPGA, which performs the time digitization of the input signals. However, the board designs differ due to different connectors and mechanical constraints necessary to fit within the new Minicrate structure.

In the same way as in the demonstrator, the digitized data is transmitted via an optical link to the readout and trigger chains. However, unlike the prototype used in the demonstrator, the communication for these final OBDT prototypes is performed via two VTRx+ transceivers [93], which provide two bidirectional links for slow control and timing, as well as six transmitter links. These transmitter links are capable of outputting data at rates of up to 10.24 Gbps to the backend system using the lpGBT protocol [70].

The 40 MHz scouting system described in the previous sections could serve as an invaluable tool for verifying the correct functionality of these OBDTs. By continuously receiving and storing all hits, the scouting system can ensure the correct behavior of the boards by monitoring data quality in real-time as a standalone system, independently from standard trigger and DAQ chains.

The new hardware for the scouting system is based on the Xilinx VCU118 Evaluation Kit, which hosts a Xilinx Virtex Ultrascale+ FPGA. The architecture of the firmware developed for this board can be divided into two logical blocks: one responsible for receiving and pre-processing hits from the front-end boards, and the other dedicated to transmitting the processed data to a server [94]. A schematic view of the system is shown in Figure 3.16.

The receiver block is responsible for collecting hits from a set of OBDTs and preparing them for transmission by the sender block. The VCU118 board is equipped with two on-board QSFP ports; one of these ports is dedicated to transmission, while the other can be used to receive up to four lpGBT links, corresponding to hit streams from four OBDTs.

To accommodate a higher number of OBDTs, with the aim of eventually collecting data from an entire DT sector, where approximately  $\approx 5$  MHz of hits is expected as measured from the slice-test, the VCU118 has been equipped with an FMC+ module that can host up to six QSFP transceivers. This configuration allows the system to receive up to 24 lpGBT links, corresponding to the 24 OBDTs required for full sector coverage.

Moreover, in contrast to the scouting system used for the preliminary version of the OBDT, where the clock was reconstructed from the data, this new setup

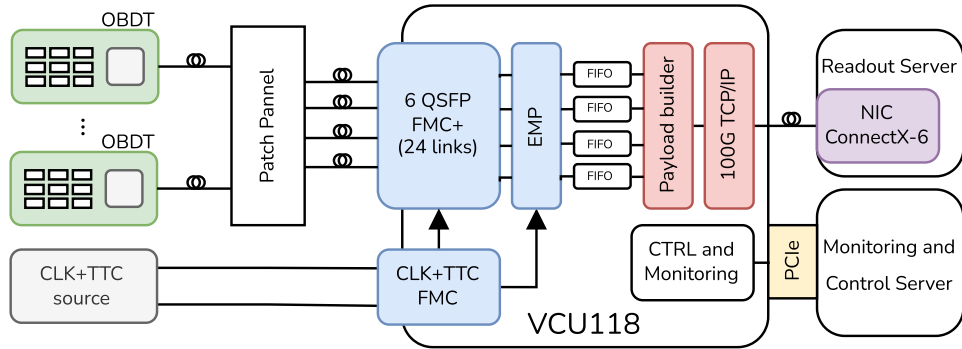


Figure 3.16: Schematic representation of the new 40 MHz readout system. From left to right: the connection to the CMS TCDS and lpGBT links from the OBDTs are received by a VCU118 using two FMC modules (FPGA Mezzanine Card). A simplified illustration of its main components is shown in the middle box. The firmware responsible for receiving the links is based on the CMS EMP framework. This is connected via a set of FIFOs to the transmitter part, which includes a payload builder process and a 100G TCP/IP module. The board is mounted on the PCIe of a control server, which is responsible for monitoring and control. A second server is used to receive and process the TCP/IP stream using a commercial NIC. The receiver components are shown in blue, while the transmitter components are shown in red.

uses an external clock, synchronized with the one driving the OBDTs. For this purpose, a second FMC module was installed on the VCU118, dedicated to receiving the external 40 MHz clock and CMS fast control signals from the CMS Timing and Control Distribution System (TCDS) [95].

The receiver firmware for this task, in blue on Figure 3.16, is built on top of the CMS EMP Framework [96], which is currently under development for the Phase-2 upgrade of the Level-1 trigger. This framework offers the advantage of implementing common functionalities that are necessary across various subsystems, ensuring full compliance with the protocols and standards that will be adopted in CMS. The EMP framework provides essential infrastructural firmware components, including the lpGBT-FPGA firmware [97], which is used to decode the input links from the frontend. Additionally, the framework includes control and monitoring software, with functionalities that can be easily extended as needed. Within this framework, the existing lpGBT module has been instantiated along with the payload formatter and spy buffers utilities that capture the data stream for debugging purposes. The TCDS decoder included

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR

in the framework is designed for the new TCDS system from the Phase-2 upgrade. However, to maintain compatibility with the current Phase-1 upgrade TCDS command stream, a module implementing the current TCDS decoder was incorporated. The payload formatter module is responsible for reformatting the data received from the OBDTs, extracting relevant TDC information such as TDC channel and bunch crossing.

Finally, the data are zero-suppressed and written into Clock Domain Crossing (CDC) FIFOs, effectively decoupling the LHC clock-synchronous readout of the lpGBT links from the asynchronous data transmission to the server.

The transmitter block, in red on Figure 3.16, is responsible for sending the TDC hits to a server, where they can be stored and analyzed. Data from the CDC FIFOs are read by a payload builder module, which merges data from all input links, producing a single stream that is fed to the transmission module.

This is implemented using a 100G TCP/IP module developed by the system group at ETH Zurich [98]. This network stack supports TCP/IP, RoCEv2, and UDP/IP at 10-100 Gbit/s, where both the choice of the network protocol and the width of the data path are parameterized at compile time. For this work, the full TCP/IP stack was used, including a retransmission module in case of lost or malformed packets at the receiver, where data are buffered on an external DDR4 memory available in the VCU118. While the library handles packet formation, final communication with the external network is handled by the Xilinx 100-Gigabit Ethernet MAC core [99].

The choice of this library was driven by multiple factors. First, it provides a ready-to-use, high-performance network stack, ideal for transmitting hits to a server. Second, this library includes an implementation of the RoCEv2 (RDMA Over Converged Ethernet) protocol, which is crucial for the ultimate goal of implementing a system using RDMA, as discussed in the FEROCe project section (section 3.4). Using TCP/IP is a first step in this direction, allowing for experience to be gained with the library.

The data stream is collected by a Readout Server, a Dell PowerEdge R750 equipped with 2xIntel Xeon Gold 5318Y CPUs @ 2.10GHz and 16x16GB DIMM DDR4 memory, hosting an NVIDIA Mellanox ConnectX-6 Network Interface Card (NIC). For the purpose of the OBDT readout demonstrator, a point-to-point connection is established between the ConnectX NIC and one of the VCU118 on-board QSFPs. The data stream is received by the Readout Server via a standard TCP socket and temporarily stored on a RamDisk. The data are then transferred

to storage for long-term preservation and offline analysis. At this point, this system could easily be integrated into one of the processing pipelines described in the previous sections.

The VCU board is mounted through an extender on the PCIe bus of a Control Server, which is responsible for the monitoring and slow control of the board. These operations consist of actions such as resetting the links, checking their status, and programming the clocks. These are implemented using IPBus registers. The user interface is provided via the EMP software libraries, included in a custom application that implements all the necessary operations to configure, monitor, and control the board.

### Results from the prototype

The system was developed, tested, and first deployed at the LNL, where a mockup system of the Phase-2 DT upgrade was constructed for the development and testing of the OBDT- $\phi$  boards. The setup at LNL includes a pre-production batch of the OBDT- $\phi$  boards mounted on a mockup chamber, along with their control and timing distribution systems compliant to the specific of CMS. This system also provides the clock to the VCU118, as described earlier.

A data-taking campaign was conducted under two scenarios: collecting pure electronics noise from the frontend, and injecting periodic signals into the frontend.

In the first data-taking scenario, the threshold of the chamber's front-end discriminator was set to zero. This configuration allows all signals generated at the front-end level to be passed to the OBDTs and digitized. Since the trigger-less readout collects the entire set of TDC hits, the data acquired in this way is dominated by electronics-induced noise rather than genuine hits produced by muons crossing the detectors.

In the case of pure white noise, a flat distribution in the TDC bins is expected. This allows for checking the differential nonlinearity (DNL) of the TDCs generated by the OBDTs by measuring the distance of each bin from its ideal width of 25/32 ns. The latter is due to the TDC implementation, which samples the signal 32 times per LHC bunch cycle of 25 ns. The results of this study, presented for a randomly selected OBDT channel in Figure 3.17a, show that the DNL is within the desired range of  $\pm 10\%$  for each channel.

The OBDT boards are designed to be capable of generating periodic signals

### 3.3. DT PHASE-2 UPGRADE DEMONSTRATOR

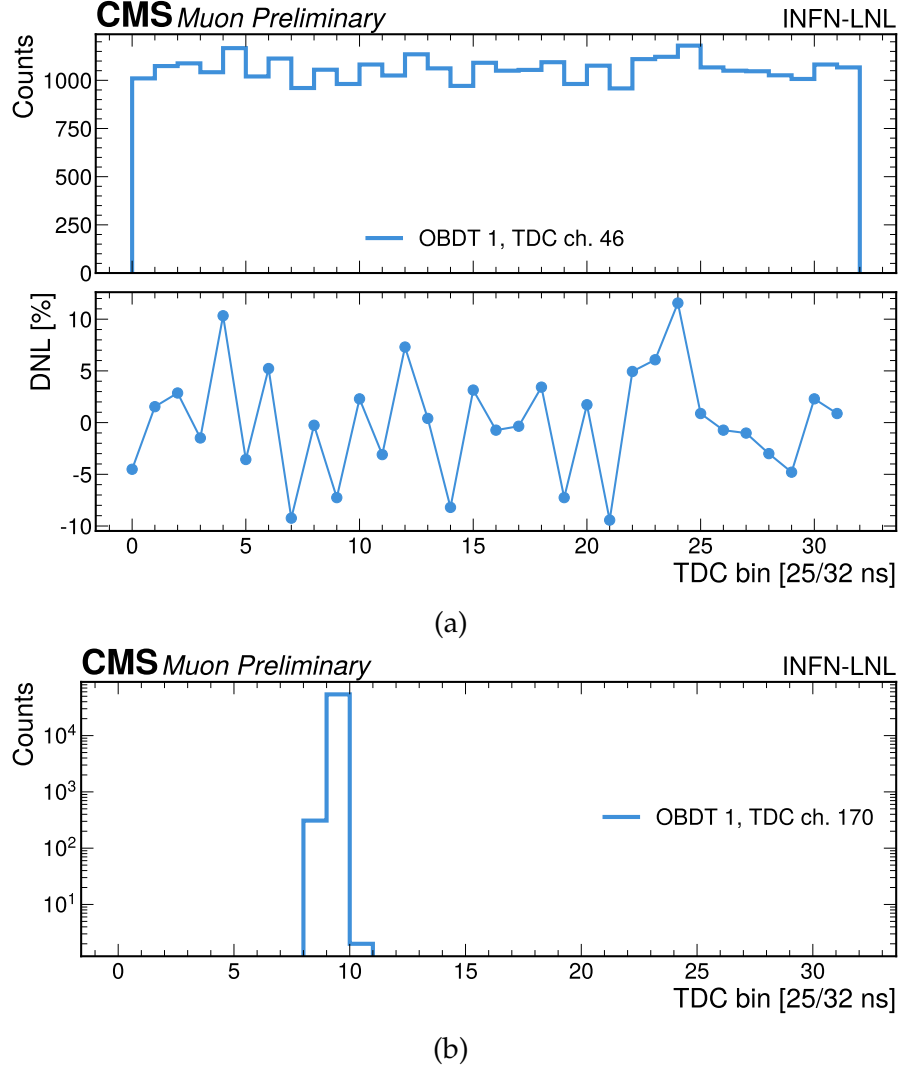


Figure 3.17: Results with data collected from the OBDTs installed at INFN Legnaro National Laboratories (LNL). (a) Distribution and differential nonlinearity (DNL) of the TDC hits produced by one of the OBDTs, expected to be within the range of  $\pm 10\%$ . (b) TDC distribution for a single channel from the periodic Test Pulse signal.



that can stimulate the chamber's front-end output, called Test Pulses (TPs). Enabling TPs on the OBDTs generates a train of signals with a user-defined time period, but always produced in correspondence with the same TDC bin for every TP. Ideally, collecting a train of TPs will result in a histogram spanning exactly one TDC bin. However, the TDC histogram of a TP may be broader due to the convolution of the TP generation and TDC uncertainties. Figure 3.17b shows the resulting TDC histogram produced in a TP run for a randomly chosen OBDT channel. The majority of hits land within two TDC bins, which is consistent with expectations.

### 3.4 FEROCE project

The ever-increasing demand for larger datasets is a common trend across many fields of experimental physics. This need is particularly pronounced in experiments that, to minimize inefficiencies or avoid being constrained by strict latency requirements, choose to operate without a hardware trigger system, thereby collecting all the data produced by the detector frontends. In such scenarios, the processing power of the computing farm tasked with analyzing the dataset becomes crucial, but this potential can be significantly hampered by inefficient data movement.

When using networking systems based on protocols such as TCP/IP, a substantial portion of CPU resources is often consumed by networking tasks rather than data processing itself. This is because conventional network stacks, responsible for managing network traffic, typically handle data through multiple copies and perform operations that are expensive in terms of both latency and throughput.

The challenge of efficiently transferring large volumes of data can be addressed by employing network protocols based on the concept of Remote Direct Memory Access (RDMA). RDMA offloads the burden of low-level data movement and protocol handling from the processor, allowing it to focus on data processing. However, this approach requires more intelligent hardware on the frontend side to initiate the RDMA transfers.

This concept is the foundation of the FEROCE (Front End ROCE) project[61], which aims to enable front-end electronics to initiate RDMA transfers directly to the computing farm. By eliminating the point-to-point connection between the frontend and backend, FEROCE introduces the flexibility to dynamically

### 3.4. FEROCÉ PROJECT

route data to computing nodes based on their processing availability. This approach not only simplifies the architecture by reducing the need for custom hardware but also leverages industry-standard Ethernet-based RDMA solutions like RoCE (RDMA Over Converged Ethernet), which is widely supported and well understood in the industry. This concept is illustrated in Figure 3.18, which compares a typical DAQ system Figure 3.18a with the FEROCÉ architecture. In the FEROCÉ schema, the network protocol is implemented directly in the front-end, allowing the use of commercial switches immediately after the front-end electronics, rather than custom concentrator boards.




Figure 3.18: Point-to-point connection between front-end and back-end, prior to transmission to a server (left). In this scenario, the network switch is introduced after the packets reach the server. In the FEROCÉ schema (right), the network protocol is implemented directly in the front-end, allowing the use of commercial switches immediately after the front-end electronics rather than custom concentrator boards.

While the primary focus of the FEROCÉ project is on front-end electronics, the same concepts and outcomes can be applied to backend boards that receive frontend links. These backend boards could perform an initial layer of data aggregation from the front-end and then use RoCE to transfer data more efficiently to server memory. Therefore, the technologies and strategies developed in this project can be utilized in various scenarios.

The starting point for this work is the FPGA network stack developed by the system group at ETH Zurich, which was introduced in the previous section. This stack includes implementations of TCP/IP, UDP, and RoCEv2 protocols, making it suitable for an initial exploration of RDMA-based data transfer. This stack has been used to understand the protocol and implement a first prototype before moving toward a custom implementation. Many features in the library, such as the RoCE receiver module, are not essential for the FEROCÉ project, as the boards are primarily designed to send data rather than receive it, allowing for efficient use of FPGA resources. Additionally, while the library is written




Figure 3.19: Diagram of the communication between the RTL simulation using Synopsys VCS [101] and Soft-RoCE. A TAP device is used to connect the RTL simulation with the Linux network stack.

in Vivado HLS (High Level Synthesis), which ties it to Xilinx devices, an HDL (Hardware Description Language) implementation would enable broader compatibility across different vendors' devices, such as the PolarFire FPGA used in the DT front-end boards.

In the following sections, a first implementation based on the ETH network stack will be described, followed by the introduction of a light-ROCE library implemented in a HDL language.

### 3.4.1 A 100G RoCEv2 Implementation

The initial phase of the FEROCe project involved developing a dynamic simulation environment for the network stack, a crucial step given the complexity of debugging network protocols. Traditional firmware simulations often rely on static test benches, which limit interaction with external systems. In contrast, a dynamic simulation model enables active interaction with its environment by exchanging network packets with the host operating system running the simulator. This means the simulated RoCE stack can, for instance, write a packet directly into the memory of a software application via an RDMA WRITE operation.

To enable this, a TAP device, a virtual network device in the Linux kernel, is employed to exchange data between the RTL simulation and the Linux network stack. Operating at the link layer (layer 2 of the OSI model[100]), the TAP device simulates an Ethernet connection, effectively bridging the RTL simulation and the software application. This setup is illustrated on the left side of Figure 3.19.

In the simulation of the ETH network stack, operating at 10 Gbps, the XG-

### 3.4. FEROCES PROJECT




Figure 3.20: Steps involved in a Reliable Connection (RC) RDMA WRITE operation: First, the two endpoints synchronize out-of-band by exchanging the information needed for the transfer. The transmitter then sends the data along with the address, remote key, and QP number. The receiver sends an acknowledgment after each packet if received correctly, which can be used to provide a measurement of latency and throughput.

MII (10 gigabit media-independent interface) stream from the MAC is used to interface with the dynamic simulation instead of targeting a gigabit transceiver. An application based on Soft-RoCE, a software-based implementation of the RoCE protocol, handles the data exchange. Unlike hardware-based RoCE solutions, Soft-RoCE leverages standard Ethernet hardware and a software stack to achieve RDMA capabilities, making it a flexible choice for simulations. The application allocates a memory buffer where data will be written by the simulation, and sharing the necessary RDMA transfer details such as Queue Pair (QP) information, remote keys, virtual address, and IP address with the simulation.

Once these details are input into the hardware simulation via a dedicated interface, the RDMA transfer can be initiated, for instance by performing an RDMA WRITE operation. The steps for this operation are outlined in Figure 3.20, with the initial synchronization step just described represented as the first line at the top. This exchange of information “occurs out-of-band”, typically via an initial UDP packet, before the RDMA transfer begins.

At this point, the data transfer can begin. In the Infiniband standard[102], the operation may involve multiple data transfers, depending on the payload size. If the message size is smaller than the Maximum Transmission Unit (MTU), it will be sent as a single RDMA WRITE ONLY packet. For larger messages, the data is segmented into packets, each with an MTU size, marked as RDMA




Figure 3.21: Wireshark capture of a test. The RDMA WRITE ONLY operation is visible in the top row, followed by the corresponding ACK. Relevant parameters such as the QP number, PSN, remote key, and IP addresses are displayed on the left panel.

WRITE FIRST, MIDDLE, or LAST. As illustrated in Figure 3.20, the receiver sends an acknowledgment (ACK) for correctly received packets; otherwise, a not-acknowledge (NAK) is sent, prompting retransmission of the packet. Finally, an out-of-band message can notify the receiver that the transfer is complete, allowing verification of the received data. The firmware includes features to measure latency and throughput based on the received ACKs, which, while not particularly useful in simulation, will be critical to measure the performance in a real hardware implementation.

In the simulation, a simple payload containing a counter was sent. The network traffic passing through the TAP device was monitored using Wireshark, with a capture of a test shown in Figure 3.21. This test involved an RDMA WRITE ONLY operation, visible in the top row, followed by the corresponding ACK. The bottom left panel displays all relevant transfer parameters, such as the QP number (0x11), PSN (10337885), remote key (0x22d), and IP addresses (Tx: 22.1.212.209, Rx: 22.1.212.210). The buffer in the application was then analyzed to verify the correct counter sequence was transmitted.

This dynamic simulation setup was invaluable during development, enabling a faster iteration cycle and allowing issues in the library to be identified and resolved efficiently.

Once all desired features were implemented and validated in simulation, the project moved to hardware implementation for the first performance measurements. The same VCU118 board used for the DT readout in the previous section was employed, targeting the 100 Gbps RoCEv2 implementation. This allowed

### 3.4. FEROCÉ PROJECT

the first set of performance measurements, presented in Table 3.1.

Transfer size	Throughput [Gbps]	Latency [ $\mu$ s]
8 kB	30.7	1.6
32 kB	62.8	2.3
64 kB	73.3	3.1
128 kB	80.5	3.3
512 kB	92.9	12.4
5.12 MB	95.8	12.4
51.2 MB	96.0	12.4
512 MB	96.0	12.4

Table 3.1: Throughput and round-trip latency as a function of the transfer sizes.

Element	RoCEv2 stack	TX only estimate
LUTs [k]	38	$\approx 15$
Registers [k]	49	$\approx 15$
BRAMs	93	$\approx 30$

Table 3.2: Resource utilization of the full RoCEv2 module and an estimate of the TX only implementation.

In this test, data was received by a server equipped with a NVIDIA ConnectX-6 NIC, running the same application used in the Soft-RoCE test.

When sufficiently large messages were transmitted, the system achieved the maximum available throughput. Similarly, the latency, as defined in Figure 3.20, reached a saturation point and did not increase further. However, it is important to note that these measurements were obtained using a direct point-to-point connection between the VCU118 and the NIC. Introducing a network switch into the setup would likely increase latency due to additional processing delays.

Resource utilization, as summarized in Table 3.2, remains manageable. The resource usage estimates for a TX-only implementation, where the receiver functions are omitted, indicate even more efficient resource utilization, which is promising for deployment on resource-constrained FPGAs.

This work represents an initial step toward developing a lightweight RoCEv2 firmware block that can be deployed on FPGAs with limited resources. The successful setup of a working simulation environment and the experience gained from implementing the RDMA module using the ETH library will greatly aid in the continued development of this project.

### 3.4.2 Towards a light-RoCE HDL Implementation

Building on the experience gained from the ETH network stack's RoCEv2 implementation, the next step is to develop a lightweight version tailored specifically to our use case. The primary objective is to create a streamlined RoCEv2 implementation that focuses exclusively on transmitting data from FPGAs to servers. To achieve this, only the essential components necessary for RDMA operations will be implemented, specifically the RDMA WRITE and WRITE with IMMEDIATE operations.

A key consideration in this development is portability across different vendors' hardware. Unlike the previous implementation, which relied on High-Level Synthesis tools specific to a particular vendor, this new approach will utilize a hardware description language. Using HDL allows the design to remain vendor-agnostic, making it easier to port the module to various FPGA platforms without being tied to the toolchains or hardware architectures of a single vendor.

Scalability is also an important requirement, as the system should ideally support various data rates (e.g., 10/25/40 Gbps) without requiring significant changes to its core structure.

By stripping down the functionality to focus solely on transmission operations, the core of RoCEv2 becomes relatively simple. RoCEv2 leverages UDP as its transport protocol, meaning that RoCEv2 messages are essentially UDP packets with specific payloads designed for RDMA operations. The structure of such a message is illustrated in Figure 3.22.

Following this structure, the development of a lightweight RoCEv2 module is simplified. Existing UDP implementations can be reused to handle the transport layer, allowing the focus to be on the RDMA-specific components within the payload. At a high level, the primary tasks involve managing the splitting of frames based on the MTU and generating the RoCE-specific headers, while the rest is handled by the UDP stack.

To achieve this, a repository containing a set of Ethernet network components implemented in Verilog[103] has been identified as a suitable foundation. This repository offers a collection of Ethernet-related components for Gigabit, 10G, and 25G packet processing, including modules for handling Ethernet frames, as well as IP, UDP, and ARP protocols. It also includes MAC modules for Gigabit and 10G/25G networks. Consequently, all the necessary components for this use



### 3.4. FEROCCE PROJECT




Figure 3.22: Structure of a RoCEv2 packet. The outermost layer is the Ethernet frame, responsible for link-layer communication over the network, with the EtherType field indicating that the packet carries an IP payload. Following this is the IP header, which specifies the protocol as UDP. The UDP header includes the port number, where a special port (4791) reserved for RoCE indicates that the next header is the InfiniBand BTH (Base Transport Header). The BTH+ header, along with RETH (RDMA Extension Transport Header) and IMMDET (Immediate Extension Transport Header), contains crucial routing information, such as the destination queue pair (QP) number, packet sequence number (PSN), and operation codes (opcodes) for RDMA operations (e.g., RDMA WRITE). The payload, typically memory data in an RDMA WRITE operation, follows the BTH. Finally, the packet includes the RoCE-specific iCRC for data integrity and the Ethernet FCS (Frame Check Sequence). The sizes of each header are noted above the structure, with the overall packet size bounded by the MTU.

case are already available, providing a strong starting point for the lightweight RoCEv2 implementation.

An illustration of the firmware structure components for the transmitter is shown in Figure 3.23.

All necessary components, except for the RoCE and iCRC modules, are sourced from the previously mentioned repository. The system has two available implementations, one operating at 10 Gbps and the other at 25 Gbps. Both implementations utilize a 64-bit datapath. The 10 Gbps design runs with a clock frequency of 156.25 MHz, while the 25 Gbps design uses a 390 MHz clock. The




Figure 3.23: Schematic representation of the RoCEv2 transmitter module. The data stream enters the RoCE module, where it is split into frames based on the MTU size. The relevant RoCE headers, BTH, RETH, and IMMDET, are added to form the IB packet encapsulated in the UDP datagram. These are combined with the payload and passed to the UDP module, where the UDP headers are added, followed by the rest of the chain that forms the Ethernet packet for final transmission.



AXI4-Stream (AXIS) interface is employed to connect the RoCE module down to the MAC module. Meanwhile, the output from the MAC module is formatted in XGMII and is sent to the Gigabit transceiver, utilizing a 66-bit datapath for data transmission over the physical layer.

The module responsible for computing the iCRC, as specified by the InfiniBand Architecture specification, is placed after the IP module. This placement is necessary because some of the fields required for the iCRC computation are inserted at the UDP and IP layers, making them available only in the later stages. Therefore, a placeholder is initially inserted by the RoCE module, which is then updated before delivering the frame to the MAC.

The data stream enters the RoCE module, where it is split into frames of the MTU size and the headers containing the metadata necessary for the operation, such as the rKey and QP number, are inserted. These form the headers of the IB packet encapsulated in the UDP datagram, as illustrated in Figure 3.22, specifically the BTH, RETH, and IMMDT. These are then combined with the payload and passed to the UDP module, where UDP headers are added, and the packet continues through the rest of the chain, ultimately leading to the formation and transmission of the Ethernet packet.

On the RX side of this system, a RoCE RX module is connected to the output of the UDP receiver, responsible for decoding the AETH (ACK Extended Transport Header), which contains the PSN, QP local number, and syndrome. If the syndrome indicates that the packet is an ACK (i.e., not an NAK), this confirms that the transmission was successful. This information can be used to compute latency and throughput, as previously described.

