



NATIONAL TECHNICAL
UNIVERSITY OF ATHENS

SCHOOL OF APPLIED
MATHEMATICAL AND
PHYSICAL SCIENCES

SCHOOL OF MECHANICAL
ENGINEERING

NATIONAL CENTRE FOR
SCIENTIFIC RESEARCH
"DEMOKRITOS"

INSTITUTE OF NANOSCIENCE
AND NANOTECHNOLOGY

INSTITUTE OF NUCLEAR AND
PARTICLE PHYSICS



Postgraduate Thesis of the M.Sc. program
"Physics and Technological Applications"

Optimization of the NSW Pad Trigger efficiency
with targeted masking in the ATLAS
experiment at CERN

Βελτιστοποίηση της απόδοσης του σκανδαλισμού
των πλακιδίων/pad του NSW με στοχευμένη
παραλλαγή στο πείραμα ATLAS στο CERN

MASTER THESIS
of Eleni Kanellaki

Supervisor: Theodoros Gerasis

Athens, October, 2024

I would like to express my deepest gratitude to my wonderful supervisor, Theodoros Gerasis, who believed in me, trusted me with this project, and supported my journey to CERN. He was always by my side, offering guidance and encouragement, and our collaboration was nothing short of perfect. He is like family to me, and I am truly thankful for his mentorship and unwavering support.

I would also like to thank the many people at CERN who have supported me, each in their own unique way: Stefano, Ludo, Paolo, Estel, Foteini, Stergios, Nikos, Yiannis, Simone, Callum, Leesa, Jared, Michelle.

Finally, I would like to express my heartfelt gratitude to my family. Even though they miss me and it is not at all easy for them to be apart, they always find a way to support me no matter how difficult it may be. I would also like to thank my friends in different cities across Greece who continue to love and tolerate me despite the distance that separates us.

Abstract

This master thesis was prepared within the framework of the interdepartmental postgraduate program: "Physics and Technological Applications" of the National Technical University of Athens, in collaboration with the National Centre for Scientific Research "Demokritos". It was carried out entirely at CERN during the 2023-2024 academic year.

The thesis presents the development and implementation of a tool responsible for the targeted masking in the New Small Wheel pad trigger system of the ATLAS experiment. The aim of this tool was to mitigate detector defects and restore the efficiency of all the sectors, ensuring their inclusion in the coincidence of the ATLAS central trigger. The tool initially focused on a stricter sTGC pad coincidence logic, but then it was expanded to include a looser coincidence as well. Throughout the project, various improvements and tests were performed to refine the tool, resulting in a robust masking strategy capable of dynamically handling complex sector conditions and improving overall system reliability.

Περίληψη

Η παρούσα μεταπτυχιακή εργασία εκπονήθηκε στο πλαίσιο του διατμηματικού μεταπτυχιακού προγράμματος: «Φυσική και Τεχνολογικές Εφαρμογές» του Εθνικού Μετσόβιου Πολυτεχνείου, σε συνεργασία με το Εθνικό Κέντρο Έρευνας Φυσικών Επιστημών «Δημόκριτος». Πραγματοποιήθηκε εξ ολοκλήρου στο CERN κατά το ακαδημαϊκό έτος 2023-2024.

Η διπλωματική εργασία παρουσιάζει την ανάπτυξη και εφαρμογή ενός εργαλείου υπεύθυνου για τη στοχευμένη παραλλαγή στο σύστημα σκανδαλισμού πλακιδίων/ rad του Νέου Μικρού Τροχού του πειράματος ATLAS. Ο στόχος αυτού του εργαλείου ήταν να μετριάσει τα ελαττώματα των ανιχνευτών και να αποκαταστήσει την απόδοση όλων των «τομέων», διασφαλίζοντας τη συμπερίληψή τους στη σύμπτωση του κεντρικού συστήματος σκανδαλισμού του ATLAS. Το εργαλείο αρχικά επικεντρώθηκε σε μια αυστηρότερη sTGC rad λογική σύμπτωσης, αλλά στη συνέχεια επεκτάθηκε ώστε να συμπεριλάβει επίσης και μια πιο χαλαρή σύμπτωση. Καθ' όλη τη διάρκεια της εργασίας, πραγματοποιήθηκαν διάφορες βελτιώσεις και δοκιμές για τη βελτίωση του εργαλείου, με αποτέλεσμα μια ισχυρή στρατηγική παραλλαγής ικανή να χειρίζεται δυναμικά πολύπλοκες συνθήκες «τομέων» και να βελτιώνει τη συνολική αξιοπιστία του συστήματος.

Contents

1	Introduction	1
1.1	Theoretical framework	1
1.2	Interaction of particles with matter	2
1.3	Interaction of photons with matter	4
1.4	Interaction of neutrons with matter	4
1.5	The quark and the parton model	5
1.6	The Higgs boson and the Standard Model	7
2	The Large Hadron Collider and the ATLAS experiment	8
2.1	The Large Hadron Collider	8
2.1.1	The particle beams of the LHC and its experiments	8
2.1.2	The LHC nominal cycle	11
2.1.3	The High Luminosity upgrade plan	13
2.2	The ATLAS experiment	15
2.2.1	Detector subsystems	17
2.2.2	Trigger and Data-Acquisition system	27
3	The New Small Wheel	30
3.1	Motivation	30
3.1.1	Precision tracking performance	30
3.1.2	Trigger selection	31
3.2	The NSW layout	32
3.3	The detector technologies of the NSW	33
3.3.1	Small-strip Thin Gap Chambers	33
3.3.2	Micromegas	35
3.4	NSW electronics and Data Acquisition dataflow	36
4	Targeted masking in the ATLAS experiment	38
4.1	Motivation	38
4.2	Analysis in the first steps	38
4.3	Targeted masking with 3/4 & 3/4 coincidence	42
4.4	Implementation of targeted masking with 3/4 & 3/4 coincidence	49
4.5	Initial manual implementation of targeted masking with 2/4 & 3/4 coincidence	56
4.6	Full development and automation of targeted masking for 2/4 & 3/4 coincidence	57
4.7	Improvements	62
5	Summary and conclusions	69

Chapter 1

Introduction

1.1 Theoretical framework

Particle physics, also known as high-energy physics, explores the nature and behavior of subatomic particles—the fundamental constituents of matter. Through the investigation of these smallest building blocks of the universe and the forces governing their interactions, particle physicists endeavor to unveil the mysteries of the cosmos. At the core of this pursuit lies the Standard Model (SM) [2], a mathematical framework that elegantly describes elementary particles and their interactions. Renowned for its remarkable accuracy, the SM has been rigorously tested and stands as a cornerstone of modern physics. Drawing upon principles from quantum field theory and gauge symmetries, it provides a comprehensive understanding of many features of nature observed in experiments. However, despite its successes, the SM does not offer a complete picture; there are still important questions it does not answer.

An illustration of the SM is shown in the figure 1.1. The particles are classified into two main categories: fermions and bosons. Fermions, which have the property of half-integer intrinsic spin, serve as the building blocks for matter in the universe and come in two types: quarks and leptons. Each type consists of six particles organized into three pairs or generations. The quark pairs are up and down, charm and strange, and top and bottom (or beauty). Quarks also exhibit three colours, which are the equivalent of charge for strong interactions. Due to a property known as colour confinement, quarks can only combine to form colourless objects. As a result, "colour" is never directly observed in nature, and all strongly interacting particles appear colour-neutral to the outside world. The lepton pairs consist of the electron and electron neutrino, muon and muon neutrino, and tau and tau neutrino. All these particles interact and are governed by the four fundamental interactions: the strong interaction, electromagnetic interaction, weak interaction, and gravity. These interactions are mediated by five particles known as bosons, which have the property of integer intrinsic spin. The Higgs boson, on the other hand, is thought to provide mass to other particles through the Higgs field, a mechanism known as the Higgs mechanism.




Figure 1.1: Overview of the Standard Model particles: the two types of elementary particles, quarks (top left) and leptons (bottom left), the mediators of the fundamental interactions (right) and at the center the Higgs boson [1].

In modern particle physics, fundamental interactions are described by Quantum Field Theory (QFT), corresponding to the exchange of the interaction-carrying particles. QFT is based on mathematical frameworks involving fields and operators that create and annihilate particles, providing a formal description of interactions at the quantum level. Quantum Electrodynamics (QED) describes phenomena involving electrically charged particles and the exchange of photons between them. It extends the classical theory described by Maxwell's equations to the relativistic limit. Similarly, Quantum Chromodynamics (QCD) governs the strong interaction, binding quarks together inside hadrons such as protons and neutrons, through the exchange of gluons. The weak interaction is described by the Electroweak Theory, which unifies the electromagnetic and weak interactions into a single theoretical framework at high energies, involving the exchange of the W and Z bosons. While QFT has been incredibly successful in explaining experimental observations, it also presents theoretical challenges, such as divergences in certain calculations. However, through techniques like renormalization, physicists have been able to overcome these challenges and achieve impressive agreement between theory and experiment.

1.2 Interaction of particles with matter

The strengths of a particle's interactions and its corresponding lifetime are determined by its couplings to the fundamental interactions, each characterized by distinct coupling constants. When the strength of the strong interaction is taken as 1, the electromagnetic interaction is about 10^{-2} times weaker, and the weak interaction is even weaker, approximately 10^{-7} times than the strong interaction. As a result, the strong interaction has the highest coupling constant, making it the most powerful interaction. Consequently, if a particle can undergo scattering via the strong interaction, this process will almost always dominate over any possible electromagnetic or weak scattering. Similarly, electromagnetic

scattering modes typically dominate over those involving the weak interaction. This hierarchy also applies to particle decays: particles decay predominantly through the strongest available interaction, provided that there is a final state with a lower total rest mass that can be reached. For any decay process to occur, it must be represented by a Feynman diagram using the interaction vertices defined by the SM. Particle decay is a fundamental process in particle physics, with most particles exhibiting very short lifetimes. The decay rates of particles are influenced by various factors, with the type of fundamental interaction involved being the most significant [3].

In describing the collision or interaction of two particles, the concept of the cross section is crucial. This quantity essentially measures the probability of a reaction occurring. Another important concept is the mean free path, which refers to the average distance a particle travels before undergoing a collision. The passage of charged particles through matter is in general characterized by two principal features, energy loss and deflection from the incident direction. These effects primarily result from two processes, inelastic collisions with atomic electrons and elastic scattering from nuclei. However, other processes, such as the emission of Cherenkov radiation, nuclear reactions, and bremsstrahlung, can also occur when particles reach hyper-relativistic speeds. Cherenkov radiation is emitted when charged particles exceed the speed of light in a medium, bremsstrahlung occurs as high-energy electrons are decelerated by electric fields, and nuclear reactions can take place during collisions between high-energy particles and atomic nuclei. The likelihood of these processes occurring depends significantly on the energy of the interacting particles. It is important to categorize charged particles into two classes: electrons and positrons, and heavier particles, which include muons, pions, protons, alpha particles, and other nuclei. This distinction is significant because the interactions and resulting behaviors of these two classes in matter can differ substantially.

For heavy particles, inelastic collisions are the primary mechanism responsible for their energy loss in matter. During these collisions, energy is transferred from the particle to the atom. Depending on the outcome, these collisions are categorized as either soft, where only excitation occurs, or hard, where sufficient energy transfer leads to ionization. In some hard collisions, enough energy is transferred that the ejected electron itself can cause substantial secondary ionization. These high energy recoil electrons are sometimes referred to as δ -rays or knock-on electrons. Elastic scattering from nuclei also occurs, though less frequently than collisions with electrons. Typically, these interactions result in very little energy transfer because the nuclei of most materials have significantly larger masses compared to the incident particles.

Thus, heavy charged particles undergo atomic collisions that are statistical in nature, occurring with a certain quantum mechanical probability. The average energy loss per unit path length, known as stopping power or simply dE/dx , describes this process. The stopping power is quantitatively described by the Bethe-Bloch formula, which provides a detailed calculation of the energy loss experienced by these particles as they pass through matter. An important exception to the applicability of this formula arises in the case of channeling in materials with a spatially symmetric atomic structure, such as crystals. In such materials, particles can experience a series of correlated small-angle scatterings, causing them to follow a slowly oscillating trajectory within an open channel for relatively long distances. Another phenomenon associated with heavy charged particles is Cherenkov radiation. This occurs when a charged particle travels through a medium at a speed greater than the speed of light in that medium, resulting in the emission of Cherenkov radiation.

Electrons and positrons experience collisional energy loss when passing through matter, similar to heavy charged particles. However, due to their small mass, they also incur additional energy loss through a process known as bremsstrahlung. This occurs because electromagnetic radiation is emitted when these particles are scattered in the electric field of the nucleus. Classically, this can be understood as radiation resulting from the acceleration of the electron or positron as it deviates from a straight-line path due to the electrical attraction of the nucleus. The collision energy loss for electrons and positrons is described by a modified Bethe-Bloch formula. This modification accounts for their small mass, which invalidates the assumption that the incident particle remains largely undeflected during the collision process. Additionally, since the collisions occur between identical particles, their indistinguishability must be considered in the calculations. Beyond inelastic collisions with atomic electrons, charged particles also experience repeated elastic Coulomb scatterings from nuclei, which are described by the Rutherford formula. Most of these interactions result in a net deflection from the original particle direction. Electrons and positrons are particularly susceptible to large-angle deflections when scattering from nuclei. In some cases, multiply scattered electrons may even reverse direction entirely, leading to backscattering out of the absorber.

1.3 Interaction of photons with matter

The interactions of x-rays and γ -rays in matter differ significantly from those of charged particles due to the absence of electric charge, which prevents frequent inelastic collisions with atomic electrons. Consequently, photons and x-rays are more penetrating in matter, and a photon beam does not degrade in energy as it passes through a material, but rather, it is attenuated in intensity. The primary interactions of photons with matter include the photoelectric effect, where a photon is absorbed by an atomic electron, resulting in the ejection of the electron from the atom; Compton scattering, where photons scatter on free electrons; and pair production, where a photon is transformed into an electron-positron pair. The dominance of each interaction depends strongly on the photon energy. At lower energies, the photoelectric effect is more prevalent, while Compton scattering dominates at intermediate energies, and pair production becomes significant at high photon energies. One of the most remarkable phenomena resulting from the combination of pair production and bremsstrahlung emission is the formation of electron-photon showers. In this process, a high-energy photon interacts with matter to produce an electron-positron pair. These particles then emit energetic bremsstrahlung photons, which can further convert into additional electron-positron pairs. This cascade or shower of photons, electrons, and positrons continues, propagating the interaction. The cascade persists until the energy of the produced electrons and positrons falls below a critical value, at which point they predominantly lose their energy through atomic collisions rather than bremsstrahlung emission, thus halting the cascade.

1.4 Interaction of neutrons with matter

Neutrons, lacking electric charge, too, do not experience Coulomb interactions with electrons and nuclei in matter. Their principal means of interaction is through the strong force with nuclei. Depending on their energy, neutrons can undergo various nuclear processes. These include elastic scattering from nuclei, which is the principal mechanism of

energy loss for neutrons in the MeV region; inelastic scattering, where the nucleus is left in an excited state and may later emit γ -rays or other forms of radiation; radiative neutron capture; other nuclear reactions where the neutron is captured and charged particles are emitted; fission; and high-energy hadron shower production. Neutrons are classified based on their energy into high-energy, fast, epithermal, thermal (or slow), and cold (or ultra-cold) neutrons. The process of slowing down fast neutrons is known as moderation. A fast neutron entering a material will scatter both elastically and inelastically off nuclei, losing energy until it reaches thermal equilibrium with the surrounding atoms. Once at thermal energy, the neutron diffuses through the material until it is either captured by a nucleus or undergoes another type of nuclear reaction. Neutrons may also participate in nuclear reactions or be captured before reaching thermal energies, especially if resonance conditions are present [4].

1.5 The quark and the parton model

In 1964, Murray Gell-Mann [5] and George Zweig [6], [7] independently proposed the quark model, a simplified model at low energies that describes mesons as bound states of a quark-antiquark pair and baryons as bound states of three quarks [8]. Later, in 1969, Richard Feynman introduced the parton model, which has been pivotal in the analysis of high-energy physics experiments. Feynman argued that the constituents of the nucleons, when incoherently scattered by an incident electron, behave as point-like particles, which he named partons [9]. Today, it is known that a nucleon consists of three "valence" quarks which carry the nucleon's quantum numbers and a "sea" of quark-antiquark pairs transiently created by gluon energy. Consequently, partons are now recognized as the quarks, antiquarks and gluons inside hadrons.




Figure 1.2: The internal structure of the proton from the perspectives of (a) the traditional quark model and (b) quantum chromodynamics. At low energies, the scattering particle primarily interacts with the valence partons. At higher energies, it also detects the sea partons [10].

Scattering processes at high-energy hadron colliders can be classified as either hard or soft interactions. These terms reflect how the strong interaction operates at different energy and distance scales. Hard interactions are high-energy processes occurring over very short distances, while soft interactions are low-energy processes occurring over longer distances. Although QCD is the underlying theory for all such interactions, the theoretical approach differs significantly between the two cases. In hard processes, such as W

boson or high transverse energy (E_T) jet production, the rates and event properties can be predicted with reasonable precision using perturbative QCD, leveraging the concept of asymptotic freedom. Asymptotic freedom describes the phenomenon where the force between quarks becomes weaker as they get closer to each other, allowing them to behave almost as free particles. In contrast, soft processes, such as total cross-section measurements or diffractive processes, are dominated by non-perturbative QCD effects, which are less well understood. In these interactions, the strong coupling between quarks remains significant, and perturbative techniques become less reliable, making precise predictions challenging [11].

Figure 1.3 illustrates a hadron-hadron collision event, such as a proton-proton (pp) collision. The central red blob represents the hard process of interest, where the protons interact at high momentum scales, producing outgoing particles like partons, leptons and/or gauge bosons. These particles are perturbatively computable. The partons, along with all particles carrying color charge, emit virtual gluons, which can further emit gluons or produce quark-antiquark pairs. This leads to the formation of particle showers, depicted as light red blobs. The transition to colorless hadrons, through the hadronization process (illustrated as light green blobs), involves the formation of hadrons from the colored partons. These hadrons then decay into stable particles that can be detected in experimental setups, with hadron decays shown in dark green. For momenta around 1 GeV, partons are confined, and non-perturbative interactions occur to form final-state hadrons through soft processes. The non-perturbative nature of these soft processes means they are typically modeled using numerical methods. In some collisions, multiple pairs of partons may interact simultaneously, a phenomenon known as multiple parton interactions (MPI). These are represented as purple blobs in the figure. Other beam remnants are shown in cyan. Additionally, the radiation of soft photons is illustrated with the yellow lines [13].




Figure 1.3: The evolution of a hadron-hadron collision as simulated by a Monte-Carlo event generator [12].

1.6 The Higgs boson and the Standard Model

The SM of particle physics is a highly successful theoretical framework that describes the fundamental particles and their interactions, with the exception of gravity. It incorporates electromagnetic, weak, and strong interactions, and has been remarkably successful in predicting experimental results. However, one crucial component had remained elusive for decades—the Higgs boson, the particle responsible for giving mass to other elementary particles through the Higgs mechanism. The discovery of the Higgs boson in 2012 by the ATLAS and CMS collaborations at the Large Hadron Collider (LHC) filled this missing piece, marking the completion of the SM. Despite this milestone, the SM still leaves many questions unanswered, such as the nature of dark matter, the matter-antimatter asymmetry in the universe, and the origin of neutrino masses. The LHC continues to study the properties of the Higgs boson in greater detail to search for deviations from the SM predictions, which could hint at new physics beyond the current framework and help address these unresolved issues. Two major research goals include precisely measuring the Higgs boson’s self-interactions, which are crucial for understanding the Higgs field and the early universe, and searching for signs of an extended Higgs sector, which could indicate physics beyond the SM, such as the multiple Higgs bosons predicted by supersymmetry (SUSY) and other beyond the Standard Model (BSM) theories.

Chapter 2

The Large Hadron Collider and the ATLAS experiment

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the world's largest and most powerful particle accelerator, operated by the European Organisation for Nuclear Research (CERN). It is the latest addition to CERN's acceleration complex and is installed in a circular tunnel that was originally built for the Large Electron-Positron Collider (LEP). Located about 100 meters below the surface, the tunnel has a circumference of 27 kilometers, spanning the boarder between Switzerland and France, near Geneva. As a particle-particle collider, the LHC machine is designed with two separate rings that allow for counter-rotating proton beams. These beams are accelerated to energies of up to 6.8 TeV per proton, resulting in collisions at a nominal center-of-mass energy of 13.6 TeV. The primary goal of these high-energy collisions is to explore fundamental physics, including the properties of the Higgs boson, and to search for phenomena beyond the SM, such as supersymmetry, dark matter candidates, and extra dimensions [14].

