université PARIS-SACLAY

Mapping the Noise - Development of a new luminosity measurement technique for the ATLAS experiment and a search for long-lived particles with a delayed jets signature

Développement d'une nouvelle technique de mesure de luminosité pour l'expérience ATLAS et recherche de particules à longue durée de vie avec une signature de jets retardés

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 576, Particules, Hadrons, Energie, Noyau, Instrumentation, Imagerie, Cosmos et Simulation (PHENIICS)
Spécialité de doctorat : Physique des Particules
Graduate School : Physique. Référent : Faculté des sciences d'Orsay

Thèse préparée dans le Département de Physique des Particules (Université Paris-Saclay, CEA), sous la direction de **Prof. Philippe Schwemling**.

Thèse soutenue à Paris-Saclay, le 12 Novembre 2026, par

Ruben POZZI

Composition du jury

Membres du jury avec voix délibérative

Olivier Arnaez Professeur, Université Savoie Mont Blanc	Président
Tetiana Hryn'ova-Berger Directrice de recherche, Université Savoie Mont Blanc	Rapporteur & Examinatrice
Eric Torrence Professeur, University of Oregon	Rapporteur & Examinateur
Emmanuel Monnier Directeur de recherche, Aix Marseille Université	Examinateur
Jose Feliciano Benitez Professeur, Universidad de Sonora	Examinateur

THESE DE DOCTORAT

NNT : 2026UPASP139

université PARIS-SACLAY

ÉCOLE DOCTORALE
Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)

Titre : Développement d'une nouvelle technique de mesure de luminosité pour l'expérience ATLAS et recherche de particules à longue durée de vie avec une signature de jets retardés
Mots clés : ATLAS, LHC, luminosité, calorimètre, argon liquide, particules à longue durée de vie, jets retardés

Résumé : Cette thèse présente le développement d'une nouvelle technique de mesure de la luminosité pour l'expérience ATLAS et une recherche de physique au-delà du Modèle Standard, toutes deux réalisées avec les données de collisions proton–proton du Run 3 du LHC enregistrées entre 2024 et 2026.

ATLAS mesure la luminosité relative avec plusieurs détecteurs et méthodes, et fixe l'échelle absolue par des balayages de van der Meer (vdM) dédiés. L'empilement prévu au LHC à haute luminosité rend de plus en plus précieuse la mise au point de méthodes relatives complémentaires, en particulier celles susceptibles de fournir une information paquet par paquet de protons.

Le système de déclenchement numérique du calorimètre à argon liquide (LAr), chaîne de lecture installée lors de l'upgrade de Phase-I, numérise les échantillons des Super Cells (SC) à chaque croisement de faisceaux de protons. La possibilité d'exploiter ces donnnées pour obtenir une mesure de luminosité est établie dans le présent travail par trois techniques complémentaires. La première utilise la dispersion des échantillons numérisés, qui est reliée à l'empilement et varie donc avec la luminosité instantanée ; sa résolution relative à la luminosité en ligne d'ATLAS est de $0,04\%$ dans un run nominal du Run 3 et de $0,21\%$ à $\langle\mu\rangle = 120,6$. La deuxième méthode mesure la luminosité paquet par paquet en déconvoluant les variations du piédestal moyen en fonction du numéro de paquet (Bunch Crossing ID, BCID) à l'aide de la forme bipolaire du signal LAr. Cette méthode retrouve la luminosité mesurée paquet par paquet par le luminomètre LUCID avec un résidu quadratique moyen de $3,4\%$, en utilisant pour la déconvolution une forme de signal extraite des données. La troisième et dernière méthode utilise le moniteur de taux d'impulsions. Il s'agit ici d'exploiter une fonction du micrologiciel back-end du calorimètre à Argon Liquide, qui compte les impulsions au-dessus d'un seuil en énergie configurable. Cette méthode atteint une résolution meilleure que le pour cent et reste sensible aux luminosités faibles présentes lors des balayage de Van der Meer. Prises ensemble, les trois méthodes couvrent la mesure de la luminosité intégrée sur les paquets, paquet par paquet et durant les balayages vdM avec un seul détecteur. Elles reposent sur un micrologiciel et un système d'acquisition que le calorimètre à Argon Liquide conservera pour le Run 4. Elles peuvent donc être développées plus avant pendant le long arrêt 3 (LS3) et déployées sans instrumentation supplémentaire.

La seconde partie introduit une recherche de particules à longue durée de vie produisant des jets hadroniques retardés, sélectionnés par des déclenchements dédiés, exploitant le temps d'arrivée des jets, dans environ 210 fb^{-1} de données de collisions proton-proton acquises à 13.6 TeV. Le signal recherché est la production électrofaible de neutralinos à longue durée de vie se désintégrant en triplets de quarks via un faible couplage violant la R-parité, pour des masses de $1,0$–$2,0$ TeV et des temps de vie de 1–30 ns. Les contributions présentées ici sont la caractérisation au niveau vérité des simulations Monte-Carlo des échantillons de signal et l'optimisation de trois régions de signal fondées sur la significativité du retard normalisée à la résolution. Le nombre d'évènements attendus atteint 86 au point de grille le plus favorable, avec un maximum à un temps de vie propre de 10 ns, domaine où les recherches par vertex perdent en acceptance. L'analyse est en aveugle et les estimations de bruit de fond sont en cours ; aucune limite n'est citée.

Title: Mapping the Noise - development of a new luminosity measurement technique for the ATLAS experiment and a search for long-lived particles with a delayed jets signature

Keywords: ATLAS, LHC, luminosity, calorimeter, liquid argon, long-lived particles, delayed jets

Abstract: This thesis presents the development of a new luminosity measurement technique for the ATLAS experiment and a search for physics beyond the Standard Model, both performed with proton–proton collision data recorded during Run 3 of the Large Hadron Collider (LHC) in the years 2024–2026.

ATLAS measures the relative luminosity with several detectors and methods, and fixes the absolute luminosity scale with dedicated van der Meer (vdM) scans. The pileup foreseen at the High-Luminosity LHC makes additional relative luminosity methods, in particular if capable of measuring the luminosity per-bunch, increasingly valuable. The readout chain installed by the Phase-I upgrade of the Liquid Argon (LAr) calorimeter, also known as Digital Trigger, is established here as a luminosity detector through three complementary techniques. The first one uses the dispersion of its digitized samples, which is driven by pileup and therefore scales with the instantaneous luminosity; its resolution relative to the ATLAS online luminosity is 0.04% in nominal Run 3 conditions and 0.21% at $\langle\mu\rangle = 120.6$. The second method measures luminosity bunch by bunch by deconvolving the mean ADC pedestal per bunch crossing identifier (BCID), using the bipolar LAr pulse shape as kernel. This method reproduces the per-bunch luminosity of the LUCID luminometer with a 3.4% root-mean-square residual, once a data-driven kernel is used. The third uses the Pulse Rate Monitor, a function implemented in the LAr back-end firmware which counts pulses above a configurable energy threshold. This method reaches sub-percent resolution and stays sensitive at vdM luminosities. Together, the three cover the bunch-integrated, the per-bunch and the vdM-scan regimes from one detector, and all are obtained from firmware and readout that the LAr system will carry into Run 4, which means they can be developed further during Long Shutdown 3 (LS3) and deployed without additional instrumentation.

The second part of the thesis introduces a search for long-lived particles producing delayed hadronic jets, selected by dedicated triggers on the jet arrival time, in a dataset of about $210\,\text{fb}^{-1}$. The targeted signal model is electroweak production of long-lived neutralinos decaying via a small R-parity-violating coupling into quark triplets, for masses of 1.0–2.0 TeV and lifetimes of 1–30 ns. The contributions are the truth-level characterization of the signal samples and the optimization of three signal regions built on a resolution-scaled delay significance. The expected yields reach about 86 events at the most favorable grid point, peaking at a proper lifetime of 10 ns, in the range where vertex-based searches lose acceptance. The analysis is blinded and the background estimates are still in progress, so no limits are quoted.

Acknowledgments

There are many people that I want to thank, but first and foremost, I want to thank my supervisors Philippe and Nikiforos. This thesis was prepared at CERN, in the frame of the CERN doctoral program, with Nikiforos Nikiforou acting as CERN thesis supervisor, and Philippe Schwemling as academic supervisor at CEA-Saclay. Philippe has been an incredible supervisor, always understanding and helpful, always present, although physically in another country for most of the project. He guided me and encouraged me from the beginning and his contribution was critical. Then, I want to thank Nikiforos. He took me under his wing three years ago, and taught me all he could about this beautiful machine called the Liquid Argon calorimeter. He transmitted his passion and love for physics to me, and always took the time and effort to solve whatever technical, bureaucratic and physics-related hassle I was encountering at any moment.

I also want to thank everyone in the LAr operations team. Almost like a family, it was a wonderful environment not only to work in, but also to have fun chats, dinners, BBQs. It was the place where I met new incredible friends. Despite the recurring pressure and stress of operations, from emergencies, through last-minute ROD board replacements went wrong, to 3 AM calls, everyone in the team always went out of their way to help each other out, and make them feel included. Believe it or not, I will almost miss operations.

The completion of this PhD would not have taken place without the endless support of my family. Mamma e Papà, from guiding me in my early years, to supporting me throughout my academic career, no matter how incomprehensible my work got, you never stopped showing your love and never-ending support. I also could not have done this without Pier and Leuz, the two brothers I could never live without. From childhood to adulthood, throughout the good times and the not so good ones, you two have always been my anchor and my best friends.

