Alma Mater Studiorum · Università di Bologna

Dipartimento di Fisica e Astronomia "Augusto Righi"
Corso di Laurea in Fisica

RPC Timing Calibration and Sensitivity Study for Heavy Stable Charged Particles in ATLAS

Relatore:
Prof. Iacopo Vivarelli

Correlatore:
Dott. Yuchen Cai

Presentata da:
Bianca Neretti

Anno Accademico 2025/2026

Abstract

The Standard Model (SM) of particle physics provides a remarkably successful description of elementary particles and their interactions, but several experimental and theoretical observations motivate the search for physics Beyond the Standard Model (BSM). Among the possible signatures of BSM models are Heavy Stable Charged Particles (HSCPs), such as long-lived staus in supersymmetric scenarios. Due to their large masses, they are expected to propagate through the detector with velocities significantly lower than those of relativistic SM particles. Their identification therefore relies critically on precise measurements of their time of flight (ToF).

This thesis investigates the timing performance of the Resistive Plate Chambers (RPCs) in the ATLAS Muon Spectrometer (MS). The study focuses on the calibration of the RPC timing response and evaluates how its precision affects the measurement of the ToF and the subsequent reconstruction of the particle velocity, β. This impact is then assessed through a sensitivity study targeting supersymmetric staus with masses between 300 and 800 GeV. The results demonstrate that precise RPC timing calibration is essential for reliable velocity reconstruction, representing a crucial tool for enhancing the sensitivity of ATLAS searches for long-lived charged particles.

Contents

List of Figures

List of Tables

Chapter 1

The Standard Model and Beyond the Standard Model

The Standard Model (SM) represents the cornerstone of modern particle physics, providing a comprehensive description of the elementary constituents of matter and three of the four known fundamental forces through the framework of quantum field theories. Despite being one of the most rigorously tested and validated theories in physics, the Standard Model is affected by well-established theoretical and experimental limitations. Consequently, various extensions and superseding theories have been proposed over the past several decades.

This chapter provides a brief overview of the SM particle content and its underlying gauge interactions; the physical mechanisms leading to long lifetimes will also be briefly addressed. Furthermore, the intrinsic limitations of the SM will be discussed, motivating the necessity for Beyond the Standard Model (BSM) physics.

Throughout this work, the standard natural unit system of particle physics, $c = \hbar = 1$, is adopted.

1.1 The Standard Model

The SM encompasses all known fundamental constituents of matter and describes their interactions through three of the four fundamental forces: the electromagnetic, the weak and the strong force. A consistent description of gravity within the framework of quantum field theory remains elusive: consequently, gravitational interactions are excluded from the SM. While gravity is the dominant force on macroscopic scales, and so essential in astrophysics and cosmology, its effects at the subatomic scale and typical particle physics energies are extraordinarily suppressed compared to the other three interactions. Therefore, the omission of gravity does not undermine the remarkable predictive power of the Standard Model.

Within this framework, both the fundamental building blocks of matter and the force-mediating gauge bosons are treated as point-like objects with no internal structure, classifying them as elementary particles. Comprehensive introductions to the Standard Model can be found in Refs. [1, 2], which served as the primary basis for the discussion in this chapter.

1.1.1 Elementary particles

According to their spin, elementary particles can be split into two major groups, as in Figure 1.1: **fermions** have half-integer spin and are the building blocks of matter; **bosons** have integer spin, act as the mediators of the forces and are responsible for the particles obtaining their masses.

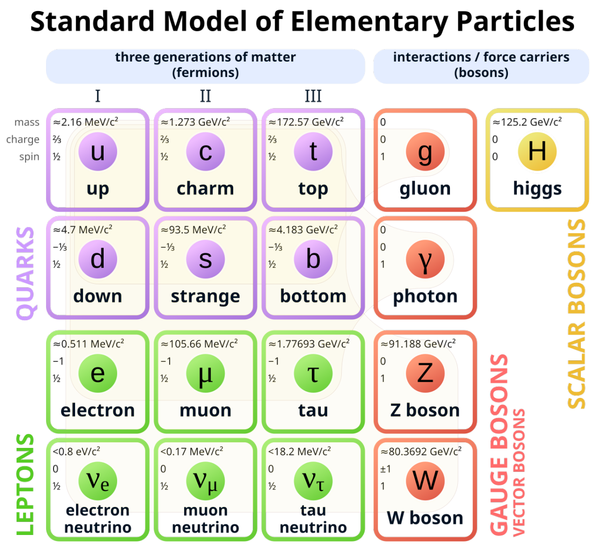

Figure 1.1: Schematic representation of the elementary particles of the SM. Figure taken from Ref. [3].

Fermions

Fermions are classified into three generations, where particles in successive generations share identical quantum numbers but exhibit progressively higher masses. Furthermore, fermions are classified according to their sensitivity to the strong interaction. Leptons do not carry colour charge and are therefore immune to the strong force, whereas quarks possess colour charge (denoted as red (r), green (g), or blue (b)) and actively participate in strong interactions.

The lepton sector is made of six distinct flavors, evenly split into two groups of three. The charged leptons carry an electric charge of $-e$ (where e[1] represents the elementary charge) and include the electron (e^-), the muon (μ^-), and the tauon (τ^-). Conversely, the neutral leptons, or neutrinos, carry no electric charge and consist of the electron neutrino (ν_e), muon neutrino (ν_μ), and tau neutrino (ν_τ). Each generation groups these particles into left-handed doublets consisting of a neutrino and its corresponding charged lepton: (ν_e, e^-) for the first generation, (ν_μ, μ^-) for the second, and (ν_τ, τ^-) for the third. This grouping stems from chirality, an intrinsic property of fermions characterized by two possible eigenstates: left-handed and right-handed. This distinction is physically paramount in the Standard Model: the weak charged-current interaction, mediated by the $W^\pm$ bosons, couples exclusively to these left-handed doublets, while completely bypassing right-handed fields (which transform as singlets). This maximal asymmetry

[1] The electric charge of a proton $e = 1.602176634 \cdot 10^{-19}$ C.

fundamentally breaks parity symmetry. A comprehensive summary of lepton properties is provided in Table 1.1.

	Particle	**Symbol**	**El. charge [e]**	**Mass [MeV]**
1st generation	electron	e^-	-1	0.511
	electron neutrino	ν_e	0	$< 27 \cdot 10^{-6}$
2nd generation	muon	μ^-	-1	106
	muon neutrino	ν_μ	0	< 0.19
3rd generation	tauon	τ^-	-1	1776.93 ± 0.09
	tauon neutrino	ν_τ	0	< 18.2

Table 1.1: Leptons of the Standard Model, ordered by generation. Masses taken from Ref. [4].

Similarly, quarks are arranged in three generations of doublets, each containing an up-type and a down-type quark. Up-type quarks, namely the up (u), charm (c), and top (t) flavors, carry a fractional electric charge of $+\frac{2}{3}e$. Down-type quarks, namely the down (d), strange (s), and bottom (b) flavors, possess an electric charge of $-\frac{1}{3}e$. These flavors are paired into generational doublets: (u, d), (c, s), and (t, b) for the first, second, and third generations, respectively. The quarks and their properties are outlined in Table 1.2.

	Particle	**Symbol**	**El. charge [e]**	**Mass [MeV]**
1st generation	up	u	$+2/3$	2.16 ± 0.04
	down	d	$-1/3$	4.70 ± 0.04
2nd generation	charm	c	$+2/3$	$(1.2729 \pm 0.0027) \cdot 10^3$
	strange	s	$-1/3$	92.9 ± 0.4
3rd generation	top	t	$+2/3$	$(172.60 \pm 0.27) \cdot 10^3$
	bottom	b	$-1/3$	$(4.1859 \pm 0.0034) \cdot 10^3$

Table 1.2: Quarks of the Standard Model, ordered by generation. Masses taken from Ref. [4].

Additionally, the Standard Model predicts an antiparticle for every fundamental fermion. Antiparticles share the same mass and spin as their matter counterparts but exhibit opposite internal quantum numbers, such as electric and colour charges. Antiparticles for quarks $(\bar{u}, \bar{d}, \bar{c}, \bar{s}, \bar{t}, \bar{b})$ and neutrinos $(\bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau)$ are denoted by a bar over the particle symbol, with antiquarks carrying anti-colours $(\bar{r}, \bar{g}, \bar{b})$. For charged leptons, however, the antiparticle is conventionally represented by explicitly denoting the positive charge (e^+, μ^+, τ^+).

Bosons

The SM incorporates two distinct categories of bosons: vector gauge bosons with spin 1, which mediate the fundamental interactions, and a single scalar boson with spin

0, the Higgs boson.

The strong interaction is described by the unbroken $SU(3)_c$ gauge symmetry and is mediated by gluons (g). Gluons interact with all colour-charged particles. Although electrically neutral, gluons carry colour charge[2], allowing them to interact with one another. The eight generators of $SU(3)_c$ correspond to eight independent gluon states, all of which are strictly massless.

In contrast, the electroweak interaction is described by the $SU(2)_L \times U(1)_Y$ gauge symmetry, which is spontaneously broken to the electromagnetic $U(1)_{\text{em}}$ symmetry through the Higgs mechanism. As a consequence, three of the four electroweak gauge bosons acquire mass, giving rise to the charged W^+ and W^- bosons and the neutral Z^0 boson. The remaining gauge boson is the photon (γ), associated with the unbroken $U(1)_{\text{em}}$ symmetry, and therefore remains strictly massless.

The Higgs boson H^0 is the scalar quantum excitation of the Higgs field [5, 6]. It is electrically neutral and colourless. Its non-zero vacuum expectation value gives rise to the masses of the W and Z bosons, while fermion masses arise through Yukawa couplings to the Higgs field. A comprehensive overview of the Standard Model bosons is presented in Table 1.3.

Interaction	**Particle**	**Symbol**	**Mass [GeV]**	**Spin**	**El. charge [e]**	**Colour**
electromagnetic	photon	γ	0	1	0	no
weak	W-boson	$W^\pm$	80.3625 ± 0.0077	1	± 1	no
	Z-boson	Z^0	91.1879 ± 0.0020	1	0	no
strong	gluon	g	0	1	0	yes
	Higgs	H^0	125.13 ± 0.11	0	0	no

Table 1.3: Bosons of the Standard Model. Masses taken from Ref. [4].

1.1.2 Interactions

The mathematical framework underlying the Standard Model is called Quantum Field Theory (QFT). Within QFT, elementary particles are described as localized excitations of their respective quantum fields, and their interactions are governed by the interaction terms in the Lagrangian density of these fields. Consequently, bound states, particle decays, and scattering events serve as primary experimental probes to investigate the dynamics of elementary constituents. By evaluating the system's dynamics through Feynman diagrams (graphical depictions of the underlying perturbation series) and computing the kinematics over the available phase space defined by the masses and momenta of the involved particles, one can analytically derive decay lifetimes and scattering cross-sections.

A schematic representation of such Feynman diagrams is illustrated in Figure 1.2. Here, the solid black external lines embody the visible, physical particles involved in the process: the left-hand lines represent the incoming initial-state particles, while the

[2] A gluon carries one colour and one anti-colour.

right-hand lines represent the outgoing final-state particles. The internal wavy line represents a virtual mediator that cannot be directly measured. Unlike real particles, virtual particles are not constrained by the energy-momentum relation ($E^2 = p^2c^2 + m^2c^4$); they can carry arbitrary mass and are said to be off-shell when their effective mass fluctuates away from their physical rest mass. Although any mass is kinematically permissible for a virtual mediator, the probability amplitude of the process suppresses significantly as the mass deviation from the shell increases. These Feynman diagrams feature processes composed of three-point interaction vertices, which constitute the dominant building blocks of higher-order perturbative diagrams. Each vertex introduces a coupling factor proportional to $\sqrt{\alpha_i}$, where α_i represents the dimensionless coupling constant of the corresponding interaction. The spatial range of a fundamental interaction is intrinsically bounded by the mass of its gauge mediator: a massless mediator gives rise to an interaction of infinite range, whereas a massive mediator severely restricts the interaction length scale.

(a) Schematic representation of an s-channel scattering process.

(b) Schematic representation of a t-channel scattering process.

Figure 1.2: Feynman diagrams for fundamental scattering channels.

Electromagnetic interaction

The quantum field theory developed to describe the electromagnetic interaction is called Quantum Electrodynamics (QED). Structurally, QED is an abelian gauge theory based on the $U(1)_{\mathrm{EM}}$ local gauge symmetry group. The mediator of this interaction is the photon (γ), which couples exclusively to electric charges.

Figure 1.3 illustrates the emission or absorption of a photon γ by a charged fermion ($f^{\pm}$), which constitutes the fundamental three-point interaction vertex of all QED processes. At low energy scales the strength of this electromagnetic coupling is governed by the fine-structure constant $\alpha_{\mathrm{em}} \approx 1/137$.

Weak interaction

As the name implies, the weak interaction is significantly weaker than the other fundamental forces described by the SM, and it is uniquely responsible for flavour-changing processes. Mathematically, the weak interaction is formulated as a gauge theory governed by the non-abelian $SU(2)_L$ local symmetry group. The gauge mediators of this force couple exclusively to the weak isospin. Within this framework, all left-handed fermi-

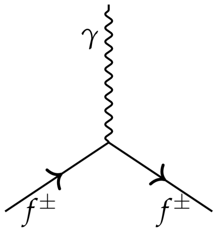

Figure 1.3: Feynman diagram showing the fundamental QED vertex.

ons and right-handed anti-fermions possess non-zero weak isospin, whereas right-handed fermions and left-handed anti-fermions carry no weak isospin charge.

The phenomena mediated by the weak force are conventionally divided into two distinct categories: charged currents and neutral currents. Charged currents are mediated by the massive $W^{\pm}$ bosons and act strictly on left-handed fermions and right-handed anti-fermions. Neutral currents are mediated by the Z^0 boson and interact with both left-handed and right-handed fermions (or anti-fermions), albeit with differing coupling strengths determined by their respective weak mixing angles.

The fundamental vertices for these interactions are illustrated in Figure 1.4.

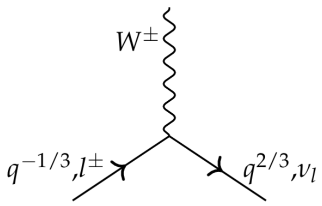

(a) Vertex of the charged current, which couples the $W^{\pm}$ bosons either to an up-type and down-type quark pair or to a charged lepton and its corresponding neutrino.

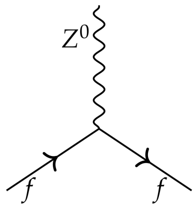

(b) Vertex of the neutral current, which couples the Z^0 boson to all known fermions.

Figure 1.4: Feynman diagram showing the two fundamental weak vertices.

In the lepton sector, the charged current vertex corresponds to the intra-generational conversion of a charged lepton into its associated neutrino, or vice versa; in the quark sector, the charged current induces transitions between up-type and down-type quarks. This cross-generational mixing arises because the weak interaction does not act upon mass eigenstates, but rather couples directly to weak eigenstates. By choosing a convenient basis, the weak eigenstates of the up-type quarks coincide with their physical mass eigenstates. In this basis, the weak eigenstates of the down-type quarks are expressed as linear combinations of the mass eigenstates. This rotation between the weak and mass bases is formalized in Equation 1.1 through the Cabibbo-Kobayashi-Maskawa (CKM) matrix. Consequently, the charged current of the weak interaction acts exclusively on left-handed doublets, leaving the right-handed singlets uncoupled, as schematically shown in Equation 1.2.

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \cdot \begin{pmatrix} d \\ s \\ b \end{pmatrix} \tag{1.1}$$

$$\begin{array}{cc} \begin{pmatrix} u \\ d' \end{pmatrix}_L \begin{pmatrix} c \\ s' \end{pmatrix}_L \begin{pmatrix} t \\ b' \end{pmatrix}_L & \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L \\ u_R\, d_R\, c_R\, s_R\, t_R\, b_R & e_R\, \mu_R\, \tau_R \end{array} \tag{1.2}$$

At low energies, the weak interaction is strongly suppressed by the large masses of the $W^\pm$ and Z^0 bosons, and is characterized by the Fermi constant G_F, related to the weak coupling g by $G_F = \frac{\sqrt{2}g^2}{8M_W^2}$. The same large mediator masses also imply a short interaction range, $\sim \mathcal{O}(10^{-18})$m.

Electroweak unification

Although the electromagnetic and weak interactions appear fundamentally distinct at low energy scales, they merge into a unified framework at energies around $\mathcal{O}(100\text{ GeV})$. Consequently, both forces can be described by a single electroweak gauge theory based on the non-abelian $SU(2)_L \otimes U(1)_Y$ local symmetry group, commonly referred to as the Glashow-Weinberg-Salam (GWS) model. Within this framework, the generator of the $U(1)_Y$ group is the weak hypercharge Y, defined via the Gell-Mann-Nishijima relation as $Y = 2(Q - I_3)$, where Q denotes the electric charge and I_3 is the third component of the weak isospin.

The gauge sector of the unbroken GWS model comprises a triplet of vector gauge fields, $W_\mu^1, W_\mu^2, W_\mu^3$, which couple exclusively to the weak isospin, and a singlet gauge field, B_μ, which couples to the weak hypercharge.

The physical, observable mediators of both the electromagnetic and weak forces emerge as linear combinations of these underlying gauge fields. This mixing is parameterized through the weak mixing angle θ_W[3] (or Weinberg angle), as formalized in Equation 1.3. This quantum-mechanical superposition explains the asymmetric behavior of chirality in the weak sector: while the Z^0 boson exhibits chiral-asymmetric couplings, the photon γ couples to left- and right-handed fields with equal strength, recovering the vector-like, parity-conserving nature of QED.

$$W^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \qquad Z^0 = W_\mu^3 \cos\theta_W - B\sin\theta_W \qquad \gamma = W_\mu^3 \sin\theta_W + B\cos\theta_W \tag{1.3}$$

Finally, to account for the large inertial masses of three out of the four physical gauge mediators ($W^\pm$ and Z^0) without violating local gauge invariance, the mechanism of Spontaneous Symmetry Breaking (SSB), commonly known as the Higgs mechanism, was introduced in 1964 and subsequently incorporated into the GWS electroweak framework.

Strong interaction

The theoretical framework describing the strong interaction is a non-abelian quantum field theory known as Quantum Chromodynamics (QCD), which is governed by the

[3] $\sin^2(\theta_W) = 0.231$ [4]

$SU(3)_C$ local gauge symmetry group. The mediators of the strong force are the gluons (g), which are strictly massless and electrically neutral vector bosons. Unlike the photons in QED, gluons couple to colour charges and carry colour charge themselves; as a consequence of this non-abelian nature, they undergo self-interactions. Within the SM, quarks and gluons are the only fields that carry colour charge, making them the exclusive participants in strong interactions.

The fundamental vertices of QCD are illustrated in Figure 1.5 with representative colour assignments.

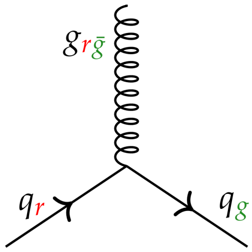

(a) Fundamental QCD vertex involving two quarks and a gluon.

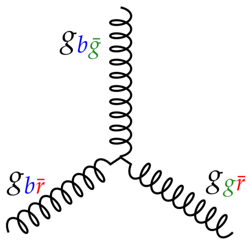

(b) Gluon three-point self-interaction vertex.

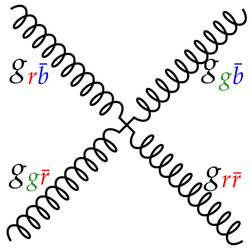

(c) Gluon four-point self-interaction vertex.

Figure 1.5: Feynman diagrams showing the three QCD vertices.

At low energy scales the strong coupling constant is $\alpha_s \sim \mathcal{O}(1)$. This implies that, unlike in the electroweak sector, higher-order perturbative diagrams do not automatically decrease in importance or contribution to the overall cross-section. However, this bare coupling is not what is directly observed in experiments. Analogous to the electric charge screening effect in a dielectric medium, quantum fluctuations introduce a running behavior for the coupling constant. In QCD, vacuum polarization is dominated by gluon loops rather than fermion loops, leading to an anti-screening effect that makes the coupling strictly distance-dependent. At short distances (or high momentum transfers), quarks experience **asymptotic freedom** and behave almost as free particles. Conversely, at long distances (or low momentum transfers), the coupling strength increases rapidly. This behavior gives rise to the phenomenon known as **colour confinement**, which dictates that colour-charged particles cannot exist as isolated states on observable macroscopic scales. Consequently, quarks and gluons are bound inside overall colour-singlet (colourless) states. Although the gluon is massless, implying an infinite theoretical range for the underlying force, confinement severely restricts its effective physical range to approximately $\mathcal{O}(10^{-15})$ m.

Therefore, quarks must organize themselves into colour-neutral composite particles called **hadrons**: this neutrality is achieved via two primary configurations[4]. The first consists of combining a colour with its corresponding anti-colour to form a colourless state; such quark-antiquark ($q\bar{q}$) bound states are classified as **mesons**. The second configuration combines three quarks into the totally antisymmetric colour-singlet representation of $SU(3)_C$, with colour wavefunction $|1\rangle = \frac{1}{\sqrt{6}}\epsilon_{ijk}|q^i q^j q^k\rangle$. These particles, com-

[4]In addition to conventional mesons ($q\bar{q}$) and baryons (qqq), exotic hadronic states containing four or five quarks have also been experimentally observed. These include tetraquarks, composed of two quarks and two antiquarks, and pentaquarks, composed of three quarks and one quark-antiquark pair [7, 8].

posed of three quarks or three antiquarks, are designated as **baryons** and anti-baryons, respectively.

Another crucial consequence of the running coupling is that the colour field between two separating colour charges behaves like a relativistic string: as the distance increases, the potential energy stored in the flux tube grows linearly. When the separation becomes large enough, it becomes energetically favorable for the field to snap, creating a new $q\bar{q}$ pair out of the vacuum, a mechanism illustrated in Figure 1.6; when colour-charged partons[5] are produced with high transverse momentum, this pair-production cascade repeats continuously until the available kinematic energy drops below the hadronization threshold. The resulting coloured partons bind into stable, colour-neutral hadrons traveling in close spatial proximity. This experimental signature is recorded in the detectors as a **jet**, as seen in Figure 1.7.