2.1.1 The particle beams of the LHC and its experiments

Before being injected into the main accelerator, the particles are prepared by a sequence of accelerators that progressively increase their energy. The process begins with the Linear particle accelerator 4 (Linac4), which accelerates negative hydrogen ions - each consisting of a proton with two electrons - to an energy of 160 MeV. These ions are then injected into the Proton Synchrotron Booster (PSB), where they are stripped of their electrons, leaving only protons. The PSB accelerates these protons further to 2 GeV before transferring them to the Proton Synchrotron (PS). The PS takes the protons and accelerates them to 26 GeV, preparing them for the next stage in the accelerator chain. Finally, the Super Proton Synchrotron (SPS) receives the particles from the PS and further accelerates them to 450 GeV. At this energy, the protons are ready to be injected into the LHC, where they will be accelerated to their maximum energy for high-energy collisions. This final acceleration is achieved through Radio Frequency (RF) cavities and the use of superconducting magnets, which guide and focus the proton beams along their circular path. These magnets operate at extremely low temperatures, close to absolute zero, using liquid helium to maintain the superconducting state required for efficient beam guidance and minimal energy loss.

The LHC physics program primarily focuses on proton-proton collisions, but it also includes heavy-ion collisions during shorter running periods. While lighter ions may be considered as well, the baseline heavy-ion scheme utilizes lead ions. The acceleration process for these ions begins with Linear accelerator 3 (Linac3), followed by storage and cooling in the Low Energy Ion Ring (LEIR). After this, the ions are further accelerated by the PS and the SPS before being injected into LHC ring. In the LHC, the lead ions

reach an energy of 2.3 TeV per nucleon, equivalent to 522 TeV per ion, surpassing the energies achieved by the Relativistic Heavy Ion Collider (RHIC). The primary objective of the heavy-ion program is to study the quark–gluon plasma, a state of matter thought to have existed in the early universe shortly after the Big Bang. By recreating and examining this extreme state of matter, physicists aim to gain insights into the strong interaction and the behavior of quarks and gluons under extreme conditions.

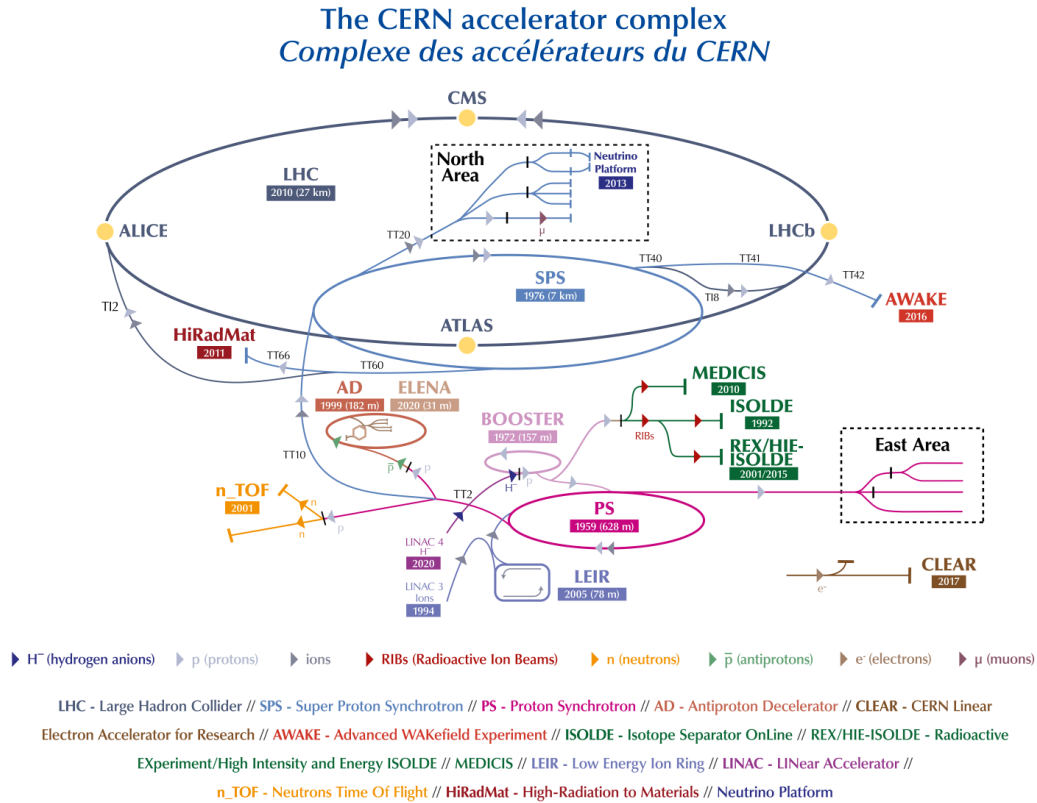


Figure 2.1: Diagram of the CERN accelerator complex, showing the various accelerators and the pathways taken by particle beams as they are prepared for injection into the LHC. The diagram also indicates the locations of major experiments, such as ATLAS, CMS, ALICE, and LHCb, which utilize the accelerated beams for high-energy physics research [15].

Beam collisions occur at four interaction points within the LHC, where the four main experiments are installed. These experiments are designed to detect and analyze the products of the high-energy collisions.

ATLAS

A Toroidal LHC Apparatus (ATLAS) [16] is a general-purpose detector at the LHC designed to investigate a wide range of particles and phenomena. ATLAS is essential for searching for the Higgs boson, supersymmetry, extra dimensions, and potential dark matter

candidates. Particle beams from the LHC collide at the center of the ATLAS detector, producing collision debris in the form of new particles, which fly out from the collision point in all directions. The detector features six different detecting subsystems arranged in layers around the collision point, which record the paths, momentum, and energy of these particles, enabling their individual identification. A huge magnet system within ATLAS bends the trajectories of charged particles, allowing for precise momentum measurements.

CMS

Compact Muon Solenoid (CMS) [17] is another general-purpose detector at the LHC, with scientific goals similar to those of ATLAS. However, CMS utilizes different technical solutions and a different magnet-system design. Unlike ATLAS, which employs a toroidal magnet system, the CMS detector is centered around a massive solenoid magnet. This variation in design affects the detector's configuration and the methods used to measure particle properties.

ALICE

The A Large Ion Collider Experiment (ALICE) [18] specializes in studying heavy-ion collisions to investigate the quark-gluon plasma, a state of matter that exists at extremely high energy densities and temperatures. When the LHC collides lead ions, it recreates conditions in the laboratory, similar to those that occurred just after the Big Bang. In these extreme conditions, protons and neutrons "melt", releasing quarks from their confinement with gluons and forming the quark-gluon plasma. Understanding the properties of this plasma is a key issue in theory of QCD, particularly regarding the phenomenon of confinement and chiral-symmetry restoration. The ALICE collaboration examines how the quark-gluon plasma evolves as it expands and cools, observing how it produces the particles that make up the matter in the universe today.

LHCb

The Large Hadron Collider beauty (LHCb) experiment [19] focuses on investigating the differences between matter and antimatter by studying the decays of particles containing b-quarks. This research is crucial for understanding the matter-antimatter asymmetry in the universe. Unlike ATLAS and CMS, which use integrated detectors that fully enclose the collision point, LHCb employs a specialized arrangement of subdetectors positioned along the direction of the particle beams. These subdetectors are placed sequentially over a distance about 20 meters, starting near the collision point. The LHCb detector is designed to detect forward particles - those emitted along the beamline after a collision. Due to the high energy of LHC collisions, a variety of quark types are produced and quickly decay into other forms. To efficiently capture and analyze b-quarks, the subdetectors of LHCb are advanced, movable tracking detectors positioned close to the beam paths.

Additionally, the LHC hosts several smaller experiments, such as TOTEM, LHCf, and MoEDAL, each with specific objectives, ranging from measuring forward particles to studying rare phenomena like magnetic monopoles. These experiments complement the larger detectors by exploring different aspects of high-energy physics.

2.1.2 The LHC nominal cycle

The LHC operates through a well-defined sequence known as the nominal cycle [20], which encompasses the stages of preparing the accelerator, delivering particle collisions, and maintaining optimal conditions for continuous operation. The time required to re-establish Stable Beams conditions after a beam dump is known as the LHC turnaround. The turnaround time is a critical factor in optimizing the collider's overall efficiency and maximizing the data collection periods for experiments. The different phases of LHC operation with the beam are:

Injection: This phase involves injecting protons or heavy ions from the injector chain into the LHC, as described above. Particles are injected in the form of bunches - tightly packed groups of particles. The LHC can accommodate up to 2808 proton bunches per beam, with each bunch containing around 1.15×10^{11} protons, spaced 25 nanoseconds apart. This configuration enables billions of proton-proton collisions per second. The injection phase is one of the most complex in the LHC cycle, involving numerous manual operations and influenced by various factors. The duration of this phase depends on the chosen filling scheme, the quality and availability of the beam from the injectors, the intensity of the circulating beam, the time required to setup the machine, and the measurement of beam parameters. Additionally, whether steering of the transfer lines is performed can also affect the injection time. Due to these variables, the time spent in this beam mode exhibits a broad distribution, with an average duration of approximately 72 minutes.

Prepare Ramp: This beam mode is declared once the injection process is completed and acts as a transition between the injection and energy ramp phases. During this stage, the beams are conditioned and stabilized in preparation for acceleration to high energies. Key tasks include ensuring that the beams are correctly aligned, optimizing machine parameters, and verifying that all systems are set for the significant increase in energy levels that occurs during the ramp phase. The average duration of this stage is approximately 10 minutes, although this can vary based on the initial conditions and any required adjustments.

Ramp: Once the bunches are injected into the LHC, the beams are accelerated to the target collision energy. The LHC's superconducting magnets gradually increase their magnetic field strength to guide and focus the beams as their energy ramps up. This process must be carefully managed to maintain beam stability and minimize losses. The ramp beam mode is declared just before the timing event is launched and terminates once arrived at flattop. Due to its well-defined nature, the duration of this phase has a very narrow time distribution, with an average of approximately 21 minutes.

Flattop: Once the energy ramp is completed, the flattop beam mode is declared to perform essential actions that prepare the beams for the next phase, the squeeze. During this stage, operations such as feedback reference changes, settings incorporation, and parameter loading are carried out. Additionally, the beam tune is adjusted to collision tunes to optimize beam dynamics for the upcoming squeeze phase. The duration of the flattop phase typically averages around 6 minutes.

Squeeze: After the beams reach their nominal energy, the squeeze phase reduces the transverse size of the beams at the collision points. This is achieved by using focusing

magnets, which tightly compress the beams to enhance the collision rate by increasing luminosity. A smaller beam size at the interaction points raises the probability of particle collisions, which is crucial for generating sufficient data for the experiments. The duration of the squeeze phase depends on the time required to execute the settings, with an average duration of approximately 16 minutes.

Adjust: During this phase, the beams are brought into collision at the designated interaction points within the LHC’s major detectors: ATLAS, CMS, ALICE, and LHCb. This phase is divided into two parts, addressing high and low luminosity regions separately to ensure optimal conditions for each detector. Once the collisions are established, luminosity is fine-tuned, and orbit feedback with reduced gain is activated to maintain beam stability. The subsequent beam mode, stable beams, is declared during the final manual adjustments. Due to the variability in the manual actions required, the duration of the adjust phase exhibits a fairly large distribution, averaging around 14 minutes.

Stable beams: The stable beams mode is the main data-taking phase of the LHC, during which experiments collect data on the resulting particle interactions from collisions. This phase is crucial as it represents the core of the LHC’s scientific output, allowing the detectors to gather extensive data for analysis. The duration of the stable beams phase can vary significantly, often lasting several hours, depending on the beam conditions and overall stability.

At the end of a collision period, or if beam conditions degrade, the beams are extracted from the LHC ring and directed into a dedicated beam dump system, which safely absorbs the energy of the particles. This process ensures that the high-energy beams are disposed of without damaging the accelerator components. Once the beams are dumped, the LHC’s magnets gradually ramp down, and the accelerator begins preparations for the next cycle. This phase includes cooling down components, recalibrating beam trajectories, and conducting comprehensive checks to ensure all systems are ready for the next injection sequence.

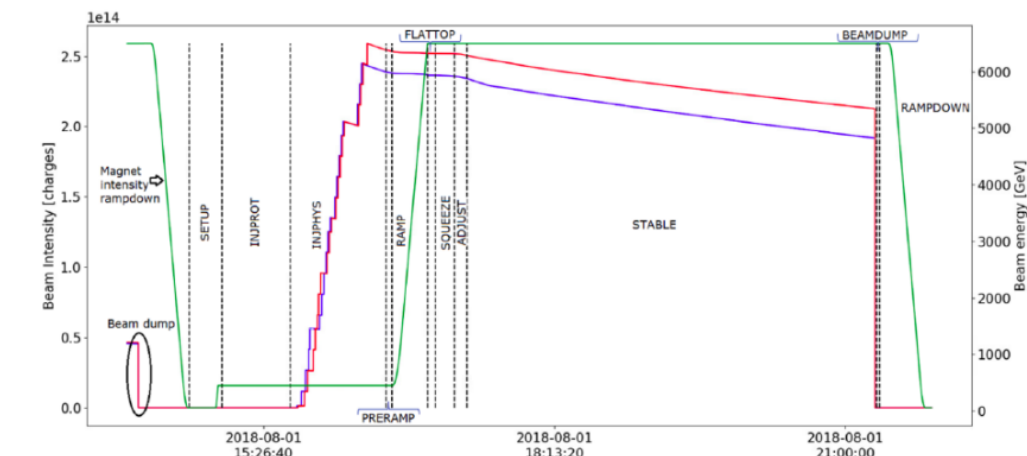


Figure 2.2: The LHC nominal cycle: the intensity of beam 1 and beam 2 is depicted in blue and red, respectively, while the energy of the beams is shown in green. The dashed black lines indicate the specific moments when the beam mode changes [21].

The LHC nominal cycle is a carefully orchestrated sequence designed to balance the need for high collision rates with the operational safety of the collider. Each phase of the cycle is meticulously monitored and adjusted by the LHC operations team to maintain beam stability and maximize the scientific output of the experiments. The organization of beams into bunches plays a critical role in achieving high luminosity, which is a key measure of the collision rate. Parameters such as bunch spacing, size, and the number of particles per bunch are optimized to maximize collision rates while adhering to the technical constraints of both the accelerator and detector systems.

Luminosity

The two most important features of a particle accelerator are its center-of-mass energy and its instantaneous luminosity \mathcal{L} [22]. The center-of-mass energy determines the types of particles that can be studied or discovered, while luminosity dictates the event rates. Luminosity measures the accelerator's capability to produce a high number of particle interactions and is defined as the number of collisions N per unit time for a given cross-section σ . High luminosity can be achieved by colliding beams with a large number of protons per bunch, within a small interaction area, and at a high collision frequency. The relationship for luminosity is expressed as:

$$\mathcal{L} = \frac{1}{\sigma} \cdot \frac{dN}{dt}. \quad (2.1)$$

The unit of luminosity is $cm^{-2} \cdot s^{-1}$. A related quantity is the integrated luminosity L_{int} , which is the integral of the luminosity over time:

$$L_{int} = \int \mathcal{L} dt. \quad (2.2)$$

Luminosity can be calculated from the knowledge of the parameters of the colliding beams, such as beam size and particle flow rate. The instantaneous luminosity of the machine can be expressed in terms of the number of particles per colliding bunch, n_1 and n_2 , the transverse beam sizes, σ_x and σ_y , the revolution frequency f at which the bunches collide in the LHC, and the number of bunches N_b . Assuming the beams have a Gaussian profile and collide head-on, the instantaneous luminosity is given by:

$$\mathcal{L}_{int} = \frac{f \cdot n_1 \cdot n_2 \cdot N_b}{4\pi \cdot \sigma_x \cdot \sigma_y}. \quad (2.3)$$

In practice, the exact properties of the colliding beams, such as their transverse profiles, are not precisely known, making it challenging to accurately calculate the instantaneous luminosity. For $f = 11.3kHz$, $n_1 = n_2 = 1.15 \cdot 10^{11}$, $N_b = 2808$ and $\sigma_x = \sigma_y = 16 \cdot 10^{-4}c$, the nominal LHC luminosity is estimated to be around $10^{34} cm^{-2} \cdot s^{-1}$.

2.1.3 The High Luminosity upgrade plan

The first proton beam in the LHC was successfully launched in September 2008. However, data-taking had to be suspended shortly after due to damage in the magnet system. The first collisions occurred in 2010, with the rest of the year to be dedicated on commissioning. Data-taking resumed in 2011 at a center-of-mass energy of 7 TeV, which was increased to 8 TeV in 2012, achieving 75% of the nominal luminosity of ($1 \times 10^{34} cm^{-2} \cdot s^{-1}$). This period is referred to as Run-1. In early 2013, the LHC underwent a two-year shutdown

(LS1) for maintenance and repair, including Phase-0 upgrades related to the experiments. During Run-2 (2015-2018), the LHC operated with proton beams at energies of 6.5 TeV and reached a peak luminosity of approximately twice the nominal value. The second long shutdown (LS2), which began in early 2019, was extended by about a year due to the COVID-19 pandemic and lasted until early 2022. This shutdown period was focused on preparations for Run-3 (2022-2025), which aims to deliver an integrated luminosity equal to the combined total of the previous runs. During Run-3, each proton beam will be accelerated to the energy of 6.8 TeV. In parallel to the accelerator activities, the Phase-1 upgrade related to the experiments was completed. The third long shutdown (LS3), scheduled from 2026 to 2028, will focus on preparing the LHC for its high-luminosity phase. This phase aims to achieve an instantaneous luminosity of 5 to 7.5 times the nominal value and reach the design center-of-mass energy of 14 TeV. The experiments will also undergo a Phase-2 upgrade during LS3. The detailed upgrade plan is illustrated in figure 2.3.



Figure 2.3: Timeline for the HL-LHC plan: the upper bold red lines indicate the energy of the collisions over the years, while the lower red lines represent the peak luminosity at different stages of the timeline [23].

2.2 The ATLAS experiment

ATLAS is the largest volume detector ever constructed for a particle collider, designed to explore a wide range of particles and interactions. During Run-3, over a billion particle interactions occur within the detector every second, with collisions happening at a center-of-mass energy of $\sqrt{s} = 13.6$ TeV. Of these interactions, only one in a million is flagged as potentially interesting and recorded for further study. The detector has a cylindrical structure, featuring a central barrel region and two end-cap regions extending on either side. Measuring 44 meters in length, about 25 meters in diameter, and weighing approximately 7000 tonnes, ATLAS operates deep underground, in a cavern 100 meters below the surface.




Figure 2.4: Cut-away view of the ATLAS detector and its subsystems [16].

The ATLAS detector is composed of six distinct detecting subsystems, wrapped concentrically in layers around the collision point. This multi-layered detection system enables the precise recording of particle trajectories, momentum, and energy, allowing for individual particle identification and measurement. As particles pass through each layer, they leave unique traces, which are analyzed to determine their properties and interactions. This comprehensive setup ensures that ATLAS can effectively capture and study the diverse array of particles produced in high-energy collisions. A huge magnet system surrounds these subsystems, bending the paths of charged particles, which allows for highly precise momentum measurements.




Figure 2.5: Diagram of particle paths in the ATLAS detector [24].

Coordinate system

ATLAS uses a right-handed coordinate system with its origin at the interaction point (IP) in the center of the detector. The z -axis is defined along the beam direction and is perpendicular to the x - y plane, referred to as the transverse plane. The x -axis points from the IP to the center of the LHC ring, while the y -axis points upward, slightly tilted with respect to the vertical. The detector is divided into two sides: A-side and C-side, reflecting the positive and negative z -values, respectively. The x - y plane can be expressed in terms of cylindrical coordinates. The azimuthal angle ϕ is measured around the beam axis, ranging from $-\pi$ to $+\pi$, while the radial component measures the distance from the beam line.

In particle collider physics, the polar angle θ is used to characterize the detected particles. The polar angle θ , ranging from 0 to π , is often reported in terms of pseudorapidity, defined as:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \quad (2.4)$$

This expression is particularly useful in the relativistic limit, with η values ranging from $-\infty$ to $+\infty$. For particles with high θ (close to perpendicular to the beam axis), η has lower values (barrel region), whereas for particles with θ closer to the beam line, η approaches infinity (end-cap regions).