Last but very much not least, I want to thank my Adri. El amor de mi vida, she met me at the beginning of my PhD, and never once stop being fascinated about what I do and curiously asking, and sometimes helping me out, in my work. She has been my constant shoulder to lean upon, celebrating me when I was succeeding, and holding me and bringing me up again when I was falling. From cute letters and presents, and delicious dinners cooked, to words of love and support while I was grinding for endless nights writing this thesis, you are the one that made all of it possible.

Résumé en français

Contexte et motivation

Le Modèle Standard laisse ouvertes des questions majeures, comme la nature de la matière noire. Dans de nombreuses extensions, notamment supersymétriques avec violation de la R-parité, les nouvelles particules peuvent acquérir un temps de vie macroscopique : ce sont les particules à longue durée de vie (*long-lived particles*, LLP), dont les signatures échappent aux techniques de reconstruction standard. Par ailleurs, la mesure de la *luminosité intégrée* délivrée par le LHC conditionne l'incertitude de toute section efficace et de toute limite d'exclusion ; elle repose dans ATLAS sur le détecteur LUCID-2, complété par les méthodes de comptage de traces et de mesure des courants de polarisation dans les intervalles d'Argon Liquide (« gaps ») des calorimètres. Les conditions attendues au HL-LHC, où le nombre moyen d'interactions par croisement atteindra $\mu \simeq 200$, rendent nécessaires des méthodes supplémentaires, indépendantes, stables et si possible résolues paquet par paquet.

Cette thèse propose trois techniques inédites de mesure de la luminosité fondées sur le système de déclenchement numérique, ou *Digital Trigger* (DT), du calorimètre à argon liquide (LAr), puis présente une contribution à la recherche de LLP se désintégrant en jets retardés dans les calorimètres. Le signal triangulaire de courant d'ionisation du calorimètre à Argon Liquide est mis en forme par une électronique bipolaire, si bien que le signal en tension obtenu possède une intégrale temporelle nulle : l'effet de l'empilement s'annule en moyenne, mais non ses fluctuations, remarque qui fonde les méthodes développées ici. Depuis la mise à niveau de l'électronique de déclenchement en vue de la Phase-I du LHC, les *Super Cellules* sont numérisées à 40 MHz, à chaque croisement de paquets, et traitées par les cartes LATOME, chaque échantillon ADC conservant son identifiant de croisement (BCID). Une partie du travail relève enfin des opérations de qualité des données du LAr en 2024–2025 (Sections A.2, A.3, A.4 et A.5).

Luminosité par la dispersion des échantillons

La contribution principale de cette thèse est la démonstration que la dispersion (écart quadratique moyen, RMS) des échantillons ADC du déclenchement numérique est un estimateur de la luminosité instantanée (Chapitre 8), le bruit électronique y fixant un plancher et le terme d'empilement croissant avec μ. La relation à la luminosité est établie empiriquement, run par run, par un ajustement linéaire à la luminosité en ligne d'ATLAS, dont l'inversion fournit l'estimateur.

Une étude de faisabilité sur les données MON[1] donne d'emblée une résolution relative de l'ordre de 2 %. L'analyse des effets instrumentaux et machine délimite le domaine de validité de la méthode.

[1] Flot de données brutes spécifique, prévu avant tout pour la vérification de la qualité des données

La réponse dépendant fortement de la couche du détecteur ainsi que de la pseudorapidité des supercellules analysées, une combinaison pondérée par l'inverse de la résolution de chaque canal améliore la résolution relative de 10 % avec une pondération par couche et de 17 % avec une pondération par pseudorapidité. Le signal est porté quasi exclusivement par les croisements actifs, et leur sélection fait passer la résolution de 2,3 % à 2,1 %. La méthode conserve une sensibilité aux collisions d'ions lourds, et sa stabilité est excellente : sur tous les runs certifiés de 2024 et 2025, la moyenne des résidus relatifs reste centrée sur zéro, sans dérive significative. Elle est en revanche insensible aux balayages d'émittance et, surtout, aux balayages van der Meer (vdM), où l'empilement devient négligeable devant le bruit électronique.

Parmi les flux officiels d'ATLAS, `physics_Main` s'avère le plus performant, la précision étant gouvernée non par la seule statistique mais par la fraction d'événements issus de croisements remplis. Les résultats obtenus avec ce flux ont été approuvés par ATLAS et publiés (Section 8.4.2) : pour un run nominal de 2025 ($\langle\mu\rangle = 58{,}6$, 2448 paquets en collision), la corrélation est linéaire et la résolution relative atteint 0,04 % ; pour un run à haut empilement ($\langle\mu\rangle = 120{,}6$), représentatif des conditions du HL-LHC, elle reste de 0,21 %. La méthode présente le rapport moyen à LUCID le plus proche de l'unité ($1{,}000 \pm 0{,}016$) parmi les méthodes d'ATLAS ; sa principale faiblesse demeure la perte de sensibilité à basse luminosité.

Luminosité paquet par paquet

Le déclenchement numérique, qui associe un BCID à chaque échantillon, se prête naturellement à une mesure résolue paquet par paquet. La méthode proposée (Chapitre 9) utilise non plus la dispersion des échantillons mais leur *valeur moyenne* en fonction du BCID. L'impulsion bipolaire du LAr s'étendant sur quelque 800 ns alors que les paquets se croisent toutes les 25 ns, le signal observé est la *convolution* discrète de la luminosité par paquet avec la forme d'impulsion prise comme noyau, et la luminosité peut donc être récupérée par *déconvolution*.

Trois estimateurs sont comparés sur une simulation Monte-Carlo simplifiée : l'inversion dans le domaine de Fourier (FFT), l'inversion par matrice inverse et l'algorithme de Richardson–Lucy, appliqué avec un noyau rendu positif par l'ajout d'une constante. La FFT et la matrice inverse, qui sont le même estimateur pour un opérateur circulant, reconstruisent la luminosité d'entrée à mieux que le pour cent (0,44 % d'erreur relative), tandis que Richardson–Lucy est d'un ordre de grandeur moins précis (6,25 %). Pour les données, une variante *contrainte* des deux premiers estimateurs est en outre introduite : elle restreint les inconnues aux BCID remplis, ajuste le piédestal comme paramètre explicite et supprime du système les BCID non mesurés.

Appliquée aux données du run 520264 avec, pour noyau, les formes d'impulsion mesurées par indice de pseudorapidité (« gabarit » 2025), la déconvolution restitue la structure en trains avec une sensibilité marquée, mais cet ordre n'est pas conservé : l'inversion non contrainte souffre d'un dépassement en début d'orbite et d'une ligne de base négative entre les trains, tandis que la matrice inverse contrainte développe une oscillation de paquet à paquet. Pour déterminer si la limitation vient du noyau ou de l'algorithme, un noyau *extrait des données* a été construit (Section 9.5) en

inversant les rôles dans la convolution : les échantillons du noyau sont ajustés par moindres carrés sur la luminosité par BCID mesurée par LUCID, indépendamment pour chacune des 218 tranches en pseudorapidité (scieta). Avec ce noyau, les deux estimateurs contraints reproduisent la référence sur toute l'orbite, avec un résidu relatif par rapport à LUCID de moyenne $-0{,}10$ % et de RMS 3,4 % pour la matrice inverse contrainte, et de $-0{,}099$ % pour 3,1 % pour la FFT contrainte, contre 19 % pour l'inversion non contrainte sur le même noyau. La matrice inverse contrainte est retenue comme la meilleure reconstruction : son résidu est plat sur toute l'orbite, alors que celui de la FFT contrainte se dégrade aux extrémités, et un biais dépendant du BCID ne se moyenne pas lors de la somme sur les paquets. C'est donc le choix du noyau, et non l'algorithme d'inversion, qui limite la mesure : obtenir une forme d'impulsion valide pour les conditions de prise de données est la principale perspective, éventuellement en *apprenant* l'inverse de la convolution au moyen d'un réseau de neurones.

Le Pulse Rate Monitor des LATOME comme luminomètre

Les deux méthodes précédentes n'exploitent que les événements acceptés par ailleurs par le système de déclenchement et perdent leur sensibilité lorsque l'empilement devient petit devant le bruit électronique. Une troisième observable, affranchie de ces deux contraintes, a été explorée (Chapitre 10) : le *Pulse Rate Monitor* (PRM), fonction du micrologiciel des LATOME qui compte en continu, pour chacune des quelque 34 000 supercellules et à chaque croisement, les impulsions dépassant un seuil en énergie transverse configurable. Le PRM est ainsi un luminomètre relatif par comptage d'impulsions, étalonné comme la méthode RMS.

Dans un run nominal de 2025, le taux moyen par supercellule est fortement corrélé à la luminosité en ligne et donne une résolution relative de 0,83 %. Les taux couvrant des échelles très différentes selon la couche, de 10 Hz environ pour les préchantillonneurs jusqu'à 10^3 Hz ailleurs, la moyenne sur tout le détecteur est contre-productive : dans le run à haut empilement ($\langle\mu\rangle = 120{,}6$), cette moyenne est bruitée, tandis que la mesure restreinte au préchantillonneur du tonneau atteint 0,79 %. Surtout, en sélectionnant la coordonnée la plus performante, le PRM montre une sensibilité claire et linéaire à la luminosité pendant un balayage van der Meer de 2025 (Section 10.6), régime où la méthode RMS échoue complètement. Sa stabilité sur les runs certifiés de 2025, avec un étalonnage unique, est raisonnable, les résidus restant majoritairement dans une plage de 1 %, mais une population secondaire apparaît, probablement liée aux évolutions du micrologiciel et des seuils.