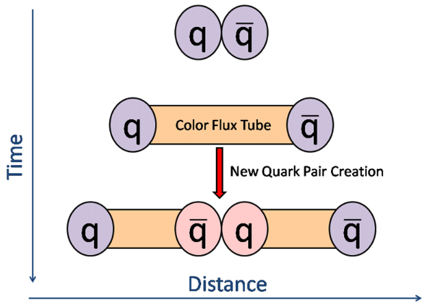

Figure 1.6: Schematic representation of the string-breaking mechanism due to colour confinement.

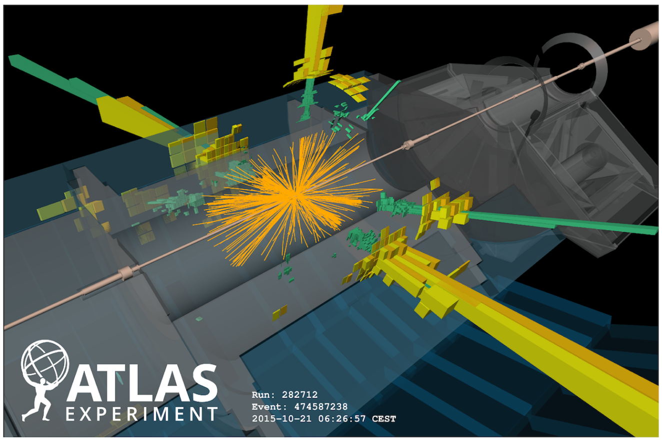

Figure 1.7: An event display from the LHC's ATLAS experiment that shows a high-energy collision producing multiple jets. The yellow lines and cones trace the trajectories of charged particles through the tracking detectors, the green and yellow blocks highlight energy deposits in the calorimeters. Figure taken from Ref. [9].

[5]In particle physics, the term parton refers to the elementary constituents of a particle, today more commonly known as quarks and gluons.

1.2 Lifetime of particles in the SM

The **mean lifetime** (τ) of a particle corresponds to the average time it exists before decaying, calculated in its proper rest frame. In the framework of QFT, the lifetime is not a discrete value but rather a statistical description of the decay probability. For a large ensemble, the exponential decay law governs the number of surviving particles $N(t)$ as a function of time t:

$$N(t) = N_0 e^{-\Gamma_{\text{tot}} t} \tag{1.4}$$

where N_0 is the initial number of particles and Γ_{tot} is the total decay rate (or total decay width). Since particles can often decay through multiple independent channels, the total decay rate is the sum of the partial widths Γ_i of all individual channels:

$$\Gamma_{\text{tot}} = \sum_{i=1}^{n} \Gamma_i \tag{1.5}$$

The mean lifetime τ is inversely proportional to this total decay rate:

$$\tau = \frac{1}{\Gamma_{\text{tot}}} \tag{1.6}$$

For each specific decay channel, the partial width Γ_i can be computed using Fermi's Golden Rule, derived from the transition amplitude (or matrix element) $\mathcal{M}$ of the process, which incorporates all the underlying dynamical properties of the interaction via Feynman calculus.

In the extreme limit where a particle is the lightest state carrying a strictly (or almost) conserved quantum number (such as the proton, which is the lightest baryon under baryon number conservation) no lighter kinematic configuration is available without violating the conservation law, rendering the particle completely stable.

In several Beyond the Standard Model (BSM) scenarios, hypothetical Long-Lived Particles (LLPs) are predicted to exhibit similar macroscopic decay lengths. Detecting and accurately reconstructing these displaced signatures is highly dependent on precise time-of-flight measurements, highlighting the critical importance of sub-nanosecond timing calibrations in systems like the ATLAS Resistive Plate Chambers (RPCs).

1.3 Limitations of the Standard Model and the need for New Physics

The SM represents one of the most successful and rigorously tested frameworks in modern science. Its theoretical predictions have been repeatedly validated by major experimental discoveries over the past decades, including the charm and top quarks, the electroweak vector bosons ($W^\pm$ and Z^0), and the Higgs boson (H^0) in 2012 by the ATLAS and CMS collaborations.

Despite these remarkable achievements, the Standard Model cannot be considered a complete theory of fundamental interactions. A wide range of astrophysical observations, cosmological data, and deep-seated theoretical inconsistencies strongly suggest that the SM is merely an effective low-energy description of a more fundamental theoretical framework. The key limitations driving the search for Physics Beyond the Standard Model (BSM) are detailed below. This section draws primarily on the sources [10, 11].

1.3.1 Experimental and cosmological shortcomings

Neutrino masses and oscillations

In its minimal formulation, the SM assumes neutrinos to be massless, with only left-handed neutrino fields included as components of the $SU(2)_L$ lepton doublets. However, the observation of neutrino oscillations, most notably through the Super-Kamiokande and SNO experiments, established that neutrino flavor eigenstates are mixtures of mass eigenstates and that at least two neutrino mass eigenstates have different, non-zero masses. These discoveries demonstrated that the minimal SM must be extended to accommodate neutrino masses and mixing, and were recognized with the 2015 Nobel Prize in Physics. Although the SM can be extended to accommodate neutrino mass terms, its minimal version fails to account for them naturally. Furthermore, while neutrinos act as hot dark matter, their low masses and cosmic abundance prevent them from accounting for the total Dark Matter relic density of the Universe.

Dark Matter

Astrophysical measurements across multiple scales provide compelling evidence for the existence of non-luminous, non-baryonic mass. Early evidence emerged in 1933 when Fritz Zwicky measured the velocity dispersion of galaxies in the Coma cluster and observed speeds far exceeding the predictions based on visible mass. Decades later, Vera Rubin confirmed this anomaly by measuring galactic rotation curves: while Newtonian dynamics predicts that the orbital velocity of outer matter should scale as $v(r) \propto 1/\sqrt{r}$ for radii $r > R$, measurements revealed flat velocity profiles ($v(r) \approx$ const). This behavior, shown in Figure 1.8, points to the presence of an extended, invisible Dark Matter halo enveloping galaxies.

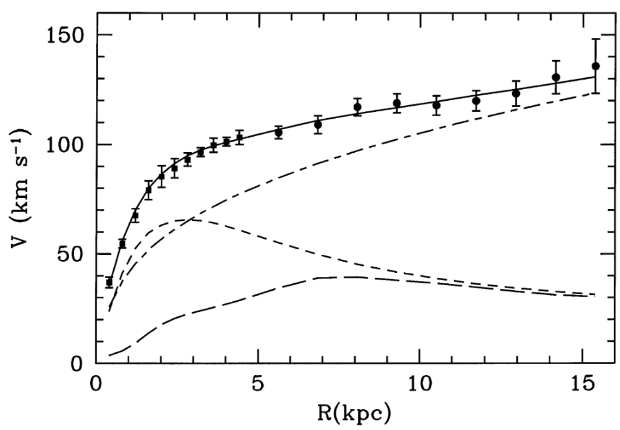

Figure 1.8: Rotation curve of the M33 galaxy. The experimental data (points) demonstrate a flat/rising velocity profile at large radii, in clear disagreement with the Keplerian decline expected from visible matter alone. The solid line represents the best-fitting model, given by the sum of the stellar disc (short-dashed line), gas (long-dashed line), and the dominant dark matter halo contribution (dot-dashed line). Figure taken from Ref. [12].

Modern cosmological surveys, most notably measurements of the Cosmic Microwave

Background (CMB) power spectrum anisotropies by the WMAP and Planck collaborations [13], determine the Universe's total mass-energy content to be:

- **Baryonic matter:** $\sim 4.9\%$
- **Dark Matter (DM):** $\sim 26.8\%$
- **Dark Energy:** $\sim 68.3\%$

To explain cosmological structure formation, Dark Matter must consist of stable, massive, electrically neutral and colour-singlet particles. Weakly Interacting Massive Particles (WIMPs) represent one of the most promising classes of candidates; however, no SM particle possesses the required properties.

Baryon asymmetry of the Universe

Standard Big Bang cosmology dictates that matter and antimatter were produced in equal amounts in the early Universe. Today's observable Universe, however, consists almost exclusively of matter. Explaining this observed asymmetry requires Charge-Parity (CP) violation processes. While CP-violating phases exist within the SM quark sector (via the CKM matrix), the resulting asymmetry is orders of magnitude smaller than what is cosmologically observed [14, 15].

The Cosmological Constant problem

Quantum field theoretical calculations of the vacuum energy density, obtained by summing zero-point energies of all SM fields up to the Planck scale, exceed the observed value of the cosmological constant Λ by over 120 orders of magnitude, constituting one of the most extreme discrepancies in theoretical physics.

1.3.2 Theoretical limitations: Naturalness and Unification

The Hierarchy/Naturalness problem

A primary theoretical motivation for BSM physics is the Hierarchy problem associated with the Higgs boson mass ($m_H \approx 125$ GeV). In QFT higher-order loop corrections modify the physical mass of a particle. For SM fermions and gauge bosons, chiral and gauge symmetries protect their masses, resulting in logarithmic quantum corrections: being logarithmic, these corrections do not significantly alter the theoretical mass value obtained from the leading order (LO) diagram. But for scalar fields like the Higgs boson no such protective symmetry exists: this results in an overestimation of the Higgs boson mass by several orders of magnitude compared to the experimentally measured value. Keeping m_H at the electroweak scale requires an extreme, fine-tuned cancellation between the bare mass parameter and quantum corrections, which is widely considered unphysical. This discrepancy between theoretical predictions and experimental results is known as the Hierarchy problem [16, 17].

GUT and Gravity

Following the evolution of the Universe from the Planck era ($t < 10^{-43}$ s) through the Grand Unified Theory (GUT) era ($t < 10^{-35}$ s) (both seen in Figure 1.9) to the electroweak scale, the fundamental gauge coupling constants $(\alpha_1, \alpha_2, \alpha_3)$ evolve with energy via the Renormalization Group Equations (RGE). Extrapolations within the SM show that while the three couplings approach each other at 10^{15}–10^{16} GeV, they fail to intersect at a single unified point, as seen in Figure 1.10. Furthermore, the SM lacks a quantum mechanical description of gravity, leaving general relativity fundamentally unintegrated at the Planck scale.

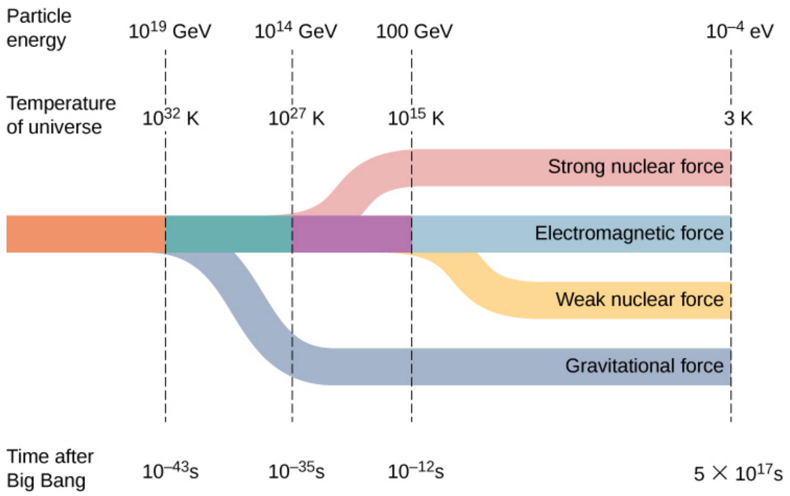

Figure 1.9: Eras of the universe's evolution from the Big Bang to the present day. Figure taken from Ref. [18].

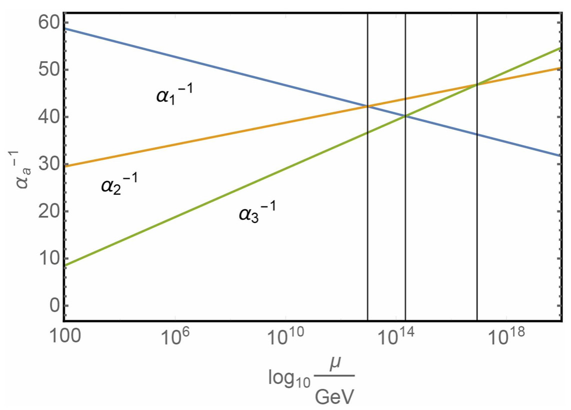

(a) SM

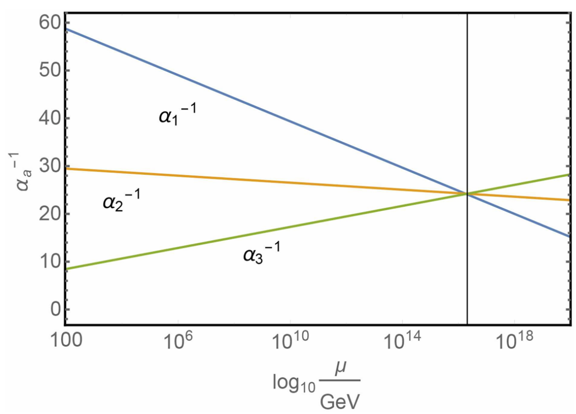

(b) MSSM

Figure 1.10: Inverse gauge couplings α_i^{-1} $(i = 1, 2, 3)$ for the electromagnetic (blue), weak (orange), and strong (green) interactions as a function of the energy scale Q. Panel (a) shows the extrapolation within the SM, where the couplings fail to unify at a single point. Panel (b) illustrates the unification achieved in the Minimal Supersymmetric Standard Model (MSSM) at $Q \sim 10^{16}$ GeV. Vertical grey lines mark the respective intersection points. Figures taken from Ref. [19].

1.3.3 Supersymmetry (SUSY) as a solution to SM limitations

Supersymmetry (SUSY) offers an elegant framework that addresses the fundamental limitations of the Standard Model [20]. Built upon Noether's theorem, which links mathematical symmetries to physical conservation laws, SUSY introduces a spacetime symmetry generated by an operator Q that transforms fermionic states into bosonic states and vice versa:

$$Q|\text{boson}\rangle = |\text{fermion}\rangle, \quad Q|\text{fermion}\rangle = |\text{boson}\rangle \tag{1.7}$$

In a supersymmetric theory, every chiral degree of freedom in the Standard Model is paired with a superpartner (visually indicated by the simbol $\tilde{}$) possessing identical quantum numbers except for its spin, which differs by a half-integer ($|\Delta S| = 1/2$). The scalar partners of SM fermions are termed *sfermions* (distinguished by the prefix *s*-). Conversely, the fermionic partners of SM gauge and Higgs bosons pick up the suffix *-ino*.

By introducing these additional scalar and fermionic degrees of freedom, SUSY provides, for example, a natural mechanism to stabilize the Higgs mass at the electroweak scale without severe fine-tuning, a unification of the electroweak and strong force [21] and dark matter candidates [22].

1.3.4 Stable Massive Particles and the stau ($\tilde{\tau}$) in SUSY

Many Beyond the Standard Model (BSM) theories naturally predict the existence of Stable Massive Particles (SMPs). A neutral SMP can provide a viable cold dark matter (DM) candidate if it is stable on cosmological timescales and sufficiently non-relativistic.[6]. In supersymmetry, imposing R-parity forbids baryon- and lepton-number-violating interactions and, as a consequence, prevents the decay of the Lightest Supersymmetric Particle (LSP), providing the stability required for it to be a potential dark matter candidate.

In R-parity-conserving SUSY models [23], a neutral[7] and weakly interacting LSP can provide a viable dark matter candidate. However, if the Next-to-Lightest Supersymmetric Particle (NLSP) is only slightly heavier than the LSP, or if its decay is strongly suppressed, it can acquire a significantly prolonged lifetime. When such a long-lived particle carries electric or colour charge, it traverses the entire detector before decaying, acting as a detectable charged SMP [24].

A prominent example of such a candidate is the scalar tau, or **stau** ($\tilde{\tau}$), the spin-0 superpartner of the Standard Model τ lepton. Present in all supersymmetric extensions containing the SM sfermion spectrum, including the Minimal Supersymmetric Standard Model (MSSM) and the Next-to-Minimal Supersymmetric Standard Model (NMSSM), the stau frequently emerges as the NLSP in various SUSY-breaking scenarios, such as Gauge-Mediated Supersymmetry Breaking (GMSB) or minimal Supergravity (mSUGRA).

[6] Otherwise, it would either decay or contribute as a hot or warm dark matter component, rather than as cold dark matter.

[7] There is no theoretical requirement in R-parity-conserving SUSY for the LSP to be electrically neutral. However, a stable charged LSP would lead to an unacceptable abundance of charged dark matter, making a neutral LSP necessary for a viable dark matter candidate.

When the stau decays into a light LSP, such as a gravitino ($\tilde{G}$), via gravitational-scale or suppressed weak couplings:

$$\tilde{\tau} \to \tau + \tilde{G} \tag{1.8}$$

it becomes long-lived on collider timescales. Staus with masses ranging from a few hundred GeV up to ~ 1 TeV thus behave as Heavy Stable Charged Particles (HSCPs) or Long-Lived Particles (LLPs) inside high-energy experiments.

1.3.5 Experimental motivation: Slow Muons and RPC Timing Calibration

Due to their large mass, long-lived staus produced in high-energy proton-proton collisions at the LHC would propagate through the detector at speeds noticeably lower than the speed of light:

$$\beta = \frac{v}{c} < 1 \tag{1.9}$$

Such particles traverse the Inner Detector (ID) and reach the outermost subdetectors, most notably the ATLAS Muon Spectrometer (MS), before decaying, leaving experimental tracks that resemble heavy, slow-moving muons [25].

Identifying these slow-moving BSM candidates requires a precise time-of-flight (ToF) measurement to cleanly separate slow particle signatures ($\beta < 1$) from ultra-relativistic SM background particles ($\beta \approx 1$) and instrumental timing jitter.

This experimental requirement provides the core motivation for the work presented in this thesis.

Chapter 2

LHC and ATLAS Detector

The European Organization for Nuclear Research (CERN)[1], founded in 1954 as one of Europe's first joint scientific ventures, is today one of the world's leading research centres in particle physics, currently supported by 23 member states [26]. Located across the Franco–Swiss border near Geneva, CERN operates an extensive complex of particle accelerators and hosts several large-scale detector experiments.

This chapter provides an overview of the experimental setup. It begins with an introduction to the main accelerator complex, the Large Hadron Collider (LHC) [27], followed by a description of the ATLAS experiment [28]. Particular emphasis is placed on the detector subsystems directly relevant to this thesis, specifically the Resistive Plate Chambers (RPCs).

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is currently the world's largest and most powerful particle accelerator. Housed in a circular underground tunnel with a circumference of approximately 26.7 km and located at an average depth of 100 m (ranging between 50 m and 175 m beneath the surface), it was originally commissioned in 1988 for the Large Electron-Positron Collider (LEP) [29]. Between 2000 and 2008, the underground infrastructure was extensively modified to accommodate the LHC. The accelerator is designed to collide hadrons (primarily protons (pp), but also heavy ions (such as Pb–Pb and p–Pb)) reaching centre-of-mass energies of up to $\sqrt{s} = 13.6$ TeV.

The strategic decision to replace the e^+e^- collider LEP with a hadron accelerator was fundamentally driven by the physics of charged particles undergoing circular acceleration. Any charged particle accelerated along a curved trajectory emits synchrotron radiation, with the energy loss per unit time governed by:

$$\frac{\mathrm{d}E}{\mathrm{d}t} \propto \frac{E^4}{m^4} \tag{2.1}$$

where E is the energy and m is the mass of the accelerated particle. Because of the small electron mass ($m_e \approx 0.511$ MeV compared to the proton mass $m_p \approx 938$ MeV), the energy loss via synchrotron radiation at LEP severely limited its maximum reachable centre-of-mass energy to $\sqrt{s} = 209$ GeV. Accelerating significantly heavier particles like

[1]Historically derived from the *Conseil Européen pour la Recherche Nucléaire.*

protons reduces radiation losses by a factor of $(m_p/m_e)^4 \approx 10^{13}$, thereby enabling the exploration of the multi-TeV energy regime within a ring of the same radius.

However, the gain in reachable energy via hadron collisions comes with distinct experimental trade-offs:

- **Lepton colliders (e^+e^-):** Electrons and positrons are elementary, point-like particles with well-defined initial kinematic states. Collisions occur in a clean background environment, making them ideal for high-precision electroweak measurements.

- **Hadron colliders (pp):** Protons are composite objects bound by strong interactions. As a result, the hard scattering process occurs between individual partons whose initial momenta are governed statistically by Parton Distribution Functions (PDFs). Hadron colliders allow probing a significantly broader spectrum of new high-energy physics phenomena, albeit at the expense of a substantially higher and more complex background driven by Quantum Chromodynamics (QCD) processes.

2.1.1 The acceleration chain and the magnet system

The LHC is a synchrotron accelerator designed to propel charged particles along circular trajectories using strong magnetic fields to bend their paths and high-frequency radio-frequency (RF) electric fields to accelerate them. Because particles cannot be accelerated directly from rest to multi-TeV energies in a single machine, proton acceleration relies on a multi-stage injection chain of pre-accelerators, as seen in Figure 2.1.

Protons are generated by ionizing hydrogen gas and undergo sequential acceleration before injection into the main LHC ring:

1. **Linac4:** A linear accelerator that accelerates protons up to an energy of 160 MeV (replacing Linac 2 since Long Shutdown 2).

2. **Proton Synchrotron Booster (PSB):** A circular accelerator that raises the beam energy to 2.0 GeV.

3. **Proton Synchrotron (PS):** Accelerates the proton bunches up to 26 GeV.

4. **Super Proton Synchrotron (SPS):** Performs the final pre-acceleration stage up to 450 GeV.

From the SPS, protons are injected into the two counter-rotating beam pipes[2] of the LHC in bunches containing approximately 1.2×10^{11} protons each, spaced by 25 ns. Final acceleration to collision energy ($\sqrt{s} = 13.6$ TeV) is accomplished within the LHC main ring by oscillating electric fields generated in superconducting RF cavities located in one of the straight insertion regions. The RF electric field is synchronized with the revolution frequency f of the particle bunches according to $f = n/T$, where T is the revolution period and n is an integer harmonic number. This synchronization ensures that proton bunches consistently experience an accelerating electric field upon each pass through the cavities.

[2] A *beam pipe* is an ultra-high vacuum metallic tube in which the particle beams travel, isolating them from atmospheric pressure to minimize collisions with residual gas molecules.