After testing and validating the system using active simulation in conjunction with Soft-RoCE, the implementation was successfully deployed on the VCU118 evaluation boards in both 10 Gbps and 25 Gbps configurations. Several tests confirmed that the target throughput was indeed achieved: 9.6 Gbps for the 10 Gbps firmware and 24.10 Gbps for the 25 Gbps firmware, with latency comparable to the values observed in previous implementations. The resource utilization remained within acceptable limits, with approximately 13k Flip-Flops (FFs), 11k Look-Up Tables (LUTs), and 7 BRAMs, aligning with the predictions in Table 3.2.

Once the firmware was validated in a point-to-point connection, both the receiver server and the VCU118 were connected to a 10G switch. A second server was then connected to the switch to act as an additional data producer alongside the VCU118, simulating two sources writing to the same server. This

### 3.4. FEROCIE PROJECT

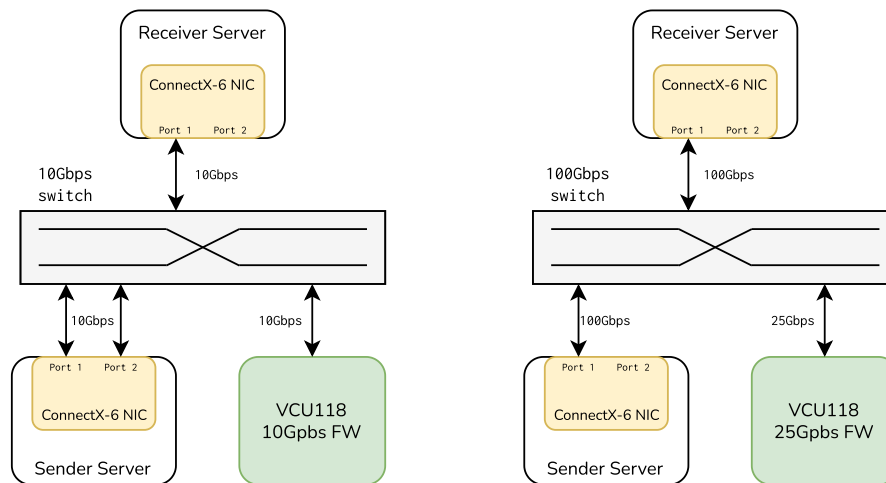


Figure 3.24: Network congestion test setup, where an additional server acts as a data source alongside the VCU118. Two scenarios were tested: one where the components were connected via a 10G switch and limited to 10 Gbps (left), and another where a 100G switch was used, with the sender and receiver set to 100 Gbps and the VCU118 using the 25 Gbps firmware (right).

setup was particularly intended to study the effects of network congestion and explore potential mitigation strategies.

In the first test, illustrated in Figure 3.24 (left), the VCU118 was running the 10G firmware, and the receiver NIC was configured with `ethtool` to autonegotiate a maximum of 10G. With just these two components, no issues were observed. However, once the additional data producer was enabled, the VCU118 firmware entered an error state, causing the system to collapse due to network congestion. This outcome was expected because the current implementation does not handle retransmission. However, even if retransmission were implemented, the throughput would likely drop to around 1 Gbps due to the congestion.

This scenario highlights the importance of Priority-Based Flow Control (PFC) [104]. PFC is a flow control technology that enables the creation of lossless networks by preventing packet loss during network congestion. In this scheme, when a device detects congestion, such as when its receive buffer becomes overloaded, it instructs the upstream device (e.g., a switch) to pause traffic transmission in that specific queue. A recovery signal is sent when the congestion eases, allowing traffic to resume.

PFC was enabled in the ConnectX NICs, the switch, and the MAC used in the firmware implementation. When the previous test was repeated with PFC

enabled, 9.6 Gbps were measured on the receiver and 4.8 Gbps on each of the transmitters, enabling the system to operate smoothly for multiple days. This stability was due to PFC effectively pausing transmission from all entities, thereby preventing network congestion.

A similar test was performed using a 100G switch, with both the receiver and additional producer set to 100G, while the VCU118 ran the 25G firmware (Figure 3.24, right). With PFC enabled, the server's additional producer operated stably at 68 Gbps, while the VCU118 maintained a throughput of 24.10 Gbps.

These tests demonstrated that, thanks to PFC, the system could withstand temporary increases in throughput without requiring retransmission. This capability ensures high performance and stability under heavy network loads, which can occur transiently even in a well-dimensioned network. Moreover, these tests highlight the potential to avoid data retransmission by adopting flow control technologies, thereby enhancing overall network efficiency.

Finally, progress is being made toward implementing this system in firmware at 100G by increasing the datapath width to 512 bits and running it at a clock speed of 322 MHz, sourced directly from the GT transmitter. However, further developments are still required.

The next steps in the FEROCÉ project involve testing this system with FPGAs from front-end boards. For example, the implementation of the 10G firmware on PolarFire FPGAs will be a key focus in the upcoming phase.



# 4

## The CMS L1-Trigger scouting system

### 4.1 Data taking strategies at the CMS experiment

In the CMS experiment, alternative data taking strategies at the high level trigger have been employed since the end of LHC Run 1 [105] to search for new physics in difficult corners of the phase space. HLT scouting and data parking strategies [49] aims at overcoming limitations in the CMS data acquisition chain, such as the finite bandwidth of data to record on disk and the finite computing power needed to reconstruct raw data as they are recorded (prompt reconstruction). An illustration of these two strategies is shown in Figure 4.1 for a typical configuration of the data flow during 2018 (Run 2).

#### HLT scouting

The data scouting technique, also referred to as HLT scouting, aims at overcoming the limitations induced by the limited bandwidth to disk available. The throughput of data to disk is the product of the selected events rate  $\mathcal{R}$  times the size of the raw data of each event  $\mathcal{S}_{\text{raw}}$ ,  $\mathcal{R} \times \mathcal{S}_{\text{raw}}$ . Therefore, in the standard data taking flow, given a fixed bandwidth  $\mathcal{B}$  and event size  $\mathcal{S}_{\text{raw}}$  the rate  $\mathcal{R}_{\text{HLT}}$  defined by the HLT menu can be optimized to fully utilize  $\mathcal{B} = \mathcal{R}_{\text{HLT}} \times \mathcal{S}_{\text{raw}}$ . The idea behind data scouting is to reduce the size of each event and consequently increase the event rate, while keeping the bandwidth constant. In other words, this allows to use lower trigger thresholds, giving access to events that would have been otherwise inaccessible. This is achieved by storing high-level physics

#### 4.1. DATA TAKING STRATEGIES AT THE CMS EXPERIMENT

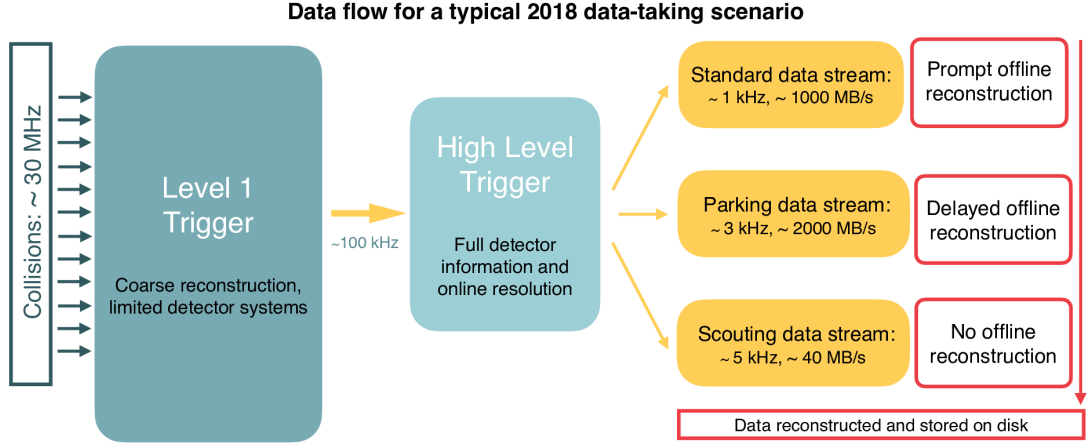


Figure 4.1: Schematic view of the typical Run-2 data flow during 2018[49]. Along with the standard data taking mode, scouting and parking data streams are shown. Numbers are related to a typical 2018 LHC fill, with an instantaneous luminosity of  $\mathcal{L}_{\text{inst}} = 1.2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$  corresponding to an average pileup of 38.

objects (such as jets or muons) reconstructed by the HLT, instead of the complete set of raw data. With this approach the event size  $\mathcal{S}_{\text{HLT-SC}}$  is reduced to the order of  $\mathcal{O}(10 \text{ kB})$ , compared to roughly 1 MB in a typical standard event. The base of the success of this strategies relies on the excellent performance of the HLT online reconstruction, closely approximating the standard offline reconstruction. A further benefit of scouting is that the raw data are processed only once: indeed, no additional computational resources are required to perform the reconstruction of the scouting events. The downside is that the objects quality depend on the reconstruction and calibrations used in the HLT, potentially limiting the accuracy and stability with respect to the one required by standard offline analyses.

This technique has been used to study the phase-space of processes which were limited by the trigger thresholds. One notable example is the search for narrow resonances in the dijet final state[106]. Previous searches were sensitive to resonances with mass above 1.2 TeV, by selecting events at the HLT with  $H_T > 1000 \text{ GeV}$  or at least a jet with  $p_T > 500 \text{ GeV}$ , while with the scouting data the sensitivity limit was pushed down to 500 GeV using scouting jets, with

trigger thresholds of approximately  $H_T > 360$  GeV and  $p_T > 180$  GeV for the single jet. Similarly, searches for low mass narrow resonances decaying to a pair of muon were performed using scouting muons in the lower end of the mass spectrum[107]. An inclusive set of results obtained using the scouting data can be found on chapter 3 of ref.[49]. The common component of all these results is that they would have not be possible without scouting data, since the event rate would have been too high to fit them in the standard data stream.

### Data parking

The parking strategy addresses the limit in the HLT output rate caused by the finite computational resources needed to perform the prompt event reconstruction. This is achieved by delaying the reconstruction, i.e. “parking” the raw data on tape storage, until the necessary computational resources are available. For instance, the reconstruction can be scheduled during long shutdown periods, or for the months during the end of the year technical stops. The main advantage of the parking strategy with respect to scouting is the fact that raw data are stored: the reconstruction is performed using the optimal set of calibrations and is less sensible to detector noise. This strategy was firstly introduced in Run1[108] with the goal of complementing the trigger, targeting for instance VBF topologies, Dark Matter searches and B physics measurements. This last point was extended during Run2, where B physics was the focus of the parking strategy. For instance, during 2018, events containing at least a displaced muon, e.g. from semileptonic B decay, and with a low  $p_T$  requirement were stored in the parking dataset allowing to record a large and unbiased sample of  $10^{10}$   $b$  hadron decays[109]. This enabled the testing of lepton flavor universality in  $B^\pm \rightarrow K^\pm \mu^+ \mu^-$  and  $B^\pm \rightarrow K^\pm e^+ e^-$  decays[110]. Previous triggers failed to provide good acceptance for final states containing electron pairs arising from rare  $b \rightarrow s e^+ e^-$  transitions due to the high  $p_T$  thresholds on single- and di-electron triggers (typically 30 and 20 GeV, respectively). Similarly, standard single-muon triggers, which typically require  $p_T > 20$  GeV and isolation from neighboring particles, decrease the acceptance for  $b \rightarrow \mu X$  decays. The data parking strategy allowed to successfully overcome these limitations.

## 4.2. THE L1-TRIGGER SCOUTING SYSTEM




Figure 4.2: Schematic view of the CMS Level-1 scouting system. Data flowing through the L1 trigger processor is intercepted before the Global Trigger final decision, and processed on a system independent from the standard DAQ and trigger chains. Trigger primitives are collected by a set of dedicated L1 Scouting boards and transmitted to a pool of computing nodes for further processing.

### 4.2 The L1-Trigger scouting system

The idea behind the L1 Trigger scouting is to develop a system capable of parasitically collecting and processing L1 Trigger objects as they flow through the trigger processor boards at the full LHC bunch crossing rate of 40 MHz. This new data-taking paradigm is conceptually similar to HLT scouting, where only the objects reconstructed at the trigger level are used. However, in this case, the objects reconstructed in the L1 Trigger are utilized. The L1 scouting system represents a parallel and independent readout and processing chain with respect to the data acquisition and trigger systems. Its independence means that it is not subject to the rate and latency constraints of the L1 and High-Level Triggers, as it does not need to operate synchronously with the central DAQ. A schematic view of such system is shown in Figure 4.2. A copy of the trigger primitives is collected using spare output links present in the trigger boards by a set of dedicated boards, using the same trigger protocol for data transmission. L1 scouting boards can eventually perform additional operations on the collected primitives, such as further processing without the latency constraints imposed



by the trigger or smart “zero-suppression”. Reprocessed data is subsequently transmitted to a pool of computing nodes, for short term storage and further analysis.

This innovative data scouting approach enables the access to physics processes that could evade the trigger chain due to, for example, L1-Trigger latency constraints or maximum available trigger rate. In addition, this system has the potential of enabling the exploration of the new exotic signatures, such as the study of correlation over several bunch-crossings, currently not possible with the L1 trigger. Moreover, alongside its physics potential, it offers the possibility of monitoring the detector and the trigger with a virtually unlimited statistics, hence enabling to monitor and studies scenarios which would be otherwise inaccessible. The main limitation for physics analysis comes from the fact that L1 trigger primitives are used for this set of studies. In fact, they are optimized to provide the best possible trigger performance while keeping under control the acceptance rate. This does not directly translate into the best possible physics performances because objects calibrations are not optimal for physics analysis. Similarly, undesirable effects could be introduced in the data by limitations of the processing hardware and algorithms or in the available bandwidth between trigger processors. These effects must be understood and studied in detail from simulations to explore the capabilities of the L1-scouting system. At the same time, trigger objects performance can be monitored to check that no evident limitations for the scouting are introduced, while keeping the best possible physics performance.

The L1 Trigger scouting system of the CMS experiment will leverage the enhanced Phase-2 upgrade of the L1 Trigger. Processes where full detector granularity is not required can potentially benefit from the large and inclusive data sample collected at the bunch crossing rate, despite the possible performance limitations of the L1 Trigger objects compared to offline objects. Some examples of such cases includes all-hadronic rare decays of Standard Model bosons, e.g.  $H \rightarrow \phi\phi$  and  $W \rightarrow \pi\pi\pi$ , have already been studied [111], proving that the L1 object resolution is sufficient for an analysis. Similarly, rare radiative  $W$  decays like  $W \rightarrow \phi\gamma$  have studied. Searches including soft hadronic final states can also benefit from the scouting approach. This includes standard dijet searches,  $X \rightarrow q\bar{q}$ , at low mass and high multiplicity unclustered hadronic final states (SUEPS) coming from the dark sector. This non-exhaustive list is currently growing in preparation of the Phase-2 upgrade, as new use-cases are frequently

## 4.2. THE L1-TRIGGER SCOUTING SYSTEM

added.

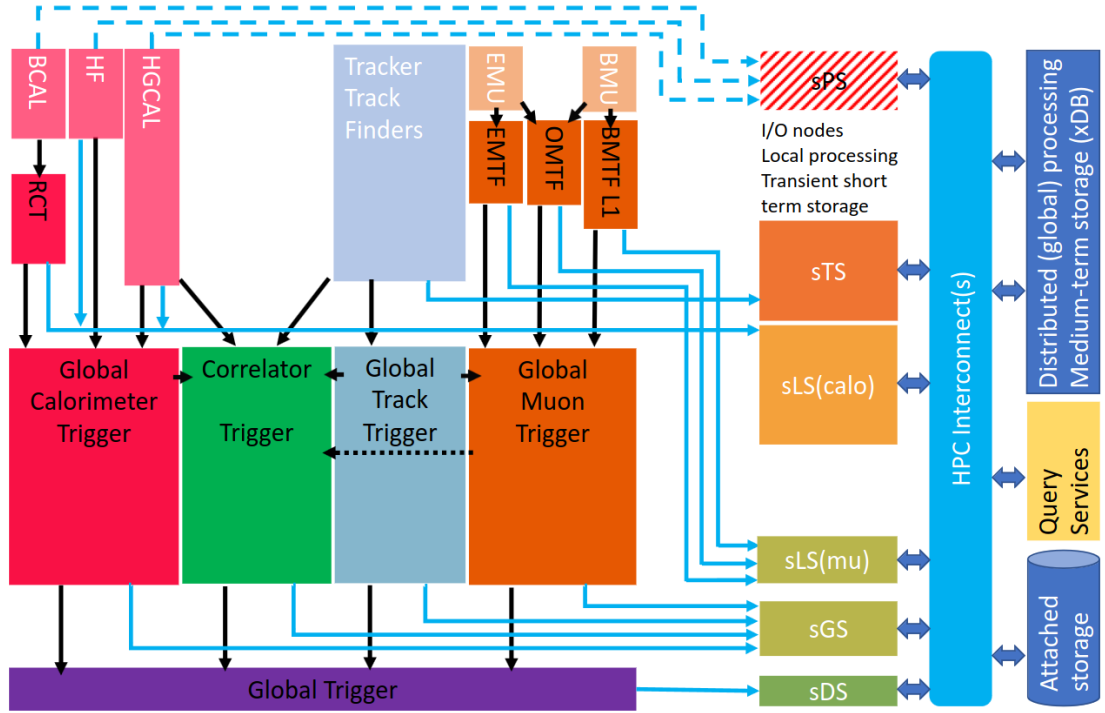


Figure 4.3: Functional architecture of the Phase-2 Level-1 trigger and scouting system. Scouting I/O will receive data collected from the Trigger processor and transmit them to a set of computing nodes via a high performance computing interconnect. Analysis results can be either transmitted to permanent storage or buffered locally for further processing. The system has a stageable design: the starting point includes the scouting of decision and global systems (sDS and sGS), while other trigger processors providing lower level information can be added in later stages.

A high level view of the proposed architecture is shown in Figure 4.3. The system is designed to be implemented in multiple stages: the baseline implementation aims at capturing data from the decision system (sDS) and global systems (sGS). The former consists of the results of the global trigger algorithms, along with information from the TCDS. On the other hand, the latter will acquire all the input objects provided to the final global trigger algorithms. This includes calorimeter and muon objects, tracker objects (such as Primary Vertex and global track quantities) and the products of the correlator layer-2 trigger particle flow-like reconstruction and pileup per particle identification (PUPPI), such as electrons and jets. This initial step could already provide trigger monitoring functionalities and good physics performance, while keeping

the throughput relatively limited thanks to the fact that higher level quantities are being used.

The system could be extended to capture the output of local systems (sLS), such as the regional muon track finders and calorimeter triggers. Similarly, provided the link availability, it could be possible to also capture the output of the tracking system (sTS). Finally, as shown in chapter 3, the raw data from the local sub-detectors could be injected into the scouting streams (sPS). The lower the information is read from the trigger chain, where the detector front-end is the limit, the larger the number of handles available in the scouting analysis. Indeed, the data underwent less processing, allowing for more complex and flexible reconstruction in the scouting analysis and enabling to work with a potentially unbiased data stream. However, this comes with the heavy price of a much larger throughput that needs to be managed from many aspects, ranging from the number of links to the increased computing power required.

In the proposed architecture, the system will collect trigger data via spare optical outputs of the L1-trigger boards, using the same 25 Gbps serial links and protocol used by the trigger itself and a set of dedicated FPGA-based boards, working as an interface between the synchronous trigger domain and the asynchronous scouting chain. Such boards are also in charge of performing pre-processing of the data, such as zero-suppression, before sending data to the attached I/O nodes. Currently, the board identified for this task is the so called DAQ800[112]. This custom aTCA board is being developed as a dedicated readout board for the CMS Phase-2 upgrade. It is capable of receiving up to  $48 \times 25$  Gbps input links, while providing an output maximum bandwidth of 800 Gbps via 100 Gbps ethernet links. In this configuration, around 7 DAQ800 boards would be sufficient for the implementation of the phase-2 scouting system, collecting data from the global and decision systems. The trigger primitives are sent from the DAQ800 to a set of commercial servers, using a custom firmware implementation of the TCP/IP protocol, where they are stored in memory before later stages of processing.

Currently many research and development projects are ongoing, with the aim of possibly improving the system. For example, more efficient protocols could be used for the communication between acquisition boards and servers, such as ROCE as introduced in chapter 3. Similarly, these boards could be based on more performant FPGA, allowing to run more complex processing on them before transmission to the servers for final processing. Given the



Demultiplexer (DEMUX in the following) and Barrel Muon Track Finder (BMTF) input muon stubs.

Two AMD Xilinx VCU128 development kits are used to receive the 10 Gbps trigger links. Both boards are mounted on a PCIe crate connected to a server, located in the CMS service cavern, used for configuration, control and monitoring. A total of 4 links are used to receive the up to 8 muons produced by the  $\mu$ GMT per each bunch crossing. The same VCU128 is used to receive the 24 links coming from the 12 BMTF boards, one per each CMS Wedge. Each BMTF processor is capable of sending to the scouting system up to 8 of the input muon stubs per “Wedge”, defined as one of the 12 azimuthal sectors into which the muon barrel is divided<sup>1</sup>. The second VCU128 receives 7 links from the Calorimeter trigger DeMux, where two of them are used to receive for each bunch crossing 12 jets, 12  $e/\gamma$ , 12  $\tau$  and one link for the full set of energy sums such as the missing transverse energy  $E_T$ . A detailed description of the primitive contents is available in the appendix (Table A.1, Table A.2 and Table A.3).

Additionally, it is possible to receive the final decision of the 512 algorithms in the  $\mu$ GT via 18 trigger links. This can be used to compare GT bits against alternative reconstructions that can be implemented in the scouting system, such as alternative usage of the BMTF stubs. This is not included during normal operations but is reserved for dedicated runs. While the trigger primitives produced by the  $\mu$ GT and Calorimeter trigger are in principle “high level” objects already viable for performing physics studies, the muon stubs collected from the BMTF provides access to “low level” quantities which can be used to test alternatives reconstruction methods. This has the goal of potentially improving the quality of the final muon, as the reconstruction can, for instance, be performed in software without the limitations imposed by the hardware.

After a zero-suppression step, data is transmitted from the VCU128 via TCP/IP to the server farm located on the surface via 100 Gbps ethernet links. Independent TCP/IP streams, one per data-type, are received and pre-processed by the scouting data acquisition software (SCDAQ), running on a set of servers called DSBU (Data Scouting Buffer Unit). These servers are acting as tempo-

---

<sup>1</sup>This correspond to a slice of barrel including the same DT sector number in all wheels

#### 4.4. DATA ACQUISITION

rary buffers for the incoming data, before being served to the Processing Units. The Processing Units (DSPU) are running the CMS reconstruction software[113] (CMSSW) and are, from a functional role point of view, equivalent to the High Level Trigger filter units. However, in the scouting system, Processing Units are also in charge of performing the “event building”. In other words, each Processing Unit aggregates raw data for the same event fetching raw fragments from all the DSBUS servers. After having merged all information from the same events, raw trigger primitives are processed and multiple data streams are produced as output of the Processing Units, following different processing criteria. These different streams will then form distinct analysis dataset for offline analysis. Products of the online processing are transferred to the CMS Tier-0 for offline storage and replication in other Tier-1 and Tier-2 sites. After this stage, users can finally access them for offline analysis.

To mark the demonstrator as successful, the full chain needs to be implemented and deployed, from the collection of trigger primitives at the full bunch crossing rate to the capability of running online processing to produce and distribute standard datasets for offline analysis. In the following sections of this chapter, the design and implementation of each component of this demonstrator system will be described in detail. Finally, the first results obtained from the data collected during the commissioning of the system will be presented.

### **4.4** Data acquisition

At its core, the L1 scouting system must be capable of collecting trigger data at the rate of 40MHz directly from the trigger board, and subsequently transfer this data to the computing nodes for further analysis. This demanding task requires robust hardware and software components to ensure seamless data acquisition and transmission.

In this section, we will describe the data acquisition components of the L1 scouting system, starting with a detailed explanation of the hardware components and firmware structure of the boards used to receive the trigger links. Following this, we will explore the software components essential for receiving the data streams and delivering them to the online processing system, ensuring that the entire data acquisition chain operates smoothly and efficiently.

### 4.4.1 Hardware

The hardware of the scouting system relies on two AMD XILINX VCU128 evaluation kits[114]. This evaluation kit was chosen because it houses the Xilinx Ultrascale+ VU37P FPGA, which is functionally equivalent to the Ultrascale+ VU35P mounted on the DAQ800 board, allowing for the reuse of firmware already developed for the DAQ800. Furthermore, the kits can be equipped with an FMC+ mezzanine hosting six QSFP slots, in addition to the four available on board, which collectively allow the reception of  $24 \times 10$  Gbps trigger links. In this configuration, the boards offer the same functionalities as half of a DAQ800 board. In other words, the two VCU128 kits together provide a similar number of links as one DAQ800 board.

The firmware structure of the boards is divided into three logical parts: the data receiver, the data processing pipeline, and the final output core.

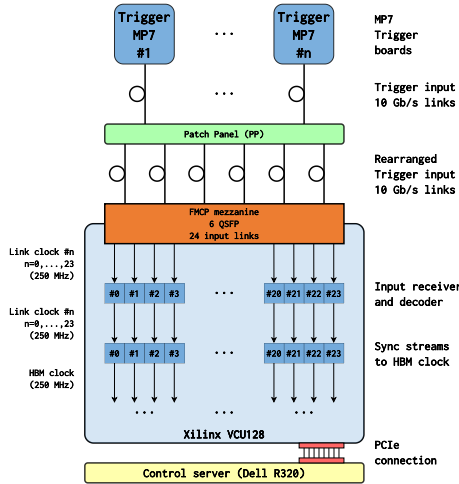


Figure 4.5: Schematic view of the input receiver logic of the L1 scouting firmware for the VCU128 and DAQ800.

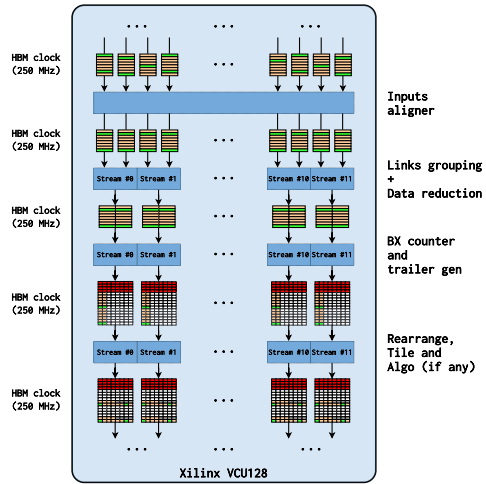


Figure 4.6: Schematic view of the processing pipeline logic of the L1 scouting firmware for the VCU128 and DAQ800.

#### Input data receiver block

The input logic of the scouting firmware is illustrated in Figure 4.5. The MP7 trigger boards [41] sends BX data to the scouting system at 250 MHz, where each BX block is a record of  $6 \times 32$ -bit frames for each input-link. Multiple BX blocks are grouped together into orbits, corresponding to a rate of around

#### 4.4. DATA ACQUISITION

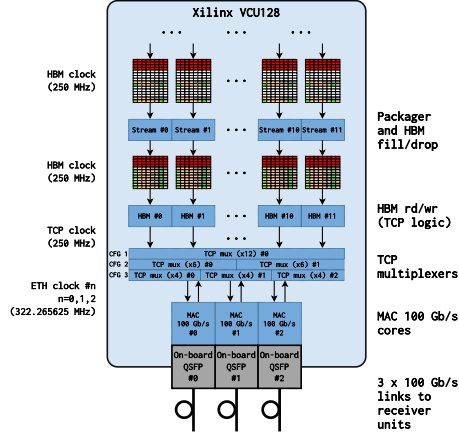


Figure 4.7: Schematic view of the orbit packager and HBM read/write logic of the L1 scouting firmware for the VCU128 and DAQ800.