Recherche de LLP avec jets retardés

La seconde partie de la thèse porte sur la recherche de LLP se désintégrant hadroniquement dans les calorimètres, où les recherches fondées sur les vertex déplacés perdent leur sensibilité. L'analyse présentée exploite une observable complémentaire et largement inexploitée, le *temps d'arrivée* des dépôts calorimétriques : une particule massive se propageant avec $\beta < 1$ puis se désintégrant à un vertex déplacé produit des jets retardés de quelques nanosecondes par rapport au temps nominal

du croisement. Le modèle de référence est la production électrofaible de neutralinos à longue durée de vie se désintégrant en trois jets via un couplage violant la R-parité, sur une grille de masses de 1000 à 2000 GeV et de temps de vie propres de 1 à 30 ns. Des chaînes de déclenchement de haut niveau dédiées, sélectionnant sur le temps du jet, ont fonctionné de juillet 2024 à la fin des opérations du LHC, pour une luminosité intégrée d'environ 210 fb^{-1}. L'analyse est aveuglée et l'estimation des quatre bruits de fond considérés, tous déterminés à partir des données, est encore en cours. Ma contribution porte sur les études au niveau de la vérité Monte-Carlo (Section 11.6) et sur la définition et l'optimisation des régions de signal (Section 11.7).

Les études menées au niveau « vérité » des signaux Monte-Carlo établissent que la position de désintégration est gouvernée par le temps de vie propre et dépend peu de la masse du neutralino, et que le retard associé atteint plusieurs nanosecondes pour les désintégrations dans les calorimètres, pleinement à l'intérieur de la fenêtre temporelle utilisée hors ligne. Elles montrent aussi que la collection standard de jets de vérité est inadaptée, son axe étant vu depuis le vertex déplacé alors qu'un jet reconstruit pointe depuis l'origine du détecteur vers le dépôt d'énergie. Des jets de vérité dédiés, construits par LLP à partir des seuls hadrons de l'état final de chaque neutralino, ont donc été définis, et une correction géométrique propageant leur axe jusqu'à la surface du calorimètre réduit la distance angulaire moyenne au jet reconstruit le plus proche de 0,152 à 0,065, très en deçà du cône d'appariement $\Delta R < 0{,}4$.

Trois régions de signal ont été définies : une région inclusive à un jet dur, destinée à un résultat indépendant du modèle, et deux régions fondées sur la multiplicité de jets retardés, à deux jets dans le même hémisphère et à trois jets. Le retard est exprimé comme une significativité mise à l'échelle de la résolution temporelle, et une optimisation bidimensionnelle simultanée du seuil en impulsion transverse et de cette significativité a été menée sur toute la grille. L'efficacité culmine autour de $\tau = 10$ ns dans les trois régions, temps de vie pour lequel la fraction de désintégrations à l'intérieur des calorimètres est maximale. Les rendements attendus atteignent environ 86 événements au point le plus favorable et tombent sous l'unité sur la majeure partie de la grille, restant partout inférieurs à un événement à 2000 GeV : la sensibilité se concentre au bas de la gamme de masse, et le choix entre les trois régions, de rendements comparables, sera dicté par leurs niveaux de bruit de fond. Le point de fonctionnement ne pouvant être fixé à partir du signal seul, la sélection reste un balayage optimisé mais non figé.

Conclusions et perspectives

Cette thèse démontre que le système de déclenchement numérique du calorimètre à argon liquide peut servir d'instrument de mesure de la luminosité, rôle pour lequel il n'avait pas été conçu. La dispersion des échantillons ADC fournit un luminomètre relatif linéaire et stable, de résolution 0,04 % dans un run nominal et 0,21 % à haut empilement ; leur valeur moyenne par BCID restitue, par déconvolution, la luminosité paquet par paquet à 3,4 % de LUCID dès lors que le noyau employé est adéquat ; le PRM, affranchi du déclenchement, atteint une résolution inférieure au pour cent et reste sensible aux échelles des balayages van der Meer. Les limitations identifiées définissent le pro-

gramme de travail à venir, dans la perspective d'une implémentation en ligne ou micrologicielle pour le Run 4.

Pour la recherche de LLP, les outils d'exploitation des échantillons simulés ont été validés et trois régions de signal définies et optimisées ; une fois les bruits de fond estimés et l'analyse désaveuglée, elle fournira les premières limites d'ATLAS obtenues à partir du temps d'arrivée des dépôts calorimétriques.

Contents

Acknowledgments **iv**

Résumé en français **vi**

Preface **xvii**
Introduction xvii
About the thesis xviii
Author's contributions xviii

I Theoretical Framework 1

1 The Standard Model of particle physics **3**
1.1 The fundamental forces and particles 3
1.2 Symmetries in the SM 6
1.3 Electroweak unification and spontaneous symmetry breaking 9
1.4 Quantum chromodynamics and asymptotic freedom 10
1.5 Cross sections and particle decays 11
1.6 Phenomenology of proton–proton collisions 14
1.7 Open questions in the Standard Model 16

2 Beyond the Standard Model **21**
2.1 The theory of supersymmetry 21
2.1.1 The Minimal Supersymmetric Standard Model 22
2.1.2 *R*-parity 22
2.2 Experimental searches for SUSY 23
2.3 Long-lived particles 25

II Experimental Setup 27

3 The Large Hadron Collider **29**
3.1 Accelerating protons into collisions 30
3.2 LHC filling scheme 32
3.3 Luminosity, pileup and HL-LHC 33

4 The ATLAS experiment **35**
4.1 ATLAS coordinate system . . . 35
4.2 Inner Detector . . . 37
4.3 Calorimeters . . . 40
4.4 Muon system . . . 41
4.5 Magnet system . . . 42
4.6 Trigger system . . . 44

5 The Liquid Argon calorimeter **49**
5.1 Calorimeters . . . 49
5.1.1 Energy measurement . . . 50
5.1.2 Electromagnetic showers . . . 51
5.2 The ATLAS calorimeters . . . 54
5.3 Electromagnetic barrel calorimeter . . . 55
5.4 End-cap calorimeters . . . 58
5.5 Forward calorimeter . . . 60
5.6 Readout electronics and trigger system . . . 61
5.6.1 Readout electrodes . . . 61
5.6.2 Front-end electronics . . . 64
5.6.3 Back-end electronics . . . 66
5.7 Noise and pileup . . . 67

6 The Digital Trigger upgrade **71**
6.1 Physics requirements for Phase-I and HL-LHC . . . 72
6.2 Phase-I upgrade and the Digital Trigger . . . 74
6.2.1 The Super Cell scheme . . . 75
6.2.2 Front-end electronics . . . 76
6.2.3 Back-end electronics: the LAr Digital Processing System . . . 76
6.3 Performance of the Digital Trigger in Run 3 . . . 79

III Luminosity Measurement **83**

7 Luminosity in high-energy physics experiments **85**
7.1 Definition and physical meaning . . . 85
7.2 Luminosity at proton collider experiments . . . 86
7.2.1 Relative luminosity measurements . . . 87
7.2.2 Absolute luminosity measurement . . . 88
7.3 Luminosity measurements in ATLAS . . . 89
7.3.1 LUCID-2 and BCM . . . 89
7.3.2 Luminometers in ATLAS . . . 91

7.4 Van der Meer scans and calibration transfer . . . 92
7.5 ATLAS Run 2 luminosity results . . . 94

8 Luminosity measurement with the Digital Trigger 97
8.1 Introduction to the methods and datasets . . . 97
8.1.1 Integration with ATLAS luminosity framework . . . 99
8.2 The RMS of Digital Trigger samples as a luminosity measurement method . . . 100
8.2.1 Theoretical foundation . . . 100
8.3 Feasibility study . . . 102
8.3.1 Data quality cuts . . . 104
8.3.2 Raw or baseline-corrected samples? . . . 104
8.3.3 Estimators per layer and per η coordinate . . . 105
8.3.4 Weighting per layer and per η coordinate . . . 106
8.3.5 Filled bunches versus empty bunches . . . 110
8.3.6 Effects of bunch structure . . . 111
8.3.7 Cross-year calibration: 2024 versus 2025 data . . . 113
8.3.8 Luminosity measurement in heavy ion runs . . . 113
8.3.9 Effect of emittance scans on the luminosity estimation . . . 115
8.3.10 Van der Meer scans . . . 116
8.3.11 Long-term stability . . . 117
8.3.12 Preliminary comparison with other luminosity methods . . . 119
8.4 Study of the SWROD data . . . 122
8.4.1 Comparison of streams and trigger biases . . . 122
8.4.2 Approved results with ATLAS physics data . . . 126
8.4.3 Comparison with ATLAS luminosity methods . . . 129

9 Bunch-by-bunch luminosity measurement 133
9.1 Idea and theoretical background . . . 133
9.2 Toy Monte Carlo simulation . . . 137
9.3 Deconvolution of the mean pedestal . . . 140
9.4 Comparison with LUCID luminosity per bunch . . . 144
9.5 Kernel extraction with LUCID . . . 148

10 LATOME Pulse Rate Monitor as a luminometer 155
10.1 Monitoring event rates with the PRM . . . 155
10.1.1 Derivation of the pulse rate in the LATOME firmware . . . 156
10.2 PRM as a luminometer . . . 157
10.3 Nominal pp physics run . . . 158
10.4 Layer and η-coordinate optimization . . . 160
10.5 High-μ run . . . 161
10.6 Sensitivity to vdM scans . . . 163