Pseudorapidity is advantageous in experimental settings because it can be easily determined from the polar emission angles of the particles. It is often preferred over the rapidity, defined as:

$$y = \frac{1}{2} \ln\left(\frac{E + p_z}{E - p_z}\right) \quad (2.5)$$

where E is the energy and p_z is the momentum component along the z -axis. The concept of pseudorapidity is particularly useful because differences in pseudorapidity between two particles are Lorentz invariant under boosts along the beam direction. The transversal momentum p_T and transverse energy E_T are also defined in terms of the polar angle as $p_T = p \cdot \sin\theta$ and $E_T = E \cdot \sin\theta$, respectively. These quantities are critical in analyzing

collisions, as they provide insights into the momentum and energy distribution in the plane perpendicular to the beam direction.




Figure 2.6: Geometric relation between pseudorapidity η , the azimuthal angle ϕ and the polar angle θ [25].

2.2.1 Detector subsystems

Building upon its layered design, each of the ATLAS's different subsystems plays a crucial role in the detection and analysis of particle interactions, working together to provide a complete picture of the collision events. In the central barrel region, these components are arranged in concentric cylinders, while in the forward regions, where the two end-caps are located, the detectors are mounted on disks perpendicular to the beam axis. This arrangement allows ATLAS to detect particles across a broad range of angles and energies, maximizing its sensitivity and enabling detailed studies of the particles produced in collisions.

Magnet System

Particles emerging from collisions normally travel in straight lines, but for precise momentum measurements, the trajectories of charged particles need to be bent. The ATLAS magnetic system achieves this by curving the paths of these particles as they pass through the various layers of the detector, allowing for the measurement of both momentum and charge. This bending is accomplished using two distinct types of superconducting magnet systems - a solenoidal magnet that surrounds the inner detector, and a set of large toroidal magnets that encases the outer sections.




Figure 2.7: (a) The field lines generated by the ATLAS magnet system, showing the influence of the solenoidal and toroidal fields on charged particles. (b) The layout of the ATLAS magnet system, highlighting the central solenoid (blue), the barrel (red) and end-cap (green) toroid magnets. The central solenoid covers a pseudorapidity range of $|\eta| < 1.4$, while the toroid magnets extend the coverage to the range of $1.4 < |\eta| < 1.6$, ensuring effective particle bending across different regions of the detector [26], [27].

The main sections of the magnet system are:

Central Solenoid Magnet

The solenoid magnet [28], located around the inner detector, plays a critical role in ATLAS by generating a strong, uniform magnetic field parallel to the beam axis. This field bends the paths of charged particles as they traverse the inner layers, enabling precise tracking and momentum measurements at lower radii. This powerful magnet spans 5.8 meters long, has a diameter of 2.56 meters, and weighs over 5 tonnes. It produces a 2 Tesla magnetic field within a remarkably thin 4.5-centimeter thickness. This impressive performance is achieved by embedding over 9 kilometers of niobium-titanium superconducting wires into strengthened, pure aluminum strips, effectively minimizing potential interactions between the magnet and the particles being studied, ensuring clear and accurate data collection.



Figure 2.8: The central solenoid magnet of the ATLAS detector, responsible for generating a strong and uniform magnetic field that bends the paths of charged particles within the inner detector, enabling precise momentum measurements [29].

Barrel Toroid Magnet

This large, cylindrical magnet [30] surrounds the calorimeters and generates a magnetic field that bends the paths of charged particles in the outer detector. Its expansive magnetic field extends over a larger volume, providing the necessary field strength to deflect particles effectively across the entire detector. At 25.3 meters in length, this is the largest magnet ever constructed. It uses more than 56 kilometers of superconducting wire and weighs about 830 tonnes, making it a critical component in guiding charged particles through ATLAS's vast detection systems and enhancing the precision of momentum measurements at larger radii.



Figure 2.9: The barrel toroid magnet of the ATLAS detector, a large superconducting magnet that surrounds the calorimeters, producing a magnetic field that bends the paths of charged particles in the outer regions of the detector, essential for momentum measurements of particles emerging from high-energy collisions [29].

End-cap Toroid Magnets

Positioned at each end of the barrel toroid, these magnets [31] extend the magnetic field into the forward regions of the detector. They ensure that particles emerging at small angles relative to the beam line are also deflected, allowing for complete momentum measurements across a broad range of angles. The end-cap toroids extend the bending capabilities to particles that leave the detector close to the beam pipe. Each end-cap toroid measures 10.7 meters in diameter and weighs 240 tonnes, providing the necessary magnetic field coverage to maintain precise momentum and charge measurements in these critical forward regions.




Figure 2.10: The end-cap toroid magnet of the ATLAS detector, positioned at each end of the barrel toroid. These magnets extend the bending capability into the forward regions, ensuring that particles emerging at small angles relative to the beam line are deflected, allowing precise momentum measurements across a wide range of angles [29].

Inner Detector

Positioned closest to the interaction point, the inner detector [32] forms the initial layer of the ATLAS detection system. It plays a critical role in measuring the direction, momentum, and charge of electrically-charged particles generated from each proton-proton collision. This compact and highly sensitive system comprises three types of sensor technologies, all immersed in a uniform magnetic field aligned with the beam axis, which aids in tracking particles with great precision. These three main components are: the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT).




Figure 2.11: The inner detector of ATLAS, extending to a radius of 1.1 meter and a length of 6.2 meters along the beam pipe, provides tracking measurements in the pseudorapidity range of $|\eta| < 2.5$ [33].

Pixel Detector

The pixel detector is the innermost component and the first point of detection in the ATLAS experiment. It consists of several layers of high-resolution silicon pixels, that provide extremely precise measurements of particle positions near the collision point, which is crucial for identifying primary and secondary vertices, such as those associated with short-lived particles. As charged particles emerge from the collision point, they deposit small amounts of energy in the pixel detector. These energy deposits are measured with a precision of almost $10\ \mu\text{m}$, allowing for accurate determination of the particle's origin and momentum. This high level of precision plays a critical role in tracking particles and studying their interactions at the heart of the detector.



Figure 2.12: The pixel detector, the innermost component of the ATLAS inner detector, featuring layers of high-resolution silicon pixels. It provides precise measurements of particle positions near the collision point, essential for identifying primary and secondary vertices [33].

Semiconductor Tracker

Surrounding the pixel detector is the SCT, which consists of multiple layers of silicon micro-strip sensors arranged in four concentric cylinders in the barrel region and nine disks in each of the two end-cap regions. The SCT adds further precision in measuring particle trajectories, especially in the intermediate radial range, by providing multiple position measurements along the track of each particle. This enhances momentum resolution and charge determination, crucial for reconstructing the paths of charged particles produced during collisions. With its fine granularity, the SCT can measure particle tracks with a precision up to $25\ \mu\text{m}$, contributing significantly to the overall tracking performance of the inner detector.




Figure 2.13: The semiconductor tracker, positioned around the pixel detector, consists of silicon micro-strip sensors arranged in concentric cylinders and disks. It enhances the precision of particle trajectory measurements in the intermediate radial range, contributing significantly to momentum resolution and charge determination [33].

Transition Radiation Tracker

The outermost layer of the inner detector is the TRT, composed of thousands of straw tubes filled with a gas mixture. Each straw, with a diameter of just 4 mm, contains a $30\ \mu\text{m}$ gold-plated tungsten wire at its center. The TRT offers continuous tracking by detecting ionization signals generated as charged particles pass through the straws, ionizing the gas to create detectable electrical signals. This setup enables the precise reconstruction of particle tracks. Additionally, the TRT has the unique capability to distinguish between electrons and other charged particles through the detection of transition radiation, emitted when high-energy charged particles cross the boundaries between materials with different refractive indices. This feature significantly aids in particle identification, enhancing the overall performance of the inner detector.




Figure 2.14: The transition radiation tracker, the outermost component of the inner detector, consists of thousands of straw tubes filled with gas. It provides continuous tracking and enhances particle identification by detecting transition radiation, particularly useful for distinguishing electrons from other charged particles [33].

Calorimeters

Following the inner detector, the calorimeter system [34] in ATLAS is responsible for detecting particles that interact via electromagnetic and strong forces. So, it comprises two main types of calorimeters: electromagnetic and hadronic. The electromagnetic calorimeters measure the energy of electrons and photons by observing their interactions with matter, while the hadronic calorimeters assess the energy of hadrons, such as protons and neutrons, through interactions with atomic nuclei. Calorimeters are designed to stop most particles, except muons and neutrinos, forcing incoming particles to deposit their energy within the detector. The ATLAS calorimeters are composed of alternating layers of high-density "absorbing" materials, which degrades the energy of incoming particles, and "active" layers that detect and measure their energy deposition as they pass through the detector. This layered structure enables precise measurement of the particles' energies and contributes to a comprehensive analysis of the collisions occurring within the detector.

The components of the ATLAS calorimetry system are: the Liquid Argon (LAr) Calorimeter and the Tile Hadronic Calorimeter (TileCal).




Figure 2.15: Cut-away view of the ATLAS calorimeter system. LAr and TileCal are illustrated along with their segmentation in the barrel and end-cap regions, providing coverage in $|\eta| < 4.9$ [35].

Liquid Argon Calorimeter

The LAr calorimeter is primarily responsible for detecting and measuring the energy of electromagnetic particles, such as electrons and photons, as well as hadrons in the forward region. This calorimeter uses liquid argon as the active detection medium due to its excellent performance characteristics, including high linearity, precise energy resolution, and strong resistance to radiation damage. It features layers of dense metal absorbers, such as lead or copper, interspersed with liquid argon. When high-energy particles strike the absorbers, they create a cascade or "shower" of secondary, lower-energy particles. As these secondary particles pass through the liquid argon, they ionize the medium, generating an

electric current. This current is then collected by electrodes, allowing the energy of the initial particle to be reconstructed from the sum of the measured signals. By combining all of the detected currents, physicists can determine the energy of the original particle that hit the detector. The central region of the calorimeter is specially designed to identify electrons and photons.

A distinctive feature of the LAr calorimeter in the central region is its accordion-like structure with a honeycomb pattern, which provides uniform coverage without gaps and ensures that all particles are efficiently detected. To maintain the argon in its liquid state, the calorimeter is kept at a cryogenic temperature of -184°C . The electronics that process the signals from the liquid argon are located outside the cold volume to avoid interference and are connected via specially designed, vacuum-sealed cylinders of cables that transmit the electronic signals from the cold liquid argon environment to the warmer readout electronics.




Figure 2.16: The accordion structure of the LAr calorimeter. Honeycomb spaces position the electrodes between the lead absorber plates [36].

Tile Hadronic Calorimeter

The TileCal is designed to measure the energy of hadrons, such as protons, neutrons, and pions, which interact strongly with the detector material and do not fully deposit their energy in the LAr calorimeter. It also provides some measurements of electromagnetic particles in the central region, complementing the LAr calorimeter's capabilities. The TileCal uses steel as the absorbing material, which slows down and absorbs the energy of incoming particles, and plastic scintillating tiles as the active medium that detects the resulting energy deposits. As the heaviest component of the ATLAS experiment, weighing almost 2900 tonnes, the TileCal is organized into a central barrel and two extended barrel sections on each side, which together cover a pseudorapidity range up to about $|\eta| < 1.7$.

When hadrons pass through the calorimeter, they excite the plastic scintillating tiles, causing them to emit light. This light is collected by wavelength-shifting fibers and then read out by photomultiplier tubes, which convert the light into electrical signals. These signals are then processed to reconstruct the energy and position of the particles that interacted with the detector. The TileCal operates in conjunction with the LAr calorimeters to provide a complete and continuous profile of the energy and position of particles, ensuring full calorimetric coverage of the detector.




Figure 2.17: The tile hadronic calorimeter in one of the end-cap regions [35].

Muon Spectrometer

Located in the outermost region of the ATLAS detector, the muon spectrometer [37] identifies and measures the momenta of muons, which are among the few particles that can traverse the entire detector without being absorbed by the calorimeters. The spectrometer covers a wide pseudorapidity range and provides excellent angular and momentum resolution, crucial for studying processes involving high-energy muons, such as those from decays of heavy particles like the W and Z bosons, or potential new particles like the Higgs boson and hypothetical supersymmetric particles. By accurately reconstructing the trajectories of muons, the spectrometer contributes significantly to the overall capability of ATLAS to explore new physics at the energy frontier.

The muon spectrometer is made up of around 4000 individual muon chambers, strategically installed in both the barrel and end-cap regions of the ATLAS detector. These chambers utilize four different gas-filled detector technologies: Resistive Plate Chambers (RPCs), Thin Gap Chambers (TGCs), Monitored Drift Tubes (MDTs), and the New Small Wheel (NSW)-which incorporates small-strip Thin Gap Chambers (sTGCs) and micromegas (MM) detectors. Each technology is tailored to specific regions of the detector, optimizing performance across various angular ranges and radiation environments.




Figure 2.18: The muon spectrometer at the ATLAS experiment and its complex system of various gas-filled chambers. The NSWs replaced the CSCs for Run-3 [38].

The MDTs and NSW provide precise position measurements, crucial for accurately tracking muons, while the RPCs and TGCs are used for fast triggering and timing information, allowing the system to rapidly identify and record interesting events. This combination of technologies enables the muon spectrometer to maintain high resolution and efficiency, even in the challenging conditions of high-energy collisions, ensuring robust muon detection and momentum measurement across the entire detector.

The MDTs are the primary detectors in the muon spectrometer and consist of cylindrical drift tubes filled with a gas mixture. As muons pass through, they ionize the gas, and the resulting electrons drift towards a central wire, allowing for precise measurement of the muon trajectory. The NSW is deployed in the forward regions, where particle flux is highest, and consists of two different technologies as is already mentioned. The sTGCs provide high spatial resolution and are designed to measure the position of muons with great accuracy, while the MM detectors enhance the readout capabilities and offer excellent time resolution. RPCs consist of two parallel resistive plates with a gas gap between them. When a muon passes through, it ionizes the gas, triggering an avalanche that is collected on external readout strips, enabling rapid detection. TGCs, employed in the forward region, are multi-wire proportional chambers with closely spaced anode wires and cathode planes. They provide fast timing and good spatial resolution, which are essential for making rapid trigger decisions.




Figure 2.19: The layout of the four primary gas-filled detector technologies in the ATLAS muon spectrometer. Top left: RPCs, which provide fast timing and triggering in the barrel region. Top right: TGCs, used for tracking and triggering in the forward regions. Bottom left: MDTs, the main precision tracking detectors covering both barrel and end-cap regions. Bottom right: NSW, offering high-resolution tracking in the high-radiation forward areas of the spectrometer [39].

The sTGCs and micromegas (MICRO MESH Gaseous Structure) detectors were a more recent addition in the muon spectrometer, designed to enhance the performance of the system, particularly in the high-radiation environment of the forward regions.

Small-strip Thin Gap Chambers (sTGCs) are designed to provide precise tracking capabilities in challenging conditions. These detectors feature thin gas gaps and are equipped with closely spaced readout strips that allow for high-resolution spatial measurements of charged particles. When a muon traverses an sTGC, it ionizes the gas, generating electron-ion pairs. The electric field across the gas gap directs the electrons towards the readout strips, where their position can be accurately determined. The sTGCs excel in providing excellent tracking efficiency and fast response times, making them especially valuable for detecting muons in high-density environments. Their integration into the ATLAS upgrade aims to significantly enhance the overall tracking performance of the muon spectrometer, ensuring that precise measurements can be obtained even in the most demanding experimental conditions.

Micromegas detectors are advanced gas-based devices that provide high spatial resolution and fast response times, making them particularly effective for precise tracking of muons. These detectors consist of a thin, micro-mesh electrode placed just above a micrometer-scale amplification gap filled with a gas mixture. When a muon passes through the detector, it ionizes the gas, and the resulting electrons are accelerated towards the mesh, creating a cascade of secondary ionizations that are then collected as a signal. This configuration allows micromegas detectors to operate with fine granularity and low noise, contributing significantly to the overall precision of the muon spectrometer. Their implementation in the ATLAS upgrade aims to improve the accuracy of muon tracking and momentum measurements, especially in the high-occupancy areas close to the beam line.

2.2.2 Trigger and Data-Acquisition system

The LHC collides protons at a frequency of around 40 MHz, meaning collisions occur every 25 nanoseconds. This high rate of particle interactions produces data that far exceeds the read-out capabilities of the detector subsystems, as well as the bandwidth and storage capacity of CERN's computing center. To manage this, the ATLAS Trigger and Data Acquisition (TDAQ) system [40] is responsible for deciding in real-time which data to record and store for offline analysis. The selective trigger system identifies events with interesting characteristics, ultimately reducing the event rate to ~ 100 Hz and recently up to 2 kHz. This significant reduction is achieved by applying selection criteria at each stage of the system, refining the decisions made at previous levels.

The TDAQ system operates through three levels of real-time event selection, designed to efficiently collect data from the detector systems, convert it digitally, and then transport it to CERN's permanent data storage for detailed offline analysis. The process begins with the first level, which rapidly examines data from specific detector components to make an initial selection. The subsequent levels apply more complex algorithms and use a broader range of detector data to refine event selection further. This multi-tiered approach ensures that only the most relevant data are kept, optimizing the use of storage and computational resources.

A diagram illustrating all the relevant components for triggering, as well as the detector read-out and data-flow used in Run-2, is shown in:

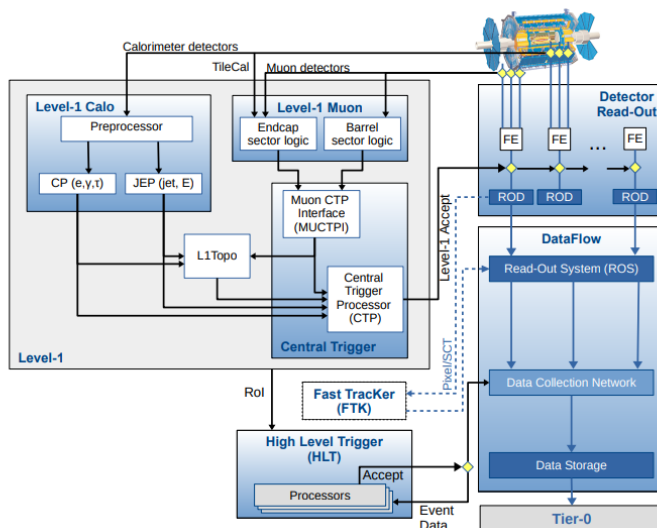


Figure 2.20: The ATLAS TDAQ system during Run-2, illustrating the key components relevant involved in event triggering, detector read-out, and data flow management [40].

The **Level-1 (L1)** trigger is a hardware-based system designed to make rapid decisions about which events to keep for further analysis. It uses custom electronics to process reduced-granularity information from the calorimeter (L1Calo) and muon detectors (L1Muon). The L1Calo trigger uses signals which are digitized and calibrated by the preprocessor and are sent in parallel to the Cluster Processor (CP), where electron, photon and τ -lepton candidates are identified above a certain threshold, and to the Jet/Energy-sum Processor (JEP), where jet candidates are identified and sums of total and missing transverse energy are produced. The L1Muon trigger relies on hit patterns from the RPCs (in the barrel) and TGCs (in the end-caps). Recently, the NSW Trigger was integrated into the L1Muon system to enhance trigger performance in the end-caps. To reduce the trigger rate from particles not originating from the interaction point, coincidence requirements are applied between the outer and inner TGC stations, as well as between the TGCs, the NSW, and the tile calorimeter. The final L1 trigger decision is made by the Central Trigger Processor (CTP), which receives inputs from the L1Calo and L1Muon through the L1Muon Central Trigger Processor Interface (MUCTPI) and the L1 topological (L1Topo) trigger, along with additional trigger signals from other detector subsystems. The CTP is also responsible for applying dead time, a mechanism to limit the number of L1 accepts to be within the $2.5 \mu\text{s}$ time window (latency) permitted by the detector electronics. As a result, the L1 trigger accepts events at a rate of up to 100 kHz, which is the maximum read-out rate permitted at this level.

In **Level-2 (L2)** trigger, each event accepted by the L1 trigger prompts the Front-End (FE) detector electronics to read out event data from all detectors. This data is first sent to the ReadOut Drivers (RODs), which handle initial processing and formatting of the information. The processed data are then transferred to the ReadOut System (ROS), where they are buffered and stored temporarily. The ROS serves as an intermediate stage, holding the data from the different subdetectors until the next trigger level requests them.

In addition, the L1 triggers identify the Regions-of-Interest (RoIs) in η and ϕ within the detector to be further investigated in the next stage. The L2 trigger system uses more refined algorithms and higher granularity data compared to L1, significantly reducing the event rate down to approximately 1 kHz. This reduction is achieved within a latency of about 40 ms per event, ensuring that only events with the most promising characteristics are passed in for further processing in the final trigger level.