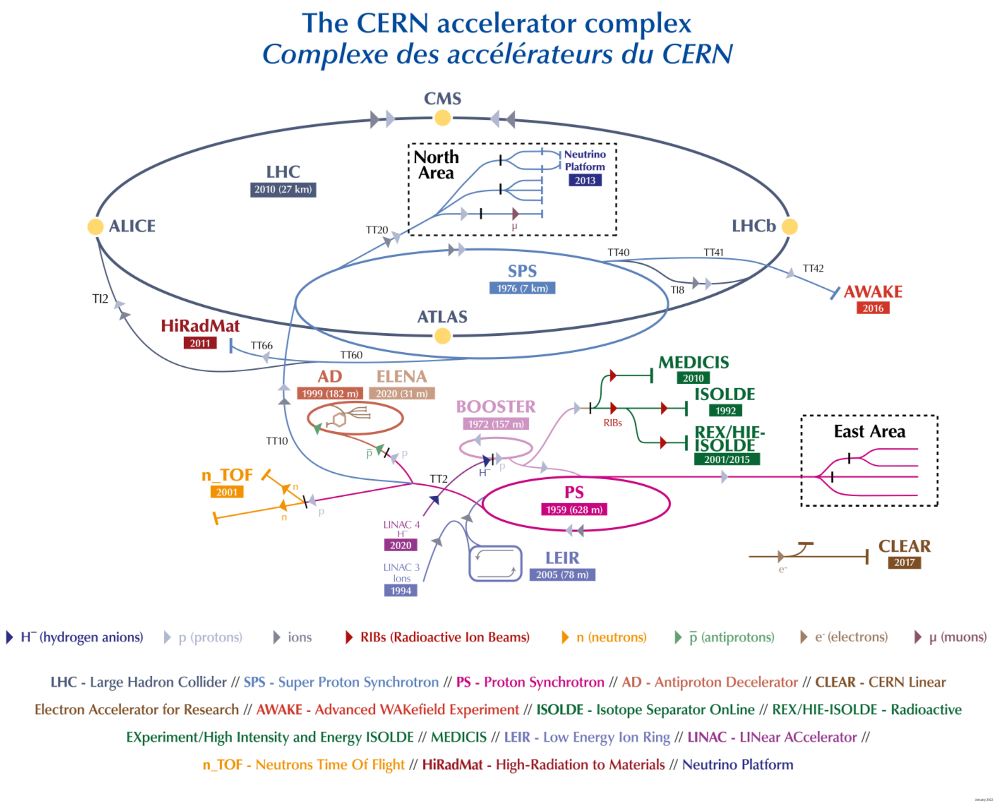

Figure 2.1: The CERN accelerator complex. Figure taken from Ref. [30].

To constrain the multi-TeV proton beams along stable circular orbits, the LHC utilizes an advanced superconducting magnetic system consisting of:

- 1232 **dipole magnets:** Each approximately 15 m long and operated at a magnetic field strength of 8.3 T, designed to bend the particle trajectories around the ring.
- 392 **quadrupole magnets:** Used to focus and collimate the beams, maintaining small beam cross-sections at the interaction regions.

To maintain superconductivity, all dipole and quadrupole electromagnets are cooled to an operating temperature of 1.9 K[3] using superfluid helium.

2.1.2 Major LHC experiments

The two counter-rotating beams cross at four main interaction points along the LHC tunnel, hosting the four primary detector experiments:

- **ATLAS** (*A Toroidal LHC ApparatuS*) [28] and **CMS** (*Compact Muon Solenoid*) [31]: general-purpose detectors sharing complementary scientific goals, including precision measurements of the Higgs boson, electroweak symmetry breaking studies, and searches for BSM physics. Although targeting similar physics programmes, ATLAS and CMS employ fundamentally distinct detector technologies and magnet configurations.

[3] $-271.3°C$

- **ALICE** (*A Large Ion Collider Experiment*) [32]: a heavy-ion specialized experiment dedicated to exploring the properties of strongly interacting matter and the Quark-Gluon Plasma (QGP).
- **LHCb** (*Large Hadron Collider beauty*[4]) [33]: A forward single-arm spectrometer specialized in heavy-flavour physics (hadrons containing b and c quarks) to investigate CP violation mechanisms and rare decays.

2.1.3 Luminosity and operational Runs

At the LHC, two proton beams[5] circulate in opposite directions. Each beam is structured into discrete packages called bunches[6], each consisting of approximately 1.15×10^{11} to 1.4×10^{11} particles, focused to transverse dimensions on the order of tens of micrometres (μm) at the interaction points.

The expected event rate $\frac{\mathrm{d}N}{\mathrm{d}t}$ for a given physical process with cross-section σ is governed by the instantaneous luminosity $\mathcal{L}$:

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \sigma \cdot \mathcal{L} \tag{2.2}$$

Under the assumption of Gaussian-shaped transverse beam profiles, $\mathcal{L}$ is defined as:

$$\mathcal{L} = \frac{n_+ n_- f}{4\pi\sigma_x\sigma_y} n_b \tag{2.3}$$

where n_+ and n_- denote the number of protons per colliding bunch in each beam, f is the revolution frequency, σ_x and σ_y characterize the transverse beam dimensions at the collision point, and n_b is the number of circulating bunches per beam.

Maximizing the instantaneous luminosity is essential for increasing the observation rate of rare physical processes. Integrating $\mathcal{L}$ over a data-taking time interval $[t_1, t_2]$ yields the total integrated luminosity L_{int}, which measures the absolute size of the recorded dataset:

$$N(t_1, t_2) = \sigma \int_{t_1}^{t_2} \mathcal{L}\,\mathrm{d}t = \sigma \cdot L_{\text{int}} \tag{2.4}$$

Integrated luminosity is expressed in units of inverse cross-section, typically inverse femtobarns (fb^{-1}), where 1 b $= 10^{-28}$ m^2 $= 10^{-24}$ cm^2.

High instantaneous luminosity causes multiple independent proton-proton collisions during a single bunch crossing[7], a phenomenon known as *pile-up*. Managing pile-up requires high spatial and timing resolutions in the subdetectors.

Table 2.1 summarizes the operational parameters and integrated luminosity collected by ATLAS during the main operational runs[8] of interest for this thesis (Run 2 and Run 3).

[4] The bottom-quark is sometimes referred to as beauty-quark.

[5] A *beam* consists of a stream of charged particles accelerated along a guided trajectory within the beam pipe.

[6] A *bunch* is a compact cluster of particles confined longitudinally by the radio-frequency fields and transversely by quadrupole magnet fields.

[7] A *bunch crossing* refers to the instant when two counter-rotating bunches overlap at an interaction

Run Period	Years	$\sqrt{s}$ [TeV]	Recorded L_{int} [fb^{-1}][9]
Run 2	2015–2018	13.0	139.0
Run 3 (2022)	2022	13.6	27.2 – 27.8
Run 3 (2023)	2023	13.6	31.0
Run 3 (2024)	2024	13.6	108.0
Run 3 (2025)	2025	13.6	115.0
Run 3 (2026)[10]	2026	13.6	28.1
Run 3	2022–2026	13.6	≈ 312.0

Table 2.1: Summary of centre-of-mass energies ($\sqrt{s}$) and recorded integrated luminosity (L_{int}) achieved during the main operational periods of the LHC. All values are certified as suitable for physics analysis according to the official Data Quality Results [35].

2.2 The ATLAS Detector

ATLAS is the largest particle detector by volume at the LHC, featuring a cylindrical geometry with a length of 44 m, a diameter of 25 m and a weight of approximately 7000 tonnes. A schematic overview of the detector and its sub-systems is shown in Figure 2.2.

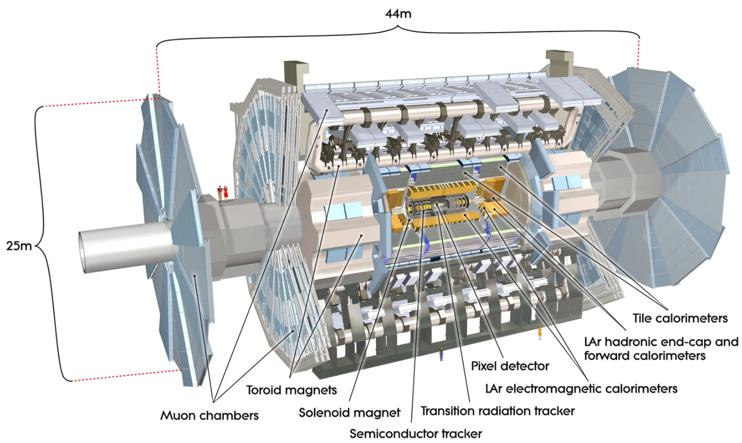

Figure 2.2: Cut-away view of the ATLAS detector. Figure taken from Ref. [28].

point, occurring every 25 ns (at LHC).

[8]A *Run* in particle accelerator terminology refers to a multi-year operational period during which the machine delivers particle collisions continuously at a target centre-of-mass energy, separated by long shutdown periods (*Long Shutdowns*, or LS) dedicated to major detector upgrades and accelerator maintenance. To date, the LHC schedule includes: *Run 1* (2009-2013, with pp collisions at $\sqrt{s}$ = 7–8 TeV), *LS1* (2013-2015), *Run 2* (2015-2018, at $\sqrt{s} = 13$ TeV), *LS2* (2019-2022), *Run 3* (2022-2026, at $\sqrt{s} = 13.6$ TeV), and *LS3* (started on June 29, 2026 [34], to prepare for the High-Luminosity LHC era, which will begin in 2030 with Run 4).

[9]All values represent the integrated luminosity recorded by ATLAS and certified as suitable for physics analysis (*Good for Physics*), meaning all subdetectors were operational and fully calibrated.

[10]The 2026 dataset corresponds to standard physics operation at nominal pile-up: nominal operation optimizes the total integrated luminosity by colliding high-density bunches, incurring higher overlapping inelastic pp interactions per bunch crossing.

2.2.1 Coordinate system and kinematic variables

ATLAS adopts a right-handed coordinate system (shown in Figure 2.3) with its origin at the nominal interaction point (IP). The x-axis points horizontally towards the centre of the LHC ring, the y-axis points vertically upwards, and the z-axis is aligned along the beam direction. In data analysis, a cylindrical coordinate system (r, ϕ, θ) is

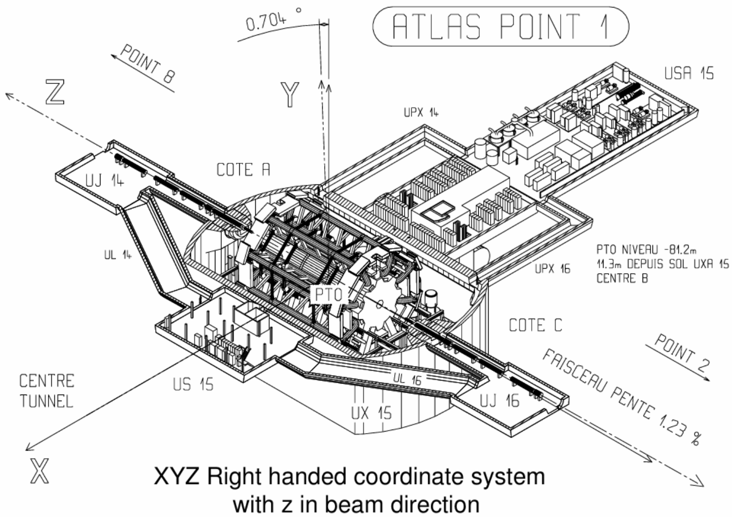

Figure 2.3: Illustration of the ATLAS Detector oriented in the global coordinate system taken from Ref. [36].

preferred, where $r = \sqrt{x^2 + y^2}$ is the radial distance, the azimuthal angle ϕ is measured in the transverse (x–y) plane relative to the positive x-axis, and the polar angle θ is measured from the z-axis. Due to the composite nature of protons, the colliding partons carry unequal and unknown fractions of the initial momentum, resulting in an unknown Lorentz boost along the z-axis for the centre-of-mass frame of each collision. To mitigate this effect, the polar angle θ is routinely replaced by the pseudorapidity η, defined as:

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

Pseudorapidity approximates the rapidity $y = \frac{1}{2}\ln\left(\frac{E+p_z}{E-p_z}\right)$ in the ultra-relativistic limit ($m \ll E$): differences in rapidity (and thus in η for light particles) are invariant under Lorentz boosts along the longitudinal beam axis. It follows that $\eta = 0$ corresponds to the plane transverse to the beam direction, while $\eta \to +\infty$ corresponds to the direction parallel to the beam. The angular distance between two particles is consequently defined in the (η, ϕ) plane as:

$$\Delta R := \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \tag{2.6}$$

Furthermore, since the incoming protons collide head-on with zero initial momentum in the transverse plane, momentum conservation applies strictly in the (x–y) plane. This motivates the use of transverse kinematic variables, such as the transverse momentum p_{T} and transverse energy E_{T}:

$$p_{\mathrm{T}} = \sqrt{p_x^2 + p_y^2}, \qquad E_{\mathrm{T}} = \sqrt{m^2 + p_{\mathrm{T}}^2} \tag{2.7}$$

Any imbalance in the vector sum of transverse momenta or energies may indicate the presence of undetectable particles and is quantified as the missing transverse momentum $\vec{p}_{\mathrm{T}}^{\mathrm{miss}}$ or energy $E_{\mathrm{T}}^{\mathrm{miss}}$:

$$\vec{p}_{\mathrm{T}}^{\mathrm{miss}} = -\sum_i \vec{p}_{\mathrm{T}}^{\,i}, \qquad E_{\mathrm{T}}^{\mathrm{miss}} = \left|\vec{p}_{\mathrm{T}}^{\mathrm{miss}}\right| \tag{2.8}$$

2.2.2 Detector architecture and subsystems

To ensure full coverage, ATLAS exhibits an "onion-like" layered structure around the beam pipe, combining a central cylindrical *barrel* region (identified by a pseudorapidity of $|\eta| < 1.05$) with two *end-cap* disks (identified by $|\eta| > 1.4$) oriented perpendicularly to the beam axis, providing an almost full 4π solid angle coverage.

Each layer serves a specific role in particle identification and reconstruction, based on how different particles interact with matter (Figure 2.4).

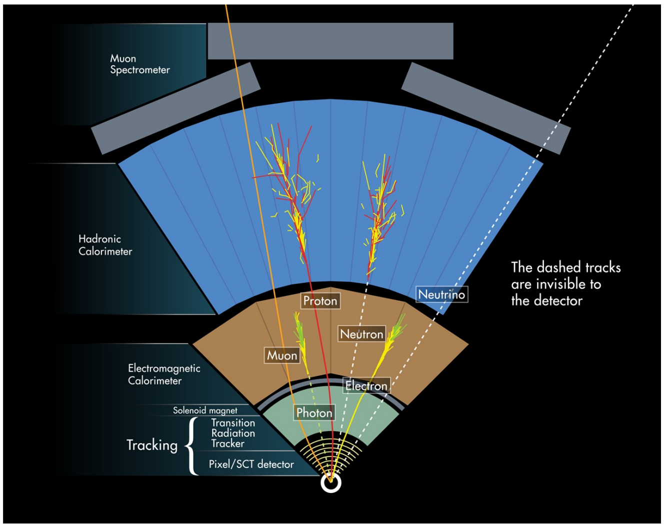

Figure 2.4: Section of the transverse plane in ATLAS showing how different particles interact with the subdetectors. Figure taken from Ref. [37].

The Inner Detector

The primary role of the Inner Detector (ID) is to measure the trajectory, momentum, and charge sign of traversing charged particles, as well as to reconstruct primary and secondary interaction vertices. At design luminosity, around 1000 particles emerge from the collision point every 25 ns, creating a high track-density environment: to maintain the vertex and momentum resolution required for high-precision physics measurements, the ID combines high-granularity semiconductor tracking with continuous straw-tube drift detection.

Enclosed in a cylindrical envelope with a length of 6.2 m and a diameter of 2.1 m, the entire ID is embedded in a 2 T axial magnetic field provided by a central superconducting solenoid aligned parallel to the beam axis. It is used to curve the trajectories of all charged particles and to measure their transverse momentum according to the relation:

$$p_{\mathrm{T}} = q|\vec{B}|\rho \tag{2.9}$$

where the symbol ρ denotes the radius of curvature of the charged particle's trajectory in the plane transverse to the magnetic field.

Covering the pseudorapidity range $|\eta| < 2.5$, the ID consists of three complementary sub-detector systems arranged concentrically around the beam pipe, spanning both the central barrel and the forward end-cap regions (as shown in Figure 2.5):

- **Pixel Detector:** The innermost tracking layer. It comprises 1744 n-type silicon pixel sensors ($250\,\mu$m thick) for a total of 87.2×10^6 readout channels. The innermost layer, known as the *Insertable B-Layer* (IBL), is positioned extremely close to the beam pipe and plays a critical role in secondary vertex reconstruction. Incident charged particles ionize the silicon, creating electron-hole pairs collected under an applied bias voltage. Beyond spatial tracking, the charge collection allows for a measurement of the specific energy loss ($\mathrm{d}E/\mathrm{d}x$), enabling particle identification at low momenta. Positioned closest to the interaction point, the Pixel Detector suffers the highest radiation damage and operational exposure.

- **Semiconductor Tracker (SCT):** Located directly outside the Pixel Detector, the SCT is a silicon microstrip system consisting of 4088 sensors. Operating on the same electron-hole pair collection principle as the Pixel Detector, the SCT provides four precision space-points per track.

- **Transition Radiation Tracker (TRT):** The outermost sub-detector layer, consisting of approximately $400,000$ straw tubes (4 mm diameter). The TRT is used to perform high-purity electron identification by distinguishing light electrons from heavier hadrons.

So, in addition to precise momentum reconstruction through trajectory curvature in the magnetic field, the ID is fundamental for primary[11] and secondary vertex detection.

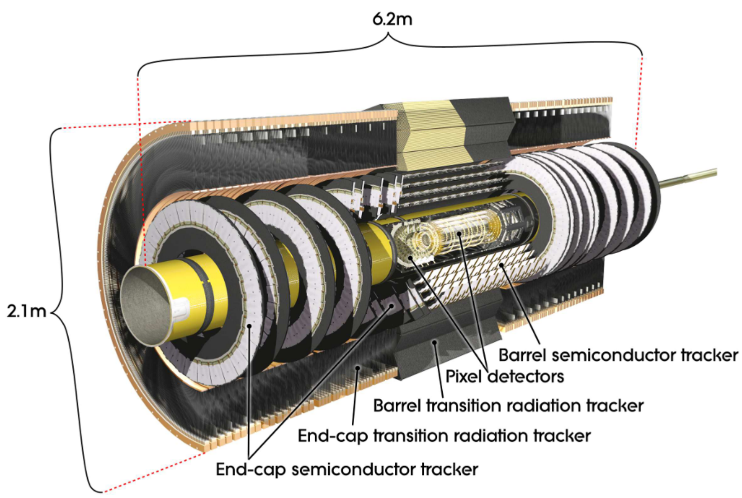

Figure 2.5: Cut-away view of the ATLAS inner detector taken from Ref. [28].

[11]The *primary vertex* corresponds to the hard-scattering proton-proton collision.

The Calorimetric System

Located outside the Inner Detector, the ATLAS calorimetric system is designed to measure the energy of both charged and neutral particles (electrons, photons, and hadrons) produced in proton-proton collisions. The system relies on *sampling calorimeters*, which consist of alternating layers of high-density passive absorber materials where traversing particles undergo hard interactions and initiate particle showers[12], and active media that detect the energy deposited by secondary shower particles.

The overall thickness of the calorimetric system corresponds to approximately 9.7 nuclear interaction lengths[13]. This depth ensures full longitudinal containment for high-energy hadronic jets and minimizes the phenomenon of *punch-through*, where energetic hadrons escape into the muon spectrometer without being fully absorbed. Furthermore, the broad coverage in pseudorapidity (η) provides high hermeticity, allowing precise measurements of the missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$).

Depending on the dominant interaction mechanism and the nature of the particle shower, the system is divided (Figure 2.6) into:

- **Electromagnetic Calorimeter:** Designed to measure electrons and photons, which interact purely via electromagnetic processes (primarily bremsstrahlung for electrons and e^+e^- pair production for photons). It employs lead plates as the passive absorber and Liquid Argon (LAr) as the active detection medium.
- **Hadronic Calorimeter:** Positioned behind the electromagnetic calorimeter, it measures hadrons that interact predominantly via the strong nuclear force. Typically, hadrons deposit roughly 60% of their energy in the electromagnetic section, while the remaining 40% develops within the hadronic calorimeter.

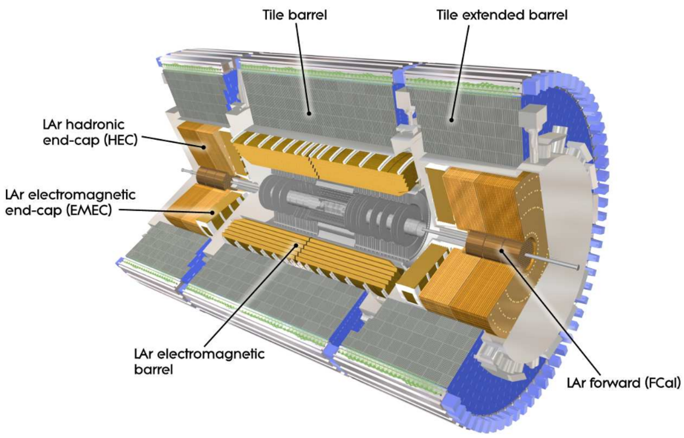

Figure 2.6: Cut-away view of the ATLAS calorimeter system taken from Ref. [28].

[12]A particle shower is a cascade of secondary particles generated when a high-energy primary particle interacts with dense absorber material, developing either via electromagnetic mechanisms (bremsstrahlung and pair production) or strong nuclear interactions.

[13]The nuclear interaction length (λ_I) is the mean path length required for a high-energy relativistic hadron to reduce its energy or probability of undergoing an inelastic nuclear interaction by a factor of $1/e$.

The Muon Spectrometer

The Muon Spectrometer (MS) forms the outermost system of the ATLAS experiment. Muons are the only charged SM particles capable of traversing the inner detector and the dense material of the electromagnetic and hadronic calorimeters without being stopped[14]. The primary functions of the MS are to provide fast and efficient Level-1 (L1) triggering for traversing muons (discussed later), measure their trajectory curvature to precisely reconstruct their transverse momentum (p_{T}), and resolve the bunch-crossing time.