10.7 Long-term stability . . . 166

IV Search for Long-Lived Particles 167

11 Search for long-lived particles with a delayed-jet signature 169
11.1 Analysis overview . . . 170
11.2 Data and simulated event samples . . . 172
11.2.1 Collision data . . . 172
11.2.2 Simulated signal samples . . . 172
11.3 Event reconstruction . . . 175
11.4 Dedicated HLT triggers for delayed jets . . . 176
11.5 Backgrounds estimation . . . 176
11.6 Truth studies . . . 177
11.6.1 LLP decay kinematics . . . 177
11.6.2 LLP decay products and truth jets . . . 179
11.6.3 Truth jets in displaced decays . . . 181
11.6.4 Truth-to-reconstruction jet matching . . . 185
11.7 Signal region definition . . . 189
11.7.1 Signal acceptance scans . . . 190
11.7.2 Expected yields . . . 195
11.8 Next steps . . . 197

Concluding remarks 199
Relevance for the HL-LHC . . . 200
Future work . . . 201

V Appendices 203

A Data quality and performance studies on the Digital Trigger 205
A.1 LAr data quality operations . . . 205
A.2 DT OFCs optimization . . . 208
A.3 Processing of MON data . . . 211
A.4 FCal correlations . . . 212
A.5 Pedestal jumps study . . . 214

B Two-dimensional signal-region optimization scans 223

Bibliography 226

Preface

Introduction

Particle physics studies the elementary particles that make up nature, and in doing so addresses one of the oldest questions humans have asked: *what are we, and the universe we live in, made of?*

The Standard Model describes these particles and their interactions with remarkable accuracy, yet leaves several questions open. Both the Standard Model and what may lie *beyond* it are tested at particle colliders such as the Large Hadron Collider (LHC) at CERN, where two beams of protons, grouped into bunches, collide at high energy and the products of the collisions are recorded by large detectors such as ATLAS.

Turning what a detector records into a physics result requires knowing how many collisions the accelerator delivered, a quantity measured by the *luminosity*. Its precise knowledge is needed to determine how often a process occurs, or to set a limit on one that is not observed. The precision of its measurement relies on a combination of instruments: dedicated calibration runs fix the absolute scale, and several independent detectors then follow the luminosity through a full year of data taking, their mutual agreement setting the final uncertainty. Every additional independent method strengthens that combination, especially one able to follow the luminosity of individual proton bunches.

The need grows with the conditions ahead. The LHC already operates well beyond its original design, with around sixty proton–proton collisions taking place every time two bunches cross, and its high-luminosity upgrade, the High-Luminosity LHC (HL-LHC), will raise this to about two hundred. These overlapping collisions, known as *pileup*, complicate most measurements, but they are exactly what a luminosity measurement relies on. The first part of this thesis exploits them with the Digital Trigger of the ATLAS Liquid Argon (LAr) calorimeter, an electronics system installed in the Phase-I upgrade of the detector to improve the real-time selection of collisions. It was never designed to measure luminosity; this thesis asks whether the signals it already records can nonetheless do so, and how well.

The second part turns the same calorimeter to a different use. Many theories that extend the Standard Model, supersymmetry among them, predict new particles that live long enough to travel a measurable distance inside the detector before decaying. Most searches for such particles rely on the tracks left by their decay products, and lose sensitivity when the decay happens beyond the tracking detectors. A heavy, slow particle that decays inside the calorimeter instead produces collimated sprays of particles,

known as *jets*, that arrive a few nanoseconds late, a largely unexplored signature that complements the tracking-based approach. The second part of this thesis presents the ATLAS search built on it.

About the thesis

The thesis is divided in four parts. Part I presents the theoretical framework, where Chapter 1 summarizes the theory of the Standard Model. Chapter 2, instead, introduces new physics beyond it, in particular supersymmetry, R-parity violation, and long-lived particles, which are the target of the physics analysis presented in this thesis.

Part II describes the experimental setup progressively, from the acceleration and collision of protons in the LHC, through the ATLAS detector (Chapter 4) and the Liquid Argon calorimeter (Chapter 5), to the Phase-I Digital Trigger upgrade (Chapter 6).

Part III contains the luminosity work. Chapter 7 serves as a literature review, defining the concept of luminosity and its physical meaning, and the relative and absolute measurement techniques and the methods currently used in ATLAS. Chapter 8 then introduces the main body of work developed in this thesis, that is, the first new luminosity measurement technique based on the root mean square (RMS) of the Digital Trigger samples. Chapter 9 describes the second innovative method, i.e., the retrieval of the luminosity per bunch using a deconvolution of the mean analog-to-digital converter (ADC) pedestal per bunch crossing identifier (BCID) and the Liquid Argon pulse shape as a kernel. Chapter 10 shows the third and final method, that is, the potential application of the Pulse Rate Monitor of the LAr Trigger prOcessing MEzzanine (LATOME) as a hit-counting relative luminometer. Finally, part IV presents the search for long-lived particles with a delayed jets signature (Chapter 11).

Two appendices complete the document: Appendix A collects the data quality and performance studies carried out on the Digital Trigger as a qualification task to achieve authorship at the ATLAS Collaboration, which also provided part of the datasets and knowledge on which the luminosity work described in Part III rests. Appendix B, instead, contains the remaining two-dimensional signal-region optimization scans not shown in Chapter 11.

Author's contributions

Experimental high-energy physics at colliders relies on infrastructure so large that it can only be built and operated by international collaborations. ATLAS counts thousands of physicists, engineers, technicians, students and support staff from around the world, which makes individual contributions difficult to isolate. The collaboration is organized into physics groups, each covering a specific part of its physics program. Every result belongs to the whole collaboration and undergoes a thorough internal review before

release, whether as a journal paper, a conference note or a set of public plots, and journal papers carry the full author list.

Alongside the physics groups, dedicated groups are responsible for building, maintaining and operating the detector, for data preparation and reconstruction, and for simulation. This detector and performance work underpins every ATLAS result, and all members are expected to contribute to it.

Throughout my PhD I was a member of the LAr calorimeter operations team. I served as on-call Run Coordinator and carried out data-quality and performance studies related to the LAr Digital Trigger upgrade, documented in Appendix A. This work provided me with the detailed knowledge of the LAr readout and data on which the luminosity study relies.

The luminosity study presented in Part III is my own work, carried out under the guidance of my supervisors, Philippe Schwemling and Nikiforos Nikiforou. It uses three methods. For the first, based on the RMS of the digitized samples, I built the infrastructure to process the MON data (see Section A.3) as part of my ATLAS qualification task. This infrastructure also served as the basis of a preliminary study of the method in Section 8.3. For the second, per-bunch method, the data were recorded with changes to the detector readout that I applied during a Run Coordinator shift, and I later processed them offline. The third method uses the LATOME Pulse Rate Monitor, which was developed and is maintained by the LAr Firmware working group. My contribution here was limited to retrieving the archived data and processing them offline.

In the last two years of my PhD I joined the Delayed Jets analysis team (Chapter 11). I set up the analysis framework used for event processing and ntuple production, and implemented the truth-level information in it. I then performed truth-level studies of the Monte Carlo signal samples, defined and optimized the signal regions, and computed the expected signal yields. At the time of writing, the analysis is blinded, and the background estimation is being carried out by other members of the analysis team.

Part I

Theoretical Framework

Chapter 1

The Standard Model of particle physics

The most accurate description of Nature available today at the smallest accessible scales is the Standard Model (SM) of particle physics. The theory describes how the world around us is built from a small number of elementary particles, and how these interact through fundamental forces to form more complex particles and structures. Its construction spans the second half of the twentieth century: the gauge field theories of the electroweak interaction were developed between 1961 and 1968 [1–3], the Brout–Englert–Higgs mechanism was proposed in 1964 [4–6], and the framework was completed in the 1970s with the formulation of quantum chromodynamics (QCD).

A large body of experimental results has tested the model ever since, confirming its predictions with increasing precision. The SM is nevertheless far from complete, as several established phenomena fall outside its scope. Despite its validity at the energy scales probed up to this day, the model is therefore widely regarded as the low-energy limit of a more fundamental theory that has yet to be established.

1.1 The fundamental forces and particles

Four fundamental interactions govern our Universe: gravity, electromagnetism, the weak interaction and the strong interaction. The first two are widely known in our everyday lives, and it is no coincidence that their framework was conceived far before the discovery of the other two. The weak force appears in radioactive decays, and in the nuclear reactions fueling the Sun and the stars, while the strong force is responsible for keeping protons and neutrons together.

The SM provides a valid and unified description of all these forces, with the exception of gravity, which finds its only description in the theory of general relativity and has no established quantum formulation.

According to the SM, these forces are mediated by spin-1 gauge bosons which carry four-momentum between the particles interacting through them. These bosons, also shown in Figure 1.1, are:

- Photon: electrically neutral, mediates the electromagnetic force.
- $W^{\pm}$ & Z: the former charged and the latter neutral, they mediate the weak force.
- Gluon: mediates the strong force.