Figure 4.8: Floorplaning on a VU37P FPGA for the BMTF stub receive firmware project.

11 kHz. A first receiver core decodes the data from each input link, received in the CMS Standard trigger protocol, recovering the clock independently for each data stream. A second step is needed to synchronize the data with the HBM clock, which is also working at 250 MHz.

#### Data processing block

In this step, schematized in Figure 4.6, data from the input block are firstly aligned BX-wise and arranged into groups. The groups are then zero-suppress by suppressing groups without any valid trigger primitive inside. Streams are subsequently padded to accommodate the 256-bit alignment of the HBM banks and, at the same time, BX blocks in an orbit are numbered. Finally, a trailer of 3564-bit is appended to the end of the orbit, carrying the information of the non-suppressed BXs. In the last stage of the pipeline, an algorithm can be applied on the trigger primitives, before rearranging the output and preparing it for the final firmware block.

#### Output block

The logic of the final step is illustrated on Figure 4.7. Data from the processing pipeline are encapsulated in a “scouting block”, which contains data from a configurable number of orbits. The occupancy of the HBM is monitored



Resource	Available	Utilization	Utilization[%]
LUT	1303680	199478	15.30
FF	2607360	328338	12.59
BRAM	2016	490	24.31
URAM	960	48	5.00
DSP	9024	6	0.07

Table 4.1: Firmware resource utilization for a VU37P FPGA.

by the packager module: in case of high occupancy, backpressure is handled by discarding the newly incoming orbits until the occupancy drops below a programmable threshold. Finally packages are written in the HBM, along with a scouting block header. The TCP/IP logic controls reading scouting blocks from the HBM and it allows to multiplex 4, 6 or 12 TCP/IP streams into a single 100 Gb Ethernet core, which is in charge transmitting the stream to the computing nodes.

### Resource utilization

The resource utilization for a firmware design with 24 trigger links receivers, 12 HBM banks, 12 TCP/IP streams and 3×100 Gb ethernet cores is lower than 30% of the available resources of a VU37P FPGA. The breakdown of the contributions is reported in Table 4.1 and floorplanning shown in Figure 4.8.

The reduced resource utilization leaves room for additional algorithms in the processing pipeline. For instance, given the high availability of DSPs, a neural network model deployed with `hls4ml`[115] could be used for tasks such as fitting the muon stubs for recalibration of the L1 muons. A preliminary study is described in [116].

#### 4.4.2 Software

The data acquisition software for the L1 scouting system, called SCDAQ (Scouting DAQ), is running on the Buffer Units with the task of receiving raw data from the scouting boards and storing it temporarily before online processing. The structure, illustrating components and their interactions, is depicted in Figure 4.9

SCDAQ is built upon the Intel Threading Building Blocks (TBB) library[117], which provides a robust framework for parallel programming. This allows

#### 4.4. DATA ACQUISITION




Figure 4.9: Illustration of the structure of the scouting DAQ software. It is logically divided into pipelines, each responsible for processing a single TCP/IP stream. The raw data processed by each pipeline is written as a file in the server’s ramdisk before being consumed by the Processing Units.

SCDAQ to efficiently handle multiple streams of data simultaneously, leveraging the power of parallel multithreaded pipelines. Each pipeline is tasked with receiving a single TCP/IP stream of data, processing it, and preparing it for further analysis.

The operation of each pipeline within SCDAQ can be logically divided into three main components. First, data from a block of orbits are written into memory slices and placed in a First-In-First-Out (FIFO) queue. This step ensures that incoming data is organized and ready for processing. Next, the processor assigned to each stream undertakes the task of reformatting and rearranging the data. The goal here is to simplify the complex data format produced by the hardware, making it easier to handle in subsequent stages of processing and more memory efficient. Finally, the processed data is handed off to an output module. This module is responsible for writing the results of each memory slice to the server’s ramdisk as raw files inside the ongoing run directory, which will then be consumed by the Processing Units.

To facilitate processing in subsequent steps, the raw data stream is divided into multiple files with a specific structure. A common counter across all dif-

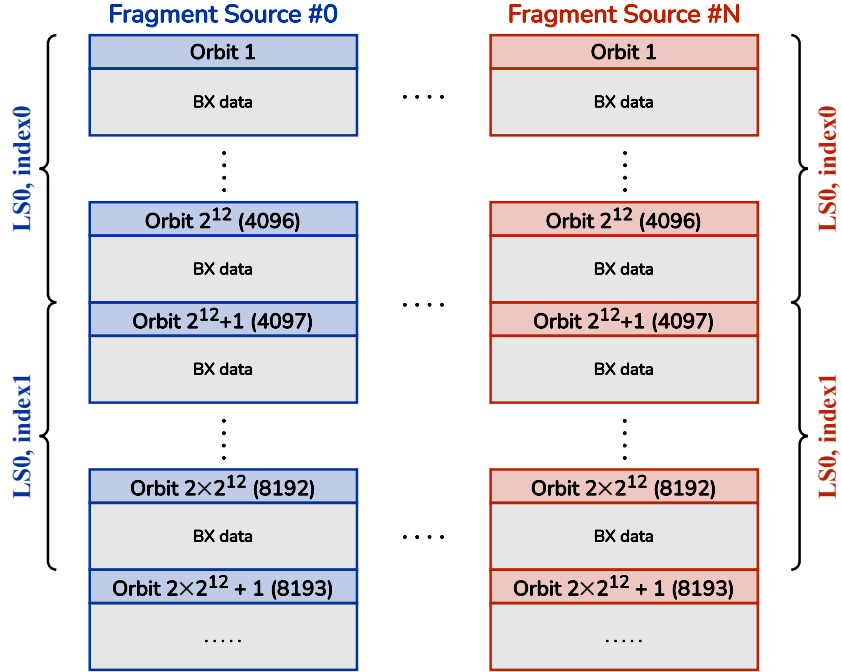


Figure 4.10: Illustration of the raw file structure. For each source, the data stream is split into multiple files based on the orbit counter.

ferent streams is the orbit counter, set to 1 at the beginning of the CMS run and incremented by 1 every 3564 BXs (i.e. every  $90 \mu\text{s}$ ). The orbit counter is the same in every trigger board, as they are synchronized, and is included in the packets sent to the scouting system. This consistency allows the data streams to be split uniformly amongst all the sources. Consequently, given a certain index, files from all the sources will contain data with the same range of orbit counters. In other words, a given range of orbit counters uniquely identifies a data taking time frame.

The time interval of a Lumisection is further divided into smaller chunks for more manageable data handling. Each chunk within a LS is identified by a unique index, resulting in files named using a pair consisting of the LS number and the index. For example, if the chunk size is set to 4096 orbits, a single LS will be divided into 64 blocks. Each file will then be named according to its block within the LS. Specifically, `LSX_index0` will contain data for orbits 1 to 4096, `LSX_index1` will contain data for orbits 4097 to 8192, and so on. This naming convention and structuring are applied consistently across all raw data fragments from all sources, ensuring a uniform approach to data segmentation.

## 4.5. ONLINE PROCESSING SYSTEM

This structured approach significantly aids in the efficient handling and processing of the raw data, allowing for easy indexing and retrieval based on the orbit counter and LS identifiers. This procedure is illustrated in Figure 4.10.

Each file contains a sequence of “Orbit headers”, each carrying information on the source which produced the data and the size of the orbits. Following each header, the file includes a raw data block for every BX within that orbit. The structure of the file is the same as the one produced by the central DAQ builder units, with the main difference being that a full orbit is treated as an “event” instead of the single BX. Moreover, data relative to a single stream is present, rather than a fully built event from multiple streams. Merging all the scouting fragments is the first step of the online processing system.

### 4.5 Online processing system

The goal of the online processing is to combine raw fragments from all the scouting sources, similarly to what the CMS Central DAQ (cDAQ) event builder does, and run processors and analyzers on them to produce new data streams. As introduced in the previous chapter, all these steps are performed *per orbit* at the fixed rate of 11 kHz, instead of a *per event* logic where data from a single bunch crossing is considered. From the scouting Buffer Units, data are read by a set of Processing Units in charge of merging and processing the scouting data, producing a new set of files in the format used by the cDAQ. From here, data are transferred to the CMS Tier-0 for repacking, i.e., data are converted from a raw format to ROOT files for persistence, and datasets are created based on the online data stream, before finally being made available to users for offline analysis. In the following section, these components will be discussed in detail.

#### 4.5.1 Orbit Building

As introduced in the previous section, the stream of trigger primitives is split consistently across all sources thanks to the orbit counter, which is common to every trigger board. This ensures that files from different sources with the same index contain raw data from the same range of orbits. In the system, one Buffer Unit is identified as the “main Buffer Unit”, where the HLT daemon (`h1td`) and `file-broker` services are running[112]. The `h1td` manages the lifecycle of a run, from detecting its start when run configuration files appear in the Buffer Unit

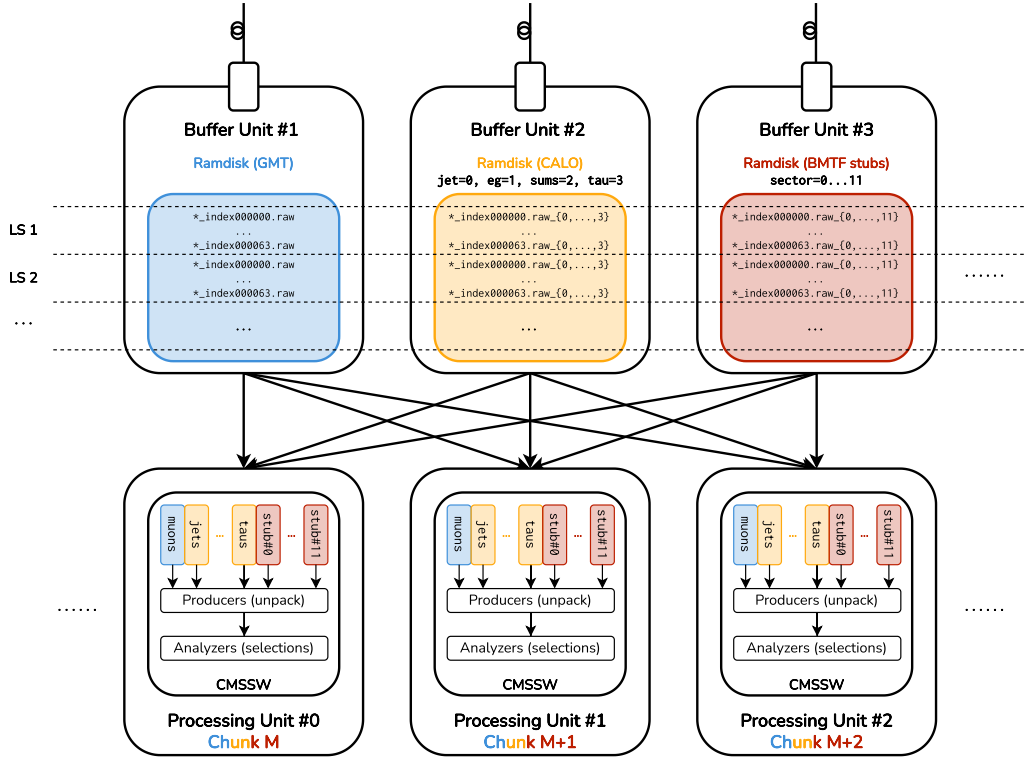


Figure 4.11: Illustration of the orbit building process. Each Buffer Unit is holding on its ramdisk subset of raw data files. Files with the same (LS, index) pair of every Buffer Unit are read through a network file system by the same Processing Unit into the CMSSW framework.

ramdisk to starting and monitoring the software processes on the filter unit (FU), based on the CMS software framework CMSSW. The file-broker monitors the buffer unit ramdisk for new raw files. Once a new file appears, it assigns it to one of the filter units connected to that buffer units for processing. All processes of the scouting Processing Units are connected to a single file-broker and notify it when they are available to process a new set of data. When a process is available, the file-broker sends it the name of one of the files that appeared in its ramdisk.

The process of orbit building is illustrated in Figure 4.11. Once a Processing Unit receives a file name, for example, `run123456_ls0001_index0000000.raw`, it must collect all the files with the same (LS, index) pair from all the Buffer Unit ramdisks. To achieve this, the ramdisks of the Buffer Units are accessible from the Processing Units through a network file system (NFS) remote mount. The CMSSW module responsible for reading raw data, the input source, needed to be extended as it was originally designed to read data from a single file produced

## 4.5. ONLINE PROCESSING SYSTEM

by a builder unit. The input source for the scouting data type is now designed to collect data from all the Buffer Units registered in the input source configuration, with filenames matching the one obtained from the file-broker.

The list of Buffer Units for the run where the file will appear is specified in the configuration of the input source, defined at the start of the run based on the scouting sources participating to the run.

Raw data are stored in a CMSSW collection called `SRDCollection`, which stands for Scouting Raw Data Collection. This naming convention follows the `cDAQ Fed Raw Data Collection (FRDCollection)`. This structure contains one buffer for each scouting data source stream. During the orbit building phase, raw data read from files are written into the corresponding slot of the `SRDCollection`, based on the source ID number stored in the event header, and are then made available for the next processing steps (Figure 4.12). Once the processing of a file index is complete, the Processing Unit will delete all files with the same (LS, index) pair from all the Buffer Units, thereby making room for new data in the ramdisk.

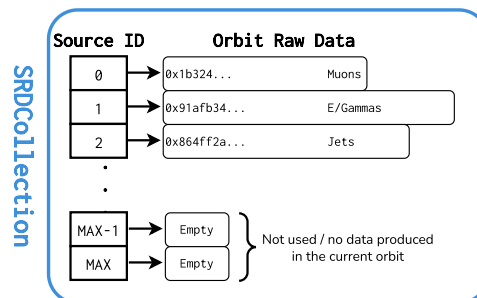


Figure 4.12: Illustration of Scouting Raw Data Collection. One buffer is available for each stream, and during the orbit building it is resized to accommodate raw data for from the source with the corresponding ID. If ID is not used by any source, or if a particular stream is not included in the run, the buffer will remain empty.

### 4.5.2 Online data streams

From the previous stage, the data reside in a collection within the standard CMS software framework. This allows for the implementation of plugins to perform various operations, starting from unpacking raw data to more complex selections and processing. The main difference compared to code running in

any other CMSSW job is that, although the processing in the framework revolves around the concept of events, the content represents data from an entire LHC orbit. Therefore, an “inner loop” on the individual bunch crossing is needed for each CMSSW-event.

The first set of modules in the dataflow is responsible for unpacking the raw data stored inside the `SRDCollection`, producing a new collection of C++ objects which can be used for analysis and persistence on disk. Each of these modules, based on a CMSSW Producer, retrieves the raw data for a set of Source IDs from the raw data collection, containing data for a full orbit. Then, the content of each BX is unpacked, and a new collection is produced and added to the event content. In the CMS software, all collections are based on the concept of events, so they are not suitable for holding the data of potentially all 3564 BX in the orbit. Therefore, a new data type called `OrbitCollection` was designed for this application. Since one collection of this type will be produced for each data type (e.g., Muon and Jet), it is essential that retrieving data from an arbitrary BX from each of them is quick, enabling correlation of arbitrary BX from different collections. Conceptually, this would be equivalent to having a collection containing two nested arrays, where the outermost contains one entry for each BX (i.e., 3564 entries), and the innermost contains the data for a specific BX. However, the same result can be achieved using two flat vectors: one holding all data objects and another containing a set of indexes allowing retrieval of objects from the data vector for a specific BX, called the `Index` vector. This option, although less intuitive, is preferable from a data-persistency point of view. The idea behind this is as follows: the `Index` vector has one entry per BX and contains the starting position index of the objects related to that BX stored in the data vector. For example, as illustrated in Figure 4.13 for the muons, if the entry for  $BX_1$  in the index vector has the index counter 0, and the adjacent  $BX_2$  contains the index counter 2, this means that the 1st and 2nd objects in the data vector are from  $BX_1$ . Generally, the content of  $BX_N$  will be stored in the indices of the data vector ranging from `Index[BXN]` to `Index[BXN+1]`.

Each unpacker module produces an `OrbitCollection<T>`, where T can be a scouting Muon, Jet, or other scouting data formats. Although existing data types for trigger objects, such as `l1t::Muon`, could have been used, new data types were designed for each scouting source. The primary reason for this was that persisting objects from the existing trigger collections resulted in larger file sizes. This was not unexpected, as these types were not designed for persistence; they

#### 4.5. ONLINE PROCESSING SYSTEM

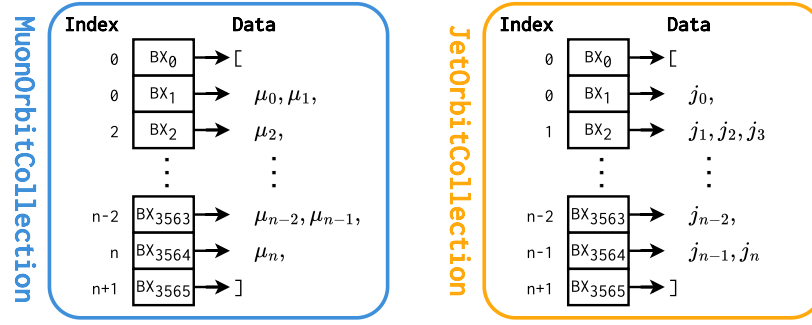


Figure 4.13: Illustration of the `OrbitCollection` for scouting Jets and Muons. Two vectors (Index and Data) are used to organize the unpacked objects.

included many quantities as both integer hardware quantities and floats, along with others that were never filled or used. This led to the design of custom data types, such as `11ScoutingRun3::Muon`, which store only the relevant information as hardware types. This new data format allowed for savings of between 30% and 40% of disk space compared to the default collections. Additionally, objects in the standard L1 trigger format can be easily obtained during analysis, thanks to a set of functions implementing the conversion between hardware and physical quantities.

After the unpacking process, multiple data streams can be formed. The first one is called the `ZeroBias` stream, which contains a prescaled fraction of all the Orbit Collections produced by the unpackers. For example, a prescale of  $N$  means that one out of every  $N$  full orbits is saved. An example of the event content of this collection is reported in Table 4.2.

Collection name	Average Size (kBytes/Orbit)	
	Uncompressed	Compressed
<code>TauOrbitCollection</code>	422	60
<code>EGammaOrbitCollection</code>	238	35
<code>BxSumsOrbitCollection</code>	226	30
<code>JetOrbitCollection</code>	103	15
<code>MuonOrbitCollection</code>	38	4
<code>BMTFStubOrbitCollection</code>	22	2

Table 4.2: Average uncompressed and compressed size (in kBytes/Orbit) for the scouting collection in the `ZeroBias` data stream computed with data from a set of lumisections at peak luminosity. The total compressed size for all the BXs in the orbit is 146 kB.

Since no selection is applied and the content of all bunch crossings is retained,



data from this stream is particularly useful for studying the characteristics and monitoring the L1 Trigger. This includes possible interactions among objects from neighboring bunch crossings, as a true unbiased stream of data is recorded.

However, saving all the collections for all the orbits, i.e., using a prescale of 1, would saturate the bandwidth dedicated to the L1 Scouting, which amounts to 300–350 MB/s. Indeed, storing all orbits would require saving  $147 \text{ kB} \times 11 \text{ kHz} \approx 1.6 \text{ GB/s}$  of data on disk. Furthermore, data from a set of loose selections are more manageable from an analysis point of view, as the ZeroBias stream contains a large number of BX with pure noise which will be most likely discarded.

For this reason, a second data stream, called `SelectionStream`, has been created. This data stream stores all the bunch crossings selected online by some loose criteria, without a prescale in the orbits, complementing the ZeroBias stream. This is implemented in the following way: a set of modules read the orbit collections of unpacked data and produce, for every orbit, a list of bunch crossings passing the selection run by that module. One example of such a module is the `JetBxSelector`, which selects bunch crossings based on jet conditions. An example is shown in Listing 4.1:

```

1  process.DijetEt30 = cms.EDProducer("JetBxSelector",
2      jetsTag          = cms.InputTag("l1ScCaloUnpacker", "Jet"),
3      minNJet          = cms.int32(2),
4      minJetEt         = cms.vdouble(30, 30),
5      maxJetEta        = cms.vdouble(99.9, 99.9)
6  )

```

Listing 4.1: Example of instantiation of a `JetBxSelector` module. In this example, the module is consuming jets produced by the `l1ScCaloUnpacker` module and selecting a bunch-crossing if it contains at least two jets with  $E_T > 30 \text{ GeV}$ , with no constraints on jets  $\eta$ .

At the time of writing this thesis, the following `BxSelector` modules are running in the scouting high-level processor:

- `DijetEt30`: requires at least two jets with  $E_T \geq 30 \text{ GeV}$
- `HMJetMult4Et20`: requires at least four jets with  $E_T \geq 20 \text{ GeV}$
- `SingleMuPt0BMTF`: requires at least one muon produced by the BMTF
- `DoubleMuPt0Qual8`: requires at least two muons with quality  $Q \geq 8$
- `MuTagJetEt30Dr0p4`: requires at least one “muon-tagged” jet, i.e., a jet and a muon such that  $\Delta R(\mu, j) \leq 0.4$  and jet  $E_T \geq 30 \text{ GeV}$

#### 4.5. ONLINE PROCESSING SYSTEM

The products of the `BxSelector` modules are then processed by a module named `FinalBxSelector`, which performs the logical-OR operation of all the bunch crossings selected by each module and produces the list of bunch crossings selected by at least one module. An example of instantiation of this module is shown in Listing 4.2.

```
1 process.FinalBxSelector = cms.EDFilter("FinalBxSelector",
2   analysisLabels = cms.VInputTag(
3     cms.InputTag("DijetEt30", "SelBx"),
4     cms.InputTag("SingleMuPt0BMTF", "SelBx"),
5     ...
6   )
7 )
```

Listing 4.2: Example of instantiation of `FinalBxSelector`. The list of selected BX by each analysis module to be considered in the logical-OR operation is provided as input tag.

The list of selected bunch crossings is then used to produce the output collections, where only the selected BX are kept. This is implemented with a class called `MaskOrbitBx<T>`, which consumes an orbit collection of objects `T` and a list of BX, producing a new one where only the selected bunch crossings are filled. These are the final collections that will be persisted on disk, forming the selection stream. The event content of the selection stream for the same data considered for the `ZeroBias` stream is reported in Table 4.3.

At analysis time, both online and offline, the events selected by a particular analysis can be retrieved by processing only the bunch crossings of the final orbit collections that are present in the list of BX produced by a certain analysis, such as `DijetEt30`. This approach can also be used to produce “datasets” based on the set of selections used to create that data stream.

Despite the use of a loose set of selections, the throughput is approximately  $15 \text{ kB} \times 11 \text{ kHz} \approx 160 \text{ MB/s}$ , which is a factor of 10 less than the `ZeroBias` stream. This allows the storage of all the selected BX without a prescale. It complements the `ZeroBias` stream, as it can be used to study processes where a high number of events is needed. The prescale of the `ZeroBias` stream is then chosen to fill the available bandwidth left by the selection module. Usually, a prescale of around 15 is used and it is important to note that the throughput is not constant during the LHC fill. The numbers reported in the previous tables are computed during peak luminosity conditions, but after the luminosity leveling phase ( $\beta^*$

Collection Name	Average Size (kBytes/Orbit)	
	Uncompressed	Compressed
EGammaOrbitCollection	43	5
BxSumsOrbitCollection	35	4.1
JetOrbitCollection	31	3.3
MuonOrbitCollection	20	1.2
BMTFStubOrbitCollection	18	0.7
FinalBxSelector_SelBx	0.9	0.2
DijetEt30_SelBx	0.6	0.1
HMJetMult4Et20_SelBx	0.4	0.1
SingleMuPt0BMTF_SelBx	0.1	0.04
DoubleMuPt0Qual8_SelBx	0.1	0.04
MuTagJetEt30Dr0p4_SelBx	0.02	0.007

Table 4.3: Average uncompressed and compressed size (in kBytes/Orbit) for the scouting collection in the Selection data stream computed with data from a set of lumisections at peak luminosity. The total compressed size per orbit is 15 kB. List of selected BXs by each BxSelector module are store and can be used for the offline analysis

leveling[92]) of the LHC, where the luminosity is maintained constant before starting to decrease due to beam burn-off.

Finally, while the current implementation exclusively runs selection modules, the system is inherently designed to accommodate the integration of additional modules that generate higher-level quantities. This flexibility allows for more sophisticated analyses by combining, for example, multiple jets or muons into composite candidates. For instance, it could be foreseen the implementation of a module that identifies pairs of muons originating from the decay of a Z boson, and store only such candidate instead of the pair.

### 4.5.3 Online data handling

Given that an architecture similar to the HLT was used, the infrastructure needed to manage the online data flow from this point was inherited from the central DAQ system. In a manner equivalent to the HLT, data from each CMSSW process and stream are written to separate files, one per input file index, including the objects and collections specified in the output module of each stream. Objects in the scouting orbit are stored on the local disk of the Processing Unit in a custom streaming format, called “streamer files”, which is dependent on the software release used to produce it and hence not meant

#### 4.5. ONLINE PROCESSING SYSTEM

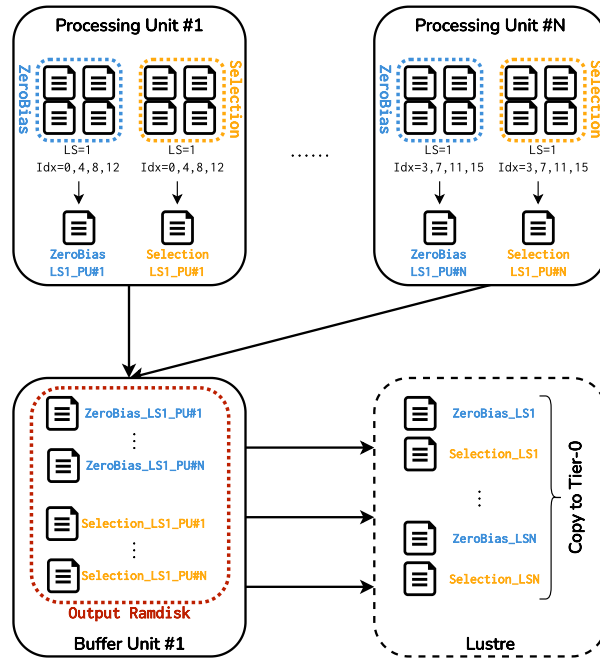


Figure 4.14: Illustration of the output files merging procedure. Local products from a Processing Unit are merged based on their Lumisection (LS) and output stream. These merged files are then transferred to the output ramdisk of a Buffer Unit. At this stage, files from different Processing Units but associated with the same stream and LS are further merged and copied to a Lustre storage area. In the final step, one file per stream and Lumisection is created and transferred to Tier-0 for additional processing.

for long-term storage. The main idea behind this data format was to provide a convenient way of concatenating files from multiple processes into a single file without the need for potentially costly merging by a ROOT-based executable.