A schematic 3D view of the overall system and its sub-detectors is shown in Figure 2.7a, while Figure 2.7b illustrates the spatial arrangement scheme.

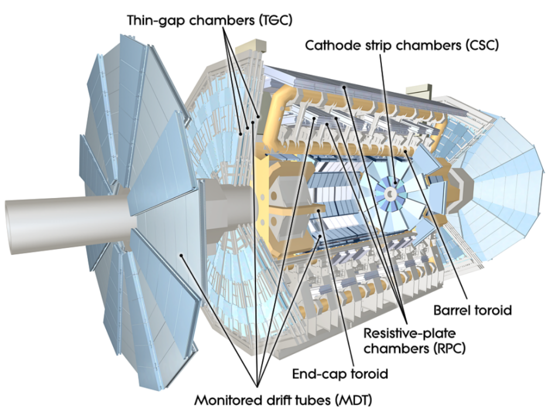

(a) Schematic cut-away drawing of the ATLAS Muon Spectrometer including its main subsystems, taken from ref. [38].

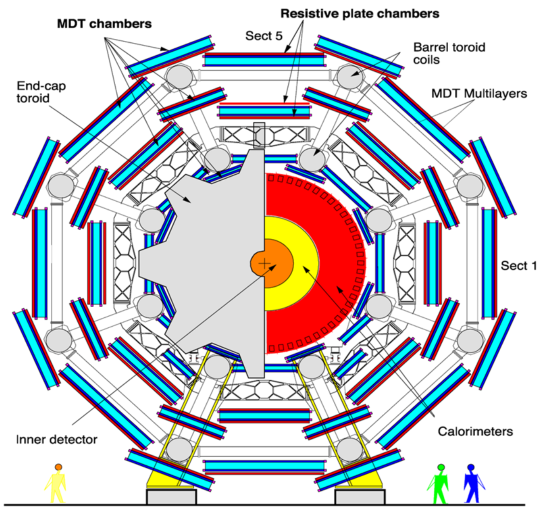

(b) Schematic representation of ATLAS in the x-y projection, taken from Ref. [39].

Figure 2.7: Overview of the ATLAS MS.

To perform high-precision momentum measurements via trajectory bending, the MS relies on a dedicated magnetic field system composed of large superconducting air-core toroids, completing[15] the ATLAS magnet system alongside the ID solenoid, as shown in Figure 2.8. Unlike this one, the toroids produce a magnetic field that is essentially orthogonal to the muon trajectories, maximizing the deflection in the bending plane: the **Barrel Toroid** ($|\eta| < 1.0$) consists of eight rectangular superconducting coils arranged symmetrically in a petal-like structure around the beam pipe (z-axis) and provides a magnetic field strength ranging between 0.5 T and 4 T in the central region; the **End-Cap Toroids** ($1.0 < |\eta| < 2.7$) each contain eight coils housed in a single cryostat, and they generate a peak magnetic field of up to 4 T to bend muons produced at forward and backward pseudorapidities.

The MS features a cylindrical symmetry divided into a central barrel and two end-caps, covering an overall pseudorapidity range of $|\eta| < 2.7$:

[14] As Minimum Ionizing Particles (MIPs), muons do not undergo strong interactions, and bremsstrahlung is strongly suppressed by their mass ($m_\mu \approx 207\, m_e$). Thus, they lose minimal energy via ionization and penetrate the calorimeters.

[15] All three magnetic systems in ATLAS operate **independently** at cryogenic temperatures maintained by a liquid helium cooling system.

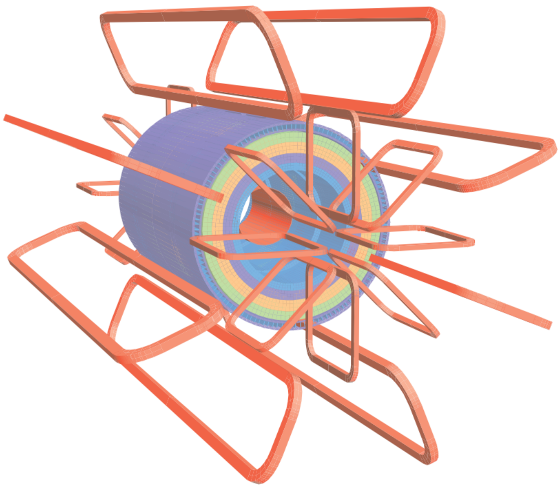

Figure 2.8: Geometry of the ATLAS magnet system, showing the end-cap toroids and barrel toroids in light red and the solenoid in the centre. Figure taken from Ref. [28].

- **Barrel Region ($|\eta| < 1.05$):** Detector chambers are arranged concentrically around the beam axis in three cylindrical layers: Barrel Inner (BI), Barrel Middle (BM), and Barrel Outer (BO). Azimuthally (ϕ), the barrel is segmented into 16 sectors. Even-numbered sectors accommodate "large" chambers (BIL, BML, BOL) mounted between the toroid coils, whereas odd-numbered sectors host "small" chambers (BIS, BMS, BOS) installed within and around the coils.
- **End-cap System ($1.05 < |\eta| < 2.7$):** The end-caps consist of four vertical disk layers oriented perpendicularly to the z-axis: End-cap Inner (EI), End-cap Extra (EE), End-cap Middle (EM), and End-cap Outer (EO). Like the barrel, these disks are divided azimuthally into large and small sectors (e.g., EIL/EIS, EEM/EES).

Finally, to fulfill the requirements of high-precision spatial resolution and ultrafast trigger response under varying background rates, four main gaseous detector technologies are employed:

1. **Monitored Drift Tubes (MDT):** High-precision tracking across almost the entire volume ($|\eta| < 2.7$, except in the innermost end-cap layer at $2.0 < |\eta| < 2.7$). When a muon ionizes the gas inside, primary electrons drift toward the anode wire with a maximum drift time of up to 700 ns. MDTs are assembled into chambers consisting of two *multilayers* separated by a mechanical spacer frame. Each multilayer contains 3 or 4 layers of drift tubes, staggered by one tube radius. Inner layer chambers feature 4-layer multilayers to cope with higher particle rates closer to the interaction point, while middle and outer layers use 3-layer multilayers. In total, 1,171 MDT chambers comprising 354,240 tubes are installed in ATLAS [40].
2. **Cathode Strip Chambers (CSC) / New Small Wheels (NSW):** In the high-pseudorapidity region ($2.0 < |\eta| < 2.7$) of the innermost end-cap layer, CSCs were historically used due to their higher granularity and ability to handle high background rates. In recent upgrades (Run 3), the inner end-cap region was replaced by the NSWs, to ensure robust tracking and fast trigger capability at high luminosity.

3. **Thin Gap Chambers (TGC):** Multi-wire proportional chambers operating in the end-cap region ($1.05 < |\eta| < 2.4$). Featuring a very small wire-to-cathode gap, TGCs provide fast signals for L1 trigger generation as well as measuring the second coordinate (ϕ) orthogonal to the bending direction.

4. **Resistive Plate Chambers (RPC):** Fast gaseous detectors dedicated to L1 triggering and second-coordinate (ϕ) measurement in the barrel region ($|\eta| < 1.05$), discussed more in the subsection 2.2.3.

2.2.3 Overview of the RPC system

Due to their fast response and excellent time resolution, RPCs are a critical component of the ATLAS trigger and offline reconstruction logic in the barrel. The RPC system is configured in three concentric layers located in the barrel region, as shown in Figure 2.7b:

- Two RPC layers are mounted directly onto both outer faces of the MDT chambers in the Middle layer (BM).

- The third RPC layer is attached to the Outer MDT chambers (BO), installed on the inner side of the MDT frame in even sectors and on the outer side in odd sectors.

An RPC unit consists of two independent gas gaps. Each gas gap is formed by two parallel rectangular plates made of resistive phenolic-melaminic plastic laminate, separated by a 2 mm distance. The gap is filled with a non-flammable gas mixture composed of 94.7 % $C_2H_2F_4$, 5 % Iso-C_4H_{10}, and 0.3 % SF_6, and operates under an electric field gradient of approximately 4.9 kV/mm. Primary ionization produced by a passing muon initiates an electron avalanche in the gas, inducing a signal on metallic readout strips located on either side of the gas gap.

Each RPC chamber contains two detector layers equipped with orthogonal strip arrays to independently measure the η (bending) and ϕ (azimuthal) coordinates. Readout strips range from 2.3 to 3.5 cm in width and up to 3.2 m in length; longer chambers are composed of two adjacent gas-gap segments joined together.

The spatial resolution provided by each strip is approximately 10 mm, with a nominal timing resolution of 1.5 ns. The internal readout electronics operate with a 320 MHz clock. Since LHC *pp* collisions occur at a frequency of 40 MHz (corresponding to a 25 ns bunch-spacing), each bunch-crossing interval is divided into eight 3.125 ns time ticks. This intrinsic digital time granularity is directly observed in RPC timing measurements and plays a central role in offline calibration and track reconstruction. Overall, the barrel RPC system comprises 606 chambers, containing roughly 3,600 gas volumes and over 370,000 individual readout channels [41].

2.2.4 Trigger System

As discussed above, at nominal LHC luminosity, bunch crossings occur every 25 ns, corresponding to an event production rate of roughly 40 MHz. Because recording and storing data at this volume is technically and economically unfeasible, ATLAS employs a two-level trigger and data acquisition system to select events of scientific interest in real time:

- **Level-1 Trigger (L1):** A hardware-based system that makes a decision within $2.5\,\mu$s. Utilizing fast custom electronics, it processes coarse-granularity signals from specific trigger detectors (e.g. calorimeters, RPCs) to identify high-p_T muons, electrons, photons, jets, and τ leptons decaying hadronically, as well as large missing transverse energy (E_T^{miss}). L1 reduces the event rate from 40 MHz to approximately 100 kHz. Furthermore, it defines one or more Regions of Interest (RoIs): geographic coordinates in terms of η and ϕ where interesting features are detected which guide the subsequent processing step.

- **High-Level Trigger (HLT):** A software-based system that accesses full-detector granularity, primarily executing complex algorithms within the RoIs passed by L1. These algorithms mirror offline reconstruction procedures to refine event selection, ultimately reducing the final storage output rate to roughly 1-3 kHz across over 1,500 dedicated trigger menus designed for a wide range of physics processes.

Chapter 3

Heavy Stable Charged Particles at ATLAS

This chapter (based on Ref. [24]) provides an overview of Heavy Stable Charged Particles (HSCPs), starting from their theoretical motivation (as briefly introduced in Section 1.3.4 of Chapter 1) to their experimental signatures and reconstruction within the ATLAS detector.

3.1 Production mechanisms and theoretical overview

In this context, the term *stable* is used in an experimental sense: it refers to particles with a lifetime long enough to traverse the relevant detector volume before decaying. The terminology used in the literature varies depending on the experimental and theoretical context. Such particles are commonly referred to as Long-Lived Particles (LLPs), while searches specifically targeting massive charged states are often described as Heavy Stable Charged Particle (HSCP) searches: in this chapter, the term HSCP is used consistently to denote the class of long-lived, massive, electrically charged particles considered in the analysis.

Consequently, HSCPs can interact with the active material of the detector, leaving experimental signatures that closely resemble those of SM muons. However, due to their significantly larger predicted masses ($m_{\mathrm{HSCP}} \gg m_\mu$), their velocity $\beta = v/c$ is substantially lower than the speed of light. This characteristic low velocity provides a clear distinction from ultra-relativistic SM muons, providing a largely model-independent probe of New Physics.

It's important to specify that the experimental signature of an HSCP is largely independent of the specific BSM model in which it arises. Regardless of the underlying production mechanism, a charged HSCP is expected to leave a reconstructed track with anomalously large ionization energy loss and a delayed time of flight due to its low velocity. The specific model primarily determines the production cross-section and the event topology, including the presence or absence of missing transverse momentum, additional particles in the final state, and possible objects suitable for triggering. These model-dependent features affect the event selection and trigger strategy, but not the basic detector signature used to identify the HSCP itself.

Direct and indirect production channels

The exact production mechanism of HSCPs is strictly dependent on the underlying BSM model. Nevertheless, because most theoretical extensions incorporate a new conserved quantum number (such as R-parity in SUSY), HSCPs are predominantly predicted to be directly produced in pairs. Prominent candidates within the Minimal Supersymmetric Standard Model (MSSM) include staus ($\tilde{\tau}$) and charginos ($\tilde{\chi}^{\pm}$) [1], which are pair-produced at tree level via electroweak interactions.

Alternatively, HSCPs can be produced indirectly through the decay chain of a heavier BSM particle. In this case, the observable production rate depends on the production cross-section of the parent particle and on the relevant branching fractions. Such cascade decays can therefore provide an important contribution to the HSCP signal when the parent particle is produced with a sufficiently large rate and has a sizeable branching fraction into the HSCP.

Single production and associated production

Single production of a HSCP is also possible in specific BSM models. However, its rate is strongly model-dependent and is determined by the couplings and quantum numbers involved in the production process. There is therefore no general relation between the single-production cross-section and the HSCP lifetime. In particular, a long lifetime does not by itself imply a negligible production cross-section, although the same small couplings or other suppressions that lead to a long lifetime can also affect the production rate in specific models.

HSCPs may also be produced in association with a neutral stable massive particle. This scenario can arise when the charged and neutral states are close in mass, with the neutral state acting as the Lightest Supersymmetric Particle (LSP). A classic example in the MSSM is the associated production of the lightest chargino ($\tilde{\chi}_1^{\pm}$) and neutralino ($\tilde{\chi}_1^0$) [2]. In such topologies, the neutral state escapes the detector unobserved, leaving a characteristic missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$) signature alongside the high-p_{T} track left by the HSCP.

3.2 Experimental observables for HSCP searches

The HSCPs' substantially lower velocity leads to characteristically high ionization energy losses in tracking detectors and delayed arrival times in outer subdetectors. Since no SM process produces relativistic charged particles with such low velocities, the background for HSCP searches is virtually devoid of physics processes and is instead dominated by detector mismeasurements and non-collision backgrounds, such as cosmic ray

[1] Charginos ($\tilde{\chi}_i^{\pm}$, with $i = 1, 2$) are mass eigenstates resulting from the mixing of the charged winos ($\tilde{W}^{\pm}$, the fermionic superpartners of the $W^{\pm}$ gauge bosons) and charged higgsinos ($\tilde{H}_u^+, \tilde{H}_d^-$, the superpartners of the Higgs doublets).

[2] Neutralinos ($\tilde{\chi}_j^0$, $j = 1, 2, 3, 4$) are neutral mass eigenstates formed by the linear combination of the bino ($\tilde{B}$, superpartner of the $U(1)_Y$ gauge boson), neutral wino ($\tilde{W}^0$, superpartner of the neutral $SU(2)_L$ gauge boson), and neutral higgsinos ($\tilde{H}_u^0, \tilde{H}_d^0$).

muons[3].

3.2.1 Ionization energy loss (dE/dx)

A charged particle traversing a detector deposits energy primarily via inelastic collisions with atomic electrons, a process quantified by the specific ionization energy loss, dE/dx. In ATLAS, dE/dx is measured by the ID using the Pixel Detector and the Transition Radiation Tracker (TRT).

The mean energy loss per distance travelled for a particle with charge z, velocity $v = \beta c$, and energy E is described by the Bethe-Bloch formula:

$$\left\langle \frac{dE}{dx} \right\rangle = 4\pi n \frac{e^4 z^2}{m_e c^2 \beta^2} \left[\frac{1}{2} \ln \left(\frac{2 m_e c^2 \beta^2 \gamma^2 T_{\max}}{I_e^2} \right) - \beta^2 - \frac{\delta}{2} \right] \tag{3.1}$$

where n is the electron density of the material, m_e and e represent the rest mass and electric charge of the electron, I_e is the mean excitation potential, and δ is a density-effect correction factor. The Lorentz factor γ is defined as:

$$\gamma \equiv \frac{1}{\sqrt{1-\beta^2}} \tag{3.2}$$

The maximum kinetic energy $T_{\max}$ transferable to a free electron in a single collision by a heavy particle ($m \gg m_e$) is given by:

$$T_{\max} = 2 m_e c^2 \beta^2 \gamma^2 \tag{3.3}$$

Because the energy loss depends primarily on the particle's $\beta\gamma$, a measurement of dE/dx provides an estimate of $\beta\gamma$ and, consequently, of the particle's velocity.

In the ATLAS Pixel Detector, four individual dE/dx measurements are collected per track (one per layer). To estimate the Most Probable Value (MPV) from the highly asymmetric Landau energy-loss distribution, a truncated mean is calculated by discarding the highest dE/dx hit to reduce tail fluctuations. An empirical parametrization is then used to relate the MPV of dE/dx to $\beta\gamma$:

$$\mathrm{MPV}_{dE/dx} = \frac{A}{(\beta\gamma)^C} + B \tag{3.4}$$

where A, B, and C are calibration constants extracted using low-momentum pions, kaons, and protons from low-luminosity runs. Combining the value of $\beta\gamma$ obtained from Equation 3.4 with the track momentum p measured by the ID allows for a direct calculation of the particle's rest mass m_0:

$$m_0 = \frac{p}{\beta\gamma c} \tag{3.5}$$

[3] Although cosmic ray muons travel at near light speed, their uncorrelated arrival time with respect to the beam collision can mimic a delayed signature, though they are eventually rejected using specific timing vetoes.

3.2.2 Time-of-Flight (ToF) measurement

A complementary approach to determine particle velocity is measuring its ToF between the interaction point (IP) and a detector element located at a path distance d. Particles traveling at velocities significantly below the speed of light ($\beta < 1$) experience a delay relative to ultra-relativistic particles ($\beta \approx 1$).

A timing offset variable t_0 is defined as the difference between the actual time of flight ToF_a and the expected time of flight $\mathrm{ToF}_c = d/c$ for a particle moving at the speed of light:

$$t_0 = \mathrm{ToF}_a - \mathrm{ToF}_c = \mathrm{ToF}_a - \frac{d}{c} \tag{3.6}$$

The velocity β can then be extracted directly from t_0:

$$\beta = \frac{d}{\mathrm{ToF}_a \cdot c} = \frac{d}{d + t_0 \cdot c} \tag{3.7}$$

In ATLAS, the RPCs provide the primary timing information used for the ToF measurement in the barrel region of the Muon Spectrometer. Their location at large distances from the interaction point, combined with their timing capabilities, makes them particularly well suited to identifying delayed signals from slowly moving charged particles. The MDT chambers primarily provide precision tracking information and are used to reconstruct and associate the particle trajectory with the RPC measurements. Combined with the momentum p, the calculated β yields another independent determination of m_0 using Equation 3.5.

3.2.3 HSCP identification in ATLAS

The experimental signature of a HSCP is characterized by the combination of a charged-particle track, anomalously large ionization energy loss, and a velocity significantly below that of relativistic Standard Model particles. The momentum measured by the ID, together with the dE/dx and ToF measurements, can therefore be used to reconstruct the mass of the candidate through:

$$m_0 = \frac{p}{\beta\gamma c} \tag{3.8}$$

The dE/dx and ToF observables provide complementary information. The former is sensitive to the particle's $\beta\gamma$ through its ionization energy loss, while the latter provides a direct measurement of its velocity. Their combination is therefore particularly powerful for distinguishing massive, slowly moving charged particles from relativistic Standard Model backgrounds.

In the present thesis, particular emphasis is placed on the ToF measurement provided by the RPCs in the ATLAS Muon Spectrometer. As discussed in the following chapter, the accuracy of the reconstructed velocity depends on the timing response of the RPC detector. A dedicated strip-by-strip timing calibration is therefore required before the calibrated β measurement can be used in the subsequent analysis.

Chapter 4

Analysis and Results

As established in the previous chapters, searches for HSCPs are characterized by an extremely small irreducible SM background. The sensitivity of such searches is therefore strongly influenced by the detector response and, in particular, by the accuracy with which the particle velocity is reconstructed. A precise determination of the velocity β is essential both for the identification of slowly moving charged particles and for the discrimination of potential signal candidates from SM backgrounds.

In the ATLAS MS, the ToF information provided by the RPCs plays a central role in the reconstruction of β. Since the measured timing is affected by detector-dependent offsets and finite time resolution, a dedicated RPC timing calibration is required. High-energy SM muons provide a suitable reference sample for this purpose, as their velocity is expected to be very close to the speed of light, $\beta \simeq 1$.

This chapter is divided into two closely related aspects of the analysis. First, the timing response of the RPC barrel is calibrated on a strip-by-strip basis using relativistic muons. The resulting calibration constants are then applied to the associated RPC hits and used to reconstruct calibrated β values. Since the RPC detector provides two independent measurement views, corresponding to the η and ϕ coordinates, the strip identifier is stored together with the view information. The calibration is therefore performed independently for the η and ϕ strips, allowing different timing offsets to be assigned to the two readout views. The performance of the calibration is assessed by comparing the reconstructed velocity distributions before and after the timing corrections.

Second, the impact of the calibrated velocity measurement is investigated in a sensitivity study targeting supersymmetric staus ($\tilde{\tau}$). In this study, the calibrated β information is combined with the ionization energy loss (dE/dx) measured in the ID to enhance the separation between signal and background. Signal samples corresponding to different stau mass hypotheses are compared with a data-driven estimate of the background, and the β selection is optimized to maximize the expected sensitivity.

4.1 RPC Timing Calibration

4.1.1 Data samples

The timing calibration and the associated studies presented in this section are based on proton-proton (pp) collision data collected by the ATLAS experiment in 2022 at a center-of-mass energy of $\sqrt{s} = 13.6$ TeV, corresponding to an integrated luminosity of 31.4 fb^{-1} [42].

The analysis starts from preprocessed and skimmed ntuples. Therefore, the event reconstruction, derivation, skimming, thinning, and slimming stages are not performed as part of the present work. However, specific selection criteria are applied during these upstream stages to provide an input file containing the required reconstructed quantities for the muon candidates and their associated RPC (and MDT) hits. At the derivation level, events are selected by requiring them to pass a single-muon trigger. In addition, muons are required to have a transverse momentum $p_{\mathrm{T}} > 20$ GeV, and the event must contain a dimuon pair with an invariant mass in the range $70 < m_{\mu\mu} < 110$ GeV. Subsequently, at the skimming level, tighter selections are enforced. The single-muon trigger requirement is maintained, but the transverse momentum threshold is raised to $p_{\mathrm{T}} > 25$ GeV. Furthermore, events must contain at least two muons satisfying the *Medium* identification working point, while the invariant mass requirement of $70 < m_{\mu\mu} < 110$ GeV is preserved. Finally, to ensure the event is suitable for a Tag-and-Probe study (see section 4.1.2), the selection requires the presence of at least one valid probe muon candidate. At this stage, this requirement is kept loose, demanding only that the probe candidate has at least one associated MDT hit and one RPC hit.