Each of these force carriers is allowed to interact only with particles exhibiting a corresponding charge: the photon couples to all electrically charged particles, while the weak carriers couple only to combinations of *weak isospin* and *weak hypercharge*. The strong interaction, on the other hand, couples to *color charge*. All the forces are characterized by a coupling constant α that establishes the relative strength of the interaction. The coupling strength depends on the energy at which the interaction is probed, and hence on the squared four-momentum transfer Q^2, but there are differences between the forces. For the electromagnetic force, α increases at large momentum transfer, or, correspondingly, at short distances, while the strong coupling α_S behaves in the opposite way, decreasing at large momentum transfer, a property known as *asymptotic freedom* and discussed in Section 1.4.

The range of the interaction is inversely proportional to the mass of the mediating particle. This means that the electromagnetic force, mediated by the massless photon, has infinite range, while the weak force, being mediated by massive bosons, has a short range. The strong interaction is an exception to this rule, and is discussed in Section 1.4.

While the bosons mediate interactions, matter itself is made up of spin-1/2 particles called *fermions*, which are further subdivided into quarks and leptons. All such particles exist in three generations, varying in terms of lifetime and mass, as well as the corresponding antiparticles, exhibiting the same mass and spin but opposite charge. An overview of the particles described below is shown in Figure 1.1.

The lepton generations, or *flavors*, correspond to three types of charged leptons: the electron e, the muon μ and the tau τ. They carry electric charge -1 in units of the elementary charge e. Similarly, corresponding electrically neutral leptons called neutrinos exist in the same generations. The SM contains no right-handed neutrino field, so they are approximated to be massless. Neutrino flavor oscillations, however, establish that the differences between the squared masses of the neutrino mass eigenstates are non-zero, and therefore that at least two of the three masses are non-zero. Oscillations are insensitive to the absolute mass scale, which is at present constrained only from above. This is one of the main shortcomings of the model, as described in more detail in Section 1.7. Charged leptons interact via both the electromagnetic and the weak force, while neutrinos interact only weakly.

Quarks are another type of fermion that exist in six *flavors*, consisting, for each generation, of an *up-type* quark of electric charge $+2/3$ and a *down-type* quark with charge $-1/3$. Namely, there are up u and down d quarks in the first generation, charm c and strange s in the second generation, and finally top t and bottom b in the third. Contrary to leptons, quarks carry color charge, and can therefore interact via the strong force. Each quark carries one of three different colors (red, green and blue), and similarly each antiquark can have a corresponding anticolor. Quarks do not exist individually in nature, but rather are always confined within color-neutral composite particles known as hadrons. These bound states exist in a color–anticolor combination, or in a combination

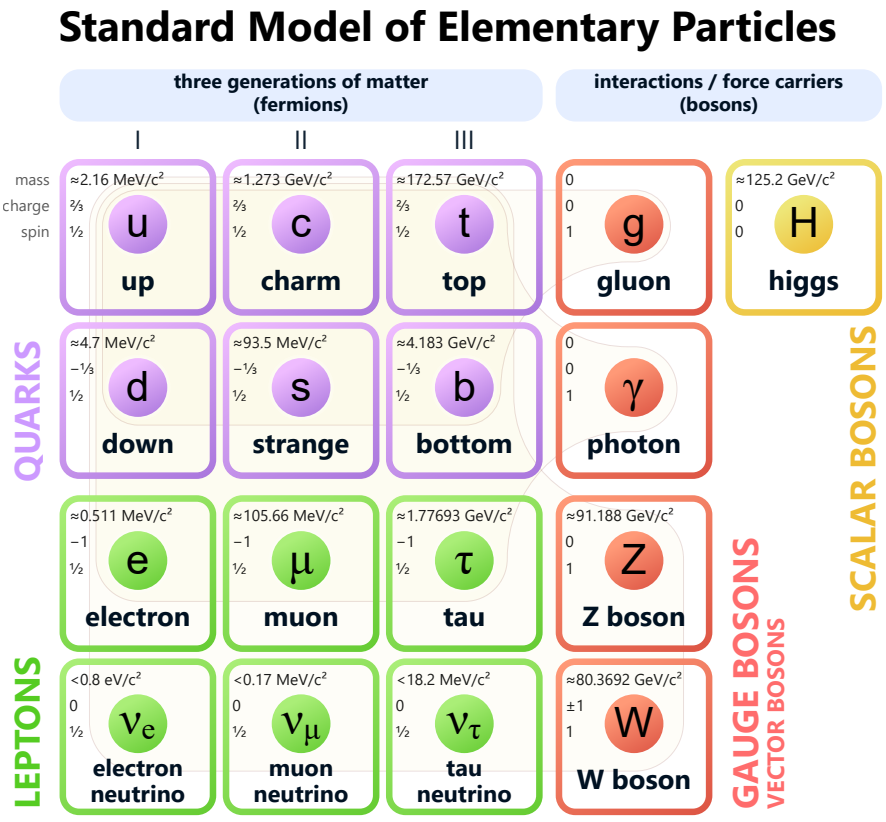

Figure 1.1: An overview of the fundamental particles in the Standard Model [7].

of all three colors. There are thus two main types of hadrons: mesons, which are made of a quark and an antiquark, and baryons (or antibaryons) consisting of three quarks (or antiquarks). The proton p and the neutron n are such particles, as they consist of the combination of uud and udd quarks, respectively. However, exotic states such as *tetraquarks* and *pentaquarks* have also been observed [8].

All fermions have a common intrinsic property called *chirality*, which is critical in weak interactions: in the massless limit, chirality corresponds to helicity, which defines the alignment of the spin relative to the direction of motion of the particle, that is, particles can be right- or left-handed, coinciding with parallel or antiparallel alignments. In this context, the weak force couples only to fermions in left-handed chiral states. Quarks and charged leptons exhibit both left- and right-handed states, while only left-handed neutrinos exist in the SM.

The last piece of the puzzle of the SM, as well as the most recent fundamental particle to be discovered, is the Higgs boson. It is a massive, electrically neutral spin-0 boson with a measured mass of $m_h \approx 125\,\text{GeV}$ [9], and is the physical excitation of the field responsible for the mechanism that gives mass to the weak gauge bosons and, through the Yukawa interactions, to the charged fermions. This mechanism accounts for the masses of the elementary particles, as described in Section 1.3.

1.2 Symmetries in the SM

The SM is formulated as a quantum field theory (QFT), where all particles are defined as quanta of various fields. Fermions are described by spin-1/2 fields and gauge bosons by spin-1 fields, also called vector fields. The SM Lagrangian controls the dynamics of the system, from which all equations of motion are derived. This is established to be *invariant* under certain transformations, and, by Noether's theorem [10], each continuous *symmetry* relates to a conserved quantity within the system. From these symmetries, a set of fundamental conditions is derived that rules the interactions allowed between the different fields.

The first condition is that the QFT must be invariant under the global transformations defined by the Poincaré group. These are the fundamental transformations in space and time, whose invariance leads to the conservation of energy and linear momentum, together with the rotations and boosts of the Lorentz group, rotational invariance being related to the conservation of angular momentum. Moreover, the fields must be invariant under a set of local phase transformations, also known as *gauge transformations*.

The gauge principle can be exemplified by the case of quantum electrodynamics (QED), which is the QFT describing electromagnetic interactions. We introduce the Lagrangian of a free spin-1/2 field ψ with mass m [11]:

$$\mathcal{L} = \bar{\psi}(x)(i\gamma^\mu \partial_\mu - m)\psi(x) \tag{1.1}$$

where γ^μ, with $\mu = 0, 1, 2, 3$, are the Dirac matrices. The Lagrangian is invariant under a global $U(1)$ transformation, where the field transforms as

$$\psi \to e^{iq\theta}\psi(x) \quad \bar{\psi} \to e^{-iq\theta}\bar{\psi}(x) \tag{1.2}$$

Nevertheless, it is not invariant a priori under a local phase transformation such that θ depends on the space-time coordinate and the field transforms as

$$\psi \to e^{iq\theta(x)}\psi(x) \quad \bar{\psi} \to e^{-iq\theta(x)}\bar{\psi}(x) \tag{1.3}$$

Invariance under this type of local transformation is restored by replacing the derivative in the Lagrangian with the covariant derivative

$$D_\mu \equiv \partial_\mu + iqA_\mu(x) \tag{1.4}$$

where $A_\mu(x)$ is a vector field which transforms according to

$$A_\mu(x) \to A_\mu(x) - \partial_\mu\theta(x) \tag{1.5}$$

Thus, the QED Lagrangian can be rewritten, including the new vector field, as

$$\mathcal{L}_{\text{QED}} = [i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi] - q\bar{\psi}\gamma^\mu\psi A_\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{1.6}$$

where

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{1.7}$$

The first expression within brackets in Eq. (1.6) relates to the free massive fermionic field, where the first two terms represent the kinetic energy and mass, respectively. The third term in the equation relates to the interaction between the fermionic field and the vector field A_μ, which is identified as the photon field introduced in Eq. (1.4). The last term in the equation is the free-field Lagrangian for a massless vector field, introducing the kinematics of the photon. All these terms are now invariant under the combined gauge transformations outlined in Eqs. (1.3) and (1.5).

A mass term for the vector field, of the form $\frac{1}{2}m^2A_\mu A^\mu$, is not invariant, since it is not left unchanged by the transformation of Eq. (1.5). Local gauge invariance therefore requires the photon to be massless, in agreement with the infinite range of the electromagnetic interaction. The same argument applied to the weak sector, whose mediators are known experimentally to be massive, is what motivates the discussion of Section 1.3.