The **Level-3 (L3)** trigger, also referred to as the High-Level Trigger (HLT), is a software based system that operates in two main stages to further refine the event selection process. Initially, dedicated fast trigger algorithms provide early rejection of events based on a quick analysis of data, focusing on efficiently discarding uninteresting events. Following this, more precise and computationally intensive algorithms, similar to those used in offline reconstruction, are applied to make the final decision on whether to keep an event. The HLT operates with a latency of approximately 4 seconds per event, as it leverages the full detector granularity and employs complex reconstruction techniques, including precise tracking, vertex reconstruction, and particle identification. It performs a comprehensive examination of the event data, which is crucial for reducing the event rate down to approximately 100 Hz and recently up to 2 kHz, the final output rate that the storage and offline analysis systems can handle. Once an event is accepted by the HLT, the data is stored in RAW (bytestream) format and sent to permanent storage for offline reconstruction. These data are then exported to the Tier-0 facility at CERN's computing center.

Chapter 3

The New Small Wheel

The New Small Wheel (NSW) [41] is a critical upgrade to the ATLAS detector’s muon spectrometer, designed to cope with the increased collision rates and challenging conditions expected in future runs of the LHC, particularly during the High-Luminosity LHC (HL-LHC) era. Replacing the original small wheel, the NSW provides enhanced capabilities for muon tracking and triggering in the forward region of the detector, where particle flux is highest and backgrounds are more severe. The NSW combines two advanced gas-filled detector technologies: small-strip Thin Gap Chambers (sTGC) and micromegas detectors (MM), which together offer excellent spatial resolution, fast response times, and robust operation under high radiation levels. This upgrade significantly improves the ATLAS experiment’s ability to maintain precise muon measurements and high trigger efficiency, ensuring continued success in probing the fundamental nature of particle interactions at unprecedented energy scales.

3.1 Motivation

The motivation for the NSW upgrade arises from the increased challenges posed by the HL-LHC environment. As the LHC pushes towards higher luminosities, the rates of proton-proton collisions will significantly increase, leading to a much higher density of particles and more complex event topologies. In the forward regions of the ATLAS detector, where the original small wheel was located, the increased particle flux results in higher background noise, reduced precision in tracking, and a decrease in trigger efficiency. These conditions threaten the performance of the ATLAS muon spectrometer, especially its ability to accurately measure the momentum of muons and effectively distinguish between signal and background events.

3.1.1 Precision tracking performance

The muon spectrometer in the ATLAS detector is affected by an unexpected high background originating from hits in the forward shielding, which are projected radially into the end-cap toroid cryostat. This background can produce track segments in the Big Wheels that falsely appear to originate from the interaction point, leading to an influx of fake muon triggers. At high luminosity, this could overwhelm the L1 Muon trigger bandwidth, primarily with these spurious triggers, risking the saturation of the system and potentially missing genuine low- p_T muon events. The NSW was specifically designed and constructed to address this issue by enhancing trigger confirmation and track reconstruction capabilities in the inner end-cap region. The NSW effectively filters out the muon-like background, thereby negating the necessity to prescale low- p_T muon triggers. Equipped with advanced detector technologies, the NSW can rapidly and precisely identify track segments, sending the trigger data within a latency of $1.10 \mu\text{s}$ after a collision to be cross-checked with the Big Wheel triggers. By matching track segments between the NSW and the Big Wheel, the

system can veto Big Wheel tracks that lack corresponding segments in the NSW, significantly reducing false triggers. Additionally, the track angle provided by the NSW allows for rejection of tracks that did not originate from the interaction point, thus refining the overall muon trigger selection and preserving the spectrometer's performance even under the high-luminosity conditions expected at the HL-LHC.




Figure 3.1: An illustration of the Muon end-cap trigger system. The existing Big Wheel trigger accepts all three tracks shown (A, B, and C). With the NSW enhancement of the muon end-cap trigger, only track "A", the desired one, which is confirmed by both the Big Wheel and the NSW, will be accepted. Track "B" will be rejected because the NSW does not detect a corresponding track originating from the interaction point that matches the Big Wheel candidate. Track "C" will be rejected because the NSW track does not point to the interaction point, indicating it is likely a background track.

3.1.2 Trigger selection

Performance studies using collision data revealed unexpectedly high rates of fake triggers in the end-cap region of the muon spectrometer. The L1 Muon trigger for the end-cap region relies on signals from the TGC chambers at the Big Wheel. The transverse momentum p_T of muons is determined by the angle of the detected signal segment relative to the direction pointing to the interaction point. However, low-energy particles, such as protons produced in the material between the Small Wheel and the electromagnetic calorimeter, can produce false triggers when they traverse the end-cap chambers at angles similar to those of genuine high- p_T muons. These fake triggers account for approximately 90% of the total triggers in the end-cap region, leading to the L1 Muon trigger rate being 8 to 9 times higher than that in the barrel region. This issue is illustrated in Figure 3.2, which shows that a significant portion of the reconstructed muons do not correspond to genuine muon candidates identified by the inner detector. The figure depicts the η distribution of muon candidates selected by the ATLAS L1 trigger with $p_T > 10$ GeV. It compares the distribution of those candidates that indeed have an offline reconstructed muon track with the distribution of reconstructed muons with $p_T > 10$ GeV. The majority (over 80%) of the muon trigger rate originates from the end-cap regions ($|\eta| > 1.0$) and is not reconstructed offline.

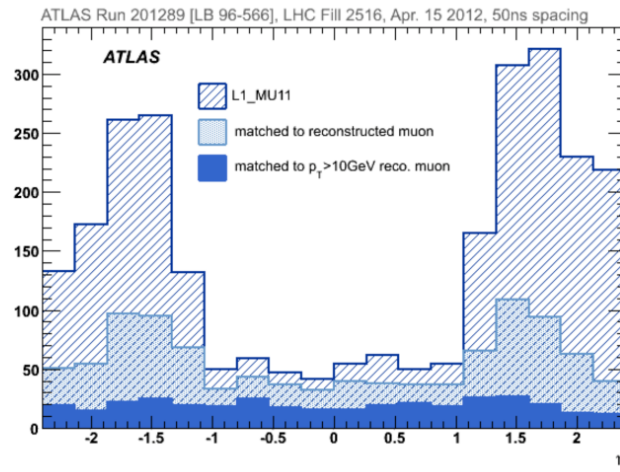


Figure 3.2: η distribution of L1 muon signal ($p_T > 10$ GeV) with the distribution of the subset with matched muon candidate to an offline well reconstructed muon, and offline reconstructed muons with $p_T > 10$ GeV.

3.2 The NSW layout

To ensure compatibility with the existing tracking detectors and the end-cap alignment system, the NSW layout adheres to the design principles of the current Small Wheels, while introducing advanced detector technologies for improved performance. The NSW is installed on each end of the ATLAS detector, consisting of a total of 16 sectors per wheel, arranged to provide comprehensive coverage in the azimuthal angle ϕ . The layout includes eight "large" and eight "small" sectors. The small sectors are oriented facing the interaction point (IP), while the large sectors are positioned closer to the confirm point (HO), optimizing the coverage and alignment with the muon spectrometer's existing components.




Figure 3.3: The sector arrangement of the NSW, where each sector consists of two wedges of sTGC (green) and MM (brown) [42].

Each sector in the NSW contains 16 detector planes arranged in a specific configuration to maximize performance. The detector planes are divided into four multi-layers, "wedges", each consisting of different detector technologies. The arrangement of these wedges is as follows: sTGC-MM-MM-sTGC. This configuration ensures that the distance between the sTGC layers on opposite wedges is maximized, which enhances spatial resolution and minimizes potential noise. Each sTGC wedge is composed of three modules or quadruplets and the MM wedges are composed of two quadruplets, with each quadruplet having four detector layers.

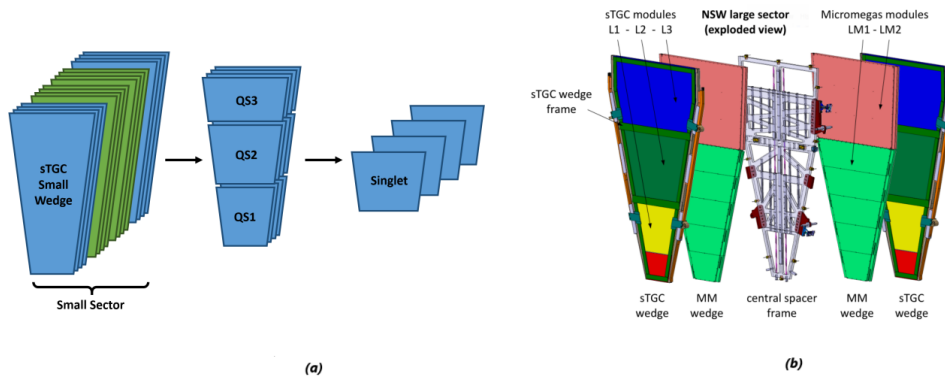


Figure 3.4: (a) Diagram of a small sector structure, showing sTGC wedges in blue and MM wedges in green. (b) Schematic representation of the detector technology layers in a NSW sector, illustrating the arrangement of sTGC and MM wedges and their respective quadruplet structures [43], [44].

3.3 The detector technologies of the NSW

The New Small Wheel (NSW) incorporates advanced detector technologies specifically selected to overcome the limitations of the previous Small Wheel, with a focus on reducing the high rates of fake triggers and enhancing spatial resolution for muon tracking in the forward region, as previously mentioned. Below, the two primary technologies are described.

3.3.1 Small-strip Thin Gap Chambers

sTGCs are multiwire ionization chambers utilized as the primary triggering detectors in the NSW. Each chamber features a central wire plane composed of $50 \mu\text{m}$ diameter gold-coated tungsten wires, spaced 1.8 mm apart, laid between two cathode planes at a 1.4 mm distance from the wire plane. These cathode planes are coated with resistive graphite. On one side, the readout planes are divided into fine strips with a 3.2 mm pitch, oriented perpendicular to the wires, and on the other side, into large rectangular pads of varying sizes. The detector operates with a gas mixture of carbon dioxide (CO_2) and n-pentane (C_5H_{12}) in a 55:45 ratio at one atmospheric pressure, with an applied voltage of 2.8 kV. This setup enables precise detection of muons, contributing to the NSW's improved trigger capabilities and spatial resolution.

The pads in the sTGCs are used to create a 3-out-of-4 (3/4) coincidence within each sTGC quadruplet, helping to identify muon tracks that approximately point to the interaction point. These pads also determine which strips need to be readout for a precise measurement in the bending coordinate, defining the region of interest for the online muon candidate selection. The strips and wires contribute to the precision of muon track reconstruction in the η direction. Charge clusters formed on the strips are used to calculate a centroid within each quadruplet (both pivot and confirm layers), which are then used to construct a track segment, enhancing the overall accuracy of muon tracking and triggering.

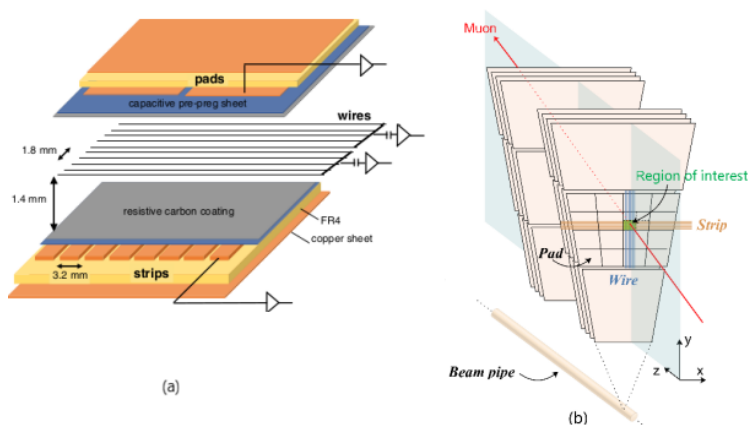


Figure 3.5: (a) Schematic illustration of the structure of sTGC detectors, showing the arrangement of the wire, strip, and pad readout components. (b) Orientation of the sTGC wire, strip and pad readout electrodes within a NSW sector, displayed in relation to the ATLAS coordinate system [45], [46].

As muons and other charged particles pass through the sTGCs, they ionize the gas within the chamber, creating electron-ion pairs. The free electrons drift towards the anode wires at the center of the chamber, where they undergo amplification through avalanche multiplication, resulting in a current signal. This process induces signals on the adjacent strips and pads, allowing for precise position measurements. However, in high-rate environments, the accumulation of drifting ions can cause long signal tails due to their slower movement compared to electrons. To mitigate this effect, sTGCs are designed with thin drift gaps, which reduce the maximum drift time and help minimize the build-up of ions, thus preventing constant current and maintaining a manageable readout voltage.

The electric potential, and consequently the electric field, is significantly lower in the midpoint between two adjacent wires compared to the regions closer to the wires. When a muon passes exactly between two wires, the ionization electrons experience a delayed drift into the higher-field region before being pulled towards a wire, where they generate an avalanche and subsequent signal. This results in a delayed signal compared to a typical detection event. To minimize this delay, the wires are staggered by 0.25 wire-spacing between layers, reducing the likelihood that a muon will pass directly between wires in multiple layers. This staggered arrangement ensures more consistent signal timing across the sTGC layers.

3.3.2 Micromegas

The MM technology enables the construction of thin, wireless gaseous particle detectors with excellent precision tracking capabilities. A MM detector consists of a planar drift electrode, a gas gap a few millimeters thick that serves as the conversion and drift region, and a thin metallic mesh positioned typically 100-150 μm above the readout electrode, forming the amplification region. The detectors are filled with a gas mixture of 93% argon, 5% carbon dioxide (CO_2), and 2% isobutane, optimized for efficient detection and fast response. MM detectors were chosen for the NSW due to their ability to deliver fast and precise tracking across a wide range of track angles.




Figure 3.6: Sketch of the layout and operating principle of a MM detector. The diagram illustrates the key components of the MM detector: the planar drift cathode, the gas gap acting as the conversion and drift region, and the metallic mesh creating the amplification region. As charged particles pass through the detector, they ionize the gas, and the free electrons drift towards the Micro Mesh, where they produce avalanches above the readout electrodes.

The original design of MM detectors was vulnerable to sparking, which occurs when the avalanche of electrons reaches a critical size, typically around a few 10^7 electrons. To mitigate this issue in the NSW, the readout electrodes were designed using a spark-protected resistive strip scheme. Additionally, a carefully configured high voltage (HV) scheme was implemented to further reduce the risk of sparking, a common problem in other MM detectors. In this scheme, a positive HV of 490 to 520 V is applied to the resistive strips, while the amplification mesh is kept at ground potential (0 V). A negative HV of -300 V is applied to the drift electrode. These voltage settings create an electric field of a few hundred V/cm in the drift region and approximately 40-50 kV/cm in the amplification region. As charged particles pass through the drift region, they ionize the gas, liberating electrons that drift towards the grounded mesh, while positive ions move towards the drift electrode. Upon reaching the mesh, electrons undergo avalanches, amplifying the signal directly above the readout electrodes. This process happens swiftly; all electrons created in the drift gap reach the mesh within approximately 100 ns, while the amplification occurs in a fraction of a nanosecond, resulting in fast pulses on the readout strips. Positive ions produced in the avalanche are cleared in about 100 ns, ensuring quick charge evacuation. This rapid clearing of ions makes MM detectors well-suited for operation in the high particle flux environment anticipated during LHC Run-3 and the HL-LHC.

3.4 NSW electronics and Data Acquisition dataflow

The NSW trigger system relies on track segments generated in real-time by the sTGC and MM detectors. To manage this, the NSW trigger electronics and data acquisition (DAQ) dataflow encompass 128 detectors and approximately 2.4 million readout channels. The system design includes distinct paths for trigger signals, data readout, and configuration/monitoring for each type of detector. The electronics chain is divided into two main sections: on-detector electronics, which handle immediate data processing and signal transmission directly from the detectors, and off-detector electronics, located in the underground service area, which are responsible for further data processing and system control. Both sTGC and MM detectors share this electronics architecture, ensuring streamlined operations and efficient data management.

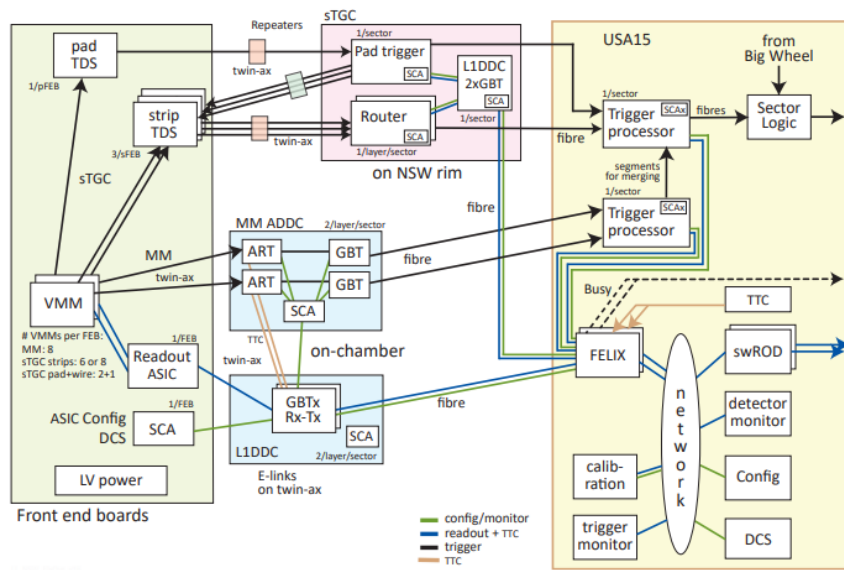


Figure 3.7: Overview of the NSW trigger and DAQ electronics [47].

The on-detector electronics include the Front-End Boards (FEBs), which are electronic cards equipped with radiation-tolerant Application Specific Integrated Circuits (ASICs). These FEBs play a crucial role in processing and transmitting data and trigger information from the NSW detectors to the subsequent stages of the electronics system. For the sTGC detectors, the FEBs are connected to the Pad Trigger and Routers for triggering and to the Level-1 Data Driver Card (L1DDC) for configuring and monitoring, while for the MM detectors, they are connected to ART Data Driver Card (ADDC) for triggering and similarly, as for the sTGC, to the L1DDC. These cards are responsible for aggregating Level-1 data - such as time, charge, and strip addresses corresponding to individual hits - from multiple FEBs. This aggregated data is then transmitted to the Front-End Link eXchange (FELIX) system, which serves as the interface between the detector electronics and the off-detector processing systems.

The off-detector electronics comprise high-throughput Back-End components that are critical for processing and managing data from the NSW detectors. Key elements include

the sTGC Rim Crate, which hosts various boards such as the Pad Trigger Board, Rim L1DDC, Router Boards, the Trigger Processor, the Sector Logic, the ALTI. Another key element is the FELIX interface that connects the detector readout systems to the data acquisition framework. Additional services like Read Out Drivers (ROD) and the Detector Control System (DCS) are also part of the off-detector setup. Communication between these electronics components is facilitated through mini-Serial Attached Small Computer System Interface (SAS SCSI) cables and optical fibers, ensuring efficient data transfer and robust connectivity throughout the system.

The Pad Trigger receives pad hits from the pad Trigger Data Serializer (pad-TDS) for all eight sTGC layers per sector and makes trigger decisions based on 3-out-of-4 layer coincidences in each quadruplet, processed independently. This coincidence logic, running parallel to the strip data collection, sends trigger information such as bunch crossing identification (BCID), strip-band ID, and a second coordinate (ϕ) ID back to the strip Trigger Data Serializer (strip-TDS). This allows for the selection of relevant charge data from the band of strips in each layer that intersect the tower producing the pad coincidence. The Router then collects data packets from the active strip-TDS and forwards them to the sTGC Trigger Processor, which computes centroids and track segments that align with the Big Wheel. The NSW muon candidates from the MM and the sTGC Trigger Processors are then merged and sent to the Sector Logic, where they are combined with Big Wheel candidates. Finally, the Sector Logic sends the Level-1 trigger candidates to the ATLAS Muon Central Trigger Processor. This entire process repeats every bunch crossing with a fixed latency.




Figure 3.8: A band of strips in each sTGC layer is selected by a particle making a 3-out-of-4 hit coincidence within a pointing tower of sTGC pads in each quadruplet. To enhance resolution, the pads in half of the layers are shifted by half a pad in both directions. Eight-layer towers pointing to the interaction point are defined by the overlapping physical pads (shown in grey), which collectively identify a logical pad (highlighted in red) in each layer [47].

Chapter 4

Targeted masking in the ATLAS experiment

This chapter will describe the development of a tool for the implementation of individual pad masking in the NSW pad trigger system. Initially, the focus was on 3/4 and 3/4 coincidence logic, but the scope expanded to include 2/4 and 3/4 coincidence as well. Various improvements of this tool followed until it took a final form. However, the dynamic nature of the system means that as conditions change, new parameters need to be considered, and further enhancements are continually needed.