The study makes use of two muon collections: standard track candidates and slow-muon candidates. The latter are associated with the RPC measurements used in the timing study. The RPC timing calibration and the subsequent β reconstruction are performed using the RPC hits associated with the slow-muon candidates stored in the input ntuple. The standard track collection is used separately for the reconstruction cross-checks and invariant-mass studies described above. The available hit information includes the RPC strip identifier, detector coordinates, measured time-of-flight, detector station information, and the detector technology associated with each hit, allowing RPC and MDT hits to be distinguished.

4.1.2 Muon pair reconstruction and the $Z \to \mu\mu$ resonance

A sample of relativistic SM muons is required for the RPC timing calibration, since their velocity is expected to be close to the speed of light, $\beta \simeq 1$. The $Z \to \mu\mu$ process provides a well-defined, data-driven source of such muons. The corresponding control sample is defined according to a Tag-and-Probe strategy. The Tag-and-Probe selection is performed at the pre-processing stage and is therefore not repeated in the present analysis, which starts from a preprocessed and skimmed ntuple. In the resulting sample, the tag muon is required to satisfy the appropriate trigger and quality requirements, while the probe provides the muon candidate used to study the detector response with reduced trigger bias.

In the Tag-and-Probe approach, one of the two muons is identified as the *tag* and is required to satisfy tight selection criteria and to be associated with the event trigger. This provides a high-purity indication that the event is consistent with a $Z \to \mu\mu$ decay. The second muon, referred to as the *probe*, is used to study the reconstruction and detector response without imposing the same trigger requirement. Consequently, by reducing this trigger bias, the probe can be investigated with a lower transverse momentum (p_{T}) threshold, avoiding the need for the p_{T} to be above the muon trigger efficiency plateau, which significantly increases the available statistics. If multiple probe candidates exist in the same event, the one yielding a dimuon invariant mass, $m_{\mu\mu}$, closest to the nominal Z-boson mass is selected. This data-driven strategy is particularly useful for studying reconstructed muons and comparing different reconstruction approaches without relying on simulation.

For each event, the available reconstructed muon candidates are combined into pairs. The invariant mass is calculated under the ultra-relativistic approximation, neglecting the muon mass ($m_\mu \simeq 0$). For two reconstructed objects with transverse momenta $p_{\mathrm{T},1}$ and $p_{\mathrm{T},2}$, pseudorapidity difference $\Delta\eta$, and azimuthal separation $\Delta\phi$, it is given by (see more in Appendix A):

$$m_{\mu\mu} = \sqrt{2p_{\mathrm{T}_1}p_{\mathrm{T}_2}\left[\cosh(\Delta\eta) - \cos(\Delta\phi)\right]} \tag{4.1}$$

The presence of a peak around the nominal Z-boson mass, which corresponds to $m_Z = 91.1879 \pm 0.0020$ GeV, provides a validation of the reconstructed muon sample and of the selected control region, as seen in Figure 4.1. The validated control sample provides the relativistic reference underlying the calibration, while the RPC hits used to determine the strip-dependent timing constants are taken from the RPC measurements associated with the slow-muon candidates stored in the input ntuples.

4.1.3 RPC hit association and quality requirements

Before reconstructing the muon velocity, quality requirements are applied to the slow-muon candidates. A candidate is retained if it is associated with at least two RPC hits and eight MDT hits. The RPC hit requirement ensures that a sufficient number of timing measurements is available for the velocity reconstruction, while the MDT requirement guarantees adequate tracking information.

It is important to note that these candidate-level selections are enforced only during the β reconstruction. The strip-timing calibration itself is performed beforehand and utilizes all available RPC hits to populate the strip-by-strip t_0^{uncal} distributions, without applying the $N_{\mathrm{RPC}} \geq 2$ and $N_{\mathrm{MDT}} \geq 8$ requirements.

For each selected slow-muon candidate, the associated RPC hits are considered individually. Each hit is identified by its detector station, detector coordinates, measurement view, and strip identifier, which together define the detector element used in the subsequent strip-by-strip timing calibration. The measured time of flight and the hit position are also retrieved for each RPC hit.

The corresponding distance from the primary vertex is calculated from the hit coordinates. The geometry of the measurement is illustrated schematically in Figure 4.2. For a relativistic muon, the expected time of flight to an RPC strip located at a distance L from the primary vertex can be compared with the measured time to determine the corresponding timing offset. These quantities form the basis for the subsequent hit-by-hit velocity reconstruction and RPC timing calibration.

4.1.4 RPC Strip Timing Calibration

The RPC timing calibration is performed on a strip-by-strip basis using the relativistic muon control sample. To fully identify the detector element for each selected RPC hit, its strip identifier is passed to the `MuonIdHelper` tool within the `Athena` software

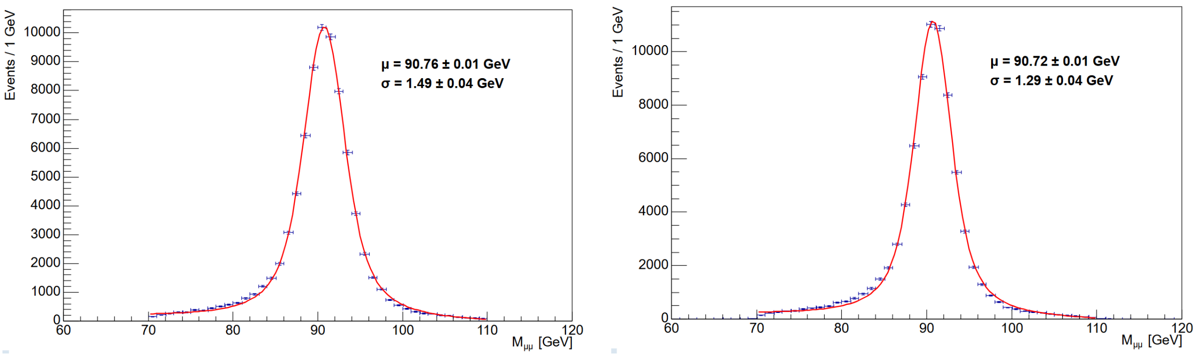

(a) SM track candidates.

(b) Slow-muon candidates.

Figure 4.1: Invariant-mass distributions of the selected opposite-charge muon pairs for (a) standard track candidates and (b) slow-muon candidates. The data points are fitted with a signal-plus-background model (solid red line) in the range $70\,\text{GeV} < m_{\mu\mu} < 110\,\text{GeV}$. The resonant $Z \to \mu\mu$ signal is described by a Voigtian profile (the convolution of a Gaussian resolution function and a Lorentzian resonance lineshape, as detailed in Appendix B), while the combinatorial background is modeled with a linear polynomial. The extracted mean mass μ and Gaussian resolution σ are displayed on the plots, demonstrating the performance of the reconstruction algorithms and the purity of the sample.

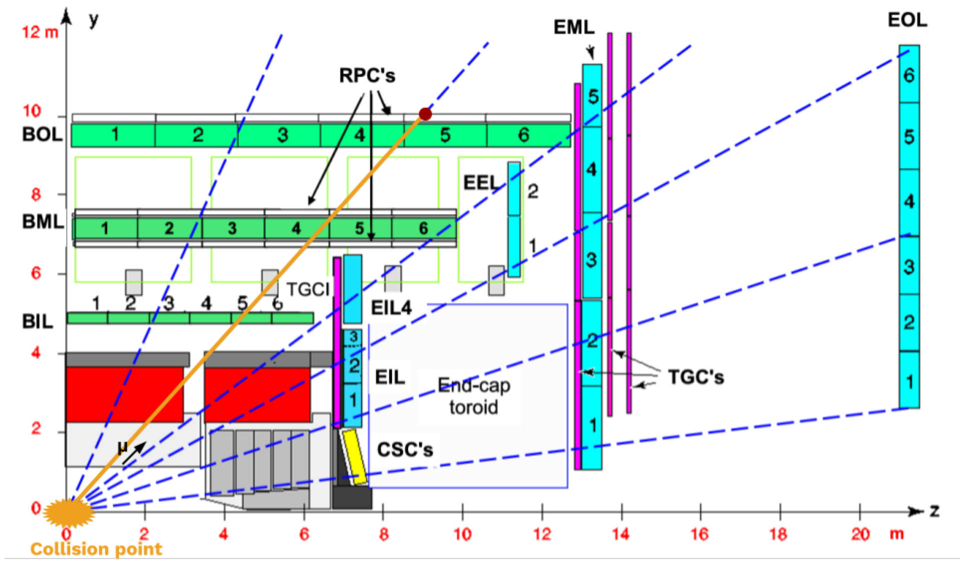

Figure 4.2: Schematic view of the ATLAS Muon Spectrometer layout in the $y-z$ plane, illustrating the geometry of the timing calibration. A relativistic probe muon (solid orange line) originating from the primary vertex (collision point) traverses the detector barrel. The theoretical time of flight of this muon is used as a reference to calibrate the timing response of a specific RPC element (indicated by the red dot) located at a flight distance L from the vertex. Drawing made on a Figure taken from Ref. [28].

framework. This allows the retrieval of the corresponding station, detector coordinates, and measurement view. The calibration procedure therefore follows the sequence of selecting the relevant RPC elements, collecting the associated hits, constructing the corresponding timing distributions, and extracting the strip-dependent timing offset.

For each RPC hit, the uncalibrated time-of-flight residual is defined as:

$$t_{0_i}^{\text{uncal}} = \text{ToF}_i - \frac{L_i}{c} \tag{4.2}$$

where ToF_i is the measured time of the i-th hit, L_i is the distance between the primary vertex and the hit position, and c is the speed of light. The corresponding distance is calculated from the three-dimensional hit coordinates according to:

$$L_i = \sqrt{x_i^2 + y_i^2 + z_i^2} \tag{4.3}$$

The relation between the measured time, the distance travelled, and the particle velocity follows from:

$$L_i = \beta_i c \text{ToF}_i \tag{4.4}$$

which gives:

$$\text{ToF}_i = \frac{L_i}{\beta_i c} \tag{4.5}$$

Combining this expression with Eq. 4.2 gives:

$$t_{0_i} = \frac{L_i}{c}\left(\frac{1}{\beta_i} - 1\right) \tag{4.6}$$

or, equivalently,

$$\frac{1}{\beta_i} = 1 + \frac{c t_{0_i}}{L_i} \tag{4.7}$$

For the relativistic control sample, $\beta \simeq 1$, and therefore the expected time-of-flight residual is close to zero in the absence of detector-dependent timing offsets. A non-zero displacement of the t_0 distribution for a given RPC strip is consequently interpreted as a strip-dependent timing offset.

For each selected strip, a separate t_0^{uncal} distribution is filled using the associated RPC hits from the control sample. A minimum statistical requirement is first imposed to ensure that the strip contains a sufficient number of entries. In order to determine the strip-dependent timing offset from the central part of the distribution while reducing the influence of isolated entries and non-Gaussian tails, a threshold based on the maximum bin content is then applied. In particular, the central region $\mathcal{R}$ is defined by retaining only those histogram bins whose content is at least 30% of the maximum bin content:

$$N_k \geq 0.3 N_{\max}, \qquad k \in \mathcal{R} \tag{4.8}$$

where N_k denotes the number of entries in bin k and $N_{\max}$ is the content of the most populated bin. This choice provides a simple and robust definition of the central peak based directly on the observed shape of the distribution, without relying on the mean or RMS of the full distribution, which can be affected by non-Gaussian tails or isolated outliers. Moreover, for an approximately Gaussian distribution with mean μ and standard deviation σ, the condition:

$$f(x) \geq 0.3 f_{\max} \tag{4.9}$$

corresponds to:

$$|x - \mu| \leq \sqrt{2 \ln \left(\frac{1}{0.3} \right)} \sigma \simeq 1.55\sigma \tag{4.10}$$

Thus, for a Gaussian-like core, the 30% maximum-content criterion selects a region approximately equivalent to a $\pm 1.5\sigma$ interval around the peak: this provides a core definition that is close to the ± 1.5 RMS criterion while remaining less sensitive to the influence of tails on the mean and RMS of the full distribution. Examples of the resulting distributions for an RPC η strip and an RPC ϕ strip are shown in Figure 4.3.

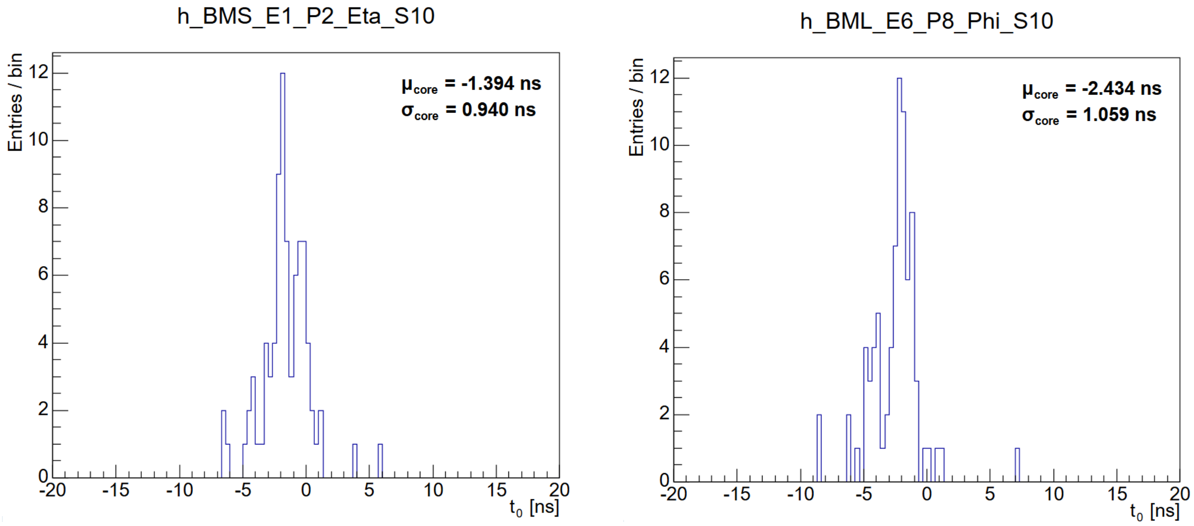

(a) Example of an RPC η strip for a BMS (Barrel Medium Small) station (Eta index 1, Phi index 2, Strip 10).

(b) Example of an RPC ϕ strip for a BML (Barrel Medium Large) station (Eta index 6, Phi index 8, Strip 10).

Figure 4.3: Examples of the t_0 distributions obtained for individual RPC strips. The central region used to determine the strip-dependent timing offset is defined by retaining bins with a content of at least 30% of the maximum bin content.

The strip-dependent timing calibration constant (i.e. strip timing offset) is extracted from the selected central region as a content-weighted mean of the corresponding bin centers:

$$t_0^{\text{corr}} = \frac{\sum\limits_{k \in \mathcal{R}} N_k t_k}{\sum\limits_{k \in \mathcal{R}} N_k} \tag{4.11}$$

where t_k is the corresponding bin center. The extracted value represents the average timing offset of the strip relative to the expected time of flight of a relativistic muon.

The statistical uncertainty of the calibration constant is estimated as:

$$\delta t_0^{\text{corr}} = \frac{\sigma_t}{\sqrt{N_{\text{core}}}} \tag{4.12}$$

where N_{core} is the total number of entries contained in the selected central region and σ_t is the RMS of the t_0^{uncal} distribution restricted to the same region. The quantity

σ_t characterizes the timing spread of the central part of the strip distribution and is subsequently used as the timing uncertainty associated with the corresponding RPC hit in the propagation of the β uncertainty.

The calibration constants are stored together with the detector identifiers defining each RPC strip and are subsequently applied to the corresponding RPC hits during the velocity reconstruction.

4.1.5 Global summary and spatial distribution of the calibration constants

A global summary of the detector's timing response is presented in Figure 4.4, which shows the one-dimensional distributions of the calibration constants t_0^{corr} and the intrinsic timing resolutions σ_t for all calibrated strips.

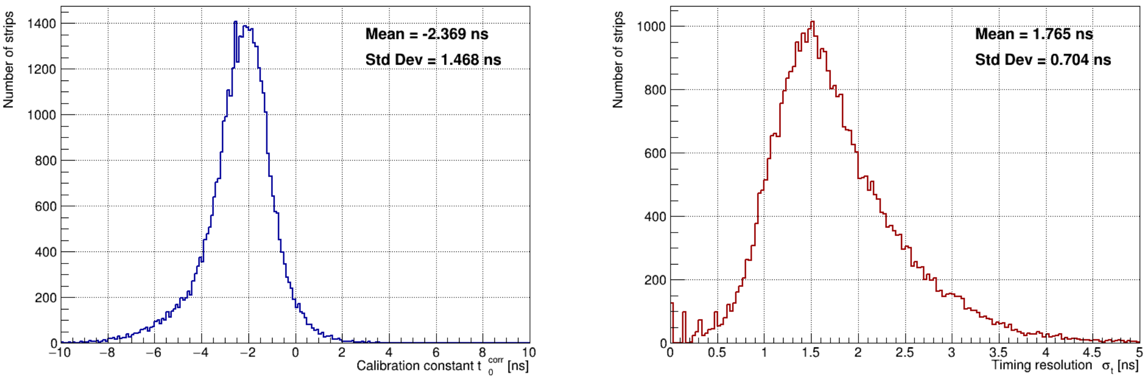

(a) Global timing offsets t_0^{corr}.

(b) Global timing resolution σ_t.

Figure 4.4: Global one-dimensional distributions of (a) the timing calibration constants t_0^{corr} and (b) the intrinsic timing resolutions σ_t for all calibrated RPC strips in the detector. The distributions summarize the overall timing response and temporal precision of the RPC system.

While the global one-dimensional distributions provide an overview of the overall detector performance, they do not retain information on the spatial dependence of the calibration parameters. To investigate the spatial uniformity of the timing response, the calibration constants are mapped onto the detector geometry using the RPC station indices `stationEta` and `stationPhi` as macroscopic coordinates. These indices identify the position of a given station within the barrel and should not be interpreted as continuous pseudorapidity and azimuthal coordinates of individual strips.

Since different RPC station types have different geometrical dimensions, strip segmentation, and identifier ranges, the spatial distributions are evaluated separately for each station type and for the two readout views corresponding to the η and ϕ measurements. The two views are distinguished through the `measuresPhi` identifier, while `strip` specifies the individual readout channel. Consequently, the resulting maps represent the spatial distribution of the calibration constants within a given station and readout view, rather than two physical layers of the chamber. This representation preserves the detector geometry and is consistent with the strip-by-strip nature of the calibration procedure.

Two main considerations motivate this representation in this analysis:

- **Inhomogeneous geometry:** Different RPC station types, even in the barrel (e.g. BML and BOS), have different physical dimensions, strip segmentations, and allowed identifier ranges. Constructing a single global spatial map would therefore mix detector elements with different geometrical and readout granularities.

- **Separate readout views:** The RPC readout provides the two spatial measurements through differently oriented strip sets. Treating the η and ϕ views separately makes it possible to identify non-uniformities associated with individual readout systems, such as channel-dependent timing offsets, electronics effects, or cabling-related delays.

The BML (Barrel Medium Large) station is presented here as a representative example. This station provides extensive coverage in the central region of the barrel and contains a large number of calibrated channels, making it suitable for illustrating the spatial structure of the calibration constants. Figure 4.5 shows the resulting two-dimensional calibration maps for the η- and ϕ-view readout strips of the BML station.

The maps of the calibration constants t_0^{corr} (Figures 4.5a and 4.5c) reveal the strip-to-strip variations in the timing response across the readout channels. The presence of coherent spatial structures in these maps reflects the underlying detector and readout organization, including channel-dependent hardware delays. The fact that these structures are reproduced by the calibration provides a useful validation of the strip-level calibration procedure.

The corresponding spatial distribution of the timing resolution σ_t is shown in Figures 4.5b and 4.5d.

Finally, all maps exhibit specific topological features that reflect the physical layout of the spectrometer. For BML, the horizontal white band at $\eta = 0$ corresponds to the physical gap separating the two halves of the ATLAS barrel, leaving the region free of RPC chambers; the randomly scattered white cells represent individual channels or chambers that were inactive (dead or masked) during the data-taking period, or for which the available statistics were insufficient to perform a reliable Gaussian core fit, leading the calibration algorithm to safely exclude them.

Together, these results confirm the validity of the strip-by-strip approach. By factoring out the structural hardware delays, the hit-level time measurements are centered back to the expected relativistic time of flight, while the resolutions provide the correct weights for the β reconstruction detailed in the subsequent section.

4.1.6 Velocity reconstruction

Once the strip-dependent timing calibration constant has been determined, it is subtracted from the corresponding measured time residual:

$$t_{0_i}^{\text{cal}} = t_{0_i}^{\text{uncal}} - t_0^{\text{corr}} \tag{4.13}$$

where the calibration constant is determined by the RPC strip associated with the i-th hit.

The calibrated time residual is then converted into a hit-level velocity measurement

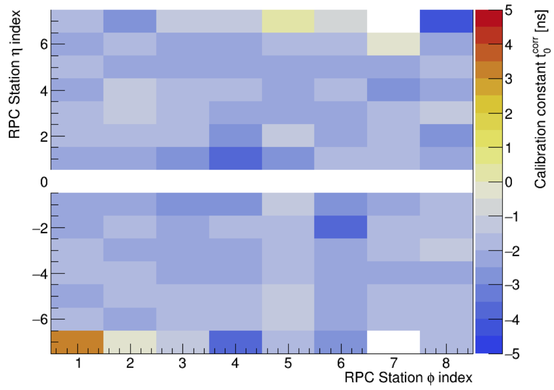

(a) Calibration constants t_0^{corr} (η-view).

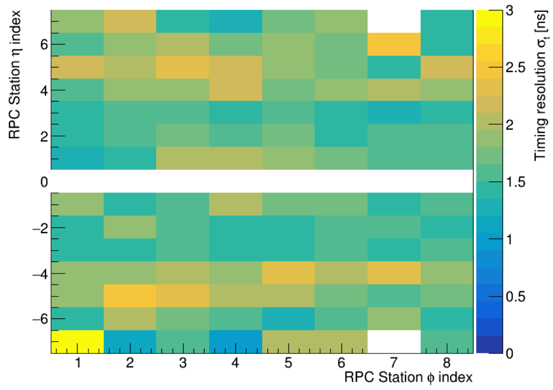

(b) Timing resolution σ_t (η-view).