This principle of local gauge invariance creating the interactions between fields, and the associated conservation of charges, can be extended to all types of interaction outlined by the SM. The corresponding symmetries are described by the gauge group

$$SU(3)_C \times SU(2)_L \times U(1)_Y \tag{1.8}$$

The strong force is obtained by requiring symmetry under color charge transformations, defined by the symmetry group $SU(3)_C$, with the associated gauge bosons being the eight gluons. On the other hand, the electromagnetic and weak interactions are combined into the unified electroweak interaction, corresponding to the invariance under the $SU(2)_L \times U(1)_Y$ symmetry. $SU(2)_L$ is the weak isospin group related to the left-handed fermions, with its associated weak isospin bosons W^1, W^2 and W^3. Finally, $U(1)_Y$ is the weak hypercharge group, with its related hypercharge boson B.

The complete field content of the SM, together with the representation under which each field transforms, is summarized in Table 1.1.

Table 1.1: Field content of the Standard Model and the representations under which each field transforms, shown for one generation of fermions [9].

Field	Spin	$SU(3)_C$	$SU(2)_L$	$U(1)_Y$
$Q_L = (u_L, d_L)^\mathsf{T}$	1/2	**3**	**2**	$+1/6$
u_R	1/2	**3**	**1**	$+2/3$
d_R	1/2	**3**	**1**	$-1/3$
$L_L = (\nu_L, e_L)^\mathsf{T}$	1/2	**1**	**2**	$-1/2$
e_R	1/2	**1**	**1**	-1
G^a_μ	1	**8**	**1**	0
W^i_μ	1	**1**	**3**	0
B_μ	1	**1**	**1**	0
ϕ	0	**1**	**2**	$+1/2$

The left-handed chiral fermion fields appear as $SU(2)_L$ doublets, pairing an up-type with a down-type quark and a charged lepton with its associated neutrino, while the right-handed fields are $SU(2)_L$ singlets. In $SU(3)_C$ the quarks are color triplets, corresponding to the three color charges, while the gluons form a color octet, the adjoint representation of the group. The weak gauge bosons form a weak isospin triplet, and the $U(1)_Y$ symmetry of weak hypercharge transformations is one-dimensional.

In addition to the conserved gauge charges, the SM possesses *accidental* global $U(1)$ symmetries, which are not imposed but arise as a consequence of requiring gauge invariance given the field content of Table 1.1. The associated conserved quantities are the baryon number B and the individual lepton numbers L_ℓ ($\ell = e, \mu, \tau$), together with the total lepton number $L = L_e + L_\mu + L_\tau$. The baryon number takes the value $+1/3$ for quarks and L_ℓ the value $+1$ for leptons, with opposite signs for the corresponding antiparticles, and both are conserved in all perturbative SM interactions [9]. For strong and electromagnetic interactions the individual quark flavors are also conserved, while flavor-changing processes are allowed by the weak interaction. The fact that B and L are accidental rather than gauged symmetries is what allows extensions of the SM to violate

them, as discussed in Chapter 2.

1.3 Electroweak unification and spontaneous symmetry breaking

The $SU(2)_L \times U(1)_Y$ sector of the SM contains four gauge fields, namely W^1_μ, W^2_μ, W^3_μ, and B_μ. Empirically, these correspond to the interaction with the photon in QED, and to the existence of two charged and one neutral massive boson mediating the weak force. The respective eigenstates are obtained by a mixing of electroweak eigenstates as follows [9]

$$W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \pm iW^2_\mu) \tag{1.9}$$

$$A_\mu = W^3_\mu \sin\theta_W + B_\mu \cos\theta_W \tag{1.10}$$

$$Z_\mu = W^3_\mu \cos\theta_W - B_\mu \sin\theta_W \tag{1.11}$$

Here, the fields denoted by $W^\pm_\mu$ and Z_μ correspond to the W and Z bosons, and A_μ is the photon field defined in Eq. (1.4). θ_W is known as the weak mixing angle and relates to the mixing between the neutral $SU(2)_L \times U(1)_Y$ gauge fields.

The theory up to this point describes massless gauge bosons, as required for the photon by the argument given in Section 1.2. Introducing mass terms, as $m^2_W W_\mu W^\mu$, violate the $SU(2)_L \times U(1)_Y$ invariance and are therefore forbidden. The same holds for the fermions: an explicit Dirac mass term $m\bar{\psi}\psi$ couples the left- and right-handed components of the field, which transform in different $SU(2)_L$ representations, as listed in Table 1.1, and is consequently not gauge invariant either. These are contradicted in experimental observations, since the weak interaction is observed to have a short range and the W and Z bosons, as well as the charged fermions, are all observed to be massive. This contradiction is resolved by introducing the *Higgs field* ϕ.

The Higgs field is a complex scalar field transforming as an $SU(2)_L$ doublet with hypercharge $Y = +1/2$, as listed in Table 1.1. The Lagrangian built from it is invariant under $SU(2)_L \times U(1)_Y$, while its ground state is not, and it is this ground state that breaks the symmetry. For such a ground state to be generated, the following scalar potential is added to the SM Lagrangian:

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda(\phi^\dagger \phi)^2 \tag{1.12}$$

The quartic coupling must satisfy $\lambda > 0$ for the potential to be bounded from below. If in addition $\mu^2 < 0$, the minimum of the potential does not lie at the origin, but on the continuously degenerate manifold defined by $\phi^\dagger\phi = v^2/2$, where $v = \sqrt{-\mu^2/\lambda} \approx 246\,\text{GeV}$ [9]. A particular point on this manifold is selected by a gauge choice, conven-

tionally taken to be

$$\phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \tag{1.13}$$

A perturbation of the Higgs field can be defined at the minimum value as follows

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \tag{1.14}$$

where $h(x)$ relates to excitations of the field, corresponding to the experimentally observed Higgs boson. The interactions with the gauge bosons, treated similarly to the Higgs field as outlined in Eq. (1.14), lead to the mass terms for the $W^{\pm}$ and Z bosons. Three out of the four scalar degrees of freedom of the Higgs field generate the masses of the $W^{\pm}$ and Z bosons, and the fourth results in the Higgs boson. The requirement of a symmetry-breaking ground state therefore gives rise to a gauge-invariant description of the masses of the electroweak gauge bosons. This process is commonly known as the Brout–Englert–Higgs mechanism [4–6].

At tree level, the resulting masses are $m_W = gv/2$ and $m_Z = m_W / \cos\theta_W$, where g is the $SU(2)_L$ gauge coupling, while the neutral combination of Eq. (1.10) remains massless and is identified with the photon. The breaking pattern is therefore $SU(2)_L \times U(1)_Y \to U(1)_{\mathrm{EM}}$, where the unbroken subgroup is generated by $Q = T^3 + Y$ and the photon is its massless gauge boson.

Finally, the fermion masses are generated by further $SU(2)_L \times U(1)_Y$ invariant interaction terms with the Higgs field, known as the Yukawa couplings. In the context of quark fields, these result in a mixing of the weak eigenstates that produces the experimentally observed mass eigenstates, which provides an explanation for the flavor-changing transitions in the decay of a quark from one generation to another. This mixing is described by a 3×3 unitary matrix known as the Cabibbo–Kobayashi–Maskawa (CKM) matrix [12, 13], which relates the weak eigenstates of the down-type quarks to their mass eigenstates.

1.4 Quantum chromodynamics and asymptotic freedom

As previously mentioned, gluons are the mediators responsible for the strong force, and they exist in eight superpositions of color–anticolor states. Hence, gluons carry a net color charge and can undergo self-interactions, unlike the photon in QED. As a consequence of renormalization, the strong coupling depends on the scale at which it is probed. This dependence is governed by the QCD beta function, defined as the

derivative of the coupling with respect to the renormalization scale μ_R,

$$\mu_R^2 \frac{\partial \alpha_S}{\partial \mu_R^2} = \beta(\alpha_S) = -\alpha_S^2 \left(b_0 + b_1 \alpha_S + b_2 \alpha_S^2 + \dots\right), \qquad b_0 = \frac{33 - 2n_f}{12\pi}, \tag{1.15}$$

where n_f is the number of quark flavors active at the scale μ_R [9]. The two terms in b_0 have opposite physical origins: the constant $33 = 11N_c$ arises from gluon loops, which antiscreen the color charge, while $2n_f$ arises from quark loops, which screen it as in QED. Since $n_f \leq 6$ in the SM, the gluon self-coupling dominates, $b_0 > 0$, and the beta function is therefore negative. This is the origin of *asymptotic freedom* [14, 15], the weakening of the strong coupling constant α_S at short distances, and, similarly, at large momentum transfers. Its measured value at the scale of the Z boson mass is $\alpha_S(m_Z^2) \approx 0.118$, and it grows as the momentum transfer approaches $\Lambda_{\text{QCD}} \approx 200\,\text{MeV}$, the scale at which the perturbative expansion breaks down [9]. The predicted and experimentally measured values of α_S are shown in Figure 1.2.

The complementary behavior at large distances, or at small momentum transfer, is known as *confinement*, establishing that color-charged particles cannot be isolated. Its consequence is that, when two particles carrying color charge move apart, the energy stored in the gluon field between them increases until it becomes energetically favorable to create a new quark–antiquark pair from the vacuum. The color field is then broken, and the newly created quarks screen the color charges of the original pair. The strong interaction therefore has a short effective range, despite gluons being massless mediators.

When quarks or gluons have enough energy to escape the confining potential, the energy of the gluon field is converted into the formation of quark–antiquark pairs, which then bind into new hadronic bound states in a process known as *hadronization*. When the energy is high enough to break out of the confinement of a hadron, a hadronization chain is started, creating a series of collimated hadrons also known as "jets".