4.1 Motivation

The coincidence 3/4 and 3/4 requires the presence of signals from all eight layers for maximum efficiency. However, several issues can hinder this performance, including High Voltage (HV) problems, failures in FEBS configuration, missing FEBS, non-connected pads, dead pads due to electronics, and also noisy pads. These detector defects impact the overall efficiency, often preventing the coincidence condition from being satisfied. To activate a sector and maintain data from this, the Sector Logic (SL) from its side requires an efficiency greater than 95%. As a result, 70% of the sectors achieved this efficiency threshold, were activated, and included in ATLAS trigger at the end of 2023 proton-proton runs.

The aim of this study was to include all the sectors with high efficiency using individual pad masking. Each pad registers either a "1" or "0", indicating whether it has detected a signal. A "dead" pad consistently outputs 0, while a noisy pad persistently outputs 1. The developed tool allows for masking of entire pFEBS or individual channels, setting the pad input to 1 for dead pads and to 0 for noisy ones. This masking has to be done in a way that maximizes the efficiency but at the same time it does not create hot spots in the pad trigger. Therefore, masking must be applied judiciously to ensure that the coincidence condition is met appropriately, avoiding to give trigger all the time incorrectly.

4.2 Analysis in the first steps

In the NSW pad trigger system, patterns are formed by projecting lines from the interaction point and determining their intersections with the detector layers. This process involves calculating the expected hit positions in each layer based on the trajectories of particles. The pad trigger system uses these patterns to identify candidate muon tracks that are consistent with originating from the interaction point. Patterns are defined to cover a wide range of possible angles and positions, ensuring that the trigger system can correctly identify tracks across the entire acceptance of the detector. When a pattern that matches the observed hits is identified, it is linked to a Lookup Table (LUT). The LUT contains precomputed information specific to each pattern, including the number of the eight pFEBS

involved and the corresponding pad channel numbers for these trajectories in the detector layers. Additionally, the LUT stores the PhiID and BandID.




Figure 4.1: Pattern formation in the NSW pad trigger system: lines projected from the interaction point intersect with detector layers, identifying candidate muon tracks.

Information about patterns is stored in the VHD files that are actually used in the FPGA of the PT. Initially, large sectors had 4260 patterns, while small sectors had 1958 patterns. The first step in the analysis involved examining the multiplicity of patterns for each combination of pFEB and pad channel. Multiplicity is defined as the number of patterns in which each pad participates for a specific pFEB. A special case exists for patterns in transition regions (these regions refer to the areas at the boundaries between different quads of the detector), where four real layers are complemented by four virtual layers. For these virtual layers, special encoding is used: pFEB 0 is assigned with pad numbers 104 and 105, which do not correspond to actual pads. Typically, each sector has 24 pFEBs, and the pad numbers vary depending on the specific pFEB, as well as whether the sector is large or small. A mapping system is employed to handle these variations, and it was crucial to precisely determine the starting point for measurements across all components. For example, pFEB numbering begins at 0 and extends up to 23.

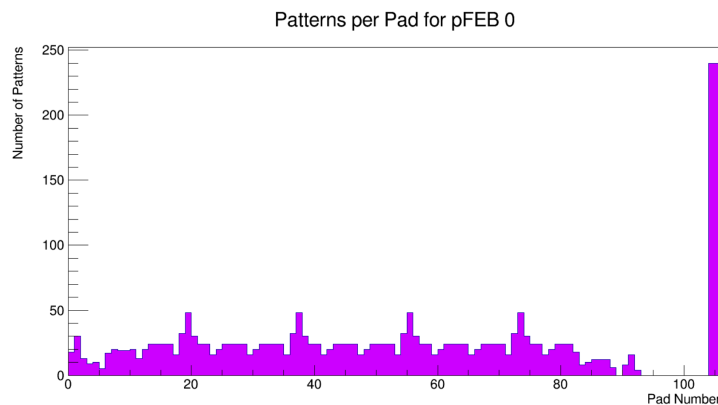


Figure 4.2: The number of the patterns in which each pad participates for the pFEB 0 of a large sector (multiplicity). The pads 104 and 105 is an encoding used for the transition regions.

Having gained an insight into the patterns, the next and very important step was to identify the dead pad channels. At the beginning, data from a specific run (R455889) were analyzed, and histograms were created to record the hits for each pFEB. Dead pad channels were then defined as those channels that registered no hits at all in the dataset. This method allowed for a straightforward identification of non-functional pad channels.

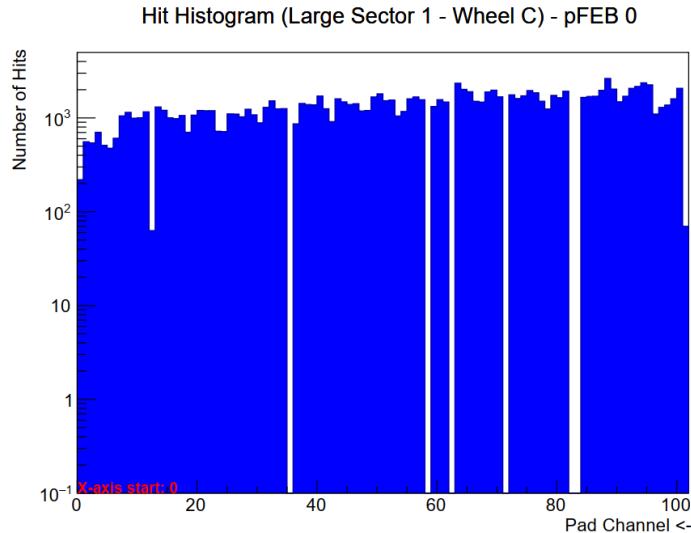


Figure 4.3: The histogram of the hits for pFEB 0 of the large sector 1 in wheel C. The arrow (<-) indicates for each pFEB whether the pad channels are in ascending or descending order. For this example they are in descending order. The pad channels that have no hits are considered as dead ones (pads: 35, 58, 62, 71, 82, 83).

The histogram above illustrates the limits used for the pads of this pFEB in the specified sector, excluding channels 104 and 105, which, as previously mentioned, are associated with transitions and serve as "flags" to indicate that such pads are either "masked to 1" or "masked to 0". A more detailed analysis involved examining the pad channel limits for each pFEB based on the pad adapter board channel mapping and comparing these to the pad channels derived from the patterns for each pFEB. This analysis revealed an additional two pad channels that were non-connected and did not correspond to real pads: channels 91 and 92, specifically for pFEB 6 in the large sectors only.

The overall goal was to identify dead pad channels using a two-pronged approach that works in tandem. First, known detector issues such as missing pFEBs or HV being off would be used to make an initial coarse determination of dead pads. This coarse assessment would then be supplemented by analyzing experimental data, specifically by filling histograms with hits, as illustrated in the example above. This combined method ensures a more comprehensive identification of dead pads.

Another useful tool for diagnosing coarse dead pad issues was the creation of 2D histograms for each sector. In these plots, the x-axis represents the detector layers, and the y-axis represents the pFEBs, with the color intensity indicating the number of hits. This visualization allows for an intuitive assessment of detector performance, making it easy to identify entire pFEBs that are dead and to determine in which quadruplet the issue resides.

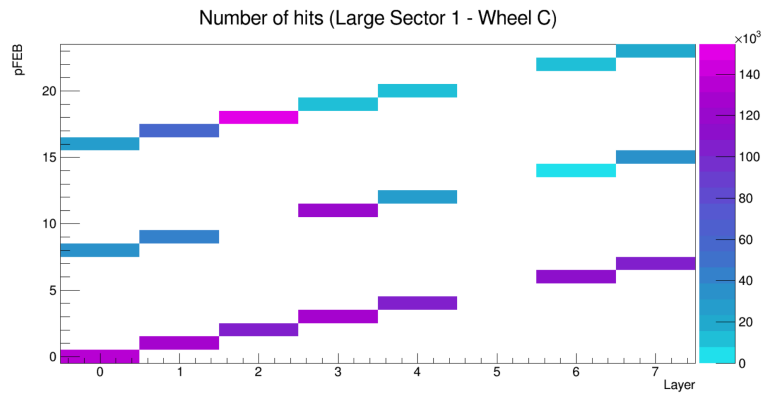


Figure 4.4: The number of the hits for each layer (x-axis) and for each pFEB (y-axis) for the large sector 1 in wheel C. It's easy to find out if a whole pFEB is dead and in which quad it is. In this example, the pFEB 10 is dead in the second quad and in the layer 5 all the pFEBs are dead in all the quads.

Finally, a 2D histogram was created for each sector, with the x-axis representing the number of dead pad channels in the IP wedge and the y-axis representing those in the HO wedge. To clarify, the IP wedge refers to the section of the detector closest to where the particle collisions occur, while the HO wedge is the section farther away from the interaction point. The color intensity in these plots reflects the number of patterns exhibiting these conditions. These histograms do not count as dead the transition layers and non-connected pads, providing a clearer view of the actual number of the dead pads per wedge for each pattern.

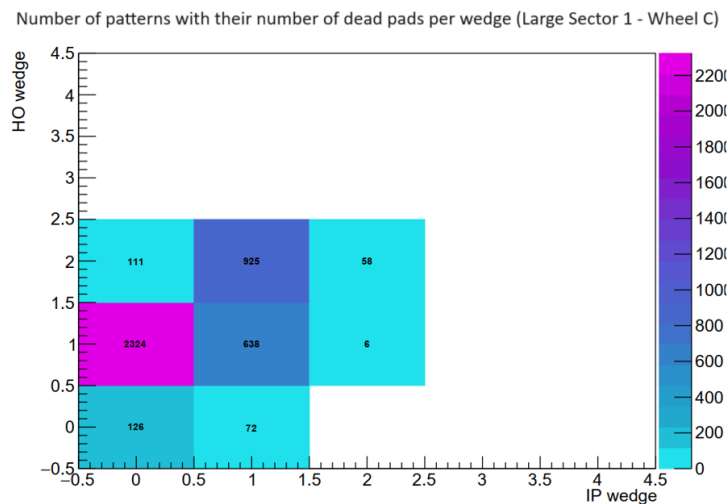


Figure 4.5: The number of the patterns that have a particular number of dead pad channels in the IP (x-axis) and in the HO wedge (y-axis) for the large sector 1 of wheel C.

4.3 Targeted masking with 3/4 & 3/4 coincidence

Having done all this study and keeping in mind the results, a new method, referred to as targeted masking, was developed to address the detector defects. The idea behind this approach was to handle each pattern individually to maximize efficiency without creating any hot spots in the detector. Initially, a 3/4 and 3/4 coincidence was supposed to be used and a pad efficiency of 95% was assumed. To implement this, a database of all individual pads with their statuses was built, taking into account the permanent and data-defined defects in the detector. Each pad's status was categorized as: okay, dead, noisy, permanently masked, dead masked to 1, or noisy masked to 0. This status was recorded with an encoding system applied, as presented in the table below:

Possible statuses	
1	okay
2	dead
3	noisy
102	permanently masked (transition - non connected)
12	dead masked to 1
13	noisy masked to 0

Table 4.1: The possible statuses of the pad channels.

After categorizing the pad statuses, the next step involved representing all pattern statuses as 8-dimensional vectors, which were then mapped onto a pattern phase space. This phase space contains all the possible pattern statuses. Within this space, a forbidden region was identified, where patterns must never appear. Patterns in this region would mean that they give a trigger on every BC (Bunch Crossing), which is undesirable. Conversely, an allowed region was defined, encompassing both patterns with zero efficiency and those with a reasonable, good efficiency. This is explained from the fact that some level of unmasking is necessary to prevent patterns from falling into the forbidden region, but this unmasking also results in patterns that will never trigger. The goal is to minimize the number of non-triggering patterns while maximizing the number of patterns that operate with optimal efficiency.




Figure 4.6: The pattern phase space showing the forbidden region and the allowed regions.

Everything described above is represented by an encoding system, as shown in the next table.

Pattern Phase Space	
0	Hot spots & zero efficiency under normal conditions
1	Regular efficiency & zero efficiency (to prevent trigger in every BC)

Table 4.2: The encoding system for the Pattern Phase Space.

To summarize, an array is programmatically constructed to record the status of each pad, for each pFEB, in every sector and wheel. This array serves as the foundation for determining the status of each pattern. Consequently, changing the status of a pad in this array will automatically update the status for every pattern that includes it. Initially, a pad status database was created by assigning values: 1 for "alive" pads, 2 for "dead" pads, and 102 for special cases. Following this, the pattern statuses were updated based on this database, and the initial efficiency of each pattern was calculated. The average efficiency before masking was then extracted for each sector of every wheel. These calculations incorporated the previously mentioned pad efficiency assumptions and the stricter coincidence requirements.

Pattern Efficiency

To calculate the efficiency for a pattern, two independent events, A and B, are considered. Event A represents the coincidence of the first four pads, while event B represents the coincidence of the other four pads. According to statistical principles, for two independent events, the probability of their intersection - the collection of all outcomes that are elements of both sets A and B - can be expressed as:

$$P(A \cap B) = P(A) \cdot P(B). \quad (4.1)$$

The probability for each of these events is calculated from the binomial distribution. In probability theory and statistics, the binomial distribution with parameters n and p describes the discrete probability distribution of the number of successes in a sequence of n independent experiments, each posing a yes-no question and yielding a Boolean-valued outcome: success (with probability p) or failure (with probability $q = 1-p$). A single success/failure experiment is known as a Bernoulli trial or Bernoulli experiment, while a sequence of outcomes is referred to as a Bernoulli process. The probability of obtaining exactly k successes in n independent Bernoulli trials (with the same rate p) is given by the probability mass function:

$$f(k, n, p) = \binom{n}{k} p^k (1-p)^{n-k} \quad (4.2)$$

where

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} \quad (4.3)$$

is the binomial coefficient.

To obtain a trigger, it is necessary to calculate the probability of achieving at least three successes in four attempts (i.e., three hits in four layers) with a probability of 95% for every wedge. This includes considering both scenarios: having three successes in four trials and having four successes in four trials. After calculating these probabilities for both wedges, the two probabilities are multiplied together to determine the overall efficiency of the pattern.

Using similar logic and the same statistical principles, all relevant probabilities can be computed, with slight variations depending on the specific situation. For example, if a single pad in one wedge is dead, a trigger requires hits from the remaining three pads, so the probability of obtaining three hits from three attempts is calculated. However, if multiple pads are dead in the same wedge, the overall efficiency drops to zero, as it becomes impossible to satisfy the coincidence condition. In cases where a pad is masked to 1 (considered permanently active), the probability is adjusted accordingly. The trigger then relies on the remaining three pads, and the probability is calculated based on obtaining either two hits from three trials or three hits from three trials.

The criteria for masking included pad participation in patterns, achieving the highest possible efficiency, and avoiding the creation of hot spots, which could lead to excessive triggering. The trigger system was originally designed to use a 3/4 and 3/4 coincidence with staggering both in η and ϕ . This staggering was expected to divide each physical pad into four logical pads. However, it was later discovered that staggering also occurs between the two wedges, adding further complexity. As a result, each physical pad could be divided into as many as sixteen logical pads, rather than just four. Due to real-world imperfections, this number could actually be even larger. This complexity made defining pad multiplicity critical, as it became a key factor in the masking process. The strategy followed was to initially mask all dead pads (setting their status to 12) and then selectively unmask some (resetting their status to 2) based on specific criteria. First, it was considered sufficient to unmask at least two pads in the same wedge. Second, the pads chosen for unmasking were those with the smallest multiplicity, meaning those that participated in the fewest patterns. After each unmasking, the pad status was updated so that changes were implemented system-wide, and the criteria were rechecked to determine whether further unmasking was necessary.

Nalve HO \ Nalve IP	0	1	2	3	4
0	x	x	x	(3, 1)	(3, 0)
1	x	x	(2, 2)	(2, 1)	(2, 0)
2	x	(2, 2)	(2, 2)	(2, 1)	(2, 0)
3	(1, 3)	(1, 2)	(1, 2)	(1, 1)	(1, 0)
4	(0, 3)	(0, 2)	(0, 2)	(0, 1)	(0, 0)


Figure 4.7: A tabular illustration of masking cases. "X" indicates no masking, and the numbers in parentheses represent the number of masked pads per wedge (N_{masked}^{IP} , N_{masked}^{HO}).

After unmasking, the new average efficiency was recalculated using the same assumptions as before. To generalize the computation of efficiency programmatically, the following formula was employed:

$$\left(\frac{4 - N_{dead}}{3 - N_{masked-to-1}} \right) p^{3 - N_{masked-to-1}} (1 - p)^{1 - N_{dead} + N_{masked-to-1}} \quad (4.4)$$

In this formula, N_{dead} represents the number of dead pads, and $N_{masked-to-1}$ indicates the number of pads that have been masked to 1. The parameters allow for a comprehensive assessment of efficiency under varying conditions.

Special attention was given to the patterns in the transition regions, where there are four real layers and four virtual layers. In these patterns, three virtual layers are masked (setting their status to 102) to enable the 3/4 and 3/4 coincidence, ensuring that a trigger is feasible. There are two options for handling the real layers. One approach is to mask one real layer, allowing triggering based on only two layers, which can yield higher efficiency but may also introduce more noise. The alternative is to leave all real layers unmasked, resulting in a drop in efficiency to approximately 70% for that pattern. Using the first approach, the expected efficiencies after applying the masking in the ATLAS experiment are presented below for each sector of each wheel.






Figure 4.9: An illustration of the efficiencies before and after masking. The errors are minimal and indistinguishable.

To sum up, this method successfully recovers most of the defective patterns, although a small number may remain inactive in particularly problematic cases. After masking, the average PT efficiency for wheel A is 97.34% with the worst-performing sector showing an efficiency of 94.0%. For wheel C, the resulting average PT efficiency is 97.26%, with the lowest sector efficiency at 95.9%.

At the final stage, the pattern phase space was computed both before and after the masking-unmasking process. Programmatically, this pattern phase space is represented as a four-dimensional array, where the dimensions are: the number of alive pads in the IP wedge, the number of alive pads in the HO wedge, the number of masked pads in the IP wedge, and the number of masked pads in the HO wedge. As previously discussed, avoiding hot spots is mandatory, even if it results in some patterns having zero efficiency. When the status of a pad changes (e.g., to masking), this change propagates across all patterns that involve that pad. Therefore, the primary goal is to eliminate hot spots, while maximizing the efficiency remains a high but secondary priority. This approach naturally divides the phase space into two key regions:

The forbidden region: Before masking, this region contains patterns with zero efficiency, due to multiple dead layers. After the masking-unmasking process, it should be empty from patterns that incorrectly trigger on every BC, but it will still contain patterns with zero efficiency caused by necessary unmasking to avoid hot spots, which affects these patterns as well.

The allowed region: Before masking, this region includes patterns with regular efficiency values. After the masking-unmasking process, it continues to contain patterns with regular efficiency, but it also includes patterns with zero efficiency to prevent them from becoming hot spots.

Some examples are provided below to illustrate these concepts and make them easier to understand. In the top left plot, the x-axis represents the number of masked pads in the IP wedge, while the y-axis corresponds to the number of masked pads in the HO wedge. This plot reflects the situation after the initial masking but before any unmasking has occurred. The color intensity indicates the number of patterns with these statuses, and the exact number is also written for each case. The top right plot follows the same format but shows the situation after the masking-unmasking process has been applied. The non-allowed area in these plots highlights the hot spots. Below these plots, a table summarizes the wheel and sector, presenting the pattern phase space before masking along with the efficiency in that state, as well as the pattern phase space and efficiency after the masking-unmasking process has been performed.