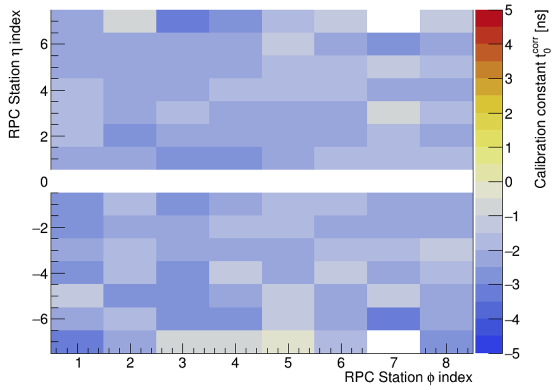

(c) Calibration constants t_0^{corr} (ϕ-view).

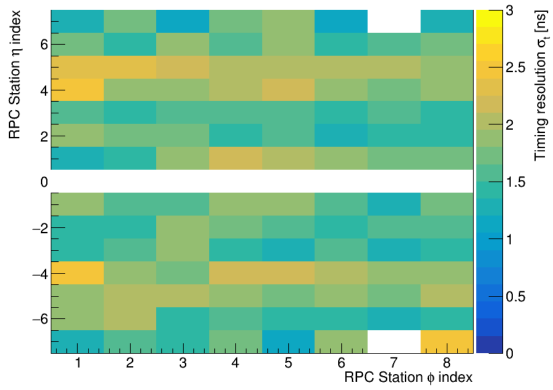

(d) Timing resolution σ_t (ϕ-view).

Figure 4.5: Two-dimensional calibration maps for the BML (Barrel Medium Large) station, separated by readout view. The top row shows (a) the extracted timing calibration constants t_0^{corr} and (b) the corresponding core timing resolution σ_t for the η-view strips. The bottom row displays the same variables, (c) t_0^{corr} and (d) σ_t, for the ϕ-view strips. The variables are plotted as a function of the station η and ϕ indices. The white horizontal band at $\eta = 0$ represents the physical separation between the two barrel halves, while isolated white cells indicate masked channels or regions with insufficient statistics.

using Eq. 4.7. The uncalibrated and calibrated values are therefore:

$$\beta_i^{\text{uncal}} = \frac{1}{1 + \dfrac{ct_{0_i}^{\text{uncal}}}{L_i}}, \qquad \beta_i^{\text{cal}} = \frac{1}{1 + \dfrac{ct_{0_i}^{\text{cal}}}{L_i}} \tag{4.14}$$

Each selected slow-muon candidate traverses several RPC detector elements, providing multiple timing measurements along the reconstructed trajectory. Although the particle has a single physical velocity, the individual hit measurements yield different values of β_i as a consequence of the finite timing resolution of the RPC strips and residual detector fluctuations. The hit-level measurements are therefore combined into a single velocity estimate for each muon candidate.

The uncertainty on each β_i measurement is obtained by propagating the timing uncertainty of the corresponding RPC strip. From Eq. 4.7:

$$\frac{\partial \beta_i}{\partial t_{0_i}} = -\frac{c}{L_i}\beta_i^2 \tag{4.15}$$

so that the corresponding uncertainty is:

$$\sigma_{\beta_i} = \frac{c}{L_i}\beta_i^2 \sigma_{t_i} \tag{4.16}$$

where σ_{t_i} is the timing resolution assigned to the strip associated with the i-th RPC hit.

The individual measurements are combined using inverse-variance weighting. The use of the inverse-variance weighted estimator follows from minimizing the corresponding χ^2 function; the derivation is given in Appendix C. For either the uncalibrated or calibrated set of measurements, the weighted average is given by:

$$\overline{\beta} = \frac{\sum_i \beta_i w_i}{\sum_i w_i}, \qquad w_i = \frac{1}{\sigma_{\beta_i}^2} \tag{4.17}$$

The same procedure is applied independently to the uncalibrated and calibrated hit measurements, resulting in one uncalibrated and one calibrated β value for each selected slow-muon candidate. These values are subsequently accumulated into the corresponding global β distributions.

4.1.7 Timing calibration performance and results

The performance of the RPC timing calibration is assessed by comparing the global β distributions obtained before and after the application of the strip-dependent timing corrections. Since the reference muons are expected to be relativistic, the calibrated distribution should be centered close to $\beta = 1$.

The effect of the timing calibration can be understood directly from the relation between the time residual and the reconstructed velocity. For a relativistic muon, the ideal time-of-flight residual is zero, so that $t_0 = 0$ corresponds to $\beta = 1$. A strip-dependent timing offset introduces a systematic displacement in the measured t_0 values and consequently shifts the reconstructed β values away from unity. Subtracting the strip-dependent calibration constant removes this systematic contribution and moves the corresponding measurements towards the expected relativistic value.

The calibration also affects the width of the global β distribution. Before calibration, the combination of measurements from different RPC strips with different timing offsets contributes an additional spread to the reconstructed velocity values. This strip-dependent contribution is not an intrinsic fluctuation of the particle velocity, but a detector-induced effect arising from differences in the timing response of the individual strips. Once the corresponding offsets are removed on a strip-by-strip basis, this contribution to the spread is reduced, while the remaining width is determined by the intrinsic timing resolution of the RPC measurements and by residual detector effects. The calibrated distribution is therefore expected to become both more closely centered around $\beta = 1$ and narrower than the uncalibrated distribution, provided that strip-to-strip timing offsets constitute a significant contribution to the pre-calibration spread. The resulting distributions and their corresponding fits are shown in Figure 4.6.

The calibrated β distribution is fitted with a Voigt profile in order to describe both its approximately Gaussian central component and its non-Gaussian tails. The same functional form and fit range are used for the uncalibrated distribution, allowing the two distributions to be compared consistently. The Gaussian parameter σ_{G} characterizes the width of the Gaussian component of the fitted resolution model, while Γ_{L} characterizes the full width at half maximum of the Lorentzian component. In the present β analysis, these parameters are used as descriptive components of an empirical model of the reconstructed distribution rather than as independent measurements of distinct physical detector effects.

To summarize the overall width of the reconstructed β distribution, the effective resolution is defined as $\sigma_{\mathrm{eff}} = \frac{\mathrm{FWHM}_{\mathrm{Voigt}}}{2\sqrt{2\ln 2}}$. The numerical factor $2\sqrt{2\ln 2} \simeq 2.35482$ follows from the relation between the FWHM and the standard deviation of a Gaussian distribution. Therefore, σ_{eff} can be interpreted as the Gaussian-equivalent standard deviation corresponding to the total FWHM of the fitted Voigt profile. It is not a directly fitted parameter, but a single-number measure of the overall width of the reconstructed β distribution that incorporates the contributions of both the Gaussian core and the Lorentzian-like tails.

The relative effective resolution is defined as $R_{\mathrm{eff}} = \frac{\sigma_{\mathrm{eff}}}{\mu_\beta}$, where μ_β is the fitted central value of the corresponding distribution. Expressed as a percentage, this dimensionless quantity measures the width of the reconstructed velocity distribution relative to its central value. A smaller R_{eff} therefore corresponds to a smaller fractional spread and hence to a more precise velocity reconstruction.

The residual displacement of the reconstructed velocity from the expected relativistic value is quantified by $\Delta\beta = \mu_\beta - 1$. A value of $\Delta\beta$ closer to zero indicates a smaller residual bias. The fitted central value and the width of the distributions can therefore be used together to quantify the effect of the timing calibration. The main fit parameters and derived quantities before and after calibration are summarized in Table 4.1.

The calibration therefore reduces the absolute displacement from the expected value $\beta = 1$ from $|\Delta\beta^{\mathrm{uncal}}| = 0.055542$ to $|\Delta\beta^{\mathrm{cal}}| = 0.007401$, corresponding to a reduction of approximately 86.67%. At the same time, the effective resolution decreases from 0.037854 to 0.026161, corresponding to a relative reduction of 30.89%. The relative effective resolution improves from 3.586% to 2.636%. These results show that the strip-dependent timing calibration substantially reduces the residual bias in the reconstructed velocity and narrows the global β distribution, bringing the measurement closer to the expected relativistic value and improving the precision of the velocity reconstruction.

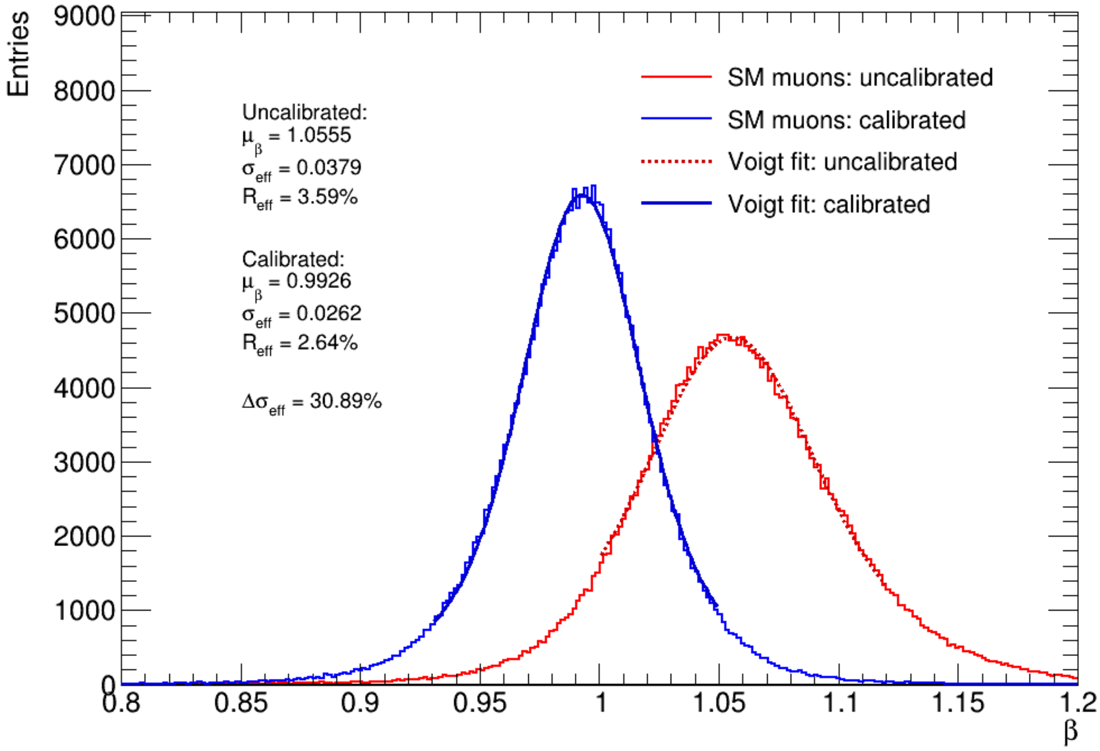

Figure 4.6: Comparison of the global β distributions obtained before and after the RPC strip timing calibration. The uncalibrated and calibrated distributions are shown together with the corresponding Voigt fits. The Gaussian parameter σ_{G} describes the width of the Gaussian component of the empirical resolution model, while Γ_{L} represents the full width at half maximum of the Lorentzian component and parametrizes the non-Gaussian tails. For the calibrated distribution, the fitted central value is $\mu_\beta = 0.9926$, close to the expected relativistic value $\beta = 1$. The fitted central values are compared quantitatively in the text to assess the residual timing bias before and after calibration.

4.2 Sensitivity Study

The timing calibration described in the previous section aims to improve the velocity reconstruction of charged particles in the RPC. To evaluate the impact of this calibrated β measurement on identifying heavy, slow-moving charged particles, a dedicated sensitivity study was performed. This analysis qualitatively assesses whether combining RPC ToF information with the ionization energy loss[1] dE/dx measured in the ID provides

[1]As a reminder, these two observables are particularly useful for HSCP searches because slowly moving charged particles are expected to exhibit both a delayed time of flight and an enhanced ionization energy loss with respect to relativistic SM particles.

Distribution	$\boldsymbol{\sigma_{\mathrm{G}}}$	$\boldsymbol{\Gamma_{\mathrm{L}}}$	$\boldsymbol{\mu_\beta}$	$\boldsymbol{\Delta\beta}$	$\boldsymbol{\sigma_{\mathrm{eff}}}$	$\boldsymbol{R_{\mathrm{eff}}}$
Uncalibrated	0.022875	0.055132	1.055542	+0.055542	0.037854	3.586%
Calibrated	0.018205	0.030707	0.992599	−0.007401	0.026161	2.636%

Table 4.1: Summary of the fitted β-distribution parameters before and after the strip-by-strip timing calibration. The Gaussian and Lorentzian contributions are characterized by σ_{G} and Γ_{L}, respectively. The fitted central value μ_β is reported together with its deviation from the nominal value, $\Delta\beta$. The effective resolution σ_{eff} combines the two contributions, while the relative effective resolution R_{eff} is expressed as a percentage of the fitted central value. The calibration substantially reduces the Lorentzian contribution and, consequently, improves the effective resolution from 3.586% to 2.636%, while shifting the fitted central value from 1.055542 to 0.992599.

sufficient discrimination between a supersymmetric stau signal and the SM background.

Rather than determining a final experimental exclusion limit, this study investigates how sensitivity scales when varying the selection criteria for the reconstructed particle velocity. The β requirement is scanned over a predefined range across several stau mass hypotheses. For each threshold, a significance estimator is calculated from the surviving signal and background yields, identifying the β cut that maximizes the expected significance as the representative working point.

The optimization is performed over the complete set of predefined β thresholds, without imposing a minimum number of retained signal MC candidates. Consequently, the estimated signal efficiency at the most restrictive thresholds may be affected by the finite size of the Monte Carlo dataset. Together with the limited statistics of the data-driven background sample, this motivates interpreting the resulting optimum as a representative working point rather than as a precisely optimized value.

4.2.1 Background and signal samples

The background estimate is obtained from a 2022 ATLAS data sample: this input file is an already preselected and event-picked[2] ntuple. The present analysis does not repeat the upstream derivation and event-picking procedure, but applies the track-level requirements needed to define the candidate sample for the sensitivity study: the dataset entries are subsequently analyzed at the track-candidate level.

The signal samples consist of simulated stau ($\tilde{\tau}$) events generated assuming the staus to be stable on the detector scale ($c\tau \to \infty$). Four mass hypotheses are considered, $m_{\tilde{\tau}} = 300, 500, 700, 800$ GeV, in order to evaluate how the expected sensitivity varies as a function of the stau mass.

The integrated luminosity used for the sensitivity study is $\mathcal{L} = 31.4\ \mathrm{fb}^{-1}$, and the production cross sections used are reported in Table 4.2.

[2]Event picking refers to the technical procedure of extracting a specific subset of collision events from primary datasets by querying a predefined list of event identifiers, such as Run and Event numbers.

Stau mass [GeV]	**Cross section [pb]**
300	6.254×10^{-3}
500	6.736×10^{-4}
700	1.235×10^{-4}
800	5.863×10^{-5}

Table 4.2: Signal cross sections used for the four stau mass hypotheses considered in the sensitivity study. Values taken from Ref. [43].

4.2.2 Data-driven background estimate with the ABCD method

The SM background in the signal region is estimated directly from data using an ABCD-type method (explained in more detail in Appendix D). Data-driven ABCD approaches exploit two observables that are approximately independent for background events. The two-dimensional phase space is divided into four regions by applying a threshold to each observable: under the factorization assumption, the background population of the signal region can be inferred from the other three regions.

In the present analysis, the two observables are the calibrated reconstructed velocity β and the pixel ionization energy loss dE/dx. The thresholds are fixed to $\beta = 1$ and $dE/dx = 1.8$.

The convention adopted in this analysis is:

$$A: \quad \beta \geq 1, \qquad dE/dx > 1.8 \tag{4.18}$$

$$B: \quad \beta \geq 1, \qquad dE/dx \leq 1.8 \tag{4.19}$$

$$C: \quad \beta < 1, \qquad dE/dx \leq 1.8 \tag{4.20}$$

$$D: \quad \beta < 1, \qquad dE/dx > 1.8 \tag{4.21}$$

Region D is the signal-enhanced region because it combines the requirement of a velocity below the relativistic value with an ionization energy loss above the chosen threshold. Regions A, B, and C are used to determine the background expectation:

$$N_D = N_C \frac{N_A}{N_B}, \tag{4.22}$$

where N_A/N_B is the transfer factor (TF).

The region C β distribution is stored as a histogram for visualization and for constructing the predicted region D shape. For the numerical β-cut scan, however, the background yield is obtained by directly counting the individual region C candidates satisfying $\beta < \beta_{\text{cut}}$, and multiplying this count by the ABCD transfer factor. This avoids introducing a dependence of the optimization on the histogram binning.

4.2.3 Signal selection and expected signal yield

The signal candidates are selected using the same baseline requirements applied to the data sample. Following these baseline selections, the surviving candidates are restricted to Region D, corresponding to the signal-enhanced region of the ABCD framework. The corresponding β distributions are stored separately for each stau mass hypothesis.

No additional β threshold is fixed at this stage, since the purpose of the sensitivity study is precisely to determine a representative working point by scanning this requirement.

For a given threshold β_{cut}, the number of surviving signal candidates, $N_S^{\text{pass}}(\beta_{\text{cut}})$, is obtained by directly counting the signal candidates in region D whose reconstructed velocity satisfies $\beta < \beta_{\text{cut}}$. The counting is performed using the individual reconstructed β values rather than the histogram binning:

$$N_S^{\text{pass}}(\beta_{\text{cut}}) = \sum_{i \in S_D} \mathrm{I}(\beta_i < \beta_{\text{cut}}) \tag{4.23}$$

where $\mathrm{I}(\cdot)$ is the indicator function, equal to 1 if the condition is satisfied and 0 otherwise. In the practical implementation, this is evaluated as the discrete integral of the unbinned candidate distribution.

To evaluate the discrimination power of the velocity requirement, the β-cut efficiency (ϵ_β) is defined as the fraction of baseline signal candidates in Region D that survive this additional velocity selection:

$$\epsilon_\beta(\beta_{\text{cut}}) = \frac{N_S^{\text{pass}}(\beta_{\text{cut}})}{N_{\text{baseline, D}}^{\text{tracks}}} \tag{4.24}$$

where $N_{\text{baseline, D}}^{\text{tracks}}$ is the number of baseline signal candidates entering Region D prior to any β cut.

Finally, the overall signal yield factor per generated event, Y_{total}, is defined relative to the total number of generated Monte Carlo events ($N_{\text{MC}}^{\text{events}}$):

$$Y_{\text{total}} = \frac{N_S^{\text{pass}}(\beta_{\text{cut}})}{N_{\text{MC}}^{\text{events}}} \tag{4.25}$$

This quantity corresponds directly to the candidate yield per input simulated event, which is then used to scale the theoretical production rate to obtain the final expected signal yield.

This construction clarifies the two distinct roles played by candidate yields and selection efficiencies in this analysis. The overall yield factor can be factorized as $Y_{\text{total}} = Y_{\text{baseline}} \times \epsilon_\beta$, where $Y_{\text{baseline}} = N_{\text{baseline, D}}^{\text{tracks}} / N_{\text{MC}}^{\text{events}}$ represents the baseline candidate yield in Region D per simulated event. Thus, the final expected physical signal yield is given by:

$$N_S^{\text{expected}} = \sigma_{\tilde{\tau}} \, \mathcal{L} \, Y_{\text{baseline}} \, \epsilon_\beta \tag{4.26}$$

where ϵ_β isolates the additional rejection power introduced by the scanned velocity requirement. Because ϵ_β strictly measures the performance of the velocity selection, it is used as the signal efficiency metric in the ROC curves presented in the following sections.

4.2.4 Scan of the β selection and sensitivity methodology

This scan provides a direct comparison between the amount of signal retained and the corresponding background contamination as the velocity requirement is progressively loosened. Since the quantities shown in the figure are cumulative yields, increasing β_{cut} includes progressively more candidates in the selected sample. At low values of the threshold, only the slowest candidates are retained and both signal and background yields are small. As the threshold is increased, the signal yield initially rises and eventually approaches a plateau when most of the signal candidates in the signal-enhanced region have already been included. The background yield, in contrast, continues to increase towards $\beta = 1$, where a much larger population of Standard Model candidates is present.

Figure 4.7 illustrates this behavior for the four stau mass hypotheses considered in the study. The background curve corresponds to the ABCD prediction from Eq. 4.22, while the signal curves represent the expectations obtained from Eq. 4.26. At each scan point, both the signal and background cumulative yields are evaluated using the reconstructed β values. Overall, these curves show how the choice of the β threshold affects signal retention and background rejection across the scanned velocity range.