First, a local merging step is performed in every filter unit, combining the output of all its processes per stream and lumisection. From the Processing Units, merged files are copied to the NFS-mounted output partition on a Buffer Unit. At this stage, additional merging steps are needed to concatenate data related to the same lumisection and stream from all the Processing Units, which are now located in the Buffer Units. These operations are performed by the Storage and Transfer System (STS) using a cluster file system based on Lustre[118], as concurrent access to a mass storage system from multiple locations is required. Once all the PUs have copied the file set to the Buffer Units, a set of workers read their content and append it into a single file in the distributed file system.

Finally, concatenated files can be transferred to Tier-0 for offline processing. An illustration of the merging chain is available in Figure 4.14. At the beginning of the 2024 data-taking period, the central STS was modified to include the L1 scouting streams. Besides merging the data, the STS is also responsible for the bookkeeping of all files produced and copied to Tier-0 from CMS. This means it needs to track files from both the cDAQ and Scouting systems. Since the two systems are independent by design, it was not possible to directly reuse the STS used in cDAQ. To address this, a common database was introduced to facilitate communication between DAQ and Tier-0. Thanks to this development, the L1 scouting data stream can now be handled by Tier-0.

#### 4.5.4 Offline data processing and distribution

The primary operation performed with scouting streamers is known as “Repacking”. Given the nature of the streamer format, there is no guarantee of backward and forward compatibility of serialization/deserialization software, making these files unsuitable for long-term storage. The repacking job addresses this issue by reading the streamer files and converting them into the ROOT-based CMSSW EDM format (Event Data Model). This conversion ensures that the data can be safely stored on disk. During the repacking process, primary datasets are created. In the CMS experiment’s computing infrastructure, Primary Datasets are used to group together events selected by the High-Level Trigger with similar physics content[119]. These datasets are then distributed to Tier-1 and Tier-2 sites for users to access the data. The ZeroBias stream is mapped into the L1Scouting dataset, while the SelectionStream is mapped into the L1ScoutingSelection dataset. The data-tier used for these datasets in the CMS computing model is L1SCOUT, which was introduced at the beginning of 2024 alongside the HLTSCOUT data tier.

Unlike other data streams, scouting data does not require additional processing steps such as “Prompt Reconstruction” or “Prompt Calibration”. Repacked data is typically available to users within 24 hours after the end of the CMS run. These files, along with their metadata, can be accessed through the CMS Data Aggregation Service (DAS)[120] and processed using the WLCG computing resources[121], similar to any other CMS dataset.

A tutorial with example code demonstrating how to access and process scouting data is available for CMS users[122]. Although this approach provides

significant flexibility, it requires writing analyzers using the CMSSW framework. Currently, this is the initial step in analyzing L1 scouting data, where custom analyzers are used to extract data and store it in ROOT Ntuples, which are easier to handle and are framework-agnostic. To mitigate the limitation of being bound to the CMSSW framework and skip this step, efforts are underway to enable the storage of L1 Scouting data in the CMS NanoAOD format[123]. In this format, data are stored in standard C++ types within a ROOT TTree, making them accessible independently of CMSSW. This initiative aims to simplify data handling and broaden accessibility for analysis.

## 4.6 Operation and performance

The L1 scouting demonstrator system described in this thesis has been successfully operating since the beginning of the 2024 data-taking period, collecting data seamlessly during a large part of this time. The ZeroBias stream was deployed from the start to debug and verify system functionalities. A test version of the selection stream was later added in May 2024 and fine-tuned by June 2024. At the time of writing this thesis, in September 2024, around  $90 \text{ fb}^{-1}$  have been collected with this system in the selection stream, comparable to the  $100 \text{ fb}^{-1}$  collected by the CMS experiment, as the demonstrator continuously operated in parallel. The corresponding number for the ZeroBias is approximately a factor of  $\sim 10$  lower, due to the prescale.

During peak luminosity, the scDAQ system received approximately 6.5 GB/s of raw data from the scouting boards. The largest fraction of this data is produced by the calorimeter trigger, with the  $\tau$  TPs being the most significant contributors. In contrast, the muon data had a much lower and manageable rate, as expected. Within the scDAQ processors, another layer of zero suppression is performed by writing only valid primitives to the raw files, along with rearranging them more efficiently. While these operations could be performed at the firmware level, the current throughput is manageable by scDAQ, making further developments unnecessary. This reduces the stream of files to be processed by CMSSW processors to around 1.5 GB/s, well below the capabilities of the processing system.

Indeed, to ensure the system's robustness, a data generator was used to stress test the system by maintaining the ramdisk of Buffer Units at a constant occupancy of 60%, with four Processing Units constantly consuming data from




Figure 4.15: Plots from the monitoring system of the “filter-farm” (F3MON[124]). Top panel: Output throughput for streamer files for each stream as a function of the lumisection. Bottom panel: Overview of the CPU time devoted to various processes, indicating a consistent resource utilization over time with approximately 75% of the time spent in the Idle state

it. The results demonstrated that each filter unit could consume nearly 2.8 GB/s, which is around 90% of the available 25 Gbps line rate. It is important to note that during this test, only the input source was tested, reading data from two input streams and discarding the output at the Processing Unit level.

Currently, with the existing scouting processing “menu”, only a small fraction of the available processing resources are utilized, between 20 – 25%, leaving significant room for the development and implementation of complex algorithms and online analysis. This is illustrated in Figure Figure 4.15, where the top panel shows the output throughput for streamer files for each data stream as a function of the lumisection. The throughput decreases over time, following the instantaneous luminosity profile. The bottom chart provides an overview of the CPU time devoted to certain processes, showing no apparent trend, which indicates that resource utilization remains constant over time. The light-green area of the plot, occupying around 75%, corresponds to the Idle state, where the process is waiting for a new input file.

The fraction of CPU used by each module to process one event in CMSSW can be visualized using the “Circle” utility[125], developed to visualize the timing of the High Level Trigger modules. One such plot is shown in Figure 4.16, focusing



#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.16: Fraction of CPU used by each CMSSW module to process an event, with a zoom on the L1Scouting section. The total time used for processing an orbit is around 1 ms, where an equal amount of 0.4 ms is occupied by the L1Scouting and Output modules.

on the contribution from the scouting modules. The total CPU time needed to process an orbit is around 1 ms, with 0.4 ms each occupied by the L1Scouting and Output modules. Among the scouting modules, the calorimeter unpacker consumes the most time due to the larger volume of raw data it processes compared to the Muon and BMTF data. Conversely, the online selection modules have an almost negligible impact on CPU usage.

### 4.7 First data from the L1 scouting demonstrator

The goal of this section is to present the characteristics of the data collected by the Level-1 Trigger Data Scouting demonstrator during the 2024 LHC Run 3 data-taking period. The main objective is to verify the quality of the trigger objects collected and processed, ensuring that their distributions match the expected ones. For instance, analyzing the distribution of bunch crossings where trigger primitives were produced and comparing it against the LHC filling scheme helps to check if the entire processing chain is functioning correctly.

Moreover, the quantities stored in the scouting collections can be compared with the same objects in the raw data of the cDAQ for triggered events. The trigger primitives used by the global trigger algorithms, specifically the muons



and calorimeter objects collected by the scouting system, are also present in the raw data sent to the High-Level Trigger for selected events. This allows to cross-verify the data between the two systems. For instance, by taking an event from the cDAQ ZeroBias dataset, it is possible to find the corresponding event in the L1Scouting dataset by matching lumisection, orbit number and BX, and checking if the hardware values of the objects match those recorded in the scouting data. This verification process has shown no mismatches in millions of events checked. Additionally, comparisons of the physical floating-point quantities have also revealed no differences, indicating that the data collected by the scouting system, as well as the scales used in conversion to physical quantities, are accurate.

With this verification, we can proceed to examine the standard distributions and features of the objects to ensure they align with expectations. Thanks to the large number of events available in the L1Scouting stream provides an excellent tool for data quality monitoring, not only for the trigger itself but also for the detector and LHC conditions. This section will present the distributions of scouting objects using data from the L1Scouting dataset, which offers a truly unbiased view of the data.

### 4.7.1 Muons

The  $\mu$ GMT system transmits up to eight muons per bunch crossing to both the  $\mu$ GT and the scouting system. The 8 muons are selected based on a sorting algorithm that combines muon quality and transverse momentum ( $p_T$ ) and integrates candidates from all three regions of the muon system: the Barrel Muon Track Finder (BMTF), Overlap Muon Track Finder (OMTF), and Endcap Muon Track Finder (EMTF). Further details on the scales and definitions of the trigger primitives can be found in Table A.1.

A first validation check is to examine the rate of bunch crossings (BX) that contain at least one muon from the BMTF. This corresponds to the rate of BXs selected online by the `SingleMuPt0BMTF` criterion. In Figure 4.17, the rate of selected BX collected by the L1 Scouting system is plotted alongside the instantaneous luminosity delivered by the LHC to CMS, as well as the rate of events selected by the L1 trigger (L1A). The rate of BXs with muons collected via L1 Scouting closely follows the profile of the instantaneous luminosity throughout the entire fill duration. During the luminosity leveling phase, where the instan-

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.17: Rate of BXs containing at least one BMTF muon (Blue), compared with the instantaneous luminosity delivered to CMS (Yellow) and the rate of events accepted by the L1 trigger (Violet). The data correspond to LHC fill 9969, which started around 14:00h on August 2nd and lasted until approximately 14:30 hours.

taneous luminosity delivered to the experiment is maintained constant, the L1A rate remains steady at approximately 100 kHz and follow the same trend. However, during the luminosity decay, the L1A rate begins to diverge from the other two rates. This divergence occurs because prescales are applied to the trigger algorithms, adjusting the thresholds to maintain a relatively constant L1A rate. For example, trigger thresholds for data parking are gradually relaxed to record parking data. It is important to note that the scouting data profile consistently mirrors the instantaneous luminosity, in contrast to the L1 trigger rate. This is because the scouting system is unaffected by trigger rules and prescales. For instance, during luminosity leveling, the rate of BXs containing BMTF muons is around 340 kHz, about three times the maximum rate of the L1 trigger.

This characteristic of scouting data could potentially be used for online measurement of instantaneous luminosity. The CMS BRIL (Beam Radiation Instrumentation and Luminosity) project is exploring the possibility of using muon trigger primitives collected from the  $\mu$ GMT as an additional online luminometer. In Figure 4.18, the rate of BXs with BMTF muons is shown during the emittance scan[91] conducted at the beginning of LHC fill 9969 (shown in Figure 4.17). During an emittance scan, the two beams are moved relative to each other in




Figure 4.18: Emittance scan performed at the beginning of LHC fill 9969 (Figure 4.17), visible as a drop in the rate during the initial phase of the fill before reaching peak luminosity. A second emittance scan was conducted at the end of the same fill.

steps along the  $x$  and  $y$  planes, resulting in the two bell-shaped distributions visible in Figure 4.18. The peak rate occurs when the beams collide head-on, i.e., with zero transverse displacement.

Muons from BMTF were considered as they exhibit a more stable behavior with respect to time than the ones produced by the EMTF. Indeed, in the endcap periodic hard-resets are required by the muon detectors, visible as sudden drops in the rate of muon trigger primitives. In Figure 4.17 they are visible in the trigger rate as periodic drops of a few kHz of the L1A rate (violet). These periodic resets are not observed in the barrel region, making the BMTF data more reliable for continuous monitoring. The larger drops in the rate are caused by backpressure or errors in other subsystems, which blocks the trigger consequently.

Similarly, the rate of GMT muons produced per bunch crossing in the orbit can be computed and analyzed. A higher rate is expected for colliding BXs compared to non-colliding ones, which are populated by cosmic muons, beam halo, and machine-induced background. These effects are visible in Figure 4.19, where the bunches colliding in CMS are marked with a black line in the bottom panel. The rate of muon primitives in these bunch crossings is multiple orders of magnitude higher than in non-colliding ones. When focusing on the BMTF muons (blue in Figure 4.19), the rate for empty bunches is flat, as expected from

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.19: Rate of GMT muons as a function of the bunch crossing in the orbit during LHC fill 9969, showing the contribution from each regional track finder of the muon system. In the bottom panel, the bunches colliding at IP5 (CMS) are marked with a black line.

cosmic muon contributions. An exception is the increase in the rate visible in the first bunches of the orbit, corresponding to muons produced by bunch trains colliding at IP2 and IP8 (corresponding to the Alice and LHCb experiments, respectively). On the other hand, the endcap muon system experiences a higher rate of particles, including beam-halo muons, well visible before and after a train of bunches. A detailed characterization of cosmic and beam-halo muons is presented in a dedicated Detector Performance Note[126].

Similarly, the distribution of muon multiplicity per BX can be analyzed, as shown in Figure 4.20. As expected, the distribution follows an exponentially decaying pattern with respect to the number of muons per BX. On average, a muon is produced in approximately 1 out of every 10 BXs. In Figure 4.21, it can be observed that the main contributors to the muon rate are the two endcaps. The muon rate remains nearly constant in the central region of the detector and increases exponentially for larger values of  $\eta$ , up to  $\eta = 2.4$ , where the muon system's coverage ends.

The distribution of muon transverse momentum  $p_T$ , shown in Figure 4.22, shows that each track-finder uses different binning schemes. For instance, the




Figure 4.20: Multiplicity of GMT muons per BX computed for LHC fill 9969. The rate in Hz is shown at the top of each bin, while the fraction of BXs with a given multiplicity is displayed at the bottom.




Figure 4.21: Distribution of the GMT muons'  $\eta$  extrapolated to the vertex, computed for LHC fill 9969. The contribution from each muon system region is shown.

OMTF adopts larger binning for higher  $p_T$  values to optimize triggering performance and FPGA resource utilization. Additionally, in all regions, the saturation value is set at 256 GeV, while the lowest possible value varies, starting from 4 GeV for the BMTF. From this distribution it is clear that the momentum cannot be used as it is for analysis purposes. Again, this is a consequence of the fact that these hardware quantities were designed with triggering purposes in mind. Attempts to improve this by recalibration L1 trigger muon were performed, leading to promising results[116].

Nevertheless, it is possible to compute physics observables, such as the invariant mass of muon pairs  $m_{\mu\mu}$ , using the muons collected by the GMT. An example is shown in Figure 4.23, where muon pairs reconstructed by the BMTF are considered. The muons are required to be produced by the Barrel Muon Track Finder, have a transverse momentum  $p_T \geq 10$  GeV, opposite charge, and an impact parameter bit  $d_{xy}^{L1} = 0$ . This implies that the standalone muons are reconstructed with a displacement  $d_{xy} \leq 32$  cm. The data are compared to a Drell-Yan (DY) Monte Carlo simulation, including the emulation of the L1 trigger reconstruction, where the muons are selected in the same way as in the data. The data are then fitted using a combination of a power-law function (p-law 3), a second-order polynomial (pol2), and a Crystal-Ball (CB) function for the resonant component. A small difference between the MC and data in the peak mean ( $\Delta\mu$ ) is related to different calibrations available in the MC. Additionally, the dif-




Figure 4.22: Distribution of the muon transverse momentum  $p_T$  of the GMT muons, with contributions from each muon track-finder region highlighted.

ference in normalization between data and MC may indicate that scale factors need to be introduced for instance to account for differences in muon reconstruction efficiencies between the two samples. Finally, the width of the resonance is broader than the one obtained by using offline objects. This is mainly due to the fact that in this range of transverse momentum, i.e.,  $p_T \approx 45$  GeV, the resolution  $\Delta p_T/p_T$  is approximately ten times worse when using standalone muons compared to when the tracker is included. This resolution will significantly improve in the Phase-2 system, where tracks from the track-trigger will be available to the muon trigger.

#### 4.7.2 Calorimeter objects

From the calorimeter trigger layer-2 demultiplexer, up to 12 jets, 12  $e/\gamma$ , 12  $\tau$ , and a set of energy sums are transmitted to the global trigger every BX for the final decision. Details on the definitions and scales used for the trigger primitives produced by the calorimeter are provided in Table A.2.




Figure 4.23: Invariant mass distribution of opposite-sign muon pairs reconstructed by the BMTF. The data are compared to a Drell-Yan Monte Carlo simulation, where the L1 muons have been produced through the L1 Trigger software simulation.

### $e/\gamma$ primitives

The effectiveness of the scouting system as a detector monitoring tool became evident during one of its first runs in early March 2024. During such run, the LHC was in a circulating mode, meaning no collisions occurred at the interaction points when the beams were present. Consequently, when observing the number of trigger primitives per BX, a flat distribution is expected, as only trigger primitives from noisy trigger towers would be produced and collected. However, when examining the distribution of  $e/\gamma$  and  $\tau$  primitives, an anomaly was found in the  $e/\gamma$  distribution, as shown in the left panel of Figure 4.24.

It can be observed that the distribution is generally flat, except for a single bin at the beginning of the orbit, corresponding to BX 24. This discrepancy is

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.24: BX occupancy of  $e/\gamma$  trigger primitives during a run without collisions (circulating run). The right panel shows only the primitives with an isolation value of 0.

enhanced when focusing only on trigger primitives with an isolation<sup>2</sup> value of 0, indicating that the isolation cut around the primitive cluster was failing. In other words, a large number of trigger towers were active in the same region simultaneously.

The source of this issue became apparent when examining the time evolution of the  $\phi$  and  $\eta$  variables of the trigger primitives produced at BX 24, shown in the top and bottom panels of Figure 4.25, respectively.

The  $\phi$  distribution reveals a clear pattern, scanning the range  $[-\pi, \pi]$ , along with constant horizontal lines produced by noisy towers. This pattern is localized in the central region of the detector and has a periodicity of about 40 minutes. This behavior is consistent with the operation of the ECAL monitoring laser[128], which periodically checks the transparency of the crystals by illuminating around 30 Towers at the same time, where each consists of  $5 \times 5$  crystals, to derive corrections. This explains both the observed periodicity and the non-isolated nature of the TPs. However, the laser is designed to fire in a specific BX located at the end of the orbit, which is then masked at the trigger level. Therefore, it should not be visible in BXs where the beam could be present. During the end-of-year shutdown between 2023 and 2024, a new calibration laser was

<sup>2</sup>The L1 isolation for  $e/\gamma$  primitives is a cut on energy in an isolation region around the cluster, considering both ECAL and HCAL towers[127]






Figure 4.25: Time evolution of the  $e/\gamma$  trigger primitives'  $\phi$  (top panel) and  $\eta$  (bottom panel) variables. The time variable is binned into lumisections, each lasting approximately 23.3 s.

installed in preparation for the HL-LHC. This new system produced reflections that illuminated unwanted BXs, which were observed and addressed during its commissioning. However, the issue observed with the scouting system went unnoticed due to the extremely low energy of the primitives, only visible in the central region of the detector where  $E_T$  is maximized and above the threshold needed to produce a primitive. Following such observations, further studies were conducted, concluding that the final event reconstruction is not impacted by such issue.

This example highlights the power of the scouting system, as it makes no assumptions or selections on the collected data, enabling the creation of an unbiased dataset suitable for anomaly detection applications.

When analyzing a collision run, the BX occupancy of the trigger primitives reflects the LHC filling scheme, as shown in Figure 4.26. Unlike the muons, the rate of trigger primitives produced in empty BXs is much higher for the calorimeter systems. This is because they are more susceptible to detector activation effects due to the large rate of particles impacting them, being closer

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.26: Rate of  $e/\gamma$  trigger primitives as a function of bunch crossing in the orbit during LHC fill 9969, showing contributions from both ECAL regions. In the bottom panel, the bunches colliding at IP5 (CMS) are marked with a black line.

to the interaction point with less material in between. No significant difference is observed between the behavior in the Barrel and Endcap regions.

When examining the trigger primitives' multiplicity per BX, it is important to differentiate between colliding and empty BXs. In the inclusive distribution shown in Figure 4.27 (left), a high rate of BXs with either one or two primitives is evident. However, when empty BXs are masked, as shown in Figure 4.27 (right), it becomes clear that these counts primarily originate from the empty BXs. In almost every colliding BX, an  $e/\gamma$  primitive is present, with an average of 5.5 TPs per colliding BX.

As visible in Figure 4.28, the majority of these trigger primitives are concentrated at the lower end of the transverse energy ( $E_T$ ) spectrum. These low-energy primitives are responsible for the increased rate in the bins corresponding to high multiplicity per BX. One notable feature in the energy distribution is a sudden increase in the number of primitives around 128 GeV. This is due to the removal of electron identification criteria for primitives with  $E_T \geq 128$  GeV[129]. A saturation bin is visible at the higher end of the spectrum, corresponding to TPs with  $E_T = 255.5$  GeV. The pseudorapidity ( $\eta$ ) distribution shown in Figure 4.29





Figure 4.27: Multiplicity of  $e/\gamma$  trigger primitives per BX, computed considering all 3564 BXs in an orbit (left) and masking the empty ones (right). In fill 9969, 2340 bunches were colliding in CMS, while the machine was filled with a total of 2352 bunches.




Figure 4.28: Transverse energy distribution of  $e/\gamma$  trigger primitives.




Figure 4.29: Pseudorapidity distribution of  $e/\gamma$  trigger primitives.

exhibits the expected shape. The majority of  $e/\gamma$  TPs are concentrated in the barrel region, while the rate increases in the endcaps as  $|\eta|$  increases, i.e., as we move closer to the beam line.

The BX occupancy distribution shown in Figure 4.26 can be further analyzed as a function of energy. In Figure 4.30, the rate of  $e/\gamma$  trigger primitives (TPs) per BX is displayed for different energy thresholds. As expected from the  $E_T$  distribution in Figure 4.28, even a small cut on the TP transverse energy significantly reduces the rate by an order of magnitude. Interestingly, the shape of the distribution changes during the transition between the last bunch of the train and the empty bunches. When no energy cut is applied, a noticeable drop in the rate occurs before stabilizing to a “baseline” value after approximately 10 BXs. However, when an  $E_T$  cut is applied, the baseline value is reached




Figure 4.30: Rate of  $e/\gamma$  trigger primitives as a function of bunch crossing in the orbit during LHC fill 9969, shown for different energy cuts used to count primitives. In the bottom panel, the bunches colliding at IP5 (CMS) are marked with a black line. This plot is zoomed in on the BX range between 50 and 180, corresponding to the first two trains of the orbit. The blue line represents all trigger primitives, while the yellow and red lines indicate the rate of TPs with  $E_T > 5$  GeV and  $E_T > 10$  GeV, respectively.

in a similar timeframe, but without the initial drop in rate. This effect is a well-known behavior of the ECAL trigger primitive generation algorithm when dealing with out-of-time (OOT) pileup and bunch trains[130]. The amplitude reconstruction is performed using digital filters (referred to as “weights”) that include pedestal subtraction. In the middle of a bunch train, the OOT pileup contribution is evenly distributed across the samples in the readout window and is correctly subtracted. However, at the end of a bunch train, the baseline in the initial part of the readout window differs, leading to a negative bias in amplitude reconstruction. This bias primarily affects low  $E_T$  TPs, which is why the effect is not observed for higher  $E_T$  TPs. On the other hand, the “staircase” descent observed at the end of the train when selecting higher energy thresholds is attributed to afterglow effects, such as detector activation.

Finally, as was done for the muons, it is possible to look at the invariant mass of  $e/\gamma$  pairs in the region where the  $Z \rightarrow e^+e^-$  peak is expected. However,




Figure 4.31: Invariant mass distribution of  $e/\gamma$  pairs identified by the calorimeter trigger in the barrel. The data are compared to a Drell-Yan Monte Carlo simulation, where the L1  $e/\gamma$  have been produced through the L1 Trigger software simulation.

given the nature of the trigger primitives, it is not possible to require them to be of opposite charge or to consist solely of electrons. Therefore, the peak we are looking for will be superimposed on a large, exponentially decreasing background of QCD photons and electrons. Nevertheless, a slightly cleaner selection can be obtained by requiring the trigger primitives to be isolated.

In Figure 4.31, the distribution of the invariant mass of  $e/\gamma$  pairs is shown. Pairs are selected by requiring them to be in the barrel ( $|\eta| < 1.44$ ), with transverse energy  $E_T > 20 \text{ GeV}$  and isolated. The data are compared with a DY simulation, the same used for the muons. Similarly, a fit to the data was performed using the same functions: a combination of a power-law and polynomial functions for the background and a crystal ball function for the resonant component. As observed for the muons, a small shift in the mean of the peak ( $\mu_{\text{Data}}$ ) with respect to MC is visible, likely due to different calibrations used in the

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.32: Multiplicity of  $\tau$  trigger primitives per BX, computed considering all 3564 BXs in an orbit (left) and masking the empty ones (right). In fill 9969, 2340 bunches were colliding in CMS, while the machine was filled with a total of 2352 bunches.

simulation. However, in this case, the standard deviation  $\sigma_{\text{Data}}$  is much smaller than that obtained for the muons, thanks to the excellent energy resolution of the CMS electromagnetic calorimeter, which is capable of precisely measuring the energy of the shower.

#### $\tau$ primitives

The BX occupancy for  $\tau$  trigger primitives is similar to that observed for  $e/\gamma$  primitives in the previous subsection, both when detecting laser reflections and during collisions. However, the rate of  $\tau$  objects is significantly higher. As shown in Figure 4.32 (left), the multiplicity per BX typically shows either no  $\tau$  or the maximum number allowed per BX. By masking the empty bunches, the multiplicity can be recalculated using only the colliding BXs. This is depicted in Figure 4.32 (right), where it becomes evident that the high rate of BXs with only 1 or 2  $\tau$  were indeed produced during the empty bunch crossings. In contrast, during collisions, a  $\tau$  trigger primitive is almost always produced, with an average of 10.5 primitives per BX, often reaching the maximum number of allowed primitives.

The distribution of the transverse energy  $E_T$  is shown in Figure 4.33. Unlike the  $e/\gamma$  distribution, a wider binning is used for primitives with energies below approximately 25 GeV, while a finer binning of 0.5 GeV is applied for higher energies. The saturation bin at 255.5 GeV is consistent with that for  $e/\gamma$  TPs, with a similar rate of approximately 1 kHz, indicating that the overall higher




Figure 4.33: Transverse energy distribution of  $\tau$  trigger primitives.




Figure 4.34: Pseudorapidity distribution of  $\tau$  trigger primitives.

rate of primitives is mainly due to low-energy TPs. The  $|\eta|$  distribution shown in Figure 4.34 appears to be a “smoothed” version of the  $e/\gamma$  distribution, especially in the transition region between barrel and endcap. This is because the  $\tau$  algorithm merges multiple neighbouring clusters, adapting the  $e/\gamma$  clustering to the selection of tau induced calorimeter signal combining the energy deposit profile in both ECAL and HCAL[131]. This is particularly helpful to recover to multi-prong  $\tau$  decays.