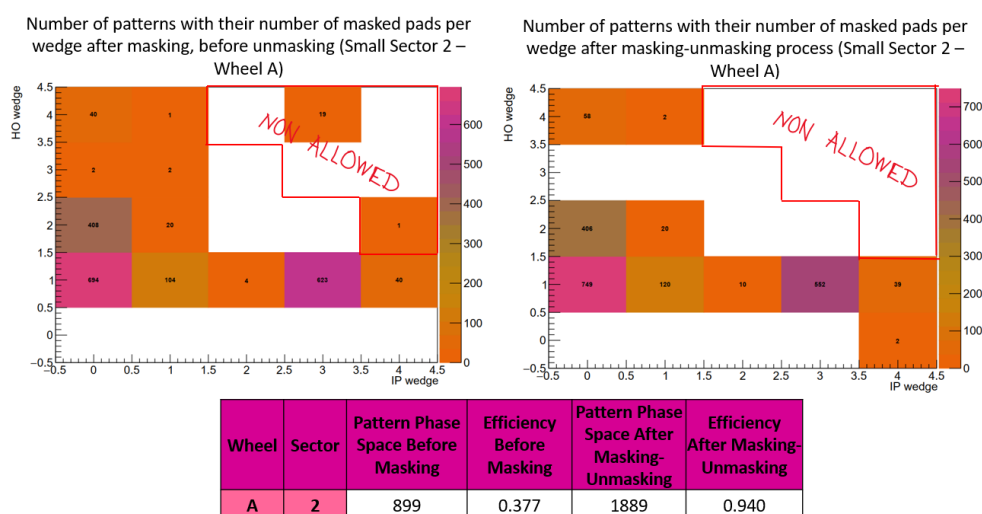


Figure 4.10: The 20 patterns shown in the top left plot represent those that trigger on every BC, requiring unmasking to remove them from the non-allowed region. This is addressed in the top right plot, where the non-allowed region is empty, indicating that the status of these patterns has changed. Before masking, only 899 patterns had regular efficiency, resulting in a very low average efficiency, while the remaining 1059 patterns had zero efficiency due to multiple dead layers. After the masking-unmasking process, the situation improved: 1889 patterns, minus the 20 patterns that must maintain zero efficiency so that they are not anymore hot spots, now contribute to the average efficiency with regular values. The remaining 69 patterns have zero efficiency as a result of the unmasking process.

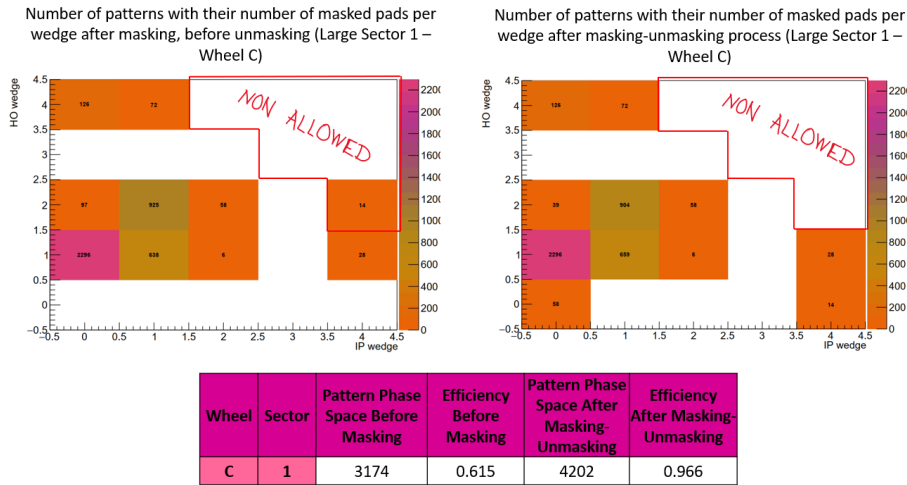


Figure 4.11: The 14 patterns shown in the top left plot were initially in the non-allowed region, but after the unmasking process, they were successfully removed, as shown in the top right plot. Before masking, 3174 patterns contributed to the average efficiency with regular values, while the remaining 1086 had zero efficiency. After the masking-unmasking process, 4202 patterns, minus the 14 that were hot spots and now must have zero efficiency, contribute to the average efficiency with reasonable values. The remaining 58 patterns necessarily have zero efficiency, as they were affected by the unmasking.

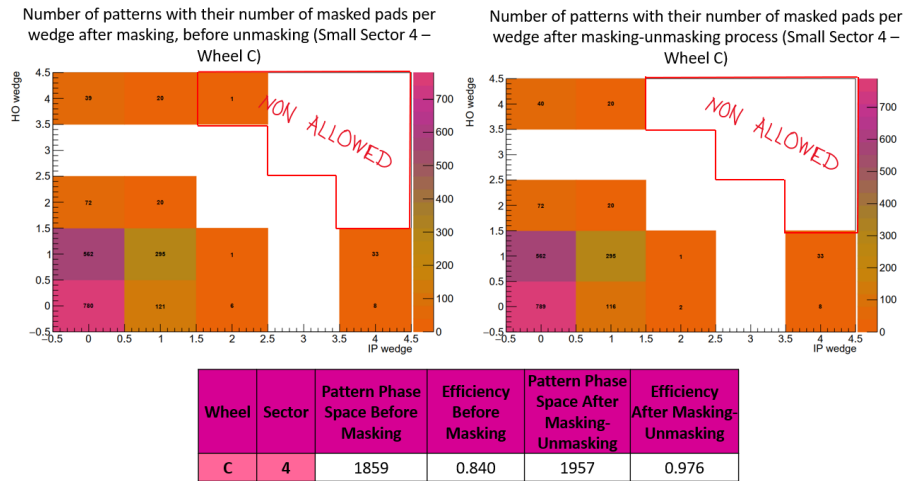


Figure 4.12: This case represents a generally good sector. Only 1 pattern is a hot spot in the top left plot, and it is successfully removed from the non-allowed region after unmasking, as shown in the top right plot. Initially, 1859 patterns had reasonable efficiency, contributing to a good average efficiency, while 99 patterns had zero efficiency. After the masking-unmasking process, 1957 patterns, minus this 1 hot spot that now has zero efficiency, contribute to the average efficiency with regular values. Only 1 additional pattern is affected by the unmasking, resulting in zero efficiency at the end.





Figure 4.13: In this example, no patterns are in the non-allowed region, meaning no unmasking was required in this case. As a result, the top left plot and the top right plot are identical. Initially, 4254 patterns contributed to the average efficiency with regular values, while only 6 patterns had zero efficiency. After masking, all patterns contribute to the efficiency with regular values.

4.4 Implementation of targeted masking with $3/4$ & $3/4$ coincidence

The implementation of the targeted masking requires the creation of masking registers, followed by the development of a software tool to generate the final JSON files for each sector. This tool reads the existing JSON files, identifies the specific block where the FPGA masking registers are defined, and replaces it with a new text that contains the correct updated values for these, to ensure that the correct configuration is applied. Once the modifications are complete, the tool generates the final JSON files for each sector, ready for use in the system. By automating this process, the tool not only minimizes human error but also enhances the speed of deployment across the entire system, confirming uniformity in the masking application. Before these JSON files are deployed, they undergo standalone testing to validate their functionality. These tests involve simulating the trigger system's behavior using the new masking configurations, checking for any unwanted side effects such as hot spots or high trigger rates. This ensures that the final configuration is stable and meets the design specifications before being integrated into the full system.

Some examples of these registers are presented below. In the JSON files, the sequence of registers is structured in three stages: first, the registers for masking to 0 are specified for the 23 pFEBs, followed by the registers for masking to 1, and finally, a register that enables the masking for specific pFEBs. The registers for masking to 0 or 1 are divided into several parts, each handling a specific range of pads for a given pFEB, which allows fine-grained control of the masking process. This structure makes sure that the masking is applied accurately and efficiently for different pads within each pFEB. The values of these registers are stored in decimal form.

pFEB_00_mask_to_0_A	30	R/W	31:8 7:0	NOT USED PADS [103:96]	CHECK TRUTH TABLE BELOW	0x00
pFEB_00_mask_to_0_B	31	R/W	31:0	PADS [95:64]		0x00000000
pFEB_00_mask_to_0_C	32	R/W	31:0	PADS [63:32]		0x00000000
pFEB_00_mask_to_0_D	33	R/W	31:0	PADS [31:0]		0x00000000



```

"030_pFeb_00_mask_to_0_A": {
  "pads[103:96]": 0
},
"031_pFeb_00_mask_to_0_B": {
  "pads[95:64]": 0
},
"032_pFeb_00_mask_to_0_C": {
  "pads[63:32]": 0
},
"033_pFeb_00_mask_to_0_D": {
  "pads[31:0]": 0
},

```




```

00000000
0000000000000000000000000000000000000000
0000000000000000000000000000000000000000
0000000000000000000000000000000000000000

```

Figure 4.14: An example of the masking registers to 0 for pFEB 0. At the top, the definition of these registers is provided, showing the number of bits, their description, and the default values. Below, the same example is displayed as it appears in the JSON file, with an arrow indicating the binary form of the values. This demonstrates how the values are represented programmatically for configuration purposes.

pFEB_00_mask_to_1_A	90	R/W	31:8 7:0	NOT USED PADS [103:96]	CHECK TRUTH TABLE BELOW	0x00
pFEB_00_mask_to_1_B	91	R/W	31:0	PADS [95:64]		0x00000000
pFEB_00_mask_to_1_C	92	R/W	31:0	PADS [63:32]		0x00000000
pFEB_00_mask_to_1_D	93	R/W	31:0	PADS [31:0]		0x00000000



```

"090_pFeb_00_mask_to_1_A": {
  "pads[103:96]": 0
},
"091_pFeb_00_mask_to_1_B": {
  "pads[95:64]": 67108863
},
"092_pFeb_00_mask_to_1_C": {
  "pads[63:32]": 4294967295
},
"093_pFeb_00_mask_to_1_D": {
  "pads[31:0]": 4294967295
},

```




```

00000000
0000001111111111111111111111111111111111
1111111111111111111111111111111111111111
1111111111111111111111111111111111111111

```

Figure 4.15: An example of the masking registers to 1 for pFEB 0. At the top, the definition of these registers is identical to that of the masking registers to 0. Below, the same example shows the values that these registers take in decimal form in the JSON file, with an arrow indicating their binary form.

PFEB_MASK_ENABLE	F0	R/W	31:24 23:0	NOT USED PFEB_MASK_ENABLE, EVERY BIT CORRESPONDS TO A PFEB		0x00000000
------------------	----	-----	---------------	--	--	------------



```

"0F0_pFeb_mask_enable": {
  "pFeb mask enable": 16777215
}

```



```

1111111111111111111111111111111111111111

```

Figure 4.16: The register for enabling masking for specific pFEBs. At the top is the definition of this register, and below is how it appears in the JSON file with a value that enables masking for all pFEBs. This is clearly indicated by the binary form, shown by the arrow, which reflects that masking is enabled for the 24 pFEBs.

Having these registers, the process of masking is determined by a truth table that outlines the logic used. This truth table is presented in a general form below, but specifically for masking pads to 1, the following logic is applied: 1) Enable masking for all pFEBs by setting the corresponding bits for each pFEB to 1. This was demonstrated in the previous

example, where the binary form of the register's value shows that there is a sequence of ones. 2) For the pads that are not masked, the "mask to 0" bits are set to zero, and the "mask to 1" bits are also set to zero. 3) For the pads that are masked to 1, the corresponding "mask to 1" bits are set to one, while the "mask to 0" bits remain to zero.

This approach guarantees that the system accurately applies the desired masking configurations and correctly interprets which pads should be masked to 1, while others remain unaffected, according to the values defined in the masking registers.

TRUTH TABLE			
PFEB MASK ENABLE	CONF MASK TO 0 [103:0] (n BIT)	CONF MASK TO 1 [103:0] (n BIT)	MASKED HITS [103:0] (n BIT)
0	X	X	UNMASKED HITS
1	0	0	UNMASKED HITS
1	0	1	MASKED TO ONE
1	1	0	MASKED TO ZERO
1	1	1	MASKED TO ONE

Figure 4.17: The truth table used to produce the desired masking configurations.

With all these preparations in place, the targeted masking was ready to be implemented in the experiment. Before proceeding to the results of this implementation, a few important remarks are worth noting. During this time, new LUTs for the patterns were introduced into the pad firmware, reflecting updated BandID definitions. Since the BandID is proportional to the radius of the pad, these adjustments led to a more homogeneous distribution of patterns across the sectors. As a consequence, the number of the patterns was updated, with 4027 patterns for the large sectors and 1869 for the small ones. Additionally, in the transition regions, it was decided not to mask any real layers, a strategy chosen to maintain lower trigger rates.

The first implementation of targeted masking was conducted for the most challenging sector, sector A10, during run R473617. Using the $3/4$ and $3/4$ coincidence without any masking, the expected average efficiency for this sector was around 28%. On the other hand, after the masking process, the anticipated average efficiency rose to approximately 96%. In practice, the experimental results revealed an actual efficiency of around 90%. While this represented a significant improvement compared to the initial efficiency without any masking, the achieved efficiency did not meet expectations. Consequently, this level of efficiency was insufficient for the sector to be activated in the coincidence by the SL. Below, several plots illustrate in more detail the first results after applying targeted masking to sector A10.

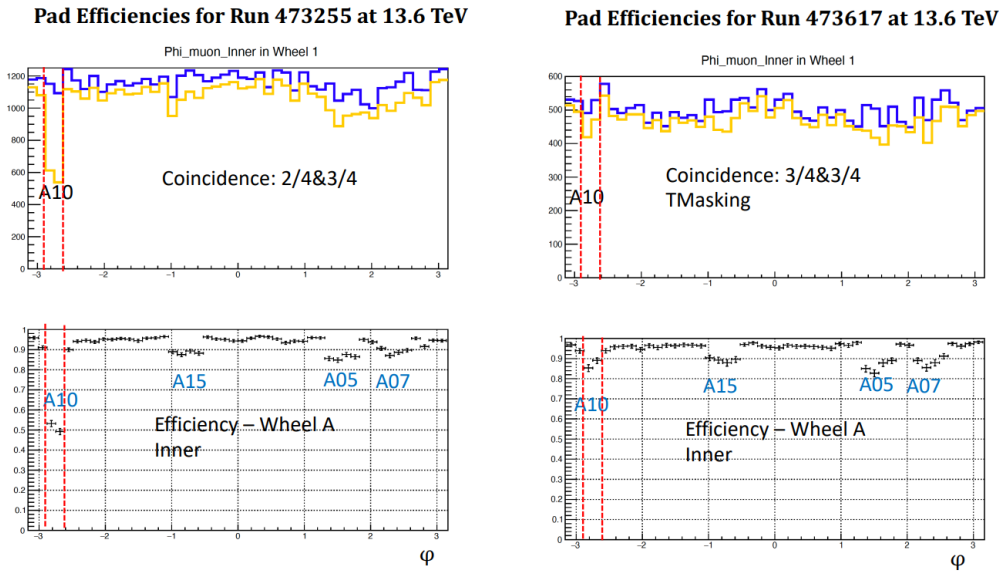


Figure 4.18: Plots of the pad efficiencies for two different runs: on the left, efficiency is shown using the looser coincidence of 2/4 and 3/4 (no masking applied), while on the right, the stricter coincidence of 3/4 and 3/4 is applied, along with the implementation of targeted masking for sector A10. These plots concern the inner part of the sectors, located on the endcaps, providing a comparison between the two different cases in this region.




Figure 4.19: Plots of the pad efficiencies, this time for the forward part of the sectors, comparing the looser coincidence of 2/4 and 3/4 with the stricter coincidence of 3/4 and 3/4 after the implementation of targeted masking for sector A10.

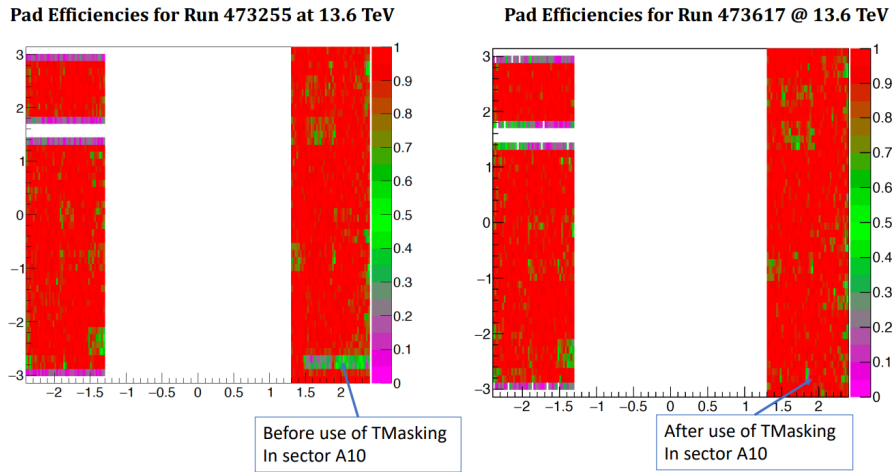


Figure 4.20: Heatmap plots of the PT efficiencies for the two different runs, highlighting the noticeable improvement in sector A10. After the implementation of targeted masking, previously inefficient areas show significant enhancement in efficiency.

Subsequently, an attempt was made to test targeted masking in additional sectors during run R473796, specifically for sectors A05, A07, and A15. Initially, using the 3/4 and 3/4 coincidence without any masking, the expected average efficiencies were approximately 87% for A05, 85% for A07, and 69% for A15. After implementing the targeted masking process, the anticipated efficiencies were expected to increase to about 97% for A05, 97% for A07, and 96% for A15. However, the actual efficiency values observed from the experiment were around 90% for A05, 80-85% for A07, and 90-95% for A15. This indicated that, once again, the actual efficiencies fell short of the expected values.




Figure 4.21: The heatmap plot of the pad efficiencies for sectors A05, A07, and A15 after the implementation of the 3/4 and 3/4 coincidence with targeted masking. The visual representation highlights the distribution of efficiencies across the pads in these sectors, illustrating the impact of the targeted masking approach on performance.

Then, an investigation was conducted to identify the causes behind the discrepancy between the expected efficiency values after targeted masking and the actual values observed during the experiment. This research led to the consideration that the initial hypothesis regarding pad efficiency might have been flawed, prompting a decision to undertake a comprehensive study on this aspect. The aim was to re-evaluate the assumptions and parameters used in the efficiency calculations, ensuring they accurately reflected the system's behavior under the applied conditions. This study was critical to refining the understanding of pad efficiency and optimizing the targeted masking approach for future implementations.

The results of this study are presented below, and they indeed confirm that the actual single pad efficiency was not the assumed 95%, but rather closer to 90%. This finding was significant as it highlighted the need for a more accurate representation of the pad efficiency in the overall efficiency calculations.

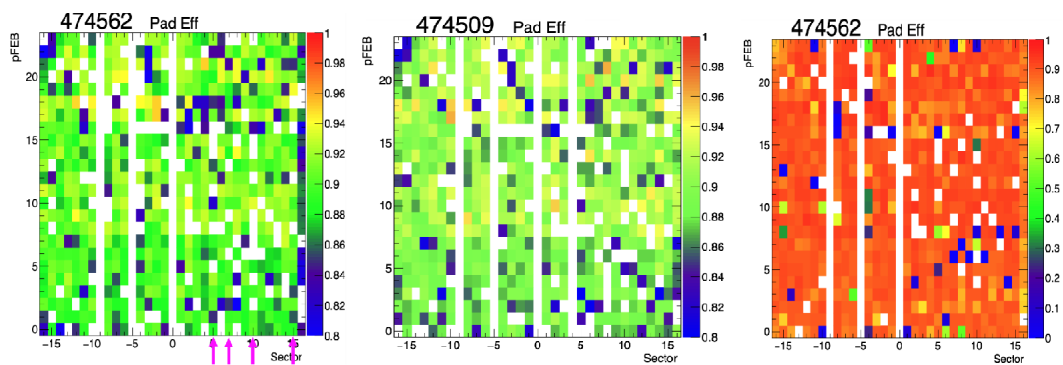


Figure 4.22: The results for pad efficiency, displayed on the left for run R474562, in the middle for another run, R474509, included for double-checking purposes, and on the right, the same plot for run R474562 but with different efficiency range. The arrays indicate the sectors A05, A07, A10, and A15. It is evident that the actual single pad efficiency is closer to 90%.

(Plots by Estel Perez Codina)

The figure 4.23 shows a plot that was constructed to display the pad trigger efficiency as a function of the pad efficiency for sector A10. When the pad efficiency is 90%, the trigger efficiency drops to around 90-92%, which further confirms the pad efficiency value, as the experimental results similarly indicated a trigger efficiency of around 90-92% as well.




Figure 4.23: The pad trigger efficiency as a function of the pad efficiency for sector A10.

Now, assuming a pad efficiency of 90%, the efficiencies for each sector and both wheels were recalculated using the 3/4 and 3/4 coincidence with targeted masking. The recalculated results are presented in the table below.

3/4 & 3/4 coincidence			
Pad Efficiency = 90 %			
Wheel C		Wheel A	
Sector	Efficiency	Sector	Efficiency
1	0.9095	1	0.9051
2	0.9007	2	0.8954
3	0.9008	3	0.9047
4	0.9091	4	0.9039
5	0.9032	5	0.9016
6	0.9015	6	0.8967
7	0.9041	7	0.9077
8	0.9016	8	0.8989
9	0.9061	9	0.9196
10	0.9103	10	0.9080
11	0.9087	11	0.9090
12	0.9069	12	0.9045
13	0.9012	13	0.8867
14	0.9066	14	0.9138
15	0.9030	15	0.8958
16	0.9030	16	0.9005

Figure 4.24: The efficiencies of all sectors and both wheels, making the assumption that the pad efficiency is 90%, for the 3/4 and 3/4 coincidence with targeted masking.