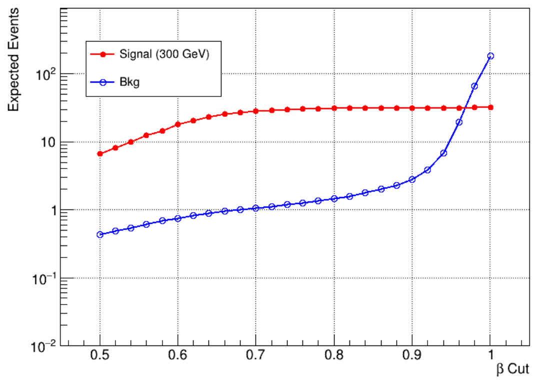

(a) $m_{\tilde{\tau}} = 300$ GeV

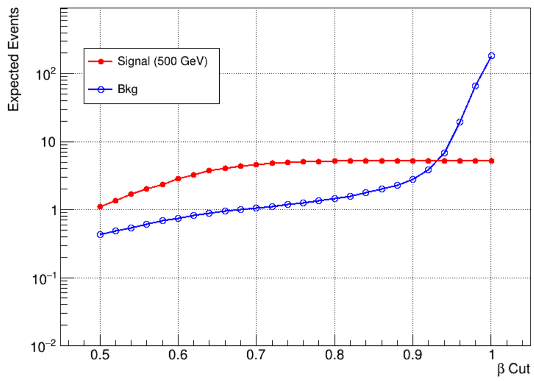

(b) $m_{\tilde{\tau}} = 500$ GeV

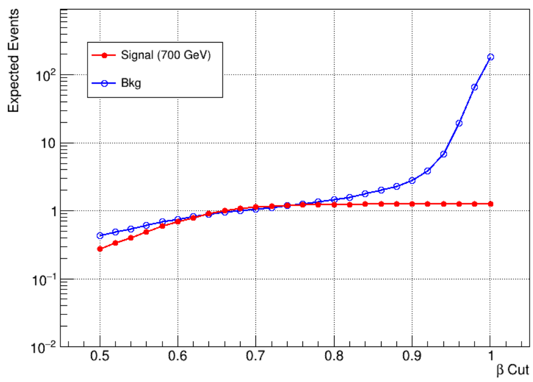

(c) $m_{\tilde{\tau}} = 700$ GeV

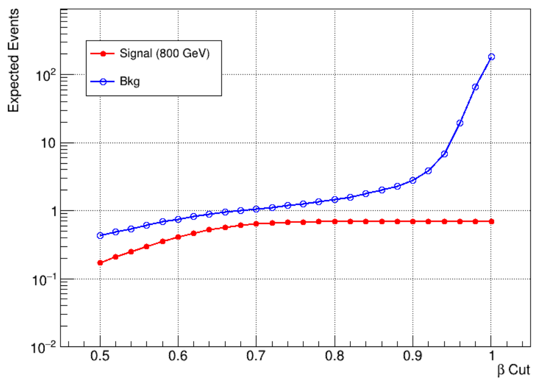

(d) $m_{\tilde{\tau}} = 800$ GeV

Figure 4.7: Expected signal and ABCD-predicted background yields as a function of the upper β selection threshold for the four stau mass hypotheses considered in the study. The scan is performed in steps of $\Delta\beta = 0.02$. The curves are cumulative, so each point represents the total number of candidates retained below the corresponding β threshold.

To quantify the discrimination power of each β threshold, a significance estimator, Σ, is calculated. It is defined using the expected signal yield N_S^{exp}, the predicted background yield N_B^{pred}, and the relative systematic uncertainty on the background prediction $\delta_B = 0.10$ (which corresponds to 10%):

$$\Sigma = \text{BinomialExpZ}\left(N_S^{\text{exp}}, N_B^{\text{pred}}, \delta_B\right) \tag{4.27}$$

This calculation is implemented through the ROOT function[3]:

`RooStats::NumberCountingUtils::BinomialExpZ`

The RooStats utility computes an expected background-only p-value for the counting experiment and converts it into the corresponding one-sided Gaussian-equivalent significance:

$$\Sigma = \Phi^{-1}(1-p) \tag{4.28}$$

where Φ denotes the cumulative distribution function of the standard normal distribution. This provides a convenient single-number metric for comparing the different β working points, which are then chosen as the β thresholds that maximize this significance:

$$\beta_{\mathrm{opt}} = \arg\max_{\beta_{\mathrm{cut}}} \Sigma(\beta_{\mathrm{cut}}) \tag{4.29}$$

Since the β requirement is scanned over a discrete set of thresholds and the background estimate is based on a finite data sample, the optimization is inherently discrete and statistically limited: for this reason, the resulting maximum should be interpreted as a qualitative representative working point rather than as a precision determination of an optimal selection threshold.

4.2.5 Signal-to-background shape discrimination and ROC analysis

Before considering the numerical sensitivity values, the separation between signal and background is investigated by comparing their normalized β distributions. For each stau mass hypothesis, the ABCD-predicted background distribution and the corresponding signal distribution in Region D are independently normalized to unit area: this normalization removes the dependence on the absolute event yields, allowing the shapes of the two distributions to be compared directly. The resulting distributions are shown on a logarithmic vertical scale in Figure 4.8.

On the other hand, a ROC-type analysis characterizes the intrinsic separation between the signal and background β distributions. In standard classification terminology, the background rejection probability is given by $1-\epsilon_B$. However, in this analysis the vertical axis of the ROC plots displays the background rejection factor, defined as $1/\epsilon_B(\beta_{\mathrm{cut}})$: this quantity quantifies the factor by which the background is suppressed[4]. To summarize the overall separation power, the Area Under the ROC Curve (AUC)[5] is evaluated in the conventional (ϵ_B, ϵ_S) plane:

$$\mathrm{AUC} = \int_0^1 \epsilon_S \, d\epsilon_B \tag{4.30}$$

[3]This function takes three arguments, $(s_{\mathrm{exp}}, b_{\mathrm{exp}}, \delta_B)$, where s_{exp} is the expected signal yield, b_{exp} is the expected background yield, and δ_B is the fractional uncertainty on the background. ROOT documents `BinomialExpZ` as the expected significance corresponding to the expected p-value for the background-only hypothesis when the background uncertainty is supplied as a fractional uncertainty [44].

[4]For instance, $R_B = 100$ corresponds to $\epsilon_B = 1\%$, meaning only one candidate in one hundred survives. Plotting ϵ_S versus $1/\epsilon_B$ on a logarithmic scale makes large background suppression factors immediately visible in the low-β regime.

[5]An AUC near 1.0 indicates strong shape separation, while an AUC of 0.5 corresponds to random classification.

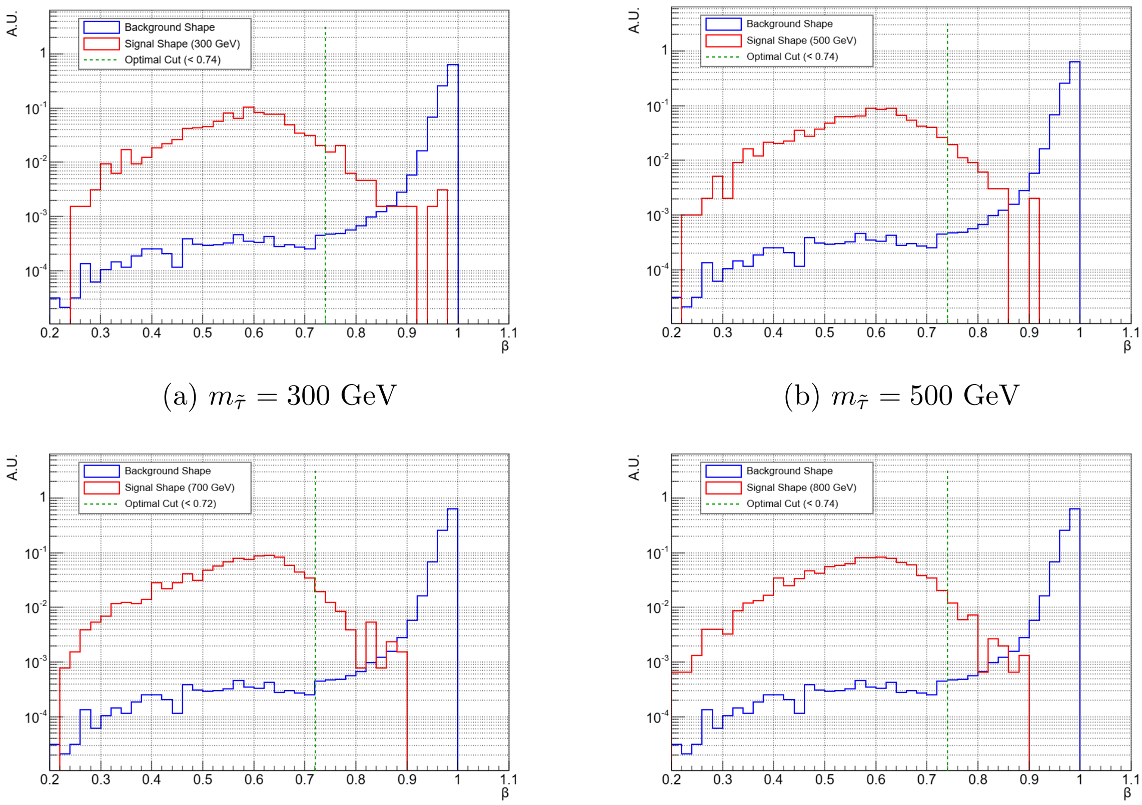

(a) $m_{\tilde{\tau}} = 300$ GeV

(b) $m_{\tilde{\tau}} = 500$ GeV

(c) $m_{\tilde{\tau}} = 700$ GeV

(d) $m_{\tilde{\tau}} = 800$ GeV

Figure 4.8: Normalized β distributions for the ABCD-predicted background and the stau signal for the four mass hypotheses. The distributions are shown on a logarithmic vertical scale. The vertical dashed line indicates the representative optimal β working point selected by the sensitivity scan.

computed numerically across discrete working points using the trapezoidal rule. The ROC curves for the four stau mass hypotheses are shown in Figure 4.9.

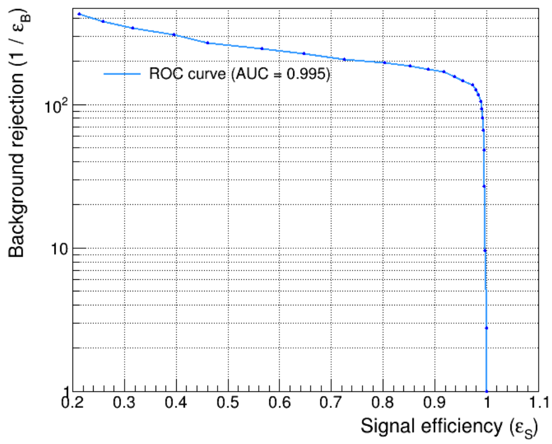

(a) $m_{\tilde{\tau}} = 300$ GeV

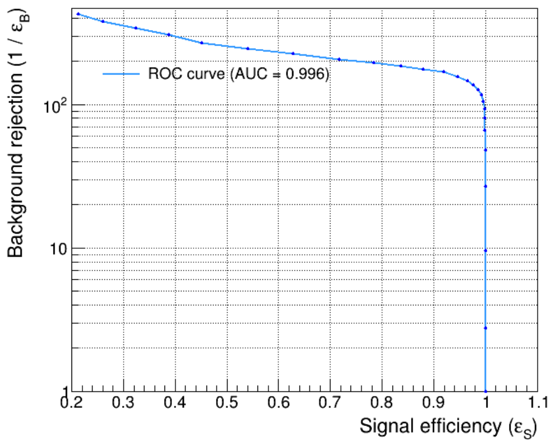

(b) $m_{\tilde{\tau}} = 500$ GeV

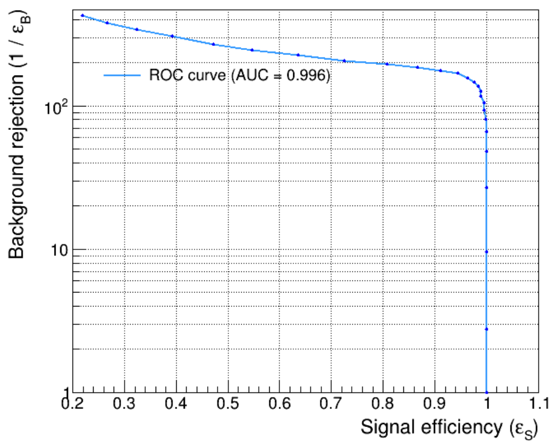

(c) $m_{\tilde{\tau}} = 700$ GeV

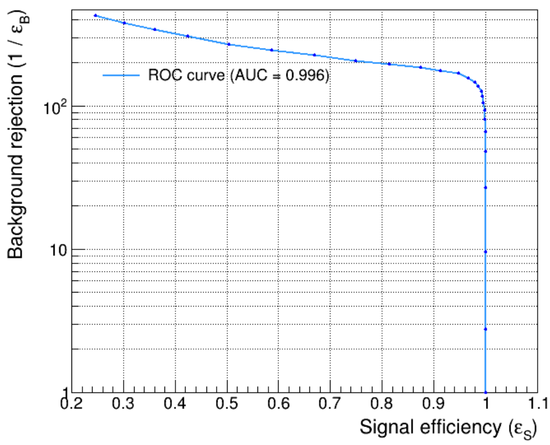

(d) $m_{\tilde{\tau}} = 800$ GeV

Figure 4.9: Signal efficiency versus background rejection factor $(1/\epsilon_B)$ for the four stau mass hypotheses. The AUC is calculated from the corresponding conventional ROC representation in the (ϵ_B, ϵ_S) plane.

4.2.6 Results and Interpretation

The resulting working points and corresponding yields are summarised in Table 4.3. The optimal β_{cut} values are found to lie in the narrow range $0.72 \leq \beta_{\text{cut}} \leq 0.74$ for all four stau mass hypotheses. This indicates that the optimal selection is relatively stable across the considered mass range.

Care is needed when comparing the different signal hypotheses. Since the background distribution is common to all signal models, whereas the signal β distributions depend inherently on the stau mass, different signal hypotheses can maximize the same significance at different discrete values of β_{cut}. However, due to the discrete binning of

$m_{\tilde{\tau}}$ **[GeV]**	β_{cut}	N_C^{real}	N_S^{real}	N_S^{exp}	N_B^{pred}	Σ	**AUC**
300	< 0.74	614	604	29.65	1.18	11.494	0.995
500	< 0.74	614	934	4.94	1.18	3.011	0.996
700	< 0.72	571	1215	1.18	1.10	0.734	0.996
800	< 0.74	614	1464	0.67	1.18	0.307	0.996

Table 4.3: Summary of the sensitivity study for the four stau mass hypotheses. N_C^{real} denotes the number of real background candidates in region C satisfying the selected β requirement, while N_S^{real} denotes the number of simulated signal candidates passing the corresponding selection. N_S^{exp} is the expected signal-candidate yield obtained after scaling the simulated sample according to the theoretical production cross section and integrated luminosity. N_B^{pred} is the corresponding ABCD background prediction. The significance Σ is evaluated with the RooStats BinomialExpZ estimator assuming a 10% relative uncertainty on the background. The AUC values quantify the overall shape-separation power of the reconstructed β observable within region D.

the scan, some of the signal hypotheses are found to reach their optimal significance at the same β_{cut} threshold. In such cases, the corresponding values of N_C^{real} and N_B^{pred} are necessarily identical, since the same background selection and the same ABCD transfer factor are applied. The repeated background yields observed for the 300 GeV, 500 GeV, and 800 GeV hypotheses are therefore an expected consequence of the common data-driven background sample, rather than an indication of identical signal behaviour.

The main variation across the mass hypotheses is observed in the absolute signal yield and, consequently, in the expected sensitivity. The expected signal yield decreases from 29.65 candidates for the 300 GeV hypothesis to 0.67 candidates for the 800 GeV hypothesis. Over the same mass range, the corresponding expected significance decreases from $\Sigma = 11.494$ to $\Sigma = 0.307$, as shown in Figure 4.10. This loss of absolute sensitivity is primarily driven by the decrease of the theoretical stau production cross section with increasing mass, together with the mass dependence of the signal acceptance and of the efficiency of the selected β requirement. Thus, even when the selection remains effective in isolating slow candidates, the rapidly decreasing signal production rate limits the number of expected signal events at high mass.

The AUC provides a complementary view of the performance of the reconstructed velocity observable. Its value remains approximately constant (AUC $\simeq$ 0.995–0.996) across all four signal hypotheses. The β observable therefore retains a very strong relative shape-separation power throughout the considered mass range. In particular, even the low-mass hypotheses, which have β values closer to the relativistic background, do not exhibit a significant degradation in the ability of the observable to distinguish signal from background at the distribution-shape level.

This apparent contrast between the nearly constant AUC and the strongly decreasing significance is a consequence of the different information encoded by the two metrics. The AUC characterises the separation between the normalised signal and background β distributions and is therefore largely insensitive to the absolute signal production rate. The expected significance, on the other hand, depends directly on the number of signal

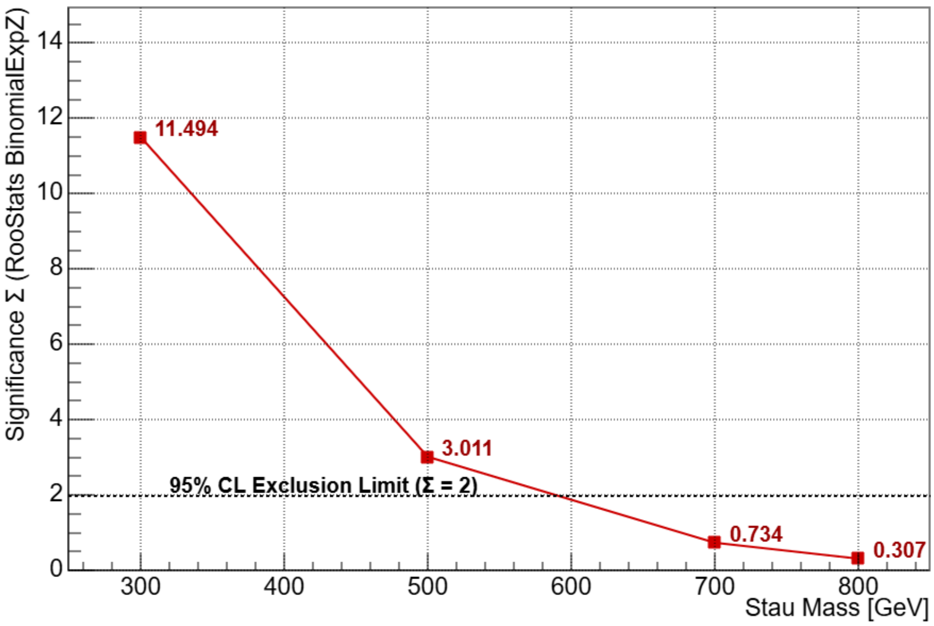

Figure 4.10: Maximum expected significance obtained in the β-cut scan as a function of the stau mass. Each point corresponds to the representative working point that maximises the RooStats BinomialExpZ estimator for the corresponding signal hypothesis. The horizontal $\Sigma = 2$ line is shown as a reference for a significance corresponding to approximately 95% confidence level.

candidates expected after selection, as well as on the corresponding background yield and its uncertainty. A high AUC can therefore coexist with a low absolute sensitivity when the signal production rate becomes small.

Conclusions

This thesis presented a dedicated strip-by-strip timing calibration of the ATLAS Resistive Plate Chambers (RPC) using 2022 pp collision data ($\sqrt{s} = 13.6$ TeV, 31.4 fb^{-1}), evaluating its performance and impact on searches for Heavy Long-Lived Charged Particles (HSCPs) via a supersymmetric stau ($\tilde{\tau}$) benchmark model.

Applying the derived strip-level timing offsets significantly enhanced the velocity (β) reconstruction quality across the detector:

- **Bias reduction:** The residual displacement $|\Delta\beta| = |\mu_\beta - 1|$ was reduced from 0.055542 to 0.007401, representing an 86.67% reduction in central timing bias.

- **Resolution improvement:** The effective resolution σ_{eff} decreased from 0.037854 to 0.026161 (a 30.89% relative reduction), improving the relative effective resolution R_{eff} from 3.586% to 2.636%.

- **Tail suppression:** The Lorentzian width Γ_{L} of the fitted Voigt profile dropped from 0.055132 to 0.030707, confirming the successful removal of dominant non-Gaussian timing spreads caused by channel-to-channel delays.

The impact of the calibrated measurement was subsequently evaluated in a sensitivity study combining RPC time-of-flight (ToF) with pixel ionization energy loss (dE/dx) within a data-driven ABCD background framework. A discrete scan over candidate β_{cut} thresholds across four stau mass hypotheses ($m_{\tilde{\tau}} = 300, 500, 700, 800$ GeV) evaluated the trade-off between signal efficiency and background rejection.

Taken together, these results demonstrate that the calibrated RPC timing information provides a powerful discriminating observable for the identification of slow, highly ionising charged-particle candidates. The reconstructed velocity β yields optimal working points in the narrow range $0.72 \leq \beta_{\text{cut}} \leq 0.74$, while maintaining Area Under the ROC Curve (AUC) values close to unity (AUC $\simeq$ 0.995–0.996) for all the investigated signal masses. The strong separation power is therefore not restricted to a specific stau mass hypothesis. Rather, the variation in sensitivity across the mass range is predominantly associated with the decreasing absolute signal yield, while the underlying discriminating capability of the calibrated timing observable remains essentially unchanged.

Specifically, three key insights emerge from the performance study:

1. **Mass-independent shape discrimination:** The consistently high AUC values (~ 0.996) prove that the shape separation between signal and background β distributions remains robust even at lower masses ($m_{\tilde{\tau}} = 300$ GeV), where particle speeds approach the relativistic regime.

2. **Decoupling of signal yield and observable power:** The steep drop in expected significance (from $\Sigma = 11.494$ at $m_{\tilde{\tau}} = 300$ GeV to $\Sigma = 0.307$ at $m_{\tilde{\tau}} = 800$ GeV) is governed by the rapid fall in theoretical production cross section (from 6.25×10^{-3} pb to 5.86×10^{-5} pb) and candidate acceptance, rather than any loss in timing performance.

3. **Stability of optimal working points:** The tight alignment of optimal selection thresholds ($\beta_{\text{cut}} \approx 0.73$) across all mass hypotheses could support the implementation of a "mass-agnostic" velocity cut in HSCP searches.

These findings establish a solid foundation for future long-lived particle searches in ATLAS:

- **Full Run 3 analysis:** Extending the strip-by-strip calibration framework to the full Run 3 dataset will boost luminosity, directly probing high-mass regimes ($m_{\tilde{\tau}} \geq 700$ GeV) limited by low event yields.

- **Multi-subsystem combination:** Integrating RPC ToF and pixel dE/dx with timing from MDT or Tile Calorimeter systems could further refine velocity resolution and maximize background rejection.

- **Preparation for HL-LHC:** High-precision time-of-flight measurements with sub-nanosecond accuracy will remain indispensable for suppressing out-of-time pileup background during High-Luminosity LHC operations.

Appendix A

Invariant mass

The invariant mass of a system of particles is a Lorentz-invariant quantity that can be reconstructed from the measured four-momenta of the final-state particles. In the case of a resonance decaying into two muons, a peak in the reconstructed invariant-mass distribution is evidence that the two muons are produced in the decay of a common parent particle. The position of the peak provides an estimate of the parent-particle mass, while its observed width receives contributions from both the intrinsic width of the resonance and the experimental detector resolution. Therefore, for a resonance with a non-negligible natural width, the observed width cannot be identified directly with the detector resolution alone. In the case of the Z boson, this distinction motivates the use of a Voigtian model, as discussed in Appendix B.