### Jet primitives

This section provides an overview of the characteristics of jet trigger primitives. A more detailed analysis of these objects will be conducted in the final chapter of this work. The BX occupancy for jet trigger primitives, shown in Figure 4.35, reveals a significantly lower rate compared to  $e/\gamma$  and  $\tau$  trigger primitives. This reduction is primarily due to the jet-finding algorithm, which effectively suppresses a larger fraction of noisy towers. Additionally, there is no major variation in the rate across different regions of the hadronic calorimeter, unlike what is observed in other subsystems, such as the muon systems.

This behavior is further reflected in the multiplicity of jet trigger primitives per bunch crossing, as shown in Figure 4.36 (left) for all BXs and Figure 4.36 (right) for colliding BXs. There is no significant difference between the two, aside from a slight variation in the rate of low-multiplicity bunch crossings, which aligns with expectations.

The pseudorapidity distribution of jet trigger primitives, shown in Figure 4.38, exhibits a pattern similar to that observed for  $e/\gamma$  primitives in Fig-

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.35: Rate of jet trigger primitives as a function of bunch crossing in the orbit during LHC fill 9969, showing contributions from all HCAL regions. In the bottom panel, the bunches colliding at IP5 (CMS) are marked with a black line.





Figure 4.36: Multiplicity of jet trigger primitives per BX, computed considering all 3564 BXs in an orbit (left) and masking the empty ones (right). In fill 9969, 2340 bunches were colliding in CMS, while the machine was filled with a total of 2352 bunches.






Figure 4.37: Transverse energy distribution of jet trigger primitives. The contribution of each region is highlighted in a different color.




Figure 4.38: Pseudorapidity distribution of jet trigger primitives. The contribution of each region is highlighted in a different color.

ure 4.29. In the central region, the rate gradually decreases before reaching the endcaps, where it increases again. The coarser segmentation of HF in the range  $3 < |\eta| < 5$  is evident in the plot as empty bins.

The transverse energy distribution of jet trigger primitives, shown in Figure 4.37, displays several notable features. Firstly, the saturation bin at the high end of the spectrum often does not correspond to actual jets with  $E_T > 1$  TeV; rather, it results from a single trigger tower saturating its energy, causing the jet and energy sums to be assigned the saturation energy. Additionally, periodic bumps appear in the energy spectrum, which should otherwise be smooth. These anomalies are due to Jet Energy Corrections (JECs) applied in the layer-2 calorimeter trigger to calibrate the trigger objects.

These features are more clearly visible in Figure 4.39, where the jet transverse energy  $E_T$  is plotted as a function of the jet pseudorapidity. In the central region, with  $|\eta| < 2$ , horizontal bands indicate an increased rate of jets with specific energy values. Furthermore, the impact of jet energy corrections is evident in two additional cases. Firstly, jets with  $E_T < 10$  GeV are shifted to the bin corresponding to  $E_T = 0$  GeV, as precisely assigning energy to such low-energy jets is not feasible. Secondly, a kinematic region where jets are not present, corresponding to energies between approximately 45 GeV and 55 GeV, indicates that these jets are either promoted or demoted in energy. This poses constraints on the possible selections applied when reconstructing events with hadronic jets. A study of this kind of event will be presented in the next chapter.




Figure 4.39: Distribution of the jet transverse energy  $E_T$  as a function of the jet  $\eta$ . While the full pseudorapidity range is shown, the transverse energy is zoomed in the range  $E_T \leq 170$  GeV.

### Energy sums

In addition to the objects described in the previous section, the calorimeter trigger also computes global energy sums for each bunch crossing, using the full detector view and trigger tower granularity. However, in the scouting system, these energy sums are only transmitted to the Buffer Units if at least one of the other calorimeter objects, such as a jet,  $e/\gamma$ , or  $\tau$  trigger primitive, is present. If none of these objects are produced, the energy sums are not read out. An example of these energy sums is shown in Figure 4.40.

The first energy sum in the list is the total  $E_T$ , which represents the scalar sum of transverse energy across all trigger towers. Similarly, the missing transverse energy,  $E_T^{\text{miss}}$ , is calculated by taking the vector sum of the transverse energies of the same trigger towers. In both cases, the full granularity available at the trigger level, the trigger tower, is utilized. Another important sum is the total scalar energy of jets, denoted as  $H_T$ , which is calculated using L1-jets with  $E_T \geq 30$  GeV and  $|\eta| < 2.5$ . A corresponding quantity,  $H_T^{\text{miss}}$ , is determined from the vector sum of these L1-jets. Additional energy sums are be computed for specific regions, such as the electromagnetic calorimeter or the forward region (HF).




Figure 4.40: Distribution of the transverse energy  $E_T$  for a set of energy sums computed by the calorimeter trigger. A wider binning is used for energy values above 512 GeV. The bin corresponding to events with saturated towers is visible in black on the right end of the spectrum.

One example are the  $ET_{miss}^{HF}$  and  $HT_{miss}^{HF}$ . At the high end of the energy spectrum, a saturation bin is visible; this bin corresponds to events containing saturated trigger primitives, such as a jet with  $E_T = 1024$  GeV. When such saturation occurs, all energy sums are capped at their maximum value.

Another sum computed for each event is the number of trigger towers, denoted as nTT. This is determined by counting the number of trigger towers with energy greater than 0.5 GeV within the eight central  $\eta$  rings of the calorimeter, i.e.,  $|\eta| \leq 0.34$ [129]. This quantity is particularly useful as a good proxy for the pileup in each event. It is employed in the pileup-correction lookup tables (LUTs), along with the trigger tower's  $\eta$ , enabling per-event pileup corrections rather than an average correction. An example of the nTT distribution is shown in Figure 4.41, where it is plotted against the bunch crossing number, along with the average pileup measured for that bunch crossing. The core of the nTT distribution closely follows the average pileup profile, which varies for each bunch in a train.




Figure 4.41: Distribution of the number of trigger towers,  $n_{TT}$ , as a function of the bunch crossing number, for bunches between 1220 and 1360. The average pileup per bunch crossing is shown in red. The core of the  $n_{TT}$  distribution closely follows the average pileup shape.

### 4.7.3 BMTF stubs

The muon stubs used by the Barrel Muon Track Finder to reconstruct muons in the barrel region ( $|\eta| < 0.83$ ) represent the lowest-level quantity collected by the scouting demonstrator. These stubs can be reprocessed to produce alternative reconstructions, rather than relying on pre-processed quantities such as final GMT muons. The input stubs to the BMTF boards are generated by the TwinMux boards, which combine DT local trigger primitives with RPC trigger primitives to form the so-called “super-primitives”. Each BMTF board collects up to 8 stubs per BX, and with 12 BMTF boards one for each CMS wedge, a horizontal slice of the CMS barrel muon system covering the same sector in all wheels a total of 96 muon stubs can be sent to the scouting system per BX. A detailed description of the quantities and information contained in each stub is provided in Table A.3 in the appendix.

Figure 4.42 shows the rate of muon stubs collected per BX, focusing on the first 500 bunch crossings of the orbit. The contribution of each stub quality type is illustrated. Except for the quality code 7, which indicates invalid tracks and is therefore not visible, the other quality codes are sorted in increasing order of




Figure 4.42: Rate of BMTF stub primitives as a function of bunch crossing and quality during LHC fill 9969. The plot focuses on the first 500 bunch crossings of the orbit. In the bottom panel, bunches colliding at IP5 (CMS) are marked with a black line.

muon stub quality. For example, the highest valid quality, code 6, corresponds to the highest-quality stubs where a trigger primitive was produced by the DT chamber with high trigger quality on both the inner and outer superlayers. In contrast, quality codes 2 and 3 represent uncorrelated triggers, meaning that a trigger primitive was produced in one superlayer but not in the other. These lower-quality stubs make up the largest fraction of the rate, especially outside the colliding bunches, likely due to detector activation effects. Additionally, a gradual increase in the rate can be observed starting around 7 BXs before the first filled bunch, likely due to a combination of factors such as beam-halo muons crossing the muon chambers and incorrect BX ID assignment. On the other hand, when examining high-quality stubs (e.g., quality code 6), a pattern similar to that observed for BMTF muons in Figure 4.19 is evident: the rate outside colliding bunch crossings is largely suppressed, as activation noise rarely produces two high-quality correlated primitives, while the train structure is more pronounced. The effects of machine-induced background are also clearly visible, such as the bump in the rate at the beginning of the orbit due to the halo of bunches not colliding in CMS, along with the non-flat rate between trains of filled bunches.

#### 4.7. FIRST DATA FROM THE L1 SCOUTING DEMONSTRATOR




Figure 4.43: Rate of BMTF stub primitives as a function of Wheel and muon station number during LHC fill 9969.




Figure 4.44: Multiplicity of BMTF stub primitives per bunch crossing during LHC fill 9969.

Figure 4.43 shows how the total rate is distributed across the 5 wheels and 4 chambers per sector. The central Wheel is WH0, while WH+2 and WH-2 represent the outermost wheels on the positive and negative sides, respectively. The chambers are numbered in increasing order from the closest to the magnet, MB1, to the furthest from the interaction point (IP), MB4. As expected, the rate decreases when moving from the MB1 chambers towards MB4, as these are further from the IP. Conversely, the rate increases when moving from the innermost to the outermost wheel, where the MB1 chambers have the highest rate. In the figure, these chambers are not visible because they are masked in the BMTF to reduce the number of duplicated (ghost) muons found in both the BMTF and OMTF (Overlap Muon Track Finder) sub-systems. Additionally, the figure represents an average across the 12 sectors; it is possible to generate similar plots to study potential asymmetries in  $\phi$ , which are expected for cosmic muons but should be uniform during collisions.

Finally, Figure 4.44 shows the multiplicity of muon stubs per BX, with values ranging from 0 to 16 out of a maximum of 96. The distribution decreases exponentially, and even at a multiplicity of 16 stubs per BX, the rate is already around  $10^{-2}$  Hz, while the rate for a single stub exceeds 1 MHz.

Data from the BMTF stubs can be used to enhance muon reconstruction and  $p_T$  assignment, with or without vertex constraints. Promising results have been achieved using a neural network that takes the muon stubs from a sector as input and uses them to regress the parameters of the muon track, including  $p_T$ ,  $\eta$ ,  $\phi$ , and charge. Comparison with muons reconstructed by the BMTF shows a visible improvement[132]. Additionally, this approach can be applied in the context of Heavy Stable Charged Particles (HSCP), which leave signals across multiple bunch crossings. This type of signature is currently challenging to trigger at the L1-trigger level, as data from multiple BXs are not available for the final decision. On the other hand, since all BXs in an orbit are available for L1 scouting processing, it is possible to develop a dedicated reconstruction that combines the products of 2 or more BXs.





# 5

## Towards a dijet analysis with L1 Scouting data

Searches for new particles that decay into pairs of jets have always played a key role in hadron collider experiments. These searches have progressively been conducted at higher energies and with increasing sensitivity, aiming to uncover physics beyond the Standard Model, motivated by numerous theoretical models predicting resonances that couple to quarks and gluons[133, 134, 135, 136, 137]. The production of such particles would manifest as resonances in the dijet invariant mass spectrum, which is otherwise dominated by QCD-induced dijet production and is expected to follow a smooth, monotonically falling distribution.

The earliest dijet searches were conducted by the UA1[138] and UA2[139] collaborations in proton-antiproton collisions at  $\sqrt{s} = 660$  GeV at the CERN Sp $\bar{p}$ S collider. These searches were later extended to higher center-of-mass energies by the CDF[140] and D $\bar{0}$ [141] collaborations, using  $\sqrt{s} = 1.96$  TeV proton-antiproton collisions at the Tevatron accelerator at Fermilab. The baton was then passed to the ATLAS and CMS experiments at the LHC, which have been performing dijet searches since Run 1, with collisions at  $\sqrt{s} = 7$  TeV and  $\sqrt{s} = 8$  TeV[142, 143]. A comprehensive summary of dijet searches at hadron colliders up to LHC Run 1 can be found in Ref.[13].

With the increased energy in LHC Run 2, at  $\sqrt{s} = 13$  TeV, the sensitivity to higher mass dijet resonances was significantly improved, setting limits across various scenarios and models[144, 145, 146, 147].

However, the enhanced energy and luminosity at the LHC have introduced challenges in probing hadronic resonances in the lower end of the mass spectrum. The increasing background cross-section from multijet events has forced tighter trigger thresholds, limiting sensitivity to lower masses. To overcome this, novel techniques, such as trigger-level analysis and searches using boosted objects, have been employed by both CMS and ATLAS to extend sensitivity to lower mass ranges.

This chapter presents an overview of the dijet search methodologies employed at CMS, with a focus on strategies for probing the low-mass region. Then, the first exploratory steps towards a dijet analysis using L1 scouting jets will be discussed, as this scouting technique presents a promising approach to bypass the trigger limitations that typically affect low-mass searches.

### 5.1 Low-mass dijet searches at CMS

Dijet searches target the production of new particles decaying to a pair of quarks and gluons, which would manifest as a resonance in the dijet invariant mass spectrum. So-called “high-mass” resonances searches typically begin from around  $\sim 1.5$  TeV and extend to a few TeV, where the LHC excels in probing previously unexplored energy scales. Events containing high- $p_T$  jets can be more easily selected, as QCD dijet production at the TeV range is more manageable. On the other hand, traditional search strategies for low-mass resonances (below a few hundred GeV) face significant challenges, primarily due to trigger limitations imposed by the overwhelming QCD background. In fact, the large QCD background at low dijet masses imposes tight trigger thresholds on jets to keep the event rate manageable. On the other hand, these high thresholds can suppress the low expected rate of BSM events in this mass region. In this region of the phase space, only loose trigger thresholds would allow for the collection of such signals, which is not compatible with the standard trigger thresholds.

In this context, HLT scouting provides an alternative data-taking strategy designed to bypass the limitations imposed by the CMS High Level Trigger, as introduced in chapter 4, by collecting objects reconstructed in the HLT. This technique has extended the sensitivity to new resonances from the offline limit of 1.2 TeV[148] down to 500 GeV[106]. Using HLT scouting, other searches were performed during Run 2 in the mass range between 0.6 – 2 TeV, as beyond this range analyses of standard data are used. By requiring an additional jet




Figure 5.1: Schematic of the decay of a boosted  $Z'$ . The two jets from the decay (blue) are merged into a single large-radius jet (red), recoiling against a high- $p_T$  jet from initial-state radiation (green).

from initial-state radiation (ISR) to overcome trigger limitations, a search was able to cover the range 350 – 700 GeV[149]: The requirement of three jets with  $p_T > 70$  GeV allowed the event to pass the L1 trigger selection while keeping the resonance mass low.

Another strategy employed to perform low-mass analyses was that of selecting events where the new particle is produced in association with a high- $p_T$  jet from initial-state radiation, e.g.  $p_T > 500$  GeV, a much higher value than what was used in the previously presented scouting analyses. Due to momentum conservation in the transverse plane, this implies that the new particle is produced with very high momentum. Consequently, the decay products are extremely collimated due to the large boost and can be clustered into a single large-radius jet, as sketched in Figure 5.1. The large-radius jet can then be cleaned of soft or large-angle particles, requiring a two-prong substructure consistent with an object decaying into a quark-antiquark pair[150, 151]. Additionally, the ISR jet can be replaced by a photon in an equivalent analysis[152]. This method helps to suppress the QCD multijet background, creating a “clean” environment for low-mass dijet searches in the 20 – 300 GeV range, at the price of a lower signal acceptance.

The result of dijet searches is typically a limit on the cross-section, which is generally model-independent since the signal simulation, combined with detector simulation, is used to extract the resonance shape. The theoretical cross-section for a particular model can then be evaluated as a function of the resonance mass and compared with the measured one to identify exclusion regions of the model’s parameters, or directly on the masses of the candidate resonances. This final step on the other hand is model-dependent.

A common benchmark model used in dijet searches is a leptophobic axial-vector  $Z'$  model, within the framework of simplified Dark Matter (DM) models[153]. In this scenario, a new boson  $Z'$  mediates interactions between Stan-

## 5.1. LOW-MASS DIJET SEARCHES AT CMS




Figure 5.2: Feynman diagram for a hypothetical  $Z'$  boson decaying into a dark matter  $\chi\bar{\chi}$  pair (left) or  $q\bar{q}$  pair (right).




Figure 5.3: Limits on the universal coupling  $g'_q$  between a leptophobic  $Z'$  boson and quarks from various CMS dijet analyses[154]. The expected limits are shown as dashed lines, while the observed limits are represented by solid lines. The hashed areas indicate the regions excluded by the observed limits.

dard Model quarks and fermionic DM particles,  $\chi$ , with couplings  $g_q$  and  $g_\chi$  respectively. The tree-level diagrams are shown in Figure 5.2.

When the decay into DM particles is kinematically suppressed ( $2m_\chi > m_{Z'}$ ), the width of the  $Z'$  is determined solely by  $g_q$ . Under these assumptions, the coupling denoted as  $g'_q$  is used to compare the results of CMS searches, as shown in Figure 5.3. As illustrated in the figure, various search techniques have enabled coverage of a wide range of  $Z'$  masses, from 10 GeV up to around 6 TeV.

In this context, the ambitious goal of performing a dijet analysis with jets collected using the L1 scouting system could follow the path of the first dijet searches conducted with HLT scouting. This approach could bridge the results from the low-mass region with the latest results obtained with HLT scouting,

focusing on the range between 200 – 500 GeV.

## 5.2 Data and Simulated samples

### 5.2.1 Data Sample

To maximize the number of events available for the search, data from the selection stream introduced in the previous chapter are utilized. Specifically, data from the `Dijet30` selection stream are used, where the only requirement is that the two leading jets satisfy  $p_T^{j_1, j_2} > 30 \text{ GeV}$ . Since the deployment of this stream, as of early September 2024, over  $50 \text{ fb}^{-1}$  of data have been collected.

In addition to this first selection, and in line with other analyses, both leading jets are required to have  $|\eta_{j_1, j_2}| < 2.5$  to ensure that the jets are well-contained within the barrel and endcap regions. No additional requirements are placed on any other jets in the events.

For this preliminary study, only a small subset of the data has been used, corresponding to approximately  $20 \text{ pb}^{-1}$ . Despite this modest sample size, the number of events selected by the  $p_T$  requirement is about  $1.5 \cdot 10^9$ , which is reduced to  $700 \cdot 10^6$  events after applying the  $|\eta| < 2.5$  criterion. Notably, this dataset corresponds to just 15 minutes of data taking. A rough estimate of the total number of dijet events satisfying these selections, assuming  $50 \text{ fb}^{-1}$  of data, is approximately  $3.7 \cdot 10^{12}$  events.

Although the relatively small event size makes it feasible to store these data, the sheer volume poses significant challenges from an analysis perspective. The first challenge lies in developing efficient methods to process and analyze such a large dataset, while the second involves determining the appropriate statistical tools to handle the data.

### 5.2.2 Signal Simulation

For this study, the focus is on understanding the L1 jets response using a simplified model. A narrow resonance is considered, where the natural resonance width is negligible compared to the L1 dijet mass resolution.

The signal model is based on the leptophobic  $Z'$  commonly used in other searches, where the  $Z'$  boson couples exclusively to quarks. The same generation procedure as in previous analyses[106, 149] is followed, but simulated with the

## 5.2. DATA AND SIMULATED SAMPLES




Figure 5.4: Left: pseudorapidity  $\Delta\eta$  and azimuthal angle  $\Delta\phi$  separation for the  $q\bar{q}$  pair produced by the decay of the  $Z'$  boson with a mass of 250 GeV. Right: Invariant mass distribution of the  $q\bar{q}$  pair. Both distributions are normalized to unit area.

CMS detector response, including pileup (PU) and trigger object reconstruction, using the 2024 CMS experiment conditions. The trigger quantities are part of the standard detector simulation and are extracted from there.

Signal samples are generated for varying  $Z'$  boson masses, ranging from 150 GeV to 350 GeV.

As shown in Figure 5.4 (left), the pseudorapidity separation  $\Delta\eta$  between the two quarks is small, and the azimuthal angle separation  $|\Delta\phi| \approx \pi$ , indicating a back-to-back topology for events where both quarks fall within  $|\eta| < 2.5$ . Applying this acceptance criterion selects approximately 60% of the events. In Figure 5.4 (right) the mass distribution computed using the quarks pairs is shown.

Next, generator-level jets are considered. These jets are clustered from generated particles resulting from the showering and hadronization of the two quarks from the  $Z'$  decay, as well as contributions from the underlying event. Clustering is performed using the anti- $k_t$  [155] algorithm with a radius parameter of 0.4. At this stage, no detector effects or PU contributions are considered.

Jets are matched to the two quarks from the  $Z'$  decay by requiring  $\Delta R(j, q) < 0.4$ . In Figure 5.5 (left), the invariant mass distribution for jets matched to quarks is shown, split by quark flavor. Heavy-flavor quarks (*charm* with PdgId = 4 and *bottom* with PdgId = 5) shift the mass toward lower values and broaden the




Figure 5.5: Left: Invariant mass distribution of generator-level jets matched to partons produced by the decay of a  $Z'$  boson with a mass of 250 GeV. The contributions from different quark flavors are shown. Right:  $p_T$  distribution of generator-level jets matched to one of the quarks from the  $Z'$  decay, as well as the unmatched jets in the event. The  $p_T$  distribution of the two quarks is shown in red. Distributions are normalized to unit area.

distribution due to the energy loss from the production of neutrinos in semileptonic decays. For lighter quarks, the distribution is sharper, but a tail toward lower mass values appears, mainly caused by final-state radiation escaping the jet cone.

In Figure 5.5 (right), the  $p_T$  distribution of matched and unmatched jets is presented. The  $p_T$  distribution of unmatched jets follows an exponential decay, while for matched jets it exhibits a peak around 125 GeV, as expected for a  $Z'$  boson with a mass of 250 GeV. The  $p_T$  distribution of the matched jets closely resembles that of the two quarks, with a slight shift toward lower  $p_T$  values. This shift is then reflected in the mass distribution's tail toward lower values.

The same procedure can be applied to jets reconstructed by the emulator of the logic implemented in the L1 calorimeter trigger, instead of those clustered at the generator level. This simulation includes the detector response and trigger primitive generation, with the contribution from pileup events also included.

Figure 5.6 (left) shows the invariant mass distribution computed from jets reconstructed by the L1 calorimeter trigger for events where generator-level quarks are matched to two distinct L1 jets. Events were selected by requiring at least two L1 jets with  $|\eta| < 2.5$  and the same for the generator-level quarks. As previously computed, 60% of the quarks are within the acceptance region, and



## 5.2. DATA AND SIMULATED SAMPLES

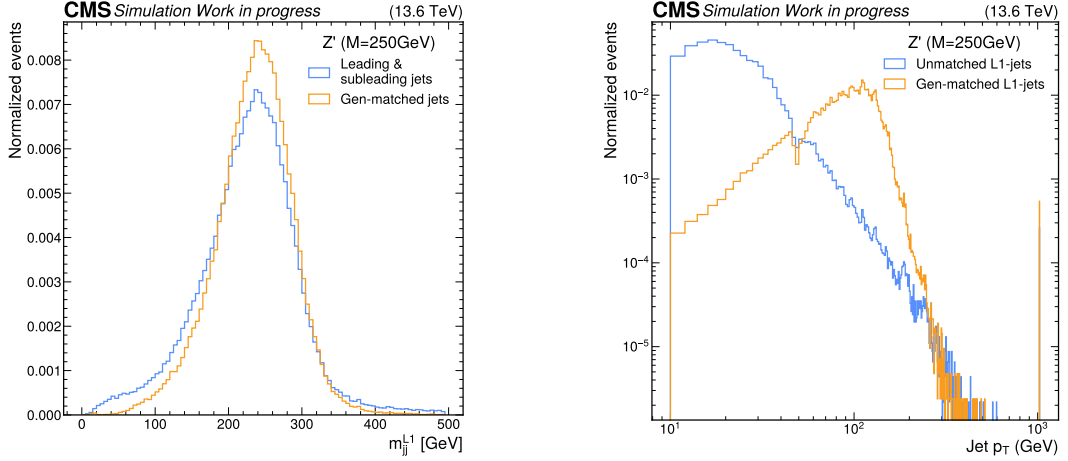


Figure 5.6: Left: Invariant mass computed using jets reconstructed by the L1 calorimeter trigger for events with a  $Z'$  boson of mass 250 GeV, calculated for events where two generator-level quarks were matched to two distinct L1 jets. The mass is computed considering two sets of jets: first, the leading and sub-leading jets in the event, as will be done in the analysis, and second, the L1 jets matched to the two quarks. Right:  $p_T$  distribution of L1 Trigger jets matched to one of the quarks produced by the  $Z'$ , and the remaining jets in the event. Distributions are normalized to unit area.

adding the trigger jet requirement does not significantly affect this percentage.

Of this subset of events, a further selection was made based on a positive match between the generator-level quarks and the L1 jets, defined as exactly two quarks matched to two distinct L1 jets with  $\Delta R < 0.4$ . This corresponds to 81% of the selected events.

With this set of simulated events, the dijet invariant mass was first computed using the leading and sub-leading jets in the event, which will be the standard procedure in subsequent steps of this analysis. This is compared to the distribution computed using the jet pair matched to the generator-level quarks to provide insight into how the distribution behaves, although the former will be used in the analysis. From the study of the MC simulation, it is observed that in approximately 90% of the events with a match, the two matching jets are the leading and sub-leading L1 jets. The remaining fraction of events, where the correct jets are not the leading and sub-leading ones, results in a broader distribution with a more pronounced tail at lower mass values.

Furthermore, the  $p_T$  distribution of jets matched to the quarks and the remaining jets in the event is shown in Figure 5.6 (right), similarly to what was






Figure 5.7: Left: Invariant mass distribution computed for signal events with different  $Z'$  masses, 150 GeV, 250 GeV, and 350 GeV. Right: Invariant mass distribution computed using generator-level jets (Gen-jets), offline Particle Flow (PF) jets, and L1 trigger jets (L1 Jets). Events were selected using the `Di jet30` selection and requiring the leading and sub-leading jets to have  $|\eta| < 2.5$ . Distributions are normalized to unit area.

done for the generator-level jets. The features observed in L1 jets in chapter 4 are clearly visible here. First, jets corresponding to saturated towers are evident at the far right end of the  $p_T$  spectrum ( $p_T = 1023.5$  GeV). Second, the discontinuity between 40 – 60 GeV is visible in both matched and unmatched jets. Matched jets exhibit a distribution similar to the expected one observed in Figure 5.5. Unmatched jets have the same  $p_T$  distribution as those collected in the zero-bias dataset, analyzed in the previous chapter (Figure 4.37), as expected.

A selection similar to that applied to the data can be used to check the invariant mass distribution for different signal masses. In Figure 5.7 (left), the invariant mass distribution was computed using the leading and sub-leading jets for events where the two leading jets have  $p_T > 30$  GeV and  $|\eta| < 2.5$ . The fraction of signal events passing this selection for different mass hypotheses is shown in Table 5.1.

$Z'$ Mass [GeV]	$p_T > 30$ GeV [%]	$p_T > 30$ GeV and $ \eta  < 2.5$ [%]
150	$71 \pm 0.1$	$45 \pm 0.1$
250	$89 \pm 0.1$	$62 \pm 0.1$
350	$95 \pm 0.1$	$69 \pm 0.1$

Table 5.1: Percentage of events passing `Di jet30` selection and the additional cut on  $|\eta|$  for different  $Z'$  mass values.