The table above clearly indicates that with a 90% pad efficiency, the trigger efficiency consistently drops to around 90-92% across all sectors. This falls of the 95% threshold required to activate the sectors in the coincidence from the SL. As a result, it was decided to re-adopt the looser 2/4 and 3/4 coincidence, while still applying targeted masking. However, this masking had to be adjusted accordingly for the new coincidence.

Nevertheless, in the end, the 3/4 and 3/4 coincidence should be used, as it is essential for trigger path involving the strips. The use of a 2/4 and 3/4 coincidence introduces a

2-fold ambiguity in determining the BandID, complicating the identification of the correct sTGC strips needed for determining the μ -direction. By implementing the 3/4 and 3/4 coincidence, a 90% efficiency can be achieved, and the rest should be complemented by the use of the MM trigger.

4.5 Initial manual implementation of targeted masking with 2/4 & 3/4 coincidence

In the previous section, the reasons that led to the adoption of the looser coincidence of 2/4 and 3/4 were explained. The following sections will focus on the development of the new masking process. Until these processes were finalized, targeted masking had to be manually applied by directly modifying the JSON files. The general logic guiding the new masking process was based on the fact that the new coincidence is looser and the primary goal is still to avoid creating hot spots. Therefore, masking was applied to a maximum of one layer per wedge to minimize unwanted effects. At this point, it's worth noting, for the sake of completeness, that sector A10 continued to use the stricter 3/4 and 3/4 coincidence, a practice that continues to this day. This sector was also used by the MM team for their tests, being the first sector where attempts were made to integrate MM into the NSW-BW coincidence. Following these changes, the SL team confirmed an efficiency increase, with the NSW sector efficiency for A10 rising from 90% to around 96%, making A10 eligible for SL activation.

Returning to the topic of masking, the new process was initially implemented for the sectors A05, A07, and A15, with masking applied very conservatively. Then, gradually, this masking strategy was extended to all sectors. The results were very encouraging, with new efficiencies around 98%. Very soon, the SL team confirmed that additional trigger sectors had stable efficiencies exceeding 95%. After including these sectors, 125 out of 144 trigger sectors (87%) were activated. Below is an image that shows these results.

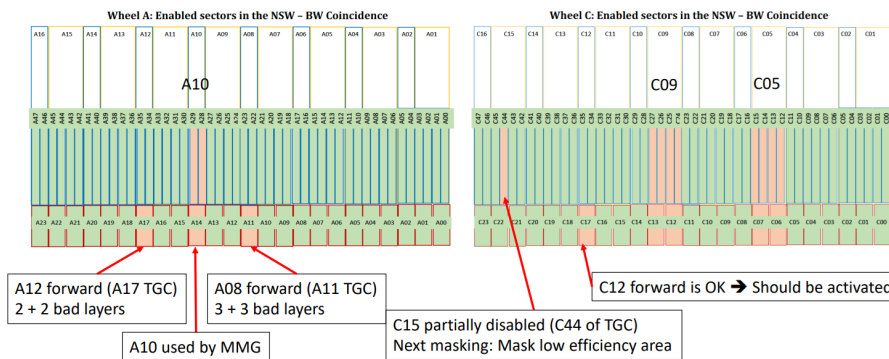


Figure 4.25: An illustration of the sectors that were enabled in the NSW-BW coincidence after the implementation of the manual targeted masking for the 2/4 and 3/4 coincidence.

At this stage, sector A10 was still being tested by the MM team and had lower efficiency. Sector C12 was left inactive by accident and was expected to be activated soon. Sectors A08 forward, A12 forward, and C15 in the endcap were showing somewhat lower efficiencies. Additionally, sector C05 had readout issues, while sector C09 experienced problems when

the toroid was on, so these sectors were not used. Further masking adjustments were planned to address inefficiencies for sectors A08, A12, and C15. The investigation into these sectors revealed that a possible cause of these inefficiencies, particularly in the forward regions, was HV trips, which primarily affected side A. There was a growing suspicion that the HV for R1 might be affecting these areas, contrary to the earlier belief that only R2 HV was responsible. Both R1 and R2 are part of Q1, with R1 being the innermost part of the detector, positioned closest to the beam.

This heatmap presented below indicates the problems in the R1 regions referred to earlier.




Figure 4.26: A heatmap for a run after the implementation of the masking, indicating that there are still some inefficiencies in the Q1 innermost part, highlighting the potential issues caused by the R1 HV trips, particularly in side A.

After implementing additional masking to cure new HV trips and inefficient areas, along with resolving the issues for sectors C05 and C09 and restoring their functionality, efficiencies for all the sectors were successfully increased above the required threshold for activation in the coincidence. A big milestone was reached on May 28, 2024, when TGC-NSW coincidence, incorporating both sTGC and MM, was integrated across the entire region of NSW!

4.6 Full development and automation of targeted masking for 2/4 & 3/4 coincidence

Programmatically, to manage the increasing complexity of this work, the code was divided into several components, and a class was constructed specifically to perform the masking. For the 2/4 and 3/4 coincidence, many things were already prepared, but some required modifications. There was an extra complexity in all the process because this coincidence is not symmetrical. The pattern phase space had to be adjusted, and the calculation of efficiency became more intricate. Two different methods were developed: a simpler approach that makes certain assumptions for convenience, and a more accurate method. Additionally, the approach to identify the critical patterns that required unmasking, along with the corresponding implementation, had to be revised. All these aspects will be examined in detail in the remainder of this section.

The first step was to separate the part of the code that handles dead pads. A combination of different processes works together to ultimately produce the array with the status of the pads, where dead pads are marked with the number "2". This optimization saves a significant amount of time, as the pad status array can be periodically updated by running this part of the code, which is time-consuming due to the data processing involved, but by isolating it, the core masking process can now run independently without delays caused by redefining pad statuses.

Finding the dead pads involves using the detector's known defects and adding mainly individual pads identified through data analysis, as previously mentioned. Initially, a part of this process is still done manually, including accounting for known detector defects. A list of HV problems, which is regularly updated, serves as a reference. An example of this list can be seen below. These defects must be input into the code, and when executed, the code flags all pads associated with the given pFEB, sector, and wheel, marking their status as "2".

HV Known Issues (Updated on 23-09-2024)

OFF or below operating voltage

Counting sectors 1-8, IP 1-4, HO 5-8

Sector	R1	R2	R3	R4	Total	Sector	R1	R2	R3	R4	Total
A01	L2, L4	L1	L6	L2	5	C01	L4	-	L7	L1, L4	4
A02	L1, L2, L3, L5, L8	L1, L5	L1, L2, L4, L5	-	11	C02	L2	L2	-	L1	3
A03	L2	-	L1, L7	L8	4	C03	L1, L3, L4	-	-	-	3
A04	L1, L2, L3, L4	-	-	-	4	C04	L3, L4, L6, L7	L8	-	L1, L5 (0 uA)	7
A05	L1, L6	L4, L6	L4, L7, L8	L1, L2	9	C05	L1, L2	-	L3	-	3
A06	L1, L2, L3, L4	L2	-	-	5	C06	L1, L2, L3, L4	-	-	L3, L7	6
A07	L1, L3, L4	L7	L2, L6	L8	7	C07	L3, L4	L4	-	L4	4
A08	L1, L2, L3, L4, L5, L6, L7	L6, L7, L8	L2, L3, L8	-	13	C08	L3, L4, L7, L8	-	-	-	4
A09	-	L5, L7	-	L2	3	C09	L1, L2, L3, L8	-	-	L5	5
A10	L3, L4, L6, L7	L8	L2	-	6	C10	L4, L6	L1	-	-	3
A11	L2	L7	L3	L5	4	C11	L1, L2	-	L5	L6	4
A12	L2, L3, L4, L5 (0 uA), L6, L8	-	L2	-	7	C12	L1, L2, L3, L4, L7	L3, L4	-	-	7
A13	L2, L5	L5	L4	L1, L6	6	C13	L2, L7	-	-	-	2
A14	L1, L2, L3, L4, L8	L5	-	-	6	C14	L1, L3, L4, L8	L4	-	-	5
A15	L1, L2, L3, L4	-	L1, L5	L1	7	C15	L1, L2, L3	-	-	L5	4
A16	L1, L2, L3, L4, L5	L2	-	-	6	C16	L1, L3	-	-	-	2
Total	55	17	21	10	103	Total	45	7	3	11	66

Figure 4.27: The list of HV known issues for the detector, regularly updated by the STGC operations team.

In this list, one can observe the R1 and R2 layers, which correspond to Q1, while R3 refers to Q2, and R4 relates to the Q3 regions of the detector. Each system uses its own method for measuring and reporting, so it is essential to verify the data carefully and use it appropriately. For example, the layers in this list are numbered from 1 to 8, but in the code, the input is indexed from 0 to 7. Notably, due to the suspicions that R1 might affect the trigger system, the R1 and R2 layers were separated and started to be handled independently. To achieve this, it was necessary to determine how many rows of the Q1 pads were associated with R1 and R2 for each layer, pFEB, and wheel, and then map these to the corresponding pads. This mapping ensured that when an issue is detected in either R1 or R2, only the pads directly associated with the affected region are marked as dead, providing more precise control over the masking process.

Another part of the code analyzes the data from a run and records in an array the hit counts for each pad across every pFEB, sector, and wheel. The code then iterates over all pads, comparing their hit counts to a predefined threshold. If the hit count for a pad falls

below this threshold, the pad is marked as dead and recorded as such. This automated step supplements the known detector defects by identifying additional problematic pads based on real-time data, ensuring a more comprehensive and accurate masking process.

These pieces of information are combined, and the final array of pad statuses is stored to be used as input in the main code that performs the masking process. Below is an image showing an example of a pad hit map for one layer. Such pad maps are often used for cross-checking and verifying that the dead pads have been accurately identified.

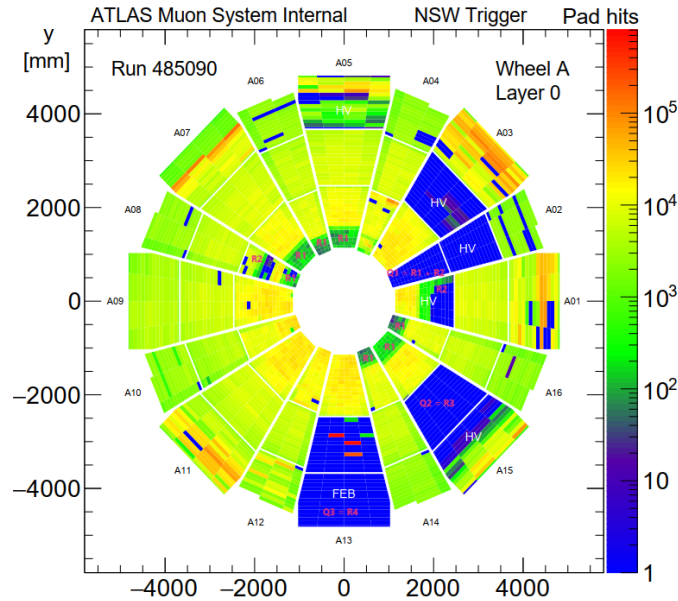


Figure 4.28: An example of a pad hit map for one layer, used for cross-checking and verifying the identification of dead pads. The blue and dark green pads are considered as dead, while those with an almost red color are classified as more noisy.

As mentioned earlier, a class was constructed to perform the masking process. At the beginning of the process, several parameters are initialized to zero, including the number of alive, dead, and masked pads in both the IP and HO wedges. Then, specific data are taken as input, among others, the PatternID, a unique identifier number assigned to each pattern, the IPattern array, a sequence of numbers where each pFEB number is followed by its corresponding pad, defining the structure of each pattern. Next, a function determines the pattern status using the array that was previously stored and prepared, which contains the current status of each pattern. All these are used to set up the patterns and determine their behavior. To do this, the code iterates through the first four or last four elements of the IPatternStatus array, which store the pad statuses for each pattern, for the IP and HO respectively. During this iteration, it checks predefined conditions to classify and count the pads in each wedge. The alive pads are identified by counting how many "1" values are present, the dead pads by the count of "2" values, and the masked pads by counting "12" or "102" values.

Another function in the class is responsible for calculating the efficiency. Statistically, these calculations were more complex for the 2/4 and 3/4 coincidence due to the asymmetry.

To simplify the process initially, the masked pads were treated as alive, resulting in slightly lower calculated efficiencies compared to the actual ones. This approximation allowed for faster computations but did not capture the true efficiency accurately. To address this discrepancy, a more precise but time-consuming approach was adopted using a pseudo-experiment method. This method essentially simulates the experiment with a random number generator to derive accurate efficiencies. During each trial, the method evaluates whether each pad fires based on its status and the pad efficiency. It then counts the number of pads that successfully fire and checks if the pattern satisfies the criteria for meeting the coincidence. The efficiency is finally computed as the ratio of successful trials to the total number of trials. This approach ensures a more accurate representation of the pattern behavior, providing reliable efficiency measurements even for complex coincidence scenarios.

The `PatternPhaseSpace` is a static variable defined differently for this coincidence. This change is necessary because the 2/4 and 3/4 coincidence is looser, allowing only one masking per wedge to prevent excessive triggers. Despite this adjustment, the overall logic remains consistent. Additionally, the `PadMultiplicityLS` and `PadMultiplicitySS` are also defined as static variables, which track how many patterns each pad participates in for the large and small sectors, respectively. This information is crucial for managing the unmasking process, as it helps identify the pads that are more suitable for unmasking due to their lower participation in multiple patterns, meaning their unmasking will not significantly affect many other patterns, particularly patterns that could give a reasonable efficiency without modifications. By leveraging these multiplicity values, the code can make informed decisions on whether unmasking a particular pad will restore performance without compromising the overall system stability.

The functions related to masking and unmasking are structured to follow a systematic approach. An initial function identifies all the dead pads and masks them by changing their status from "2" to "12". Following this, another function identifies the critical patterns - those that have more than one masking in either the IP or the HO wedge. Managing these critical patterns is essential since only a single mask is allowed per wedge to maintain the integrity of the 2/4 and 3/4 coincidence logic. A specialized function is employed to handle unmasking of transition regions if this option is selected. In the stricter 3/4 and 3/4 coincidence logic, transition regions are identified by finding exactly four occurrences of the number "102" in either the IP or HO wedge, which indicates a transition region. But now, for the 2/4 and 3/4 coincidence, a new pad status had to be introduced, the "0", to signify permanently dead pads. In this case, transition detection is performed by identifying two "102" occurrences and two "0" ones. This means that in this configuration, two virtual layers are masked to 1, and two virtual layers are masked to 0. The trigger thus requires hits from at least three real layers. In such cases, the function adjusts the other wedge by replacing every "12" with "2". At the end, there are some functions dedicated to unmasking the other patterns. Each of these functions checks specific conditions based on the number of masked pads in a wedge. The first function targets wedges with four masked pads and unmask the one that participates in the fewest patterns. Similarly, the second function addresses wedges with three masked pads, and the third handles those with two, applying the same logic each time.

In the main part of the code, instances of the class are created, and the corresponding functions are called for each instance. Initially, all the necessary data is set up, and the efficiency and pattern phase space are calculated before any masking is applied. Next,

all the dead pads are masked, and the pattern status is updated accordingly. The code then proceeds to count the number of alive, dead, and masked pads again, and the critical patterns are identified. The first step in handling these critical patterns is the unmasking of transitions. After applying the transition unmasking, the statuses are updated once more. The code then iterates over the critical patterns and applies the unmasking functions, updating the statuses each time. Once this iterative unmasking process is completed, the efficiency and pattern phase space are recalculated.

This section briefly describes the code developed to execute these processes and generate the final array with the updated pad statuses. From this array, the corresponding JSON files could then be produced. The code was written in the C++ programming language. Once the tool was fully developed, it was executed using the most recent list of HV trips, and the estimated efficiencies are summarized in the table below. It should be noted that at this stage, no unmasking was performed for the transition regions, and the separation of R1 and R2 had not yet been implemented.

2/4 & 3/4 coincidence			
Accurate			
Wheel C		Wheel A	
Sector	Efficiency	Sector	Efficiency
1	0.9882	1	0.9903
2	0.9909	2	0.8419
3	0.9880	3	0.9895
4	0.9832	4	0.9875
5	0.9885	5	0.9870
6	0.9807	6	0.9835
7	0.9866	7	0.9911
8	0.9676	8	0.9025
9	0.9887	9	0.9760
10	0.9932	10	0.9565
11	0.9899	11	0.9900
12	0.9833	12	0.9873
13	0.9882	13	0.9645
14	0.9892	14	0.9832
15	0.9820	15	0.9770
16	0.9884	16	0.9774

	L0	L1	L2	L3	L4	L5	L6	L7	IP	HO
R1	X (HV)		X (HV)						3	2
R2		X (HV)		X (HV)						1
R3	X (HV)	X (HV)							3	
R4				X (HV)						

	L0	L1	L2	L3	L4	L5	L6	L7	IP	HO
R1	X (HV)		X (HV)	X (HV)	X (HV)	X (FEB)			3	2
R2	bad pads				X (FEB)					3
R3		X (HV)	X (HV)			X (HV)	X (HV)		2	1
R4							X (HV)			

Figure 4.29: A table displaying the accurate efficiency calculations after the complete development of the tool implementing targeted masking for the 2/4 and 3/4 coincidence. To the right, two additional tables highlight the dead layers in two sectors that exhibited poor efficiency.

Two sectors exhibiting multiple issues began to show low efficiencies, highlighting the need for improvements in the masking implementation method. Throughout this period, as new HV trips continued to occur and the method was still being further developed, manual adjustments to the masking were periodically made to ensure that high efficiencies were maintained across all sectors. This manual intervention continued until the new improvements were fully integrated.

4.7 Improvements

The overall goal was to reduce dependence solely on pad multiplicity for unmasking decisions. A more strategic approach was needed to prioritize masking based on the specific conditions of the layers. For instance, when dead pads are present in multiple layers, the method should aim to mask pads from the layer where the entire layer is inactive, rather than a spurious dead channel in another layer. Similarly, if multiple layers are dead, the method should mask all pads from a single layer rather than randomly masking a few pads from each affected layer.

The newly developed method followed a specific structure. First, a four-dimensional parameter was created to systematically categorize the pads. The first dimension represents the sector type (either large or small). The second dimension corresponds to the R regions (R1, R2, R3, or R4). The third dimension identifies the layer (L0, L1, L2, L3, L4, L5, L6, or L7). Finally, the fourth dimension stores the pad details, including the lowest pad number, the highest pad number, and their difference, which provides the total number of pads for each sector type, R region, and layer. Alongside this, another set of parameters tracks the number of dead pads for each sector type, R region, and layer, and identifies fully inactive layers. This determination is made by comparing the total number of pads in the first parameter with the number of dead pads recorded in the second parameter.

Based on these parameters, the method identifies the best masking strategy. If a detector layer is completely inactive, it is prioritized for masking. If more than one fully dead layer is found, the method selects the layer with the maximum number of pads to be masked. Otherwise, if no fully dead layers are found, targeted masking is applied to individual dead pads. A special flag is used to control the masking process - marking entire layers for masking if needed or targeting specific dead pads. During the masking process, instead of masking all dead pads indiscriminately, the masking is done only for the pads that the flag indicates. Finally, after the masking is applied, a final check is performed using the same unmasking process described previously, utilizing the three separate unmasking functions. This step ensures that any wedges with more than one masked pad are adjusted, maintaining the required condition of only one masking per wedge. This final step serves as both a safeguard against cases where, for instance, there are seven real pads and an eighth pad is already masked, which could lead to multiple masked pads in a single wedge, and as a general double-check of the overall masking process.

This masking strategy was initially tested and implemented in sectors A01, A02, A03, and A04. After successful testing, the efficiency of these sectors was improved, validating the approach. Following this, the next step was to gradually implement the masking method across the other sectors to achieve similar enhancements throughout the entire system. The first phase focused on implementing this masking strategy for the remaining sectors on side A, applied during run R481510. In general, no drastic changes in efficiency were expected, however there was some improvement in some areas. The main differences stemmed from a few adjustments: 1) a correct separation between R1 and R2, which allowed for more accurate masking in these regions, 2) updated masking based on new HV trips occurred, and 3) a more systematic layer selection process in scenarios with multiple dead layers. Previously, when performed manually, the choice of which layer to mask was often arbitrary, whereas the new code employed a more structured approach to identify the optimal layer for masking, potentially resulting in some layers being masked

differently compared to previous manual methods. The efficiency results, reflecting these improvements, are shown in the heatmaps below.