A.1 Invariant mass from the four-momentum

The four-momentum of a particle is written as:

$$p^\mu = (E, \mathbf{p}) \tag{A.1}$$

where E is the particle energy and $\mathbf{p}$ is its three-momentum. In natural units, with $c = 1$, the invariant mass m is defined by the Lorentz-invariant relation:

$$p^\mu p_\mu = E^2 - \mathbf{p}^2 = m^2 \tag{A.2}$$

For a system of two particles with four-momenta p_1^μ and p_2^μ, the total four-momentum is:

$$P^\mu = p_1^\mu + p_2^\mu \tag{A.3}$$

The invariant mass of the two-particle system is therefore defined as:

$$m_{\mu\mu}^2 = P^\mu P_\mu = (p_1 + p_2)^2 \tag{A.4}$$

Expanding the four-vector product gives:

$$m_{\mu\mu}^2 = p_1^2 + p_2^2 + 2p_1 \cdot p_2 \tag{A.5}$$

Since each muon satisfies:

$$p_i^2 = m_\mu^2 \tag{A.6}$$

the exact expression is:

$$m_{\mu\mu}^2 = 2m_\mu^2 + 2\left(E_1 E_2 - \mathbf{p}_1 \cdot \mathbf{p}_2\right) \tag{A.7}$$

A.2 Ultra-relativistic approximation

For a relativistic particle:

$$E_i^2 = |\mathbf{p_i}|^2 + m_i^2 \tag{A.8}$$

When the particle energy is much larger than its rest mass, $E_i \gg m_\mu$, the mass term can be neglected with respect to the momentum term, giving:

$$E_i \simeq |\mathbf{p}_i| \tag{A.9}$$

For the muons considered in the present analysis, this approximation allows the exact expression in Eq. A.7 to be simplified to:

$$m_{\mu\mu}^2 \simeq 2\left(E_1 E_2 - \mathbf{p}_1 \cdot \mathbf{p}_2\right) \tag{A.10}$$

The three-momenta are conveniently expressed in terms of transverse and longitudinal components. In cylindrical coordinates around the beam axis:

$$\mathbf{p_i} = (p_{x_i}, p_{y_i}, p_{z_i}) \tag{A.11}$$

with

$$p_{\mathrm{T}_i} = \sqrt{p_{x_i}^2 + p_{y_i}^2} \tag{A.12}$$

and the azimuthal angle ϕ_i defined by

$$p_{x_i} = p_{\mathrm{T}_i} \cos\phi_i, \qquad p_{y_i} = p_{\mathrm{T}_i} \sin\phi_i \tag{A.13}$$

The scalar product of the transverse components is then:

$$\mathbf{p}_{\mathrm{T}_1} \cdot \mathbf{p}_{\mathrm{T}_2} = p_{\mathrm{T}_1} p_{\mathrm{T}_2} \left(\cos\phi_1 \cos\phi_2 + \sin\phi_1 \sin\phi_2\right) = p_{\mathrm{T}_1} p_{\mathrm{T}_2} \cos(\phi_1 - \phi_2) \tag{A.14}$$

where we define

$$\Delta\phi = \phi_1 - \phi_2 \tag{A.15}$$

The longitudinal contribution is described in terms of the pseudorapidity η. For a particle with three-momentum magnitude $p = |\mathbf{p}|$, the pseudorapidity is defined as:

$$\eta = \frac{1}{2} \ln\left(\frac{p + p_z}{p - p_z}\right) \tag{A.16}$$

Solving Eq. A.16 for p_z gives:

$$e^{2\eta} = \frac{p + p_z}{p - p_z} \tag{A.17}$$

and therefore

$$e^{2\eta}(p - p_z) = p + p_z,\ \ p_z(e^{2\eta} + 1) = p(e^{2\eta} - 1)$$

so that

$$\frac{p_z}{p} = \frac{e^{2\eta} - 1}{e^{2\eta} + 1} = \tanh\eta \tag{A.18}$$

Hence:

$$p_z = p \tanh\eta \tag{A.19}$$

The transverse momentum is related to the total momentum by:

$$p_{\mathrm{T}}^2 = p^2 - p_z^2 \tag{A.20}$$

Using Eq. A.19:

$$p_{\mathrm{T}}^2 = p^2 \left(1 - \tanh^2 \eta\right) \ = p^2 \mathrm{sech}^2 \eta$$

Therefore:

$$p = p_{\mathrm{T}} \cosh \eta \tag{A.21}$$

and consequently

$$p_z = p_{\mathrm{T}} \sinh \eta \tag{A.22}$$

Combining Eq. A.9 with Eq. A.21 yields the requested relation

$$E \simeq p_{\mathrm{T}} \cosh \eta \tag{A.23}$$

The three-dimensional scalar product can now be written as the sum of its transverse and longitudinal contributions:

$$\mathbf{p}_1 \cdot \mathbf{p}_2 = \mathbf{p}_{\mathrm{T}_1} \cdot \mathbf{p}_{\mathrm{T}_2} + p_{z_1} p_{z_2} \tag{A.24}$$

Using Eqs. A.14 and A.22:

$$\mathbf{p}_1 \cdot \mathbf{p}_2 = p_{\mathrm{T}_1} p_{\mathrm{T}_2} \left[\cos \Delta\phi + \sinh \eta_1 \sinh \eta_2\right] \tag{A.25}$$

Substituting the ultra-relativistic energy relation from Eq. A.23 into Eq. A.10 gives:

$$m_{\mu\mu}^2 \simeq 2 p_{\mathrm{T}_1} p_{\mathrm{T}_2} \left[\cosh \eta_1 \cosh \eta_2 - \cos \Delta\phi - \sinh \eta_1 \sinh \eta_2\right] \tag{A.26}$$

Using the hyperbolic identity

$$\cosh(a - b) = \cosh a \cosh b - \sinh a \sinh b \tag{A.27}$$

the longitudinal terms become:

$$\cosh \eta_1 \cosh \eta_2 - \sinh \eta_1 \sinh \eta_2 = \cosh(\eta_1 - \eta_2) \tag{A.28}$$

Defining

$$\Delta\eta = \eta_1 - \eta_2 \tag{A.29}$$

the invariant mass therefore takes the form

$$m_{\mu\mu}^2 = 2 p_{\mathrm{T}_1} p_{\mathrm{T}_2} \left[\cosh(\Delta\eta) - \cos(\Delta\phi)\right] \tag{A.30}$$

Taking the positive square root gives the final expression used in the analysis:

$$\boxed{m_{\mu\mu} = \sqrt{2 p_{\mathrm{T}_1} p_{\mathrm{T}_2} \left[\cosh(\Delta\eta) - \cos(\Delta\phi)\right]}} \tag{A.31}$$

This expression is particularly convenient for collider analyses because it depends only on the transverse momenta, pseudorapidities, and azimuthal separation of the two reconstructed muons. Under the ultra-relativistic approximation, the explicit dependence on the muon rest mass disappears, while the longitudinal information is encoded in the pseudorapidity difference $\Delta\eta$.

Appendix B

Voigt profile

The Voigt profile arises naturally when a physical resonance with a finite intrinsic width is measured by a detector with finite resolution. The intrinsic mass distribution of a resonance can be described by a Breit–Wigner distribution, which has the Lorentzian form:

$$L(m; m_0, \Gamma) = \frac{1}{\pi} \frac{\Gamma/2}{(m - m_0)^2 + (\Gamma/2)^2} \tag{B.1}$$

where m_0 is the central resonance mass and Γ is its full width at half maximum (FWHM).

The measured invariant mass is affected by the finite detector resolution: experimental effects such as uncertainties in the reconstructed transverse momentum, angular measurements, detector alignment, and other instrumental effects introduce a smearing of the true invariant mass. This detector response can be approximated by a Gaussian distribution:

$$G(m; m_{\text{true}}, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(m - m_{\text{true}})^2}{2\sigma^2}\right] \tag{B.2}$$

where σ represents the Gaussian detector resolution.

Consequently, the experimentally measured resonance is described by the convolution of the intrinsic Breit–Wigner lineshape with the Gaussian detector response. The resulting distribution is the Voigt profile

$$V(m; m_0, \sigma, \Gamma) = \int_{-\infty}^{+\infty} L(m'; m_0, \Gamma) G(m - m'; 0, \sigma) \mathrm{d}m' \tag{B.3}$$

or, equivalently,

$$V(m) = [L(m; m_0, \Gamma) * G(m; 0, \sigma)] \tag{B.4}$$

In this analysis, the Voigt profile is used to model both the reconstructed $Z \to \mu\mu$ invariant-mass distributions and the calibrated β distribution. The implementation used for the fit is provided by the ROOT function [45]:

`TMath::Voigt`

For the Z-boson invariant-mass distributions, the Lorentzian component represents the intrinsic resonance lineshape, while the Gaussian component describes the experimental mass resolution. For the calibrated β distribution, the same functional form is used as an empirical model of the central Gaussian resolution and the non-Gaussian Lorentzian-like tails.

Appendix C

Inverse-variance weighted estimator

The combination of the individual hit-level velocity measurements can be derived by minimizing a χ^2 function with respect to the common velocity value. For a set of measurements β_i with uncertainties σ_{β_i}, the χ^2 is defined as:

$$\chi^2(\beta) = \sum_i \frac{(\beta_i - \beta)^2}{\sigma^2_{\beta_i}} \tag{C.1}$$

where β_i is the measured value, β is the expected value and σ_{β_i} is the error on the measure.

The minimum is obtained by requiring $\frac{\partial \chi^2}{\partial \beta} = 0$. Differentiating Eq. C.1 gives:

$$\frac{\partial \chi^2}{\partial \beta} = -2 \sum_i \frac{\beta_i - \beta}{\sigma^2_{\beta_i}} \tag{C.2}$$

Therefore:

$$\sum_i \frac{\beta_i}{\sigma^2_{\beta_i}} = \beta \sum_i \frac{1}{\sigma^2_{\beta_i}} \tag{C.3}$$

from which the estimator follows:

$$\overline{\beta} = \frac{\sum\limits_i \beta_i/\sigma^2_{\beta_i}}{\sum\limits_i 1/\sigma^2_{\beta_i}} \tag{C.4}$$

The uncertainty of the estimator is obtained from the curvature of the χ^2 function. The second derivative is:

$$\frac{\partial^2 \chi^2}{\partial \beta^2} = 2 \sum_i \frac{1}{\sigma^2_{\beta_i}} \tag{C.5}$$

For a Gaussian likelihood, the variance of the estimator is therefore:

$$\sigma^2_{\overline{\beta}} = \left(\frac{1}{2} \frac{\partial^2 \chi^2}{\partial \beta^2} \right)^{-1} = \left(\sum_i \frac{1}{\sigma^2_{\beta_i}} \right)^{-1} \tag{C.6}$$

and hence

$$\sigma_{\overline{\beta}} = \left(\sum_i \frac{1}{\sigma^2_{\beta_i}} \right)^{-1/2} \tag{C.7}$$

Appendix D

The ABCD background estimation method

The ABCD method is a data-driven technique for estimating a background contribution in a signal-enriched region from background-dominated control regions. Its basic assumption is that two observables used to define the regions are approximately independent for the background. Under this assumption, the joint probability density can be factorised as:

$$P(x, y) \simeq P_x(x)P_y(y) \tag{D.1}$$

The phase space is divided into four regions by two selection thresholds. Using the convention adopted in this thesis, the four regions are:

$$A: \quad x \geq x_0, \quad y > y_0 \tag{D.2}$$
$$B: \quad x \geq x_0, \quad y \leq y_0 \tag{D.3}$$
$$C: \quad x < x_0, \quad y \leq y_0 \tag{D.4}$$
$$D: \quad x < x_0, \quad y > y_0 \tag{D.5}$$

For an independent background distribution, the number of candidates in a region can be written schematically as the product of the corresponding marginal probabilities. This gives:

$$\frac{N_A}{N_B} = \frac{N_D}{N_C} \tag{D.6}$$

and therefore

$$N_D = N_C \cdot \frac{N_A}{N_B} \tag{D.7}$$

The ratio

$$TF = \frac{N_A}{N_B} \tag{D.8}$$

acts as a transfer factor from the low-y region to the high-y region. The method is particularly useful when the background is difficult to model reliably with simulation and when the control regions provide a large data sample [46].

Bibliography

[1] David J. Griffiths. *Introduction to elementary particles; 2nd rev. version.* Physics textbook. New York, NY: Wiley, 2008. URL: https://cds.cern.ch/record/111880.

[2] Mary K. Gaillard, Paul D. Grannis and Frank J. Sciulli. 'The Standard Model of Particle Physics'. In: *Reviews of Modern Physics* 71.2 (1999), S96–S111. DOI: 10.1103/RevModPhys.71.S96. URL: https://arxiv.org/abs/hep-ph/9812285v1.

[3] Wikipedia. *Modello standard.* Wikimedia Foundation. 2025. URL: https://it.wikipedia.org/wiki/Modello_standard (visited on 06/08/2026).

[4] F. Takahashi et al. 'Review of Particle Physics'. In: *Int. J. Mod. Phys. A* 41 (2026), p. 2630011.

[5] ATLAS Collaboration. 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Physics Letters B* 716.1 (Aug. 2012). Report No. CERN-PH-EP-2012-218, arXiv:1207.7214, pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. URL: https://cds.cern.ch/record/1471031.

[6] CMS Collaboration. 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Physics Letters B* 716.1 (July 2012). Report No. CERN-PH-EP-2012-220, CMS-HIG-12-028, pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. URL: https://cds.cern.ch/record/1471016.

[7] LHCb Collaboration. 'Observation of an exotic narrow doubly charmed tetraquark'. In: *Nature Physics* 18 (2022), pp. 751–754. DOI: 10.1038/s41567-022-01614-y.

[8] LHCb Collaboration. 'Observation of pentaquark states in $\Lambda_b^0 \to J/\psi p K^-$ decays'. In: *Physical Review Letters* 115 (2015), p. 072001. DOI: 10.1103/PhysRevLett.115.072001.

[9] ATLAS Collaboration. *Event display for multi-jet production from 13 TeV proton-proton collisions.* CERN Document Server, Report No. ATLAS-DISPLAY-2015-004. 2016. URL: https://cds.cern.ch/record/2194545 (visited on 05/08/2026).

[10] Carlos S. Frenk and Simon D. M. White. 'Dark matter and cosmic structure'. In: *Annalen der Physik* 524.9-10 (2012), pp. 507–534. DOI: 10.1002/andp.201200212.

[11] ATLAS Collaboration. 'The quest to discover supersymmetry at the ATLAS experiment'. In: *Physics Reports* 1116 (2025), pp. 261–300. DOI: 10.1016/j.physrep.2024.09.010. arXiv: 2403.02455 [hep-ex].

[12] Edvige Corbelli and Paolo Salucci. 'The extended rotation curve and the dark matter halo of M33'. In: *Monthly Notices of the Royal Astronomical Society* 311.2 (2000), pp. 441–447. DOI: 10.1046/j.1365-8711.2000.03075.x.

[13] Planck Collaboration. 'Planck 2013 results. I. Overview of products and scientific results'. In: *Astronomy & Astrophysics* 571 (2014), A1. DOI: 10.1051/0004-6361/201321529.

[14] M. B. Gavela et al. 'Standard model CP violation and baryon asymmetry'. In: *Modern Physics Letters A* 9.09 (1994), pp. 795–810. DOI: 10.1142/S0217732394000629. arXiv: hep-ph/9312215.

[15] P. Huet and E. Sather. 'Electroweak baryogenesis and standard model CP violation'. In: *Physical Review D* 51.1 (1995), pp. 379–394. DOI: 10.1103/PhysRevD.51.379. arXiv: hep-ph/9404302.

[16] Michael Dine. 'Naturalness Under Stress'. In: *Annual Review of Nuclear and Particle Science* 65.1 (2015), pp. 43–62. DOI: 10.1146/annurev-nucl-102014-022053. arXiv: 1501.01035 [hep-ph].

[17] Leonard Susskind. 'Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory'. In: *Physical Review D* 20.10 (1979), pp. 2619–2625. DOI: 10.1103/PhysRevD.20.2619.

[18] Samuel J. Ling, Jeff Sanny and William Moebs. 'University Physics'. In: vol. 3. OpenStax / University of Central Florida Pressbooks, 2016. Chap. Evolution of the Early Universe. URL: https://pressbooks.online.ucf.edu/osuniversityphysics3/chapter/evolution-of-the-early-universe/ (visited on 06/08/2026).

[19] D. Croon et al. 'GUT Physics in the Era of the LHC'. In: *Frontiers in Physics* 7 (2019), p. 76. DOI: 10.3389/fphy.2019.00076.

[20] Abdus Salam and J. A. Strathdee. 'Supersymmetry and nonabelian gauges'. In: *Physics Letters B* 51.4 (1974), pp. 353–355. DOI: 10.1016/0370-2693(74)90226-3.

[21] S. Dimopoulos, S. Raby and F. Wilczek. 'Supersymmetry and the scale of unification'. In: *Physical Review D* 24.6 (1981), pp. 1681–1683. DOI: 10.1103/PhysRevD.24.1681.

[22] J. R. Ellis et al. 'Supersymmetric relics from the Big Bang'. In: *Nuclear Physics B* 238.2 (1984), pp. 453–476. DOI: 10.1016/0550-3213(84)90461-9.

[23] H. K. Dreiner. 'An introduction to explicit R-parity violation'. In: *Perspectives on Supersymmetry II*. Ed. by Gordon L. Kane. Vol. 21. Advanced Series on Directions in High Energy Physics. World Scientific, 2010, pp. 565–583. DOI: 10.1142/9789814307505_0017. arXiv: hep-ph/9707435.

[24] M. Fairbairn et al. 'Stable Massive Particles at Colliders'. In: *Physics Reports* 438.1-4 (2007), pp. 1–63. DOI: 10.1016/j.physrep.2006.10.002. arXiv: hep-ph/0611040.

[25] ATLAS Collaboration. 'Search for heavy charged long-lived particles in the ATLAS detector in 36.1 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV'. In: *Physical Review D* 99.9 (2019), p. 092007. DOI: 10.1103/PhysRevD.99.092007. arXiv: 1902.01636 [hep-ex].

[26] CERN. *Our Member States*. 2020. URL: https://home.cern/about/who-we-are/our-governance/member-states (visited on 10/08/2026).

[27] Lyndon R. Evans and Philip Bryant. 'LHC Machine'. In: *JINST* 3 (2008). Abridged version of the LHC Design Report (CERN-2004-003), S08001. DOI: 10.1088/1748-0221/3/08/S08001.

[28] ATLAS Collaboration et al. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[29] Thomas Taylor and Daniel Treille. 'The Large Electron Positron Collider (LEP): Probing the Standard Model'. In: *Technology Meets Research: 60 Years of CERN Technology Ideas and Innovation*. Vol. 27. Advanced Series on Directions in High Energy Physics. World Scientific, 2017, pp. 217–261. DOI: 10.1142/9789814749145_0007.

[30] Ewa Lopienska. *The CERN Accelerator Complex*. CERN Document Server: CERN-GRAPHICS-2022-001-1. CERN. Feb. 2022. URL: https://cds.cern.ch/images/CERN-GRAPHICS-2022-001-1 (visited on 10/08/2026).

[31] CMS Collaboration et al. 'The CMS Experiment at the CERN LHC'. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[32] ALICE Collaboration. 'The ALICE Experiment at the CERN Large Hadron Collider'. In: *JINST* 3 (2008), S08002. DOI: 10.1088/1748-0221/3/08/S08002.

[33] LHCb Collaboration. 'The LHCb Detector at the LHC'. In: *JINST* 3 (2008), S08005. DOI: 10.1088/1748-0221/3/08/S08005.

[34] CERN. *Long Shutdown 3: Preparing the LHC for High-Luminosity*. CERN. 2026. URL: https://home.cern/science/long-shutdown-3/ (visited on 11/08/2026).

[35] ATLAS Collaboration. *ATLAS Data Quality Public Results*. CERN. 2026. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/DataQualityResults (visited on 10/08/2026).

[36] K. Assamagan et al. *A Hierarchical Software Identifier Scheme for the ATLAS Muon Spectrometer*. ATLAS Communication Note ATL-COM-MUON-2002-008. Version 2.0. ATL-COM-MUON-2002-008. CERN, Aug. 2002. URL: https://cds.cern.ch/record/684070.

[37] Joao Pequenao and Paul Schaffner. *How ATLAS detects particles: diagram of particle paths in the detector*. CERN Document Server. CERN-EX-1301009. © 2013 CERN for the benefit of the ATLAS Collaboration. Jan. 2013. URL: https://cds.cern.ch/record/1505342.

[38] Edward Diehl. 'Calibration and Performance of the ATLAS Muon Spectrometer'. In: (2011). arXiv: 1109.6933 [physics.ins-det].

[39] ATLAS Collaboration. 'Commissioning of the ATLAS Muon Spectrometer with cosmic rays'. In: *The European Physical Journal C* 70.3 (Oct. 2010), pp. 875–916. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-010-1415-2. URL: http://dx.doi.org/10.1140/epjc/s10052-010-1415-2.

[40] ATLAS Collaboration. *ATLAS Muon Spectrometer: Technical Design Report*. Technical Design Report ATLAS CERN-LHCC-97-022. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[41] D. Boscherini. 'Performance and operation of the ATLAS Resistive Plate Chamber system in LHC Run-1'. In: *Journal of Instrumentation* 9.12 (Sept. 2014), p. C12039. DOI: 10.1088/1748-0221/9/12/C12039. URL: https://cds.cern.ch/record/1753173.

[42] ATLAS Collaboration. *Preliminary analysis of the luminosity calibration of the ATLAS 13.6 TeV data recorded in 2022*. ATLAS Public Note ATL-DAPR-PUB-2023-001. CERN, 2023. URL: https://cds.cern.ch/record/2853525 (visited on 26/08/2026).

[43] *SUSY Cross Sections 13 TeV: slepton-slepton*. URL: https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections13TeVslepslep (visited on 26/08/2026).

[44] ROOT Collaboration. URL: %5Curl%7Bhttps://root.cern/doc/master/namespaceRooStats_1_1NumberCountingUtils.html%7D (visited on 25/08/2026).

[45] ROOT Collaboration. *ROOT v6.38: TMath Namespace Reference*. https://root.cern.ch/doc/v638/namespaceTMath.html. Documentation for TMath::Voigt.

[46] Suyong Choi and Hayoung Oh. 'Improved extrapolation methods of data-driven background estimations in high energy physics'. In: *European Physical Journal C* 81 (2021). DOI: 10.1140/epjc/s10052-021-09404-1.