1.5 Cross sections and particle decays

Cross sections and particle decay rates are two critical concepts in experimental particle physics, as they express computable quantities for the process in which an initial particle state transitions into a final state via a set of interactions.

Cross sections quantify the probability that a certain process will take place. For a $2 \to 2$ process, such as the electron–positron annihilation whose Feynman diagrams are shown in Figure 1.3, where two initial particle states transition to two final states, the differential cross section is expressed as [9]

$$d\sigma = \frac{(2\pi)^4 \left|\mathcal{M}\right|^2}{4\sqrt{(p_1 \cdot p_2)^2 - m_1^2 m_2^2}} \times d\Phi\left(p_1 + p_2; p_3, p_4\right). \tag{1.16}$$

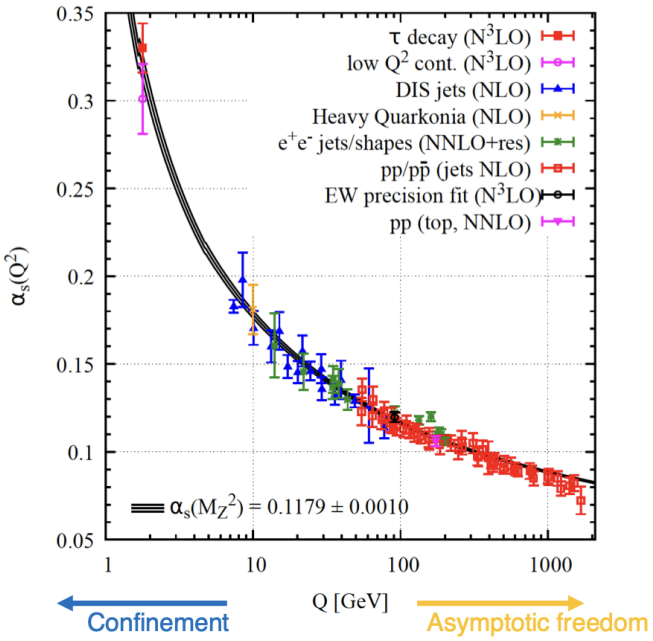

Figure 1.2: The observed values of α_S as a function of momentum transfer Q [9].

Here, p_1, p_2, m_1, m_2 are the four-momenta and masses of the initial state, and p_3, p_4 are the four-momenta of the final state particles. $d\Phi$ is the phase-space factor, representing the number of kinematically allowed states available to the final state particles, while $|\mathcal{M}|$ is the matrix element describing the dynamics of the interaction.

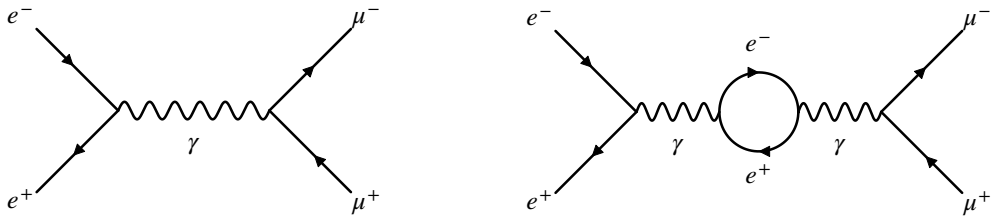

(a) (b)

Figure 1.3: The Feynman diagrams of electron–positron annihilation processes in leading order (a) and second order (b).

The matrix element is computed from the SM Lagrangian using perturbation theory, as the amplitude is expanded as a power series in the relevant coupling constant, with each order represented diagrammatically by Feynman diagrams containing an increasing number of interaction vertices. This is illustrated in the Feynman diagrams

of Figure 1.3. The leading order (LO) contribution retains only the diagrams with the minimum number of vertices, while higher orders, next-to-leading order (NLO), next-to-next-to-leading order (NNLO) and beyond, add virtual loop corrections and additional real emissions. The expansion is meaningful only where the coupling is small, which for QCD restricts perturbative calculations to large momentum transfers, as discussed in Section 1.4.

Loop diagrams introduce integrals over unconstrained internal momenta which diverge in the ultraviolet. These divergences are removed by *renormalization*, where the bare parameters of the Lagrangian are redefined so that all divergences are absorbed into a finite set of measured quantities. The couplings and masses then become functions of the scale at which they are probed, which is the origin of the running of α and α_S discussed in the previous sections.

Another property of fundamental particles is that they will normally decay into lighter particles, if allowed by the conservation laws in the SM. This is not true, for instance, for electrons, which are not allowed to decay as they are the lightest electrically charged particles in the SM.

An unstable particle decays into a given n-particle final state at a rate given by the partial width [9]

$$\Gamma_i = \frac{(2\pi)^4}{2M} \int |\mathcal{M}|^2 \, d\Phi_n(P; p_1, ..., p_n). \tag{1.17}$$

Here, M and P are the mass and four-momentum of the decaying particle, and p_i are the four-momenta of the decay products, while $|\mathcal{M}|$ and $d\Phi_n$ are the matrix element and the phase-space factor defined earlier. The total width is the sum of the partial widths over all accessible final states, and is the inverse of the mean proper lifetime τ of the particle,

$$\Gamma = \sum_i \Gamma_i = \frac{1}{\tau}, \tag{1.18}$$

while the ratio $\mathcal{B}_i = \Gamma_i/\Gamma$ defines the branching fraction of the channel i. The lifetime determines the probability that the particle survives for a time t in the laboratory frame before decaying:

$$P(t) = e^{-t/\gamma\tau} \tag{1.19}$$

where γ is the Lorentz factor.

The decay rate is governed both by the phase space available to the final state and by the structure of the matrix element. A decay is *phase-space suppressed* when the sum of the masses of the decay products approaches the mass of the parent, since $d\Phi_n$ then shrinks to zero. It is *dynamically suppressed* when the matrix element itself is small.

A similar mechanism, whether realized through a heavy mediator, a small coupling, or a near-degenerate mass spectrum, can also push the proper lifetime of a hypothetical particle into the range where it decays at a macroscopic distance from the interaction point. Such scenarios are discussed in Chapter 2.

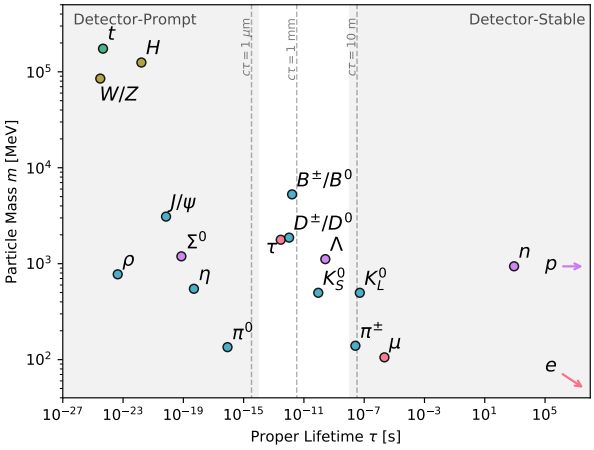

Figure 1.4: An excerpt of SM particles with their proper lifetimes and masses. The gray regions represent the detector-prompt and detector-stable sectors of lifetime space, for a particle traveling almost at the speed of light [16].

A selection of the masses and corresponding lifetimes of elementary and composite particles is shown in Figure 1.4. Particles are considered to be *long-lived* in the context of the resolution of the experimental instrumentation. For instance, a particle traveling a distance that can be resolved by the detector before it decays is said to be long-lived. More details are given in Section 2.3.

1.6 Phenomenology of proton–proton collisions

High-energy proton–proton (pp) collisions produced at the Large Hadron Collider (LHC) form the basis of the particle physics research presented in this thesis. Thus, this section describes the phenomenology of such collisions.

Starting with the protons, made of three valence quarks uud, together with *sea* quarks and gluons, in other words *partons*. The bulk of the pp interactions at the LHC are *soft*, involving a small momentum transfer between the colliding partons. In a small fraction of the collisions, two individual partons interact with a large momentum transfer, known as hard interactions, and they are normally the focus of physics analyses. Alongside the hard interaction, numerous QED and QCD processes take place. A schematic view of the evolution of an LHC collision between two protons is shown in Figure 1.5. The protons are denoted by the gray ellipses and the hard interaction by the red dot, as seen in Figure 1.5a. There, only one parton in each proton takes part in the process. The others can be seen as the black lines coming out of the proton, going on to generate the underlying event.

The partons then undergo a *parton shower*, illustrated in Figure 1.5b. Radiation emit-

ted by the incoming partons before the hard scattering is referred to as *initial state radiation* (ISR), while radiation emitted by the outgoing partons after it is referred to as *final state radiation* (FSR). Both are dominated by soft and collinear gluon emission, which generates a cascade of successive branchings ($q \to qg$, $g \to gg$, $g \to q\bar{q}$, shown as the curly lines) that evolves the system down from the hard scale toward the scale at which perturbation theory breaks down. ISR largely determines the transverse momentum distribution of the hard system, while FSR determines the internal structure of the resulting jets [17].

Finally, Figure 1.5c shows the hadronization process, where partons radiated from the shower eventually reach an energy low enough to bind with each other into observable, color-neutral hadrons. The newly formed hadrons potentially decay into further particles until only long-lived or stable particles are left. The partons not participating in the hard process (shown in brown) can themselves produce gluons (shown in blue) and participate in the hadronization process, shown in Figure 1.5d.