## 5.2. DATA AND SIMULATED SAMPLES

Finally, in Figure 5.7 (right), the dijet mass is compared for generator-level jets, offline reconstructed jets using Particle Flow (PF)[57], and jets reconstructed by the L1 trigger.

### 5.2.3 Background Simulation

The dominant background for this analysis is expected to come from QCD production of two or more jets. As an initial step in this study, a comparison between the data and QCD background predictions for dijet events is performed. The QCD samples were generated using Pythia8[156] with the CP5 tune[157], where the phase space is divided into exclusive bins based on the transverse momentum of the hard-scattered parton  $\hat{p}_T$ . These samples were centrally produced by the CMS generation group, with the lowest  $\hat{p}_T$  bin starting from 30 GeV, and include full simulation of the detector response and reconstructed quantities. The cross-sections and equivalent luminosities for the simulated QCD samples are presented in Table 5.2.

Sample $\hat{p}_T$ [GeV]	Cross Section [pb]	Equivalent Luminosity [ $\text{pb}^{-1}$ ]
$\hat{p}_T \in [30, 50]$	113300000.0	0.034
$\hat{p}_T \in [50, 80]$	16760000.0	0.084
$\hat{p}_T \in [80, 120]$	2534000.0	0.385
$\hat{p}_T \in [120, 170]$	445800.0	3.25
$\hat{p}_T \in [170, 300]$	113700.0	8.42
$\hat{p}_T \in [300, 470]$	7589.0	130
$\hat{p}_T \in [470, 600]$	626.4	1534
$\hat{p}_T \in [600, 800]$	178.6	5335
$\hat{p}_T \in [800, 1000]$	30.57	31611

Table 5.2: Cross sections and equivalent luminosities for different QCD  $\hat{p}_T$  bins.

One key observation is the large cross-section for samples with low  $\hat{p}_T$ , resulting in their low equivalent simulated luminosity. To match the low integrated luminosity used in this initial study, approximately  $2.6 \cdot 10^9$  events would need to be generated and simulated just to match  $20 \text{ pb}^{-1}$ .

This is one of the main reasons why most dijet analyses traditionally rely on data-driven methods to estimate the background. Such approaches are further justified by the inherent uncertainties in QCD calculations, making them unreliable estimators for backgrounds in dijet resonance searches. However, QCD MC samples still play a crucial role in validating the quality of the data, ensuring that no unexpected behaviors are present.

### 5.3 Dijet Data and QCD Background

The first comparison was made between data and QCD MC without any additional cuts beyond the two mentioned earlier, requiring the two leading jets to have  $p_T > 30 \text{ GeV}$  and  $|\eta| < 2.5$ . The goal is to develop a set of selections that can suppress the QCD background as much as possible and at the same time enhance the signal. The distribution of  $m_{jj}$  for QCD is expected to follow a smooth, exponentially falling shape. For this comparison, the MC samples were scaled according to their cross-sections and integrated luminosities relative to the data.

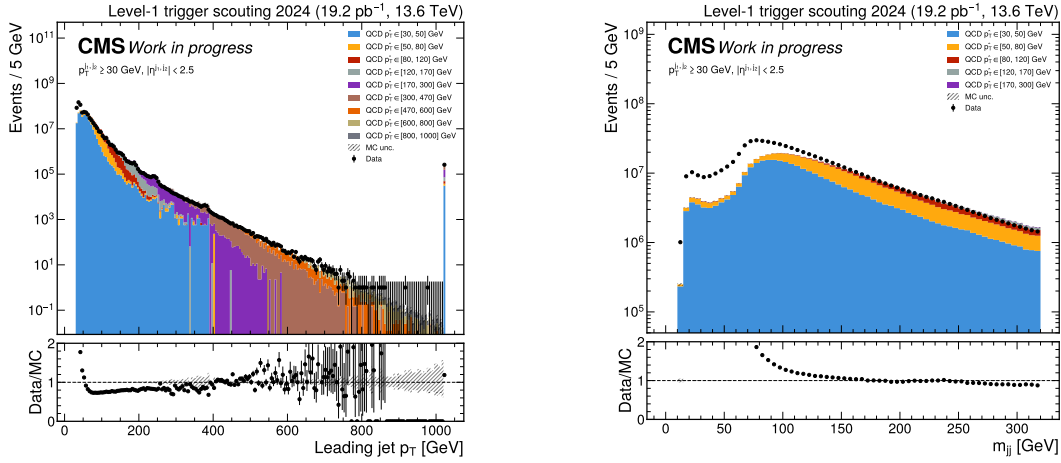


Figure 5.8: Left: Distribution of the leading L1 jet  $p_T$  for data and QCD simulation, split by  $\hat{p}_T$  bins to show the contribution of each sample. Right: Invariant mass distribution computed using the leading and subleading jets, zoomed in on the low mass range.

Figure 5.8 (left) shows the distribution of the leading jet  $p_T$  for both data and Monte Carlo simulation. Notably, the main features observed in the data are well described by the MC, from the “bumps” in the lower part of the spectrum to the appearance of saturated jets (Figure 4.39 and Figure 4.37). Additionally, despite the low  $\hat{p}_T$  of the first sample, events are present even above 200 GeV. This is likely due to pile-up (PU) events within the sample. Furthermore, there are more data than predicted MC events in the 30 – 50 GeV range. This mismatch in low- $p_T$  jets could be attributed to the absence of lower  $\hat{p}_T$  MC samples. Although a sample exists for  $\hat{p}_T \in [15, 30] \text{ GeV}$ , its large cross-section, an order of magnitude higher than the  $\hat{p}_T \in [30, 50] \text{ GeV}$  sample, would require an

### 5.3. DIJET DATA AND QCD BACKGROUND




Figure 5.9: Distribution of pseudorapidity separation  $\Delta\eta$  and azimuthal angular separation  $\Delta\phi$  between the leading and subleading jets, computed for data (left) and QCD simulation (right).

unreasonable number of events to be generated, which is beyond the scope of this study. When examining the ratio between data and MC, this mismatch is evident as a divergence at the lower end of the  $p_T$  distribution. Additionally, a slow increasing trend in the ratio can be observed.

Figure 5.8 (right) shows the invariant mass distribution computed using the leading and subleading jets. Similarly to the transverse momentum distribution, a deficit of low-mass events compared to the MC is visible, likely due to the absence of low  $p_T$  MC events. To mitigate this effect, one possible cut is on the mass, for example, requiring  $m_{jj} > 150$  GeV. This also helps to remove events with low- $p_T$  jets. However, events with low  $p_T$  but large  $\Delta\eta$  may still pass the cut, as large  $m_{jj}$  values can still be produced in such cases.

Figure 5.9 (left) shows the distribution of  $\Delta\eta$  and  $\Delta\phi$  between the leading and subleading jets in the data, while Figure 5.9 (right) presents the same for QCD simulation. The simulation reproduces all the key features observed in the data. At the center of the plot, a region devoid of jets is visible. This corresponds to the geometric area of the jet, which is computed using  $9 \times 9$  trigger towers and roughly matches a jet of radius  $R = 0.4$ , where two jets cannot be present simultaneously. However, a halo around this region is also visible, indicating a significant number of events where the two leading jets are extremely close. Since the energy computation for such jets may be biased, a decision was made to discard these events by requiring an angular separation of  $\Delta\phi > \pi/3$  between the two jets. This selection favors events where the dijet system is more “back-to-back”, i.e.,  $|\Delta\phi| \approx \pi$ , which is expected for a two-body system in the absence of

additional radiation due to conservation of transverse momentum. This feature is clearly seen in the figures, where a larger event count is observed at  $|\Delta\phi| = \pi$ . Additionally, a vertical band of  $\Delta\eta \approx 0.8$  is visible. This is caused by the pile-up (PU) subtraction algorithm introduced for the 2024 data-taking period, which uses a “ $\phi$ -ring” method. In this approach, the PU is estimated by averaging the energy in a full  $\phi$  section of the calorimeter for a given  $\eta$  value, and this estimated PU is subtracted from the jets in that ring. Inevitably, in the presence of a high  $p_T$  jet, the energy of all other softer jets in the same ring will be significantly reduced, sometimes pushing them below the detection threshold.

Connected to the angular variables, one of the most powerful handles for suppressing QCD background is the pseudorapidity separation  $|\Delta\eta|$  between the two leading jets. In almost all dijet searches, a cut on  $|\Delta\eta|$  is applied, often around  $|\Delta\eta| < 1.3$ , to suppress the QCD background. This is because QCD dijet production at high mass is dominated by  $t$ -channel processes, which peak at  $\cos\theta^* = 1$ , where  $\theta^*$  is the scattering angle in the center-of-mass frame[13]:

$$\frac{dN}{\cos\theta^*} \sim \frac{1}{\hat{t}^2} \sim \frac{1}{(1 - \cos\theta^*)^2} \quad (5.1)$$

On the other hand, resonances produced via the  $s$ -channel, such as a  $Z'$  boson, exhibit a much flatter angular distribution in  $\cos\theta^*$ . For example, for a  $Z'$ :

$$\frac{d\hat{\sigma}}{\cos\theta^*} \sim 1 + \cos\theta^{*2} \quad (5.2)$$

This suggests that a cut can be applied to optimize the signal-to-background ratio,  $\frac{S}{\sqrt{B}}$ . If Equation 5.1 is used for the background and a signal with a flat  $\cos\theta^*$  distribution is assumed, the optimal cut is found to be  $|\cos\theta_{MAX}^*| = 1/\sqrt{3} = 0.58$ . Using the approximation  $\cos\theta^* \approx \tanh(\Delta\eta/2)$ , this translates into  $\Delta\eta_{MAX} \approx 1.3$ . For the purposes of this exploratory study, this value will be adopted, though a more precise optimization can be performed in future work using detailed MC simulations of both signal and background.

Applying the selections described earlier, i.e.,  $m_{jj} > 150 \text{ GeV}$  and  $\Delta\phi > \pi/3$  to remove low-mass and potentially pathological events, leads to the pseudorapidity separation  $\Delta\eta$  shown in Figure 5.10 for both signal and QCD background. The signal sample has been scaled to match the number of events in the QCD sample for comparison purposes. It is clear that the signal peaks at lower values

### 5.3. DIJET DATA AND QCD BACKGROUND




Figure 5.10: Pseudorapidity separation between leading and subleading jets for QCD and  $Z'$  signal samples. The  $Z'$  signal is normalized to the QCD simulation for visualization.

of  $|\Delta\eta|$ , while QCD backgrounds dominate for larger values  $|\Delta\eta|$ . Additionally, the impact of the pile-up (PU) subtraction algorithm is visible as a discontinuity in the low  $|\Delta\eta|$  region for both signal and background.

Furthermore, the increase in the signal fraction at smaller  $|\Delta\eta|$  values is reflected in the lower panel of the figure, where the  $S/\sqrt{B}$  ratio rises as  $|\Delta\eta|$  decreases.

In the next step, the selections  $m_{jj} > 150 \text{ GeV}$ ,  $\Delta\phi > \pi/3$ , and  $|\Delta\eta| < 1.3$  are applied to the selected events. These selections are intended to remove the majority of lower mass events, which fall outside the scope of this study, and to further suppress the QCD background. Additionally, events containing a saturated jet are discarded, as such jets would typically be the leading jet, but their energy is essentially unknown. This additional set of cuts retains approximately 3.5% of the events that passed the initial selection based on the  $p_T$  and  $|\eta|$  of the leading jets.

Furthermore, for the purpose of comparing the shape of the distributions, the QCD simulation has been normalized to the data by multiplying it by a factor of 0.85, as was done in Ref.[144].

The distribution of the leading jet  $p_T$  for the selected events is shown in Figure 5.11 (left). A noticeable improvement in agreement between the data and the QCD simulation can be observed in the lower region of the distribution. Although no explicit cuts on  $p_T$  were applied, the shape of the distribution

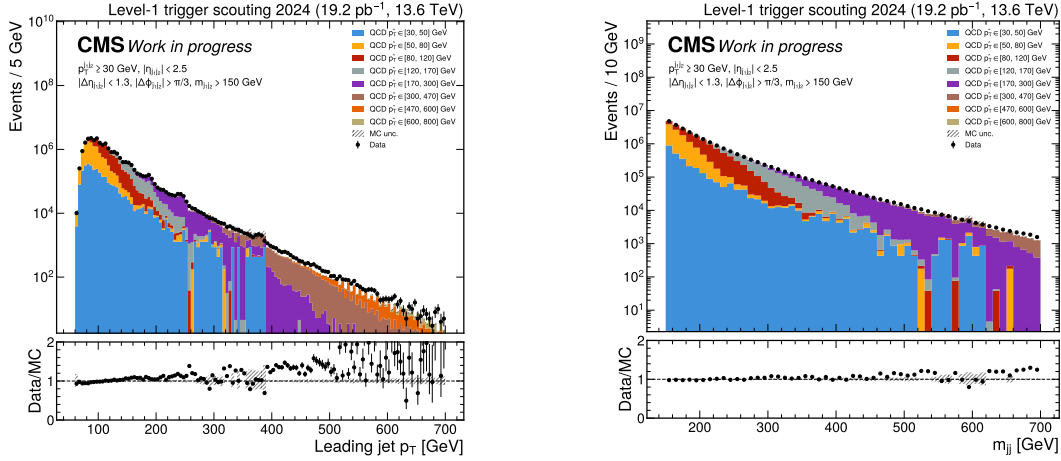


Figure 5.11: Distribution of the leading jet transverse momentum  $p_T$  (left) and the invariant mass of the two leading jets  $m_{jj}$  (right) compared to the QCD simulation normalized to the data.

changes, starting from around 55 GeV due to the simultaneous cuts on  $m_{jj}$  and  $|\Delta\eta|$ . Assuming a small pseudorapidity separation and a back-to-back system, this constrains  $p_T \approx M/2$ . Consequently, the  $p_T$  distribution now lies just above the region affected by energy correction in the calorimeter trigger, which impacted the 40 – 55 GeV range. Nonetheless, the increasing trend in the data/MC ratio remains visible.

Figure 5.11 (right) shows the invariant mass distribution of the dijet system. A similar trend in the data/MC ratio is observed, as seen in the  $p_T$  distribution. The regions with higher MC uncertainty are a result of a small number of events from high-weight samples, which produce events with high mass. Specifically, the low  $\hat{p}_T$  sample (blue) suffers from a lack of simulated events for masses above 400 GeV, failing to model the tail of the distribution. Nevertheless, the  $m_{jj}$  distribution in the data appears to be a smooth, falling function.

The azimuthal angular separation  $\Delta\phi$  and pseudorapidity separation  $\Delta\eta$  between the two leading jets are shown in Figure 5.12. In both cases, good agreement between data and simulation is observed. The azimuthal separation of the two jets shows the expected distribution, with a strong peak at  $|\Delta\phi| = \pi$ , as expected from QCD dijet production. The distribution for  $|\Delta\eta|$  is consistent with the one shown in Figure 5.10, displaying good agreement with the data.

## 5.4. MODELLING OF THE QCD BACKGROUND

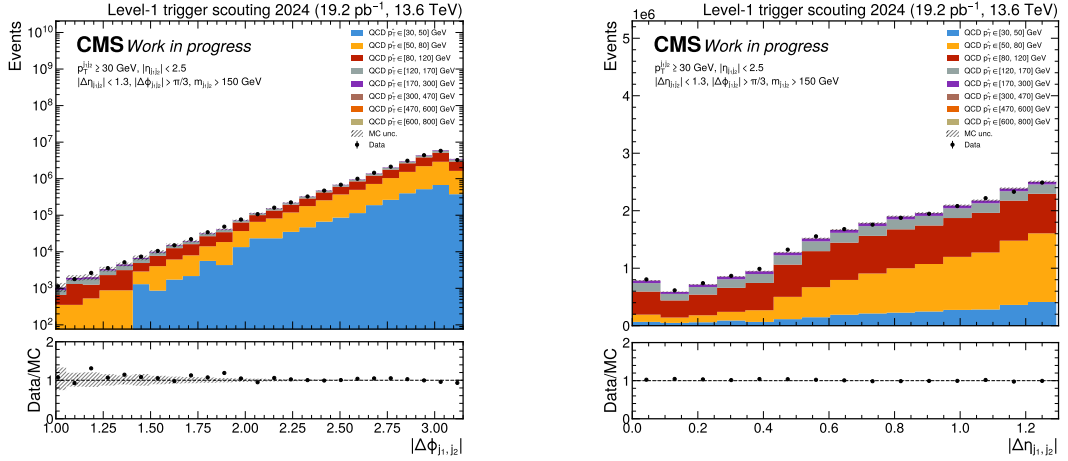


Figure 5.12: Azimuthal angular separation  $\Delta\phi$  (left) and pseudorapidity separation  $\Delta\eta$  (right) between the two leading jets compared to the QCD simulation normalized to the data for selected events with  $|\Delta\eta| < 1.3$ .

### 5.4 Modelling of the QCD background

In previous dijet searches, the main QCD background is modeled using a set of empirical functions. The signal extraction is performed via a likelihood fit of the dijet invariant mass spectrum, where the signal shape is extracted from the MC and the background modeled with an empirical function. This procedure provides an upper limit on the signal cross-section, more precisely, on  $\sigma \cdot \mathcal{B} \cdot \mathcal{A}$ , where  $\sigma$  is the cross-section,  $\mathcal{B}$  the branching ratio, and  $\mathcal{A}$  the acceptance. The focus of this section is on the fit under the background-only hypothesis, with the goal defining a strategy for the modelling of the background.

Following the approach used in previous searches[13, 148, 106, 142, 149, 145, 144], the dijet mass spectrum is defined by the following equation:

$$\frac{d\sigma}{dm_{jj}} = \frac{N_i}{L \cdot \Delta m_i} \quad (5.3)$$

where  $N_i$  represents the number of events in the  $i$ -th bin,  $\Delta m_i$  is the bin width, and  $L$  is the integrated luminosity. This formalism allows for variable bin sizes, although for the purpose of this test, a fixed bin size is used.



### 5.4.1 Parametric model

A four-parameter function is commonly used to model the QCD background, extensively utilized in dijet searches, with the following analytical form:

$$\frac{d\sigma}{dm_{jj}} = \frac{p_0(1 - m_{jj}/\sqrt{s})^{p_1}}{(m_{jj}/\sqrt{s})^{p_2+p_3} \log(m_{jj}/\sqrt{s})} \quad (5.4)$$

where,  $p_0$ ,  $p_1$ ,  $p_2$ , and  $p_3$  are the four free fit parameters, and  $\sqrt{s}$  is the  $pp$  collision energy. The model can be extended to include five free parameters, expressed as:

$$\frac{d\sigma}{dm_{jj}} = \frac{p_0(1 - m_{jj}/\sqrt{s})^{p_1}}{(m_{jj}/\sqrt{s})^{p_2+p_3} \log(m_{jj}/\sqrt{s}) + p_4 \log(m_{jj}/\sqrt{s})^2} \quad (5.5)$$

Similarly, other families of functions may be used. These functions are motivated by QCD calculations, where the term in the numerator behaves like the parton distribution functions at an average fractional momentum  $x$  of the two partons, while the term in the denominator provides a mass dependence similar to QCD matrix elements. The historical development of these functions is discussed in Ref. [13]. The background modelling function is typically selected from a pool of alternative functions using a Fisher test[158].

Examples of binned maximum likelihood fits performed using the 4-parameter function defined in Equation 5.4, selected amongst alternatives by the F-test, are shown in Figure 5.13. In both cases, the same function is adopted, as increasing the number of parameters did not improve the fit quality. In the first example, the fit range spans from the lowest available mass bin at 150 GeV up to 800 GeV. The fit results in large residuals, as seen in the bottom panel of the plots, along with a clear pattern in the distribution of the pulls. The  $\chi^2$  per degree of freedom,  $\chi^2/\text{ndf} = 547/20$ , further indicates the poor quality of the fit.

On the other hand, restricting the fit range to  $190 < m_{jj} < 600$  GeV appears to improve the fit quality, as seen in Figure 5.13 (right). In this case, the fit yields  $\chi^2/\text{ndf} = 64/16$ , with a more uniform distribution of pulls, where no evident patterns are visible. However, the  $\chi^2$  remains significantly larger than the number of degrees of freedom due to the extremely low statistical uncertainty of the data. Furthermore, as the number of events increases, the discrepancies observed in the pulls will likely become more pronounced.

## 5.4. MODELLING OF THE QCD BACKGROUND




Figure 5.13: Dijet mass spectrum (black points) compared to a fitted parametrization of the background (solid red curve). Two examples of fits are shown, differing in the fit range: the first (left) is performed over the range  $150 < m_{jj} < 800$  GeV, while the second covers  $190 < m_{jj} < 600$  GeV. The black dashed line represents the dijet mass distribution for the  $Z'$  signal with a mass of 350 GeV, assuming a cross-section of 100 pb. The bottom panel shows the difference between the data and the fitted parametrization, divided by the statistical uncertainty of the data.

This could suggest that the dijet mass distribution is not as smooth as it initially appears. This observation may arise from the fact that the distributions of jet  $p_T$  and  $\eta$  are themselves not entirely smooth. Given the large number of events being analyzed, these patterns become more evident. Specifically, differences in  $p_T$  distribution and the varying behavior of jets in different  $\eta$  regions naturally manifest in the dijet mass distribution.

### 5.4.2 Ratio method

An alternative approach to estimating the background, orthogonal to purely parametric fit of the QCD background, is the *ratio method*, introduced in Ref. [145]. The key idea behind this approach is to predict the background in the signal region (SR) using the dijet mass distribution from a control region (CR) and applying a transfer factor computed from QCD simulation. The transfer factor,  $\mathcal{R}$ , is computed from the QCD simulation as follows:

$$\mathcal{R}_i = \frac{N(m_i)_{\text{SR}}^{\text{Simulation}}}{N(m_i)_{\text{CR}}^{\text{Simulation}}} \quad (5.6)$$

where  $N(m_i)^{\text{Simulation}}$  represents the number of events in the  $i$ -th bin of the dijet mass distribution in the signal and control regions. This ratio is then expected to reflect the ratio of distributions computed from the data, allowing it to be used to map the mass distribution from the control region to the signal region:

$$N(m_i)_{\text{SR}}^{\text{Prediction}} = \mathcal{R}_i N(m_i)_{\text{CR}}^{\text{Data}} \quad (5.7)$$

This method can be extended to include multiple control regions, as shown in the original paper, where additional control regions are used to estimate a correction factor  $C$ , improving prediction quality. However, in this initial test, a single control region is used.

The signal and control regions are defined similarly to the selections described in the previous section, where it was observed that the signal is more significant for smaller pseudorapidity separation between the jets, whereas QCD events dominate for larger values. For this test, the signal region is defined as  $|\Delta\eta| < 1.3$ , and the control region is selected to be  $1.3 < |\Delta\eta| < 2$ , where the kinematics are similar to the signal region but with a lower fraction of signal events. Moreover, the MC description of the data in this region remains accurate. A careful study should be conducted to determine the best range for the control region, as minimizing signal contamination is crucial. A significant number of signal events in the control region would lead to an overestimation of the background in the signal region, reducing the signal significance.

The dijet mass distribution for the control region is shown in Figure 5.14 (left), where good agreement with the MC simulation is observed, exhibiting a similar pattern to that seen in the signal region in Figure 5.11 (right).

The dijet mass distribution, compared to the prediction from the ratio method, is presented in Figure 5.14 (right). Apart from a visible bias in the lower mass region, the overall agreement appears satisfactory, indicating that this method could be a viable alternative to the parametric fit. An advantage of this approach is its ability to naturally model the features present in the trigger objects.

However, this method relies on having an accurate simulation of the data. It is not necessary for the simulation to match the number of events in the data, as it is only used to compute a transfer factor applied to the data itself. Nevertheless, a large number of MC events is needed as the uncertainty on the

## 5.5. FUTURE STEPS

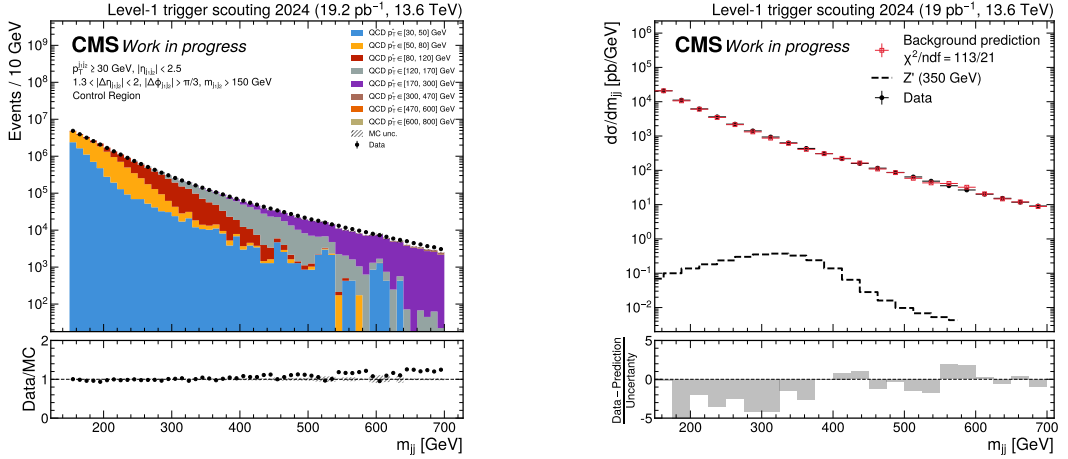


Figure 5.14: Left: Dijet mass distribution for jets with pseudorapidity separation  $1.3 < |\Delta\eta| < 2$ . Right: Mass spectrum (black points) compared to the background prediction (red points) obtained from the control region  $1.3 < |\Delta\eta| < 2$ . The black dashed line represents the dijet mass distribution for a  $Z'$  signal with mass 350 GeV, assuming a cross-section of 100 pb. The bottom panel shows the difference between the data and the background prediction, divided by the statistical uncertainty of the data, including those from the control region.

ratio  $\mathcal{R}$  impacts the background estimation. The simulation needs to accurately model the  $m_{jj}$  distribution, as regions with high MC uncertainty such as those seen in Figure 5.14 (left) for the region above 500 GeV translate into a higher uncertainty in the prediction, visible as fluctuations in the 500–600 GeV interval. This results in a systematic error that limits the benefits of the large data sample. Moreover, if differences between the signal and control regions are not well modeled by the simulation, it can introduce biases, such as the one visible at the lower end of the spectrum, potentially diminishing the sensitivity to a signal.

## 5.5 Future steps

This section has explored the initial steps toward using L1 objects for a dijet analysis using the current L1T objects available before the improvements foreseen with the Phase2 upgrade. While the virtually unlimited number of events offers the potential to enhance sensitivity to new dijet resonances in the low-mass regime, it also presents several challenges. These challenges span from managing and processing large datasets to addressing computational and methodological issues in the statistical analysis.