Figure 4.30: Heatmap plots of the pad trigger efficiency for a run before the implementation of the masking on the entire side A (left) and after the implementation (right). The right side of the heatmaps corresponds to side A, where the new masking strategy was implemented. Sector A10 is excluded from consideration due to its use of a different coincidence configuration.

In the next run, R481553, high rate issues emerged. The root cause of these problems was not clear, as the increased rates coincided with both the implementation of the new masking changes and a simultaneous adjustment in ATLAS's μ (average number of interactions per bunch crossing). Both factors could potentially contribute to the issue, but what was particularly strange was that these changes had already been implemented during R481510, whereas the high rate problems only appeared in R481553. Sector A13, in particular, exhibited the highest rate among all the sectors.

A13 rate from grafana




Figure 4.31: The pad trigger rates from grafana, illustrating the issues with high rates observed in the sectors.

In order to reduce the observed rates, a test was conducted where the transition regions were unmasked. This means that no masking was applied to the wedge containing real layers for the sectors A01, A03, A05, A07, A11, and A13 during run R481865.

The results of this test are presented below.




Figure 4.32: Heatmap plots of the pad trigger efficiency for the run where the high rate problems began (left) and for the run after the unmasking of the large sectors A01, A03, A05, A07, A11, and A13 (right). After this unmasking, some problematic regions are evident.




Figure 4.33: The pad trigger rates from grafana, illustrating the reduction of rates after the unmasking of the large sectors A01, A03, A05, A07, A11, and A13. This reduction is analytically summarized in the table on the right.

The unmasking of the transition regions led to a reduction in rates for large sectors, ranging between 5% and 14%, with A11 experiencing the smallest reduction. Despite this overall improvement, the trigger occupancy plots revealed that while many high-occupancy towers disappeared after unmasking, two specific towers showed no change in occupancy. At the time of writing this thesis, the code required a minor fix related to this. Specifically, when this coincidence was applied, adjustments were necessary for the transition regions, which included the introduction of a new status, "0", indicating permanently dead pads, as it is already mentioned. In this setup, two out of the four virtual layers are masked to 1, while two layers are masked to 0. This configuration is crucial because, in this way, to

meet the coincidence requirement, three real pads must be activated to register a hit.




Figure 4.34: Trigger occupancy plots illustrating the effects of unmasking transitions for sectors A01, A03, A05, A07, A11, and A13. The right plot highlights the high-occupancy towers that remained unchanged, while the left plot indicates the towers that showed improvement after the unmasking.

(Plots by Estel Perez Codina)

The next step involved extending the masking strategy to side C. During run R482028, the masking was applied across the entire side C, without masking the transitions. However, similar to previous implementations, there were instances where the unmasking was not applied appropriately. Consequently, there remains a need for the same fix that was mentioned earlier to address these inconsistencies.

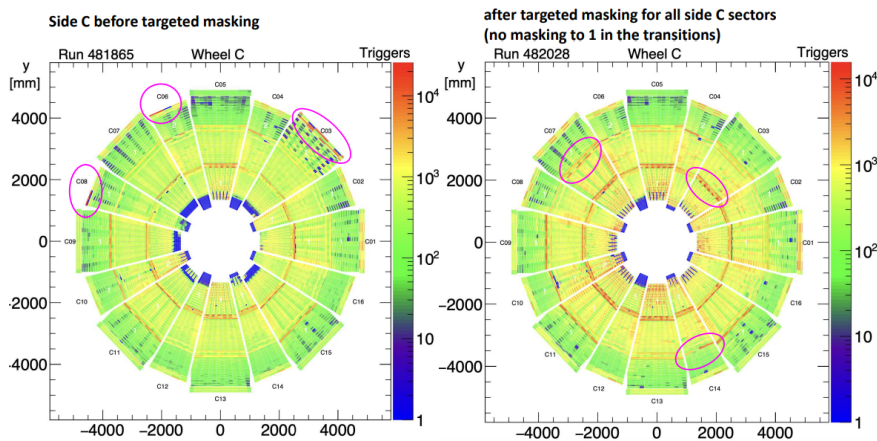


Figure 4.35: Trigger occupancy plots illustrating the effects of unmasking transitions.

The rate was reduced in some transitions, but a few transitions require fixes.

(Plots by Estel Perez Codina)

Again, alongside the implementation of the new masking strategy, ATLAS adjusted its mu. As a result, the rates exhibited varied behaviors across the sectors. Notably, sector

C13 recorded the highest rate during this period.





Figure 4.36: Heatmap plots of the pad trigger efficiency for the run before the changes implemented in the masking for side C (left) and after (right). The left side of the heatmaps corresponds to side C.



	481865	482028	increase ratio	
C01	4.84	5.79	0.20	
C03	5.29	4.96	-0.06 *	C03 rate decreases!
C05	5.74	5.92	0.03	
C07	3.17	6.01	0.90 *	C07 rate doubles!
C09	3.81	4.63	0.22 *	
C11	3.85	6.24	0.62	
C13	5.92	6.96	0.18	
C15	3.91	4.76	0.22 *	* = recovered R1 efficiency

Figure 4.37: The pad trigger rates from grafana, illustrating the changes in rates after the implementation of the masking in side C, without masking the transitions. These changes are analytically summarized in the table on the bottom.

Observing the cases of A13 and C13 revealed a common pattern in their configuration. Both sectors had one dead R1 layer in each wedge, and initially, both layers were masked. As a result, a test was conducted to evaluate the impact of modifying this setup on the trigger rate. This test was performed during run R482084, where only one of the dead R1 layers was masked, leaving the other unmasked. This adjustment aimed to see if the rates

would be further reduced. The results of this modification are presented and analyzed below.

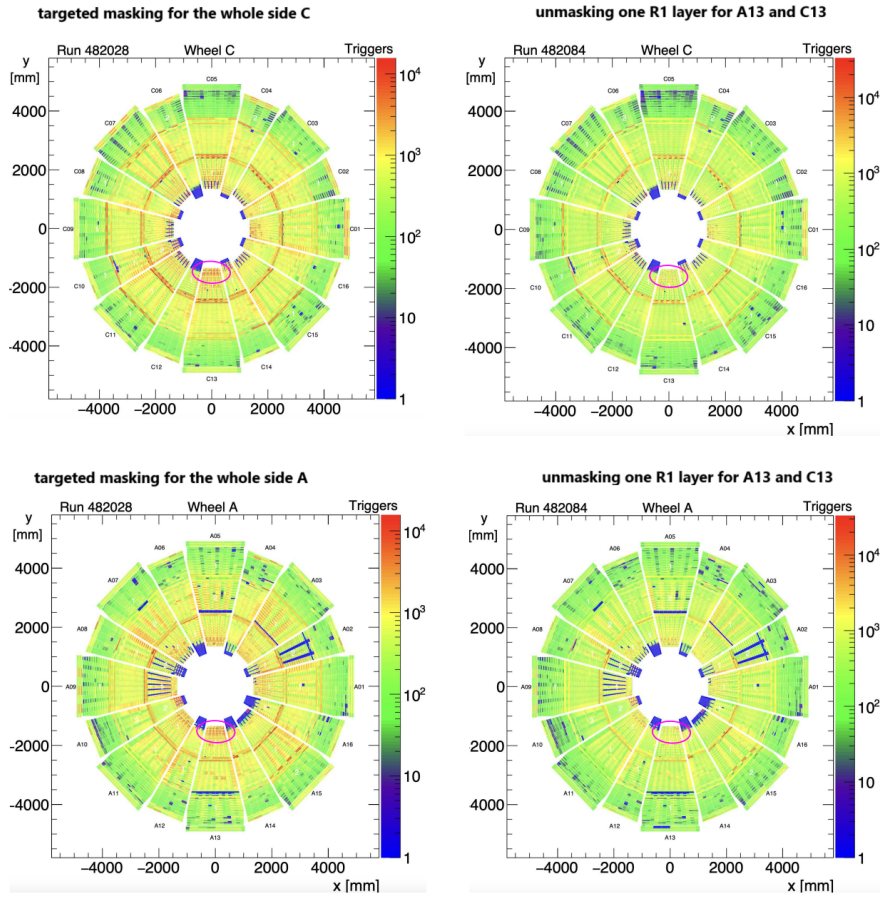


Figure 4.38: Trigger occupancy plots illustrating the effects of unmasking one R1 layer for both sectors A13 and C13. The rates were indeed reduced as it was expected.

(Plots by Estel Perez Codina)




Figure 4.39: Heatmap plots of the pad trigger efficiency before (left) and after (right) the unmasking of one out of the two dead R1 layers in A13 and C13. There is no visible impact on the efficiency.




Figure 4.40: The pad trigger rates from grafana, illustrating the changes in rates after the unmasking of one out of the two dead R1 layers in A13 and C13.

As a result of this change, the trigger rate was reduced by approximately 13% to 17%, without any noticeable loss in efficiency.

Chapter 5

Summary and conclusions

To summarize, the targeted masking strategy proved to be a highly effective tool, with a significant impact on the NSW pad trigger system of the ATLAS experiment. It successfully enhanced the efficiency of all sectors, ensuring that each was adequately included in the coincidence from the SL team. Another important outcome was the substantial reduction of fake triggers. The trigger rejection using the 2/4 and 3/4 coincidence logic in combination with the targeted masking strategy is illustrated in the scheme below.




Figure 5.1: The rejection of fake triggers without using the coincidence, i.e., without Tile and NSW (blue), the rejection with the coincidence used last year (yellow), and the rejection achieved using this year's coincidence with targeted masking (red).

Additionally, a plot showing the reduction in fake rates is presented below. The implementation of targeted masking led to a significant decrease in the rate, and with the complementary use of the MM and sTGC detectors, the rate was further reduced to 12 kHz.




Figure 5.2: On the left, the plot shows the initial fakes. A noticeable reduction occurs after implementing targeted masking, followed by an even lower level due to the full implementation of this masking strategy. After May 28th, the fake rates are further reduced to 12 kHz by the integration of the MM into the NSW-BW coincidence for all sectors.

The 2/4 and 3/4 coincidence tolerates up to one HV trip per quad while maintaining high efficiency. As a result, a JSON update is required every few days on a regular basis. On the other hand, the 3/4 and 3/4 coincidence is more vulnerable to HV trips; a single HV trip in a quad can cause an efficiency drop down to approximately 70%. Additionally, given the current pad efficiency, the pad trigger efficiency cannot exceed 92%. Despite this, the 3/4 and 3/4 coincidence is better for the strips, as the 2/4 and 3/4 coincidence introduces an ambiguity in the BandID, resulting in a 50% efficiency for the strips, which is suboptimal for their performance. Consequently, adopting the 3/4 and 3/4 coincidence configuration, along with targeted masking, is recommended for future operations. Complementary use of the MM and sTGC detectors can help maintain high trigger efficiency. In the long term, an automated tool for dynamically updating masking registers in real-time should be considered to optimize the overall performance and maintain consistent efficiency.

In general, given the dynamic nature of the operating conditions, it is essential to continuously adapt and implement improvements that consider every new situation that emerges, as has been demonstrated throughout this work. Finding the right balance between achieving higher efficiencies and maintaining lower rates through targeted masking is crucial. This approach makes sure that while maximizing detector performance, the overall stability and reliability of the system, are preserved. Thus, an ongoing effort is required to fine-tune the masking strategy in response to changing conditions, optimizing for both efficiency and rate suppression.

For the immediate future, a key priority for improving the targeted masking applied in the 2/4 and 3/4 coincidence would be to first resolve the issue with the unmasking of the

transitions and ensure that it is consistently applied across all sectors. This would help eliminate any remaining discrepancies and establish a uniform masking strategy throughout the system, enhancing its overall stability and performance. What is more, another approach that could be explored involves the patterns with the seven real layers. Specifically, the one virtual layer could be masked to 0 instead of masking it to 1. This adjustment may help mitigate the occurrence of unphysical rates that have been observed.




Figure 5.3: An illustration depicting the scenario where the virtual layers in patterns with seven real layers are masked to 1. This configuration leads to the occurrence of unphysical rates in the trigger system.

Finally, an important next step is to extend the targeted masking strategy to address the noisy channels as well. This will involve developing and optimizing the tool to effectively mask these noisy channels to 0, ensuring that they do not interfere with the overall performance of the trigger system. By implementing this enhancement, we aim to improve data quality and reliability further, ultimately contributing to more accurate measurements within the experiment.

Bibliography

- [1] D. Dominguez, "Particles of the Standard Model of particle physics", Available: <https://home.cern/science/physics/standard-model>
- [2] S. Weinberg, "The Making of the Standard Model", Eur. Phys. J. C, vol. 34, pp. 5-13, 2004.
- [3] M. Thomson, "Modern Particle Physics", Cambridge Univ. Press, 2013.
- [4] W. R. Leo, "Techniques for Nuclear and Particle Physics Experiments", Springer Berlin, Heidelberg, 1987.
- [5] M. Gell-Mann, "A Schematic Model of Baryons and Mesons", Physics Letters, vol. 8, issue 3, pp. 214-215, 1964.
- [6] G. Zweig, "An SU(3) Model for Strong Interaction Symmetry and its Breaking", Version 1, CERN Document Server, CERN-TH-401, 1964.
- [7] G. Zweig, "An SU(3) Model for Strong Interaction Symmetry and its Breaking", Version 2, CERN Document Server, CERN-TH-412, 1964.
- [8] R. L. Workman et al. (Particle Data Group), "Review of Particle Physics: Quark Model", Prog. Theor. Exp. Phys. 2022, 083C01, 2022. Available: <https://pdg.lbl.gov/2022/reviews/rpp2022-rev-quark-model.pdf>
- [9] T. M. Yan, S. D. Drell, "The parton model and its applications", International Journal of Modern Physics A, vol. 29, no. 30, 1430071, 2014.
- [10] H. X. Chen, W. Chen, X. Liu, Y. R. Liu, S. L. Zhu, "An Updated Review of the New Hadron States", Rept. Prog. Phys. 86 (2023) no.2, 026201, 2023.
- [11] F. Bechtel, "The underlying event in proton-proton collisions", CERN Document Server, CERN-THESIS-2009-215, 2009. Available: https://cds.cern.ch/record/1446555/files/TS2009_008.pdf
- [12] S. Hoche, "Introduction to parton-shower event generators", p. 40, lectures presented at TASI 2014, 2015. Available: <https://arxiv.org/pdf/1411.4085>
- [13] C. Kitsaki, "Testing the operation and response with cosmics of the Micromegas detector for the NSW upgrade of the ATLAS experiment and study of the Higgs boson production in the WW* -> l ν l ν decay channel through VBF mechanism", CERN Document Server, CERN-THESIS-2023-187, 2023. Available: <https://cds.cern.ch/record/2874937>
- [14] L. Evans, P. Bryant, "LHC Machine", Journal of Instrumentation, vol. 3, no. 8, S08 001, 2008.
- [15] "The CERN accelerator complex", Available: <http://cds.cern.ch/record/2197559>
- [16] The ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider", Journal of Instrumentation, vol. 3, no. 8, S08 003, 2008.

- [17] The CMS Collaboration, "The CMS experiment at the CERN LHC", *Jinst*, vol. 3, S08 004, 2008.
- [18] The ALICE Collaboration, "The ALICE experiment at the CERN LHC", *Journal of Instrumentation*, vol. 3, no. 8, S08 002, 2008.
- [19] The LHCb Collaboration, "The LHCb Detector at the LHC", *Journal of Instrumentation*, vol. 3, no. 8, S08 005, 2008.
- [20] M. S. Camillocci, "LHC nominal cycle", CERN Document Server, CERN, 2016. - 4 p., 2015. Available: https://indico.cern.ch/event/434129/contributions/1917195/attachments/1205096/1765722/Nominal_cycle.pdf
- [21] D. Prelicpean, "Comparison between measured radiation levels and FLUKA simulations at CHARM and in the LHC tunnel of P1-5 within the R2E project in Run 2", Technical University of Munich, 2021-07-26. - 131 p., CERN-THESIS-2021-101, 2021.
- [22] W. Herr, B. Muratori, "Concept of luminosity", CERN Accelerator School: Intermediate Course on Accelerator Physics, Zeuthen, Germany, pp. 361-378, 2003.
- [23] "High Luminosity LHC Project". Available: <https://hilumilhc.web.cern.ch/content/hl-lhc-project>
- [24] "How ATLAS detects particles: diagram of particle paths in the detector", CERN Document Server, 2013. Available: <https://cds.cern.ch/record/1505342>
- [25] T. Lenzi, "Development and Study of Different Muon Track Reconstruction Algorithms for the Level-1 Trigger for the CMS Muon Upgrade with GEM Detectors", CERN Document Server, CERN-THESIS-2013-042, 2013.
- [26] "ATLAS Detector Magnet System", CERN Document Server, ATLAS-PHOTO-2021-029-1, 2021. Available: <https://cds.cern.ch/record/2770604>
- [27] J. J. Goodson, "Search for Supersymmetry in States with Large Missing Transverse Momentum and Three Leptons including a Z-Boson", CERN Document Server, CERN-THESIS-2012-053, 2012.
- [28] The ATLAS Collaboration, "ATLAS central solenoid : Technical Design Report", CERN Document Server, CERN-LHCC-97-021, ATLAS-TDR-9, 1997.
- [29] "Magnet System". Available: <https://atlas.cern/Discover/Detector/Magnet-System>
- [30] The ATLAS Collaboration, "ATLAS barrel toroid : Technical Design Report", CERN Document Server, CERN-LHCC-97-019, ATLAS-TDR-7, 1997.
- [31] The ATLAS Collaboration, "ATLAS end-cap toroids : Technical Design Report", CERN Document Server, CERN-LHCC-97-020, ATLAS-TDR-8, 1997.
- [32] The ATLAS Collaboration, "ATLAS inner detector : Technical Design Report, 1", CERN Document Server, CERN-LHCC-97-016, ATLAS-TDR-4, 1997.
- [33] "The Inner Detector". Available: <https://atlas.cern/Discover/Detector/Inner-Detector>

-
- [34] The ATLAS Collaboration, "ATLAS calorimeter performance : Technical Design Report", CERN Document Server, CERN-LHCC-96-040, ATLAS-TDR-1, 1996.
- [35] "Calorimeter". Available: <https://atlas.cern/Discover/Detector/Calorimeter>
- [36] P. Krieger, "The ATLAS Liquid Argon Calorimeter : construction, integration, commissioning and performance from selected particle beam test results", CERN Document Server, 52nd IEEE Nuclear Science Symposium and Medical Imaging Conference, 2006. Available: https://www.physics.utoronto.ca/~krieger/procs/Krieger_NSS05_Proc.pdf
- [37] The ATLAS Collaboration, "ATLAS muon spectrometer : Technical Design Report", CERN Document Server, CERN-LHCC-97-022, ATLAS-TDR-10, 1997.
- [38] "Computer generated image of the ATLAS Muons subsystem", CERN Document Server, CERN-GE-0803017-01, 2008.
- [39] "Muon Spectrometer". Available: <https://atlas.cern/Discover/Detector/Muon-Spectrometer>
- [40] The ATLAS Collaboration, "Operation of the ATLAS Trigger System in Run 2", JINST, vol. 15, no. 10, P10004, 2020.
- [41] T. Kamawoto, S. Vlachos, L. Pontecorvo et al, "New Small Wheel Technical Design Report", CERN Document Server, CERN-LHCC-2013-006, ATLAS-TDR-020, 2013.
- [42] The ATLAS Collaboration, "The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3", CERN Document Server, CERN-EP-2022-259, 2023.
- [43] J. C. Rivera Vergara, "Cosmic ray tests for the QS1 module of th New Small Wheel in the ATLAS experiment and prospects on the search for heavy vector triplet bosons in the leptonic decay channels with the ATLAS experiment at the HL-LHC", Pontifical Catholic University of Chile, 2018. Available: <https://repositorio.uc.cl/handle/11534/22188>
- [44] P. Scholer, "Test Beam Results and Performance Studies of ATLAS Micromegas Production Modules", J. Phys.: Conf. Ser. 1498 012053, 2020.
- [45] A. Canesse, "Small-Strip Thin Gap Chambers for the Muon Spectrometer Upgrade of the ATLAS Experiment", PoS (LHCP2020) 245, 2020.
- [46] P. Miao, F. Li, L. Guan et al, "The development of the Front-End Boards for the small-strip Thin Gap Chambers detector system of the ATLAS Muon New Small Wheel upgrade", JINST, vol. 15, P11024, 2020.
- [47] G. Iakovidis, L. Levinson, Y. Afik et al, "The New Small Wheel electronics", JINST, vol. 18, P05012, 2023.