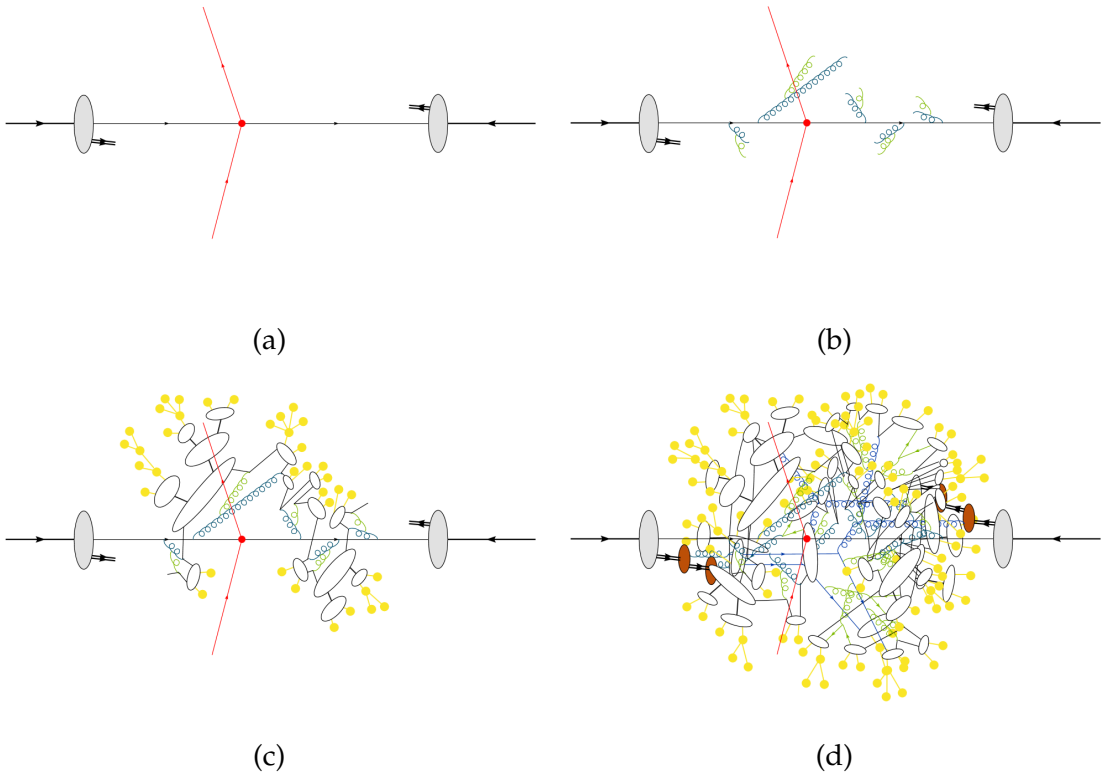

Figure 1.5: The various steps of a proton–proton collision involving one parton: hard process (a), parton shower (b), hadronization (c), fragmentation and underlying event (d) [18].

Hadronization occurs at momentum scales of $\mathcal{O}(1\,\text{GeV})$, while the hard scatter oc-

curs at the $\mathcal{O}(1\,\text{TeV})$ scale, at which QCD is perturbative and the matrix element can therefore be computed. The hadronization process, on the other hand, lies at non-perturbative scales, and can therefore only be modeled and computed phenomenologically. The QCD factorization theorem [19] allows the two to be treated as independent. The separation is defined by the collinear approximation, in which the partons entering the hard process are taken to be collinear with the parent hadron, so that the long-distance structure of the proton is absorbed into universal parton distribution functions and separated from the short-distance partonic cross section. That the partons can be treated as quasi-free over the short time scale of the hard interaction is a distinct statement, and follows from asymptotic freedom. The cross section for a proton–proton collision developing into a final state O can thus be expressed in the factorized form as follows [19]

$$\sigma_{pp\to O} = \sum_{i,j} \int_0^1 dx_i \int_0^1 dx_j \, f_i(x_i, \mu_F^2) \, f_j(x_j, \mu_F^2) \times \hat{\sigma}_{ij\to O}(x_i p_i, x_j p_j, \mu_F^2, \mu_R^2). \tag{1.20}$$

Here, $f_i(x_i, \mu_F^2)$ and $f_j(x_j, \mu_F^2)$ are known as the parton distribution functions (PDFs), $\hat{\sigma}_{ij\to O}$ is the partonic cross section, and the sum runs over all parton types i and j available in each proton that can take part in the process. The PDFs parameterize the probability of a given parton in the proton carrying a fraction x of the momentum of the proton itself, and are defined at the factorization scale μ_F, while μ_R is the renormalization scale introduced in Section 1.5. The PDFs are not derivable from first principles and therefore have to be determined from experimental data. They are measured at particular scales and then evolved to the energy scale relevant for the physical process under consideration.

1.7 Open questions in the Standard Model

The SM has been an undeniable success, providing accurate theoretical descriptions of fundamental particles and interactions, and its predictions were confirmed over the years by direct experimental observation: the W and Z bosons were discovered in 1983 by the UA1 and UA2 experiments at the CERN Super Proton Synchrotron [20–23], the top quark in 1995 by the CDF and D0 experiments at the Fermilab Tevatron [24, 25], and the Higgs boson in 2012 by the ATLAS and CMS experiments at the LHC [26, 27]. Measurements of SM observables have so far shown no confirmed deviation from the theory, although a number of tensions at the level of a few standard deviations persist [28] and are the subject of ongoing investigation.

Nevertheless, there are several critical open questions that the SM cannot explain, suggesting that extended theories are needed:

- **Dark matter:** A broad range of astrophysical and cosmological observations

shows that most of the matter in the Universe is not made of SM particles. The evidence includes the rotation curves of spiral galaxies, gravitational lensing, the dynamics of galaxy clusters and the formation of large-scale structure, and it is completed by fits to the Cosmic Microwave Background, which indicate that only roughly 15% of all matter in our Universe is conventional SM matter [29]. The nature of dark matter is still very much open to debate. What the observations do require is that it is non-luminous, that it is massive and non-relativistic, and that its lifetime is at least of the order of the age of the Universe itself. A coupling to the SM through the weak interaction is not required by these observations, but is the additional assumption that defines the weakly interacting massive particle (WIMP) hypothesis. A WIMP produced thermally in the early stages of the Universe and accounting for the observed relic density has a mass in the range of 100 to 1000 GeV. Such a particle would lie well within the electroweak scale and could potentially be observed in LHC searches. Extensive direct-detection and collider searches have nevertheless excluded much of the canonical WIMP parameter space, which motivates alternative scenarios in which dark matter resides in a hidden sector coupled only feebly to the SM. Such feeble couplings generically imply long-lived mediators, briefly discussed in Chapter 2.

- **The mass of neutrinos:** In the framework of the SM, neutrinos are assumed to be massless. However, neutrino flavor oscillations have been observed by the Super-Kamiokande and SNO experiments [30, 31], implying non-zero mass-squared differences between the neutrino mass eigenstates and therefore at least two non-zero neutrino masses. An extended framework is required to accommodate this, and a considerable neutrino physics program worldwide is under way to support it experimentally.

- **Matter–antimatter asymmetry:** The observable Universe is made almost exclusively of matter, whereas matter and antimatter are expected to have been produced in equal amounts in the Big Bang. Nearly all of it annihilated, leaving the small residual asymmetry observed today, quantified by the baryon-to-photon ratio $n_B/n_\gamma \approx 6 \times 10^{-10}$ [29].

- **The hierarchy problem:** Unlike fermion and gauge boson masses, the mass of a fundamental scalar receives quantum corrections that are quadratically sensitive to the highest scale Λ at which the theory remains valid, and at which its behavior is expected to change due to new physics. Loop-level diagrams have to be included in the computation of the physical mass of the SM Higgs boson, which is expressed with radiative corrections as

$$m_h^2 = m_{h,0}^2 + \Delta m_h^2, \tag{1.21}$$

where $m_{h,0}$ is known as the *bare* Higgs mass. The corrections are related to the cutoff scale Λ as

$$\Delta m_h^2 = -\frac{\lambda_f^2}{16\pi^2}\left[2\Lambda^2 + \mathcal{O}\left(m_f^2 \ln\left(\frac{\Lambda}{m_f}\right)\right)\right], \tag{1.22}$$

where λ_f is the Yukawa coupling parameter to the fermion field, and m_f is the respective fermion mass. The dominant correction comes from the top-quark loop seen in Figure 1.6, and if the SM is valid up to the Planck scale, $\Lambda \sim 10^{19}$ GeV, the observed value $m_h \approx 125$ GeV requires the bare mass parameter and the radiative corrections to cancel to more than thirty significant figures. No symmetry of the SM enforces this cancellation, which is widely regarded as unnatural.

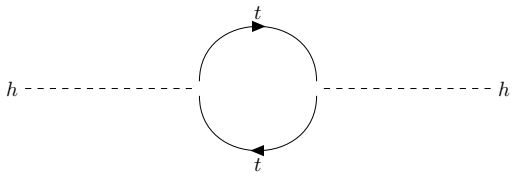

Figure 1.6: The correction to the mass of the Higgs boson via the top-quark loop.

This raises the question of *naturalness*, namely that a theoretical framework should ideally reproduce the observed scales without recourse to fine-tuning. A solution to this problem can be theoretically provided by the theory of supersymmetry, which will be described in Chapter 2.

- **Grand unification:** A Grand Unified Theory is a framework in which the electroweak and strong forces