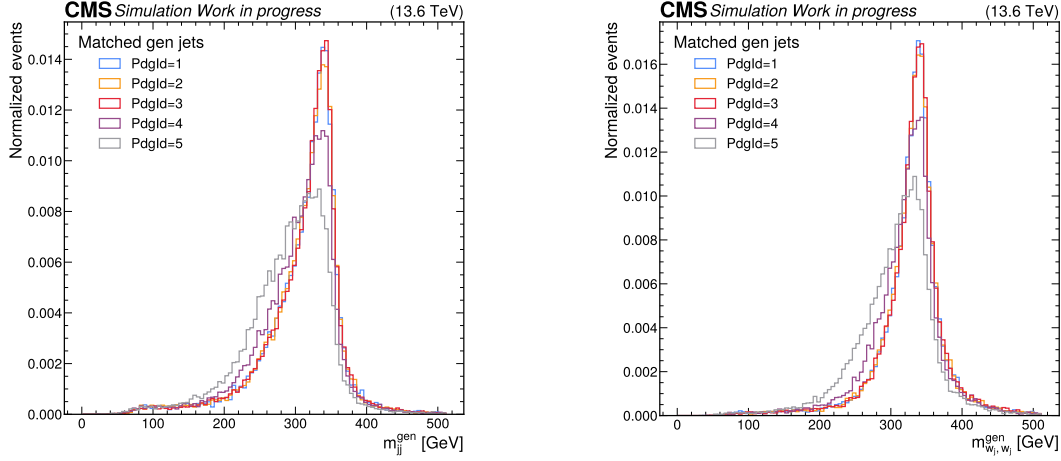


Figure 5.15: Dijet mass distribution of generator-level jets matched to the partons produced by the decay of a  $Z'$  boson with a mass of 350 GeV. On the left, jets as produced by the calorimeter trigger are used, while on the right, the same distribution is computed using wide-jets. In both cases, the contribution of each quark flavor is shown.

The current events selections have been applied based on educated guesses, informed by theoretical motivations and results from previous analyses.

A thorough study is needed to improve our understanding of their impact on signal significance. This holds true for both background prediction methods, especially for the one based on the control region, which has shown promising results but requires further refinement. Nonetheless, both approaches could be used in parallel to provide a reliable cross-check of the methods.

Besides the analysis strategy, gaining complete control and understanding over the objects being used, such as jets produced by the trigger, is essential. Jets as produced and used in the trigger were employed in this work. Their response, when compared to offline reconstruction, shows some alignment, as observed in the signal simulation in Figure 5.7 (right), which provides hope for the potential success of this analysis. However, dedicated calibrations for the jets, both in  $p_T$  and  $\eta$ , should be derived, similar to the approach used in offline and HLT scouting jets.

Before applying these techniques to the data, a comprehensive understanding of the behavior of jets at small  $\Delta R$ , as shown in Figure 5.9, is needed to avoid double-counting of jet energy due to the jet finder algorithm implemented in the trigger, along with the effects of PU subtraction.

Another potential improvement involves combining geometrically close jets

## 5.5. FUTURE STEPS

into “wide-jets”. This can be done by using the leading and subleading jets as “seeds” and incorporating all jets within  $\Delta R < 1$ . This method helps to recover final-state radiation that escapes the jet cone, sharpening the signal peak, especially for lighter quarks, as seen in Figure 5.15 comparing standard (left) and wide (right) jets.

It is also necessary to address how certain  $p_T$  regions, such as 40 – 60 GeV, contain no jets, as this will impact the dijet mass distribution. The combination of cuts on the invariant mass and pseudorapidity separation between the two jets helps to filter out many events from this region. However, in order to extend sensitivity to lower dijet masses and avoid dealing with problematic events, this correction should ideally be removed from the trigger. This is not feasible in the short term, as the trigger is carefully tuned for optimal performance. In this type of study, the primary focus is on “physics performance”, such as optimal jet response, which may not always align with the trigger’s operational goals.

One potential solution, assuming sufficient spare links are available, is to collect the raw jets instead of those processed by the calorimeter trigger’s Layer-2, where jet energy corrections and PU subtraction are applied. This would allow the derivation of a custom set of corrections that could offer improved performance since they would not be constrained by the limitations of Look-Up Tables required to fit inside an FPGA. Additionally, these corrections would be designed with physics performance in mind rather than optimizing the trigger rate.

A particularly promising possibility is to directly collect the raw energy from trigger towers and implement a custom jet-finding algorithm, which could overcome the constraints imposed by the trigger, such as latency and resource utilization. Jets could be reconstructed either within dedicated scouting boards, leveraging new technologies such as the AI Engine present in AMD Versal<sup>1</sup>, or in scouting processing units. If throughput proves too high, even collecting a fraction of the events or data from part of the detector (e.g., only the barrel) would be invaluable. This approach could not only benefit dijet searches but also enable the study of processes characterized by a large multiplicity of soft, unclustered energy patterns and the application of anomaly detection using full granularity. This concept is similar to what is being explored by the CICADA[159] project at

---

<sup>1</sup><https://www.xilinx.com/products/technology/ai-engine.html>

the CMS experiment, where a convolutional neural network is used to identify anomalies in the calorimeter.







## Conclusions

The work developed in this thesis concerns the implementation of a novel data-taking strategy in the CMS experiment, namely the development of a 40 MHz Level-1 trigger scouting system and focusing on its potential to explore uncharted regions of phase space that would otherwise be inaccessible due to traditional trigger limitations. The L1 scouting approach offers a unique opportunity to probe areas such as low-mass dijet analyses, which are often limited by trigger thresholds due to the overwhelming background from QCD multi-jet production. By bypassing these constraints, scouting opens up new possibilities for discovering physics beyond the Standard Model.

One of the major achievements presented in this thesis is the first successful implementation and operation of the 40 MHz L1 scouting demonstrator, which was deployed and operated in parallel to the CMS data acquisition system throughout the 2024 LHC run. The data collected with this system during the 2024  $pp$  run was instrumental in making the first steps toward a low-mass dijet search, where different techniques for background modeling were explored along with an initial study of a  $Z'$  signal's characteristics using L1 trigger objects. This milestone highlights the potential of using the scouting system to perform analyses that were previously limited by traditional trigger selections. Furthermore, the development of the demonstrator serves as a crucial step toward the implementation of the full scouting system envisioned for the Phase-2 upgrade of the CMS Level-1 trigger.

In addition to the L1 trigger scouting, this work also explored the feasibility of reading out front-end data in a triggerless mode, demonstrated in the context

of the DT Phase-2 demonstrator. The data collected with such strategy proved invaluable for detector monitoring and background studies due to its unbiased nature. This approach provided deeper insights into detector performance and environmental backgrounds while opening new possibilities for extending triggerless data acquisition techniques to other parts of the CMS detector. Such methods could complement the L1-Trigger scouting system and provide additional opportunities for analysis involving raw data.

Furthermore, the implementation of a new data acquisition scheme based on Remote Direct Memory Access technology has been presented. This innovative approach holds the potential to enhance both scouting and standard data acquisition systems by improving the efficiency of data transfer from the detector to server-based processing farms.



## L1 Scouting data formats

The list of hardware quantities and scales collected during Run-3 with the L1-Trigger scouting system are shown in Table A.1 for the muons, Table A.2 for calorimeter objects and Table A.3 for the muon stubs collected from the Barrel Muon Track Finder.

Parameter	Bits	Range	Notes
$\phi_{\text{ext}}$	10	$2\pi$	$\phi$ extrapolated at the vertex
$p_T$	9	$0 \dots 256 \text{ GeV}$	$p_T$ assignment with vertex constraint
quality	4	$0 \dots 15$	Muon quality code
$\eta_{\text{ext}}$	1+8	$-2.45 \dots 2.45$	Pseudorapidity extrapolated at the vertex
Isolation	2	$0 \dots 3$	Muon isolation (not used up to date)
Charge sign	1	sign	Charge of the muon 0: $\mu^+$ 1: $\mu^-$
Charge valid	1	bool	Validity of charge assignment
Index	7	$0 \dots 107$	Track finder index $\text{idx} \in [36, 70]$ : barrel $\text{idx} \in [17, 35] \cup [71, 89]$ : overlap $\text{idx} \in [0, 16] \cup [90, 107]$ : endcap
$\phi$	10	$2\pi$	$\phi$ at second muon station
$p_T^{\text{unc}}$	8	$0 \dots 256 \text{ GeV}$	$p_T$ assignment without vertex constraint
HS trigger	1	bool	Hadronic shower trigger
$d_{xy}$	2	$0 \dots 3$	Impact parameter

Table A.1: Scale and definitions of the  $\mu$ GMT muon primitives transmitted from  $\mu$ GMT to the scouting system.

Object	Parameter	Bits	Range	Notes
Jet	$E_T$	11	0...1024 GeV	
	$\eta$	7+1	-5...5	
	$\phi$	8	$2\pi$	
	Quality flags	2		Not used
$e/\gamma$	$E_T$	9	0...256 GeV	
	$\eta$	7+1	-5...5	
	$\phi$	8	$2\pi$	
	Isolation	2	0...3	01 or 11 for isolated $e/\gamma$
Tau	$E_T$	9	0...256 GeV	
	$\eta$	7+1	-5...5	
	$\phi$	8	$2\pi$	
	Isolation	2	0...3	Not used
ET	$E_T$ [ET]	12	0...2048 GeV	Total $E_T$
	$E_T$ [ETTEM]	12	0...2048 GeV	Total $E_T$ in ECAL
HT	$E_T$	12	0...2048 GeV	Total hadronic $E_T$
	Tower count	13	0...8191	Sum of firing trigger towers
ET <sub>miss</sub>	$E_T$	12	0...2048 GeV	Total missing $E_T$
	$\phi$	8	$2\pi$	$\phi$ direction of missing $E_T$
HT <sub>miss</sub>	$E_T$	12	0...2048 GeV	Total hadronic missing $E_T$
	$\phi$	8	$2\pi$	$\phi$ direction of hadronic missing $E_T$

Table A.2: Scale and definitions of the calorimeter primitives transmitted from the DeMux board to the scouting system. For the calorimetric sums, only the main ones are reported.

Parameter	Bits	Range	Notes
valid	1	flag	Keep or Zero Suppress stub
$\phi$	12	-2048...2047	Relative position of a segment inside a sector
$\phi_b$	10	-512...511	Bending angle of stub
quality	3	0...7	Encodes the number of superlayers used
$\eta$ hits	7	pattern	Each bit corresponds to one chamber area 0: no hit (less then 3 SLs hit) 1: hit (3 or more SLs hit)
$\eta$ quality	7	pattern	Each bit corresponds to one chamber area 0: 3 SLs hit 0: 4 SLs hit
station	2	0...3	Station of stub
wheel	3	-2...2	Wheel of stub
reserved	3		Space reserved

Table A.3: Scale and definitions of the super primitive parameters transmitted from BMTF to the scouting system.

# References

- [1] M. Herrero. “The Standard model”. In: *NATO Sci. Ser. C* 534 (1999). Ed. by T. Ferbel, pp. 1–59. DOI: [10.1007/978-94-011-4689-0\\_1](https://doi.org/10.1007/978-94-011-4689-0_1). arXiv: [hep-ph/9812242](https://arxiv.org/abs/hep-ph/9812242).
- [2] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [3] CMS Collaboration. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235 \[hep-ex\]](https://arxiv.org/abs/1207.7235).
- [4] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214 \[hep-ex\]](https://arxiv.org/abs/1207.7214).
- [5] John Ellis. “The need for new physics”. In: *Nature* 481.7379 (Jan. 2012), pp. 24–24. ISSN: 1476-4687. DOI: [10.1038/481024a](https://doi.org/10.1038/481024a). URL: <https://doi.org/10.1038/481024a>.
- [6] Michael E Peskin. “Beyond the Standard Model”. In: (1997). URL: <https://cds.cern.ch/record/326995>.
- [7] Pran Nath et al. “The Hunt for New Physics at the Large Hadron Collider”. In: *Nucl. Phys. B Proc. Suppl.* 200-202 (2010). Ed. by George Alverson, Pran Nath, and Brent D. Nelson, pp. 185–417. DOI: [10.1016/j.nuclphysbps.2010.03.001](https://doi.org/10.1016/j.nuclphysbps.2010.03.001). arXiv: [1001.2693 \[hep-ph\]](https://arxiv.org/abs/1001.2693).
- [8] F. Sauerburger. *LHC cross section plot*. <https://lhc-xsecs.org/>.
- [9] Henning Bahl. “BSM theory perspectives for Run 3”. In: *PoS LHCP2023* (2024), p. 213. DOI: [10.22323/1.450.0213](https://doi.org/10.22323/1.450.0213).

## REFERENCES

- [10] Aniello Spiezia. “Status and prospect of LHC BSM searches”. In: *PoS ALPS2018* (2018), p. 025. doi: [10.22323/1.330.0025](#).
- [11] Philip Bechtle et al. “What if the LHC does not find supersymmetry in the  $\sqrt{s} = 7$  TeV run?” In: *Phys. Rev. D* 84 (2011), p. 011701. doi: [10.1103/PhysRevD.84.011701](#). arXiv: [1102.4693 \[hep-ph\]](#).
- [12] Simon Knapen and Steven Lowette. “A Guide to Hunting Long-Lived Particles at the LHC”. In: *Ann. Rev. Nucl. Part. Sci.* 73 (2023), pp. 421–449. doi: [10.1146/annurev-nucl-101920-013011](#). arXiv: [2212.03883 \[hep-ph\]](#).
- [13] Robert M. Harris and Konstantinos Kousouris. “Searches for Dijet Resonances at Hadron Colliders”. In: *Int. J. Mod. Phys. A* 26 (2011), pp. 5005–5055. doi: [10.1142/S0217751X11054905](#). arXiv: [1110.5302 \[hep-ex\]](#).
- [14] Malcolm Fairbairn et al. “Constraints on  $Z'$  models from LHC dijet searches and implications for dark matter”. In: *JHEP* 09 (2016), p. 018. doi: [10.1007/JHEP09\(2016\)018](#). arXiv: [1605.07940 \[hep-ph\]](#).
- [15] Ewa Lopienska. “The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022”. In: (2022). General Photo. URL: <https://cds.cern.ch/record/2800984>.
- [16] R Bailey and Paul Collier. *Standard Filling Schemes for Various LHC Operation Modes*. Tech. rep. Geneva: CERN, 2003. URL: <https://cds.cern.ch/record/691782>.
- [17] C Wiesner et al. “Abort Gap Keeper and Abort Gap Protection”. In: (2019), pp. 141–145. URL: <https://cds.cern.ch/record/2813530>.
- [18] Arkadiusz Andrzej Gorzawski. “Luminosity control and beam orbit stability with beta star leveling at LHC and HL-LHC”. presented 30 Nov 2016. Ecole Polytechnique, Lausanne, 2016. doi: [10.5075/epfl-thesis-7338](#). URL: <https://cds.cern.ch/record/2238309>.
- [19] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. doi: [10.23731/CYRM-2020-0010](#). URL: <https://cds.cern.ch/record/2749422>.
- [20] *High Luminosity LHC Project*. <https://hilumilhc.web.cern.ch/content/hl-lhc-project> [Accessed: 2024].

- [21] *Public CMS Luminosity Information*. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults> [Accessed: 2024].
- [22] ATLAS collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08003. DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003). URL: <https://cds.cern.ch/record/1129811>.
- [23] CMS Collaboration. “The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08004. DOI: [10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004). URL: <https://cds.cern.ch/record/1129810>.
- [24] ALICE collaboration. “The ALICE experiment at the CERN LHC. A Large Ion Collider Experiment”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08002. DOI: [10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002). URL: <https://cds.cern.ch/record/1129812>.
- [25] LHCb collaboration. “The LHCb Detector at the LHC”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08005. DOI: [10.1088/1748-0221/3/08/S08005](https://doi.org/10.1088/1748-0221/3/08/S08005). URL: <https://cds.cern.ch/record/1129809>.
- [26] CMS Collaboration. *The CMS tracker system project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: <https://cds.cern.ch/record/368412>.
- [27] CMS Collaboration. *CMS Technical Design Report for the Pixel Detector Upgrade*. Tech. rep. Additional contacts: Jeffrey Spalding, Fermilab, [Jeffrey.Spalding@cern.ch](mailto:Jeffrey.Spalding@cern.ch) Didier Contardo, Universite Claude Bernard-Lyon I, [didier.claude.contardo@cern.ch](mailto:didier.claude.contardo@cern.ch). 2012. URL: <https://cds.cern.ch/record/1481838>.
- [28] CMS Collaboration. “The CMS Silicon Strip Tracker”. In: *J. Phys.: Conf. Ser.* 41 (2006). 8 pages, 8 figures, talk given at XIX EPS NPDC Conference on New Trends in Nuclear Physics Applications and Technology, September 5-9, 2005 Pavia, Italy Subj-class: Instrumentation and Detectors, pp. 127–134. DOI: [10.1088/1742-6596/41/1/011](https://doi.org/10.1088/1742-6596/41/1/011). URL: <https://cds.cern.ch/record/914891>.
- [29] CMS Collaboration. *The CMS electromagnetic calorimeter project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: <https://cds.cern.ch/record/349375>.

## REFERENCES

- [30] CMS Collaboration. “Energy resolution of the barrel of the CMS electromagnetic calorimeter”. In: *JINST* 2 (2007), P04004. DOI: [10.1088/1748-0221/2/04/P04004](https://doi.org/10.1088/1748-0221/2/04/P04004).
- [31] CMS Collaboration. “Energy Calibration and Resolution of the CMS Electromagnetic Calorimeter in  $pp$  Collisions at  $\sqrt{s} = 7$  TeV”. In: *JINST* 8 (2013), P09009. DOI: [10.1088/1748-0221/8/09/P09009](https://doi.org/10.1088/1748-0221/8/09/P09009). arXiv: [1306.2016](https://arxiv.org/abs/1306.2016) [hep-ex].
- [32] CMS Collaboration. “Performance of calorimeters at the LHC”. In: *J. Phys. Conf. Ser.* 293 (2011). Ed. by Yifang Wang, p. 012001. DOI: [10.1088/1742-6596/293/1/012001](https://doi.org/10.1088/1742-6596/293/1/012001).
- [33] CMS Collaboration. “Jet energy scale and resolution in the CMS experiment in  $pp$  collisions at 8 TeV”. In: *JINST* 12.02 (2017), P02014. DOI: [10.1088/1748-0221/12/02/P02014](https://doi.org/10.1088/1748-0221/12/02/P02014). arXiv: [1607.03663](https://arxiv.org/abs/1607.03663) [hep-ex].
- [34] CMS Collaboration. *The CMS muon project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: <https://cds.cern.ch/record/343814>.
- [35] CMS Collaboration. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at  $\sqrt{s} = 13$  TeV”. In: *JINST* 13.06 (2018), P06015. DOI: [10.1088/1748-0221/13/06/P06015](https://doi.org/10.1088/1748-0221/13/06/P06015). arXiv: [1804.04528](https://arxiv.org/abs/1804.04528) [physics.ins-det].
- [36] CMS Collaboration. “Performance of the CMS drift tube chambers with cosmic rays”. In: *Journal of Instrumentation* 5.3 (Mar. 2010), T03015. DOI: [10.1088/1748-0221/5/03/T03015](https://doi.org/10.1088/1748-0221/5/03/T03015). arXiv: [0911.4855](https://arxiv.org/abs/0911.4855) [physics.ins-det].
- [37] CMS Collaboration. “Commissioning of the CMS DT electronics under magnetic field”. In: (2009). DOI: [10.5170/CERN-2009-006.81](https://doi.org/10.5170/CERN-2009-006.81). URL: <https://cds.cern.ch/record/1234871>.
- [38] CMS Collaboration. “CMS. The TriDAS project. Technical design report, vol. 1: The trigger systems”. In: (Dec. 2000).
- [39] CMS Collaboration. “Commissioning of the CMS High Level Trigger”. In: *JINST* 4 (2009), P10005. DOI: [10.1088/1748-0221/4/10/P10005](https://doi.org/10.1088/1748-0221/4/10/P10005). arXiv: [0908.1065](https://arxiv.org/abs/0908.1065) [physics.ins-det].



- [40] CMS Collaboration. “Run-3 Commissioning of CMS Online HLT reconstruction using GPUs”. In: *EPJ Web Conf.* 295 (2024), p. 11020. doi: [10.1051/epjconf/202429511020](https://doi.org/10.1051/epjconf/202429511020).
- [41] CMS Collaboration. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Tech. rep. Additional contacts: Jeffrey Spalding, Fermilab, Jeffrey.Spalding@cern.ch Didier Contardo, Universite Claude Bernard-Lyon I, didier.claude.contardo@cern.ch. 2013. URL: <https://cds.cern.ch/record/1556311>.
- [42] CMS Collaboration. “The CMS Level-1 muon triggers for the LHC Run II”. In: *PoS ICHEP2018* (2019), p. 918. doi: [10.22323/1.340.0918](https://doi.org/10.22323/1.340.0918).
- [43] CMS Collaboration. “Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger”. In: *JINST* 11.01 (2016), p. C01051. doi: [10.1088/1748-0221/11/01/C01051](https://doi.org/10.1088/1748-0221/11/01/C01051). arXiv: [1511.05855 \[physics.ins-det\]](https://arxiv.org/abs/1511.05855).
- [44] R. Frazier et al. “A demonstration of a time multiplexed trigger for the CMS experiment”. In: *JINST* 7 (2012), p. C01060. doi: [10.1088/1748-0221/7/01/C01060](https://doi.org/10.1088/1748-0221/7/01/C01060).
- [45] Johannes Wittmann et al. “The upgrade of the CMS Global Trigger”. In: *JINST* 11.02 (2016), p. C02029. doi: [10.1088/1748-0221/11/02/C02029](https://doi.org/10.1088/1748-0221/11/02/C02029).
- [46] CMS Collaboration. “Development of the CMS detector for the CERN LHC Run 3”. In: *JINST* 19.05 (2024), P05064. doi: [10.1088/1748-0221/19/05/P05064](https://doi.org/10.1088/1748-0221/19/05/P05064). arXiv: [2309.05466 \[physics.ins-det\]](https://arxiv.org/abs/2309.05466).
- [47] CMS Collaboration. *New trigger strategies for CMS during Run 3*. Tech. rep. Geneva: CERN, 2022. doi: [10.22323/1.414.0681](https://doi.org/10.22323/1.414.0681). URL: <https://cds.cern.ch/record/2842439>.
- [48] CMS Collaboration. “Anomaly Detection in the CMS Global Trigger Test Crate for Run 3”. In: (2023). URL: <https://cds.cern.ch/record/2876546>.
- [49] CMS Collaboration. “Enriching the Physics Program of the CMS Experiment via Data Scouting and Data Parking”. In: (Mar. 2024). arXiv: [2403.16134 \[hep-ex\]](https://arxiv.org/abs/2403.16134).

## REFERENCES

- [50] CMS Collaboration. *Technical Proposal for the Phase-II Upgrade of the CMS Detector*. Tech. rep. Upgrade Project Leader Deputies: Lucia Silvestris (INFN-Bari), Jeremy Mans (University of Minnesota) Additional contacts: Lucia.Silvestris@cern.ch, Jeremy.Mans@cern.ch. Geneva, 2015. doi: [10 . 17181 / CERN . VU8I . D59J](https://doi.org/10.17181/CERN.VU8I.D59J). URL: <https://cds.cern.ch/record/2020886>.
- [51] CMS Collaboration. *The Phase-2 Upgrade of the CMS Tracker*. Tech. rep. Geneva: CERN, 2017. doi: [10 . 17181 / CERN . QZ28 . FLHW](https://doi.org/10.17181/CERN.QZ28.FLHW). URL: <https://cds.cern.ch/record/2272264>.
- [52] CMS Collaboration. *A MIP Timing Detector for the CMS Phase-2 Upgrade*. Tech. rep. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2667167>.
- [53] CMS Collaboration. *The Phase-2 Upgrade of the CMS Barrel Calorimeters*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283187>.
- [54] CMS Collaboration. *The Phase-2 Upgrade of the CMS Endcap Calorimeter*. Tech. rep. Geneva: CERN, 2017. doi: [10 . 17181 / CERN . IV8M . 1JY2](https://doi.org/10.17181/CERN.IV8M.1JY2). URL: <https://cds.cern.ch/record/2293646>.
- [55] CMS Collaboration. *The Phase-2 Upgrade of the CMS Muon Detectors*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283189>.
- [56] CMS Collaboration. *The Phase-2 Upgrade of the CMS Level-1 Trigger*. Tech. rep. Final version. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2714892>.
- [57] CMS Collaboration. “The CMS Particle Flow Algorithm”. In: *EPJ Web Conf.* 191 (2018), p. 02016. doi: [10 . 1051 / epjconf / 201819102016](https://doi.org/10.1051/epjconf/201819102016). URL: <https://cds.cern.ch/record/2678077>.
- [58] Daniele Bertolini et al. “Pileup Per Particle Identification”. In: *JHEP* 10 (2014). Comments: v1 - 23 pages, 10 figures, p. 59. doi: [10 . 1007 / JHEP10\(2014\)059](https://doi.org/10.1007/JHEP10(2014)059). arXiv: [1407 . 6013](https://arxiv.org/abs/1407.6013). URL: <https://cds.cern.ch/record/1745357>.
- [59] Matteo Migliorini et al. “A horizontally scalable online processing system for trigger-less data acquisition”. In: *Nuclear Instruments and Methods in Physics Research A* 1036, 166869 (Aug. 2022), p. 166869. doi: [10 . 1016 / j . nima . 2022 . 166869](https://doi.org/10.1016/j.nima.2022.166869). arXiv: [2111 . 05155](https://arxiv.org/abs/2111.05155) [physics.ins-det].

- [60] M. Migliorini et al. “Muon trigger with fast Neural Networks on FPGA, a demonstrator”. In: *Journal of Physics: Conference Series* 2374.1 (Nov. 2022), p. 012099. DOI: [10.1088/1742-6596/2374/1/012099](https://doi.org/10.1088/1742-6596/2374/1/012099). URL: <https://dx.doi.org/10.1088/1742-6596/2374/1/012099>.
- [61] G. Bortolato et al. “Front-end RDMA over Converged Ethernet, real-time firmware simulation”. In: *Journal of Instrumentation* 19.3, C03038 (Mar. 2024), p. C03038. DOI: [10.1088/1748-0221/19/03/C03038](https://doi.org/10.1088/1748-0221/19/03/C03038).
- [62] N. Amapane et al. “Study of muon pair production from positron annihilation at threshold energy”. In: *JINST* 15.01 (2020), P01036. DOI: [10.1088/1748-0221/15/01/P01036](https://doi.org/10.1088/1748-0221/15/01/P01036). arXiv: [1909.13716](https://arxiv.org/abs/1909.13716) [hep-ex].
- [63] F Gonella and M Pegoraro. “The MAD, a Full Custom ASIC for the CMS Barrel Muon Chambers Front End Electronics”. In: (2001). DOI: [10.5170/CERN-2001-005.204](https://doi.org/10.5170/CERN-2001-005.204). URL: <https://cds.cern.ch/record/529411>.
- [64] CMS Collaboration. “Electronics Developments for Phase-2 Upgrade of CMS Drift Tubes”. In: *PoS TWEPP2018* (2019), p. 035. DOI: [10.22323/1.343.0035](https://doi.org/10.22323/1.343.0035).
- [65] Bruce Taylor. “TTC distribution for LHC detectors”. In: *Nuclear Science, IEEE Transactions on* 45 (July 1998), pp. 821–828. DOI: [10.1109/23.682644](https://doi.org/10.1109/23.682644).
- [66] S. Baron et al. “Implementing the GBT data transmission protocol in FPGAs”. In: *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP09)*. CERN. CERN, 2009. DOI: [10.5170/CERN-2009-006.631](https://doi.org/10.5170/CERN-2009-006.631).
- [67] *DMA for PCI Express (PCIe) Subsystem*. available at <https://www.xilinx.com/products/intellectual-property/pcie-dma.html>.
- [68] *AXI DMA Controller*. available at [https://www.xilinx.com/products/intellectual-property/axi\\_dma.html](https://www.xilinx.com/products/intellectual-property/axi_dma.html).
- [69] CMS Collaboration. “The Analytical Method algorithm for trigger primitives generation at the LHC Drift Tubes detector”. In: *Nucl. Instrum. Meth. A* 1049 (2023), p. 168103. DOI: [10.1016/j.nima.2023.168103](https://doi.org/10.1016/j.nima.2023.168103). arXiv: [2302.01666](https://arxiv.org/abs/2302.01666) [hep-ex].
- [70] Paulo Moreira et al. *lpGBT documentation: release*. 2022. URL: <https://cds.cern.ch/record/2809058>.

## REFERENCES

- [71] CMS Collaboration. “The OBDT-theta board: time digitization for the theta view of Drift Tubes chambers.” In: *Journal of Instrumentation* 19.3, C03035 (Mar. 2024), p. C03035. DOI: [10.1088/1748-0221/19/03/C03035](https://doi.org/10.1088/1748-0221/19/03/C03035). arXiv: [2401.03818](https://arxiv.org/abs/2401.03818) [hep-ex].
- [72] J. L. Autran et al. “Characterization of atmospheric muons at sea level using a cosmic ray telescope”. In: *Nucl. Instrum. Meth. A* 903 (2018), pp. 77–84. DOI: [10.1016/j.nima.2018.06.038](https://doi.org/10.1016/j.nima.2018.06.038).
- [73] Gaia Grosso et al. “Triggerless data acquisition pipeline for Machine Learning based statistical anomaly detection”. In: *European Physical Journal Web of Conferences*. Vol. 295. European Physical Journal Web of Conferences. May 2024, 02033, p. 02033. DOI: [10.1051/epjconf/202429502033](https://doi.org/10.1051/epjconf/202429502033).
- [74] Gaia Grosso et al. “Fast kernel methods for data quality monitoring as a goodness-of-fit test”. In: *Machine Learning: Science and Technology* 4.3, 035029 (Sept. 2023), p. 035029. DOI: [10.1088/2632-2153/acebb7](https://doi.org/10.1088/2632-2153/acebb7). arXiv: [2303.05413](https://arxiv.org/abs/2303.05413) [hep-ex].
- [75] CMS Collaboration. *Trigger-less readout and unbiased data quality monitoring of the CMS Drift Tubes muon detector*. Tech. rep. 01. Geneva: CERN, 2023. DOI: [10.1088/1748-0221/18/01/C01003](https://doi.org/10.1088/1748-0221/18/01/C01003). URL: <https://cds.cern.ch/record/2839764>.
- [76] *Apache Kafka*. available at <https://kafka.apache.org/>.
- [77] Matthew Rocklin. “Dask: Parallel computation with blocked algorithms and task scheduling”. In: *Proceedings of the 14th python in science conference*. 130-136. Citeseer. 2015.
- [78] *Kubernetes*. available at <https://kubernetes.io/>.
- [79] Paolo Andreetto et al. “The CloudVeneto initiative: 10 years of operations to support interdisciplinary open science”. In: *European Physical Journal Web of Conferences*. Vol. 295. European Physical Journal Web of Conferences. May 2024, 07042, p. 07042. DOI: [10.1051/epjconf/202429507042](https://doi.org/10.1051/epjconf/202429507042).
- [80] *ZeroMQ, an open-source universal messaging library*. available at <https://zeromq.org/>.
- [81] CMS Collaboration. “A slice-test demonstrator for the upgrade of the CMS drift tubes at High-Luminosity LHC”. In: *Journal of Instrumentation* 17.10, C10007 (Oct. 2022), p. C10007. DOI: [10.1088/1748-0221/17/10/C10007](https://doi.org/10.1088/1748-0221/17/10/C10007).

- [82] CMS Collaboration. *The Phase-2 upgrade of the CMS Drift Tube detectors for High Luminosity LHC*. Tech. rep. Geneva: CERN, 2023. doi: [10.1016/j.nima.2022.167657](https://cds.cern.ch/record/2815418). URL: <https://cds.cern.ch/record/2815418>.
- [83] CMS Collaboration. “The Analytical Method algorithm for trigger primitives generation at the LHC Drift Tubes detector”. In: *Nuclear Instruments and Methods in Physics Research A* 1049, 168103 (Apr. 2023), p. 168103. doi: [10.1016/j.nima.2023.168103](https://arxiv.org/abs/2302.01666). arXiv: [2302.01666](https://arxiv.org/abs/2302.01666) [hep-ex].
- [84] CMS Collaboration. “Calibration of the CMS drift tube chambers and measurement of the drift velocity with cosmic rays”. In: *Journal of Instrumentation* 5.3 (Mar. 2010), T03016. doi: [10.1088/1748-0221/5/03/T03016](https://arxiv.org/abs/0911.4895). arXiv: [0911.4895](https://arxiv.org/abs/0911.4895) [physics.ins-det].
- [85] CMS Collaboration. “Background in the CMS muon detectors: simulation and measurements with pp collision data”. In: *Journal of Instrumentation* 15.2 (Feb. 2020), p. C02038. doi: [10.1088/1748-0221/15/02/C02038](https://arxiv.org/abs/2002.09060). arXiv: [2002.09060](https://arxiv.org/abs/2002.09060) [physics.ins-det].
- [86] CMS Collaboration. “Background in the CMS drift tubes: measurements with LHC collision data and implications for detector longevity at HL-LHC”. In: *Journal of Instrumentation* 17.9, C09028 (Sept. 2022), p. C09028. doi: [10.1088/1748-0221/17/09/C09028](https://arxiv.org/abs/2002.09060).
- [87] CMS Collaboration. “Background studies in the CMS muon system”. In: (2022). URL: <https://cds.cern.ch/record/2815407>.
- [88] CMS Collaboration. “Background measurements in the CMS DT chambers during LHC Run 2”. In: (2020). URL: <https://cds.cern.ch/record/2710126>.
- [89] CMS Collaboration. “Background studies in the CMS DT chambers during LHC Run 2”. In: (2020). URL: <https://cds.cern.ch/record/2717926>.
- [90] T Baer et al. “UFOs in the LHC”. In: (2011). URL: <https://cds.cern.ch/record/1379150>.
- [91] CMS Collaboration. *Emittance scans for CMS luminosity calibration in Run 2*. Tech. rep. Geneva: CERN, 2020. doi: [10.22323/1.364.0193](https://cds.cern.ch/record/2693840). URL: <https://cds.cern.ch/record/2693840>.
- [92] M Hostettler et al. “Operational  $\beta^*$  levelling at the LHC in 2022 and beyond”. In: *JACoW IPAC 2023* (2023), MOPL045. doi: [10.18429/JACoW-IPAC2023-MOPL045](https://cds.cern.ch/record/2884758). URL: <https://cds.cern.ch/record/2884758>.

## REFERENCES

- [93] Jan Troska et al. “The VTRx+, an optical link module for data transmission at HL-LHC”. In: *PoS TWEPP-17* (2017), p. 048. DOI: [10.22323/1.313.0048](https://doi.org/10.22323/1.313.0048). URL: <https://cds.cern.ch/record/2312396>.
- [94] CMS Collaboration. *40MHz Trigger-less readout of the CMS Drift Tube muon detector*. Tech. rep. Geneva: CERN, 2023. URL: <https://cds.cern.ch/record/2879316>.
- [95] Jeroen Hegeman et al. “The CMS Timing and Control Distribution System”. In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 2015, pp. 1–3. DOI: [10.1109/NSSMIC.2015.7581984](https://doi.org/10.1109/NSSMIC.2015.7581984).
- [96] *EMP Framework*. available at <https://serenity.web.cern.ch/serenity/emp-fwk/>.
- [97] J. M. Mendez et al. “New LpGBT-FPGA IP: Simulation model and first implementation”. In: *Topical Workshop on Electronics for Particle Physics. 17-21 September 2018. Antwerp*. Sept. 2018, 59, p. 59. DOI: [10.22323/1.343.0059](https://doi.org/10.22323/1.343.0059).
- [98] David Sidler et al. “StRoM: smart remote memory”. In: *EuroSys ’20: Fifteenth EuroSys Conference 2020, Heraklion, Greece, April 27-30, 2020*. ACM, 2020, 29:1–29:16. DOI: [10.1145/3342195.3387519](https://doi.org/10.1145/3342195.3387519).
- [99] AMD. *UltraScale+ Devices Integrated 100G Ethernet Subsystem Product Guide (PG203)*. AMD, May 2023. URL: <https://docs.xilinx.com/r/en-US/pg203-cmac-usplus>.
- [100] J.D. Day and H. Zimmermann. “The OSI reference model”. In: *Proceedings of the IEEE* 71.12 (1983), pp. 1334–1340. DOI: [10.1109/PROC.1983.12775](https://doi.org/10.1109/PROC.1983.12775).
- [101] *Synopsys VCS Simulator*. <https://www.synopsys.com/verification/simulation/vcs.html>.
- [102] I. T. Association. “InfiniBand architecture, Specification Volumen 1, Release 1.0.a, available from the InfiniBand Trade Association”. In: <http://www.infinibandta.org> (2001). URL: <https://cir.nii.ac.jp/crid/1570291224874626176>.
- [103] Alex Forencich. *Verilog Ethernet Components*. URL: <https://github.com/alexforencich/verilog-ethernet> (visited on 08/15/2024).



- [104] “IEEE Standard for Information technology–Local and metropolitan area networks–Specific requirements– Part 3: CSMA/CD Access Method and Physical Layer Specifications - Amendment 8: MAC Control Frame for Priority-based Flow Control”. In: *IEEE Std 802.3bd-2011 (Amendment to IEEE Std 802.3-2008)* (2011), pp. 1–31. DOI: [10.1109/IEEESTD.2011.5982058](https://doi.org/10.1109/IEEESTD.2011.5982058).
- [105] CMS Collaboration. “Data Parking and Data Scouting at the CMS Experiment”. In: (2012). URL: <https://cds.cern.ch/record/1480607>.
- [106] CMS Collaboration. “Search for narrow resonances in dijet final states at  $\sqrt{s} = 8$  TeV with the novel CMS technique of data scouting”. In: *Phys. Rev. Lett.* 117.3 (2016), p. 031802. DOI: [10.1103/PhysRevLett.117.031802](https://doi.org/10.1103/PhysRevLett.117.031802). arXiv: [1604.08907](https://arxiv.org/abs/1604.08907) [hep-ex].
- [107] CMS Collaboration. “Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at  $\sqrt{s} = 13$  TeV”. In: *Phys. Rev. Lett.* 124.13 (2020), p. 131802. DOI: [10.1103/PhysRevLett.124.131802](https://doi.org/10.1103/PhysRevLett.124.131802). arXiv: [1912.04776](https://arxiv.org/abs/1912.04776) [hep-ex].
- [108] CMS Collaboration. “The CMS trigger system”. In: *JINST* 12.01 (2017), P01020. 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].
- [109] CMS Collaboration. “Recording and reconstructing 10 billion unbiased b hadron decays in CMS”. In: (2019). URL: <https://cds.cern.ch/record/2704495>.
- [110] CMS Collaboration. “Test of lepton flavor universality in  $B^\pm \rightarrow K^\pm \mu^+ \mu^-$  and  $B^\pm \rightarrow K^\pm e^+ e^-$  decays in proton-proton collisions at  $\sqrt{s} = 13$  TeV”. In: *Rept. Prog. Phys.* 87.7 (2024), p. 077802. DOI: [10.1088/1361-6633/ad4e65](https://doi.org/10.1088/1361-6633/ad4e65). arXiv: [2401.07090](https://arxiv.org/abs/2401.07090) [hep-ex].
- [111] CMS Collaboration. “Level-1 trigger scouting in Phase-2”. In: (2024). URL: <https://cds.cern.ch/record/2916191>.
- [112] CMS Collaboration. *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*. Tech. rep. This is the final version of the document, approved by the LHCC. Geneva: CERN, 2021. URL: <https://cds.cern.ch/record/2759072>.
- [113] CMS Collaboration. *CMS Software*. <https://github.com/cms-sw/cmssw> [Accessed: 2024].

## REFERENCES

- [114] CMS Collaboration. “CMS Level-1 trigger Data Scouting firmware prototyping for LHC Run-3 and CMS Phase-2”. In: *JINST* 19.05 (2024), p. C05027. DOI: [10.1088/1748-0221/19/05/C05027](https://doi.org/10.1088/1748-0221/19/05/C05027).
- [115] FastML Team. *fastmachinelearning/hls4ml*. Version v0.8.1. 2023. DOI: [10.5281/zenodo.1201549](https://doi.org/10.5281/zenodo.1201549). URL: <https://github.com/fastmachinelearning/hls4ml>.
- [116] CMS Collaboration. “40 MHz Scouting with Deep Learning in CMS”. In: (2022). URL: <https://cds.cern.ch/record/2843741>.
- [117] Chuck Pheatt. “Intel® threading building blocks”. In: *Journal of Computing Sciences in Colleges* 23.4 (2008), pp. 298–298.
- [118] Jean-marc Olivier Andre et al. *Online data handling and storage at the CMS experiment*. Tech. rep. 8. Geneva: CERN, 2015. DOI: [10.1088/1742-6596/664/8/082009](https://doi.org/10.1088/1742-6596/664/8/082009). URL: <https://cds.cern.ch/record/2016893>.
- [119] CMS Collaboration. “Dataset definition for CMS operations and physics analyses”. In: *Nucl. Part. Phys. Proc.* 273-275 (2016), pp. 929–933. DOI: [10.1016/j.nuclphysbps.2015.09.144](https://doi.org/10.1016/j.nuclphysbps.2015.09.144).
- [120] M Giffels et al. “The CMS Data Management System”. In: *Journal of Physics: Conference Series* 513.4 (June 2014), p. 042052. DOI: [10.1088/1742-6596/513/4/042052](https://doi.org/10.1088/1742-6596/513/4/042052). URL: <https://dx.doi.org/10.1088/1742-6596/513/4/042052>.
- [121] M Mascheroni et al. “CMS distributed data analysis with CRAB3”. In: *Journal of Physics: Conference Series* 664.6 (Dec. 2015), p. 062038. DOI: [10.1088/1742-6596/664/6/062038](https://doi.org/10.1088/1742-6596/664/6/062038). URL: <https://dx.doi.org/10.1088/1742-6596/664/6/062038>.
- [122] Matteo Migliorini. *L1 Scouting example analyzer*. <https://github.com/Mmiglio/L1ScoutingAnalyzer.git> [Accessed: 2024]. 2024.
- [123] Marco Peruzzi et al. “The NanoAOD event data format in CMS”. In: *Journal of Physics: Conference Series* 1525.1 (Apr. 2020), p. 012038. DOI: [10.1088/1742-6596/1525/1/012038](https://doi.org/10.1088/1742-6596/1525/1/012038). URL: <https://dx.doi.org/10.1088/1742-6596/1525/1/012038>.
- [124] Jean-marc Olivier Andre et al. *A scalable monitoring for the CMS Filter Farm based on elasticsearch*. A scalable monitoring for the CMS Filter Farm based on elastic search. Tech. rep. Geneva: CERN, 2015. DOI: [10.1088/1742-6596/664/8/082036](https://doi.org/10.1088/1742-6596/664/8/082036). URL: <https://cds.cern.ch/record/2020877>.



- [125] *Circles*. <https://github.com/fwyzard/circles> [Accessed: 2024]. 2019.
- [126] CMS Collaboration. “40MHz Scouting Muon Studies”. In: (2023). URL: <https://cds.cern.ch/record/2859463>.
- [127] CMS Collaboration. “The Level-1 CMS electron and photon trigger for the LHC Run II”. In: *PoS LHCP2018* (2018), p. 073. DOI: [10.22323/1.321.0073](https://cds.cern.ch/record/2668901). URL: <https://cds.cern.ch/record/2668901>.
- [128] CMS Collaboration. “Calibration and Performance of the CMS Electromagnetic Calorimeter in LHC Run2”. In: *European Physical Journal Web of Conferences*. Vol. 245. European Physical Journal Web of Conferences. Nov. 2020, 02027, p. 02027. DOI: [10.1051/epjconf/202024502027](https://cds.cern.ch/record/202024502027).
- [129] CMS Collaboration. “Performance of the CMS Level-1 trigger in proton-proton collisions at  $\sqrt{s} = 13$  TeV”. In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10017. DOI: [10.1088/1748-0221/15/10/P10017](https://cds.cern.ch/record/200610165). arXiv: [2006.10165](https://arxiv.org/abs/2006.10165) [hep-ex].
- [130] CMS Collaboration. “Optimizing the Performance of the CMS ECAL Trigger for Runs 2 and 3 of the CERN LHC”. In: *arXiv e-prints*, arXiv:1910.06232 (Oct. 2019), arXiv:1910.06232. DOI: [10.48550/arXiv.1910.06232](https://cds.cern.ch/record/191006232). arXiv: [1910.06232](https://arxiv.org/abs/1910.06232) [physics.ins-det].
- [131] CMS Collaboration. *The CMS Level-1 Tau identification algorithm for the LHC Run II. The CMS Level-1 Tau algorithm for the LHC Run II*. Tech. rep. Geneva: CERN, 2016. DOI: [10.1016/j.nuclphysbps.2015.09.444](https://cds.cern.ch/record/1971159). URL: <https://cds.cern.ch/record/1971159>.
- [132] Nicolo’ Lai. “Barrel muon track reconstruction with deep learning for Level-1 trigger data scouting in the CMS experiment”. In: (2023). URL: <https://cds.cern.ch/record/2883887>.
- [133] Luis A. Anchordoqui et al. “Dijet signals for low mass strings at the LHC”. In: *Phys. Rev. Lett.* 101 (2008), p. 241803. DOI: [10.1103/PhysRevLett.101.241803](https://cds.cern.ch/record/241803). arXiv: [0808.0497](https://arxiv.org/abs/0808.0497) [hep-ph].
- [134] R. Sekhar Chivukula et al. “Hadron Collider Production of Massive Color-Octet Vector Bosons at Next-to-Leading Order”. In: *Phys. Rev. D* 87.9 (2013), p. 094011. DOI: [10.1103/PhysRevD.87.094011](https://cds.cern.ch/record/13031120). arXiv: [1303.1120](https://arxiv.org/abs/1303.1120) [hep-ph].

## REFERENCES

- [135] Tao Han, Ian Lewis, and Zhen Liu. “Colored Resonant Signals at the LHC: Largest Rate and Simplest Topology”. In: *JHEP* 12 (2010), p. 085. doi: [10.1007/JHEP12\(2010\)085](https://doi.org/10.1007/JHEP12(2010)085). arXiv: [1010.4309](https://arxiv.org/abs/1010.4309) [hep-ph].
- [136] Lisa Randall and Raman Sundrum. “An Alternative to compactification”. In: *Phys. Rev. Lett.* 83 (1999), pp. 4690–4693. doi: [10.1103/PhysRevLett.83.4690](https://doi.org/10.1103/PhysRevLett.83.4690). arXiv: [hep-th/9906064](https://arxiv.org/abs/hep-th/9906064).
- [137] E. Eichten et al. “Supercollider physics”. In: *Rev. Mod. Phys.* 56 (4 Oct. 1984), pp. 579–707. doi: [10.1103/RevModPhys.56.579](https://doi.org/10.1103/RevModPhys.56.579). URL: <https://link.aps.org/doi/10.1103/RevModPhys.56.579>.
- [138] UA1 Collaboration. “Two Jet Mass Distributions at the CERN Proton - Anti-Proton Collider”. In: *Phys. Lett. B* 209 (1988), pp. 127–134. doi: [10.1016/0370-2693\(88\)91843-6](https://doi.org/10.1016/0370-2693(88)91843-6).
- [139] UA2 Collaboration. “A Measurement of two jet decays of the W and Z bosons at the CERN  $\bar{p}p$  collider”. In: *Z. Phys. C* 49 (1991), pp. 17–28. doi: [10.1007/BF01570793](https://doi.org/10.1007/BF01570793).
- [140] CDF Collaboration. “The Two jet invariant mass distribution at  $\sqrt{s} = 1.8$  TeV”. In: *Phys. Rev. D* 41 (1990), pp. 1722–1725. doi: [10.1103/PhysRevD.41.1722](https://doi.org/10.1103/PhysRevD.41.1722).
- [141] D0 Collaboration. “Search for new particles in the two jet decay channel with the D0 detector”. In: *Phys. Rev. D* 69 (2004), p. 111101. doi: [10.1103/PhysRevD.69.111101](https://doi.org/10.1103/PhysRevD.69.111101). arXiv: [hep-ex/0308033](https://arxiv.org/abs/hep-ex/0308033).
- [142] CMS Collaboration. “Search for Dijet Resonances in 7 TeV pp Collisions at CMS”. In: *Phys. Rev. Lett.* 105 (2010), p. 211801. doi: [10.1103/PhysRevLett.105.211801](https://doi.org/10.1103/PhysRevLett.105.211801). arXiv: [1010.0203](https://arxiv.org/abs/1010.0203) [hep-ex].
- [143] ATLAS Collaboration. “Search for New Particles in Two-Jet Final States in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC”. In: *Phys. Rev. Lett.* 105 (2010), p. 161801. doi: [10.1103/PhysRevLett.105.161801](https://doi.org/10.1103/PhysRevLett.105.161801). arXiv: [1008.2461](https://arxiv.org/abs/1008.2461) [hep-ex].
- [144] CMS Collaboration. “Search for narrow and broad dijet resonances in proton-proton collisions at  $\sqrt{s} = 13$  TeV and constraints on dark matter mediators and other new particles”. In: *arXiv e-prints*, arXiv:1806.00843 (June 2018), arXiv:1806.00843. doi: [10.48550/arXiv.1806.00843](https://doi.org/10.48550/arXiv.1806.00843). arXiv: [1806.00843](https://arxiv.org/abs/1806.00843) [hep-ex].

- [145] CMS Collaboration. “Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at  $\sqrt{s} = 13$  TeV”. In: *arXiv e-prints*, arXiv:1911.03947 (Nov. 2019), arXiv:1911.03947. doi: [10.48550/arXiv.1911.03947](https://doi.org/10.48550/arXiv.1911.03947). arXiv: [1911.03947](https://arxiv.org/abs/1911.03947) [hep-ex].
- [146] ATLAS Collaboration. “Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in  $pp$  collisions at  $\sqrt{s} = 13$  TeV”. In: *Phys. Rev. Lett.* 121.8 (2018), p. 081801. doi: [10.1103/PhysRevLett.121.081801](https://doi.org/10.1103/PhysRevLett.121.081801). arXiv: [1804.03496](https://arxiv.org/abs/1804.03496) [hep-ex].
- [147] ATLAS Collaboration. “Search for new phenomena in dijet events using 37 fb<sup>-1</sup> of  $pp$  collision data collected at  $\sqrt{s} = 13$  TeV with the ATLAS detector”. In: *Phys. Rev. D* 96.5 (2017), p. 052004. doi: [10.1103/PhysRevD.96.052004](https://doi.org/10.1103/PhysRevD.96.052004). arXiv: [1703.09127](https://arxiv.org/abs/1703.09127) [hep-ex].
- [148] CMS Collaboration. “Search for resonances and quantum black holes using dijet mass spectra in proton-proton collisions at  $\sqrt{s} = 8$  TeV”. In: *Phys. Rev. D* 91.5 (2015), p. 052009. doi: [10.1103/PhysRevD.91.052009](https://doi.org/10.1103/PhysRevD.91.052009). arXiv: [1501.04198](https://arxiv.org/abs/1501.04198) [hep-ex].
- [149] CMS Collaboration. “Search for dijet resonances using events with three jets in proton-proton collisions at  $\sqrt{s} = 13$  TeV”. In: *arXiv e-prints*, arXiv:1911.03761 (Nov. 2019), arXiv:1911.03761. doi: [10.48550/arXiv.1911.03761](https://doi.org/10.48550/arXiv.1911.03761). arXiv: [1911.03761](https://arxiv.org/abs/1911.03761) [hep-ex].
- [150] CMS Collaboration. *Search for low mass vector and scalar resonances decaying into quark-antiquark pairs*. Tech. rep. Geneva: CERN, 2024. URL: <https://cds.cern.ch/record/2904945>.
- [151] ATLAS Collaboration. “Jet substructure measurements with CMS and ATLAS”. In: (2021). URL: <https://cds.cern.ch/record/2772255>.
- [152] CMS Collaboration. “Search for Low-Mass Quark-Antiquark Resonances Produced in Association with a Photon at  $\sqrt{s} = 13$  TeV”. In: *Phys. Rev. Lett.* 123.23 (2019), p. 231803. doi: [10.1103/PhysRevLett.123.231803](https://doi.org/10.1103/PhysRevLett.123.231803). arXiv: [1905.10331](https://arxiv.org/abs/1905.10331) [hep-ex].
- [153] Daniel Abercrombie et al. “Dark Matter benchmark models for early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum”. In: *Phys. Dark Univ.* 27 (2020). Ed. by Antonio Boveia et al., p. 100371. doi: [10.1016/j.dark.2019.100371](https://doi.org/10.1016/j.dark.2019.100371). arXiv: [1507.00966](https://arxiv.org/abs/1507.00966) [hep-ex].

## REFERENCES

- [154] CMS Exotica Summary plots for 13 TeV data. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEX013TeV> [Accessed: 2024].
- [155] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- $k_t$  jet clustering algorithm”. In: *JHEP* 04 (2008), p. 063. doi: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063). arXiv: [0802.1189](https://arxiv.org/abs/0802.1189) [hep-ph].
- [156] Torbjörn Sjöstrand. “The PYTHIA Event Generator: Past, Present and Future”. In: *Comput. Phys. Commun.* 246 (2020), p. 106910. doi: [10.1016/j.cpc.2019.106910](https://doi.org/10.1016/j.cpc.2019.106910). arXiv: [1907.09874](https://arxiv.org/abs/1907.09874) [hep-ph].
- [157] CMS Collaboration. “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”. In: *Eur. Phys. J. C* 80.1 (2020), p. 4. doi: [10.1140/epjc/s10052-019-7499-4](https://doi.org/10.1140/epjc/s10052-019-7499-4). arXiv: [1903.12179](https://arxiv.org/abs/1903.12179) [hep-ex].
- [158] R. A. Fisher. “On the Interpretation of 2 from Contingency Tables, and the Calculation of P”. In: *Journal of the Royal Statistical Society* 85.1 (1922), pp. 87–94. ISSN: 09528385. URL: <http://www.jstor.org/stable/2340521> (visited on 09/07/2024).
- [159] CMS Collaboration. “Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm”. In: (2023). URL: <https://cds.cern.ch/record/2879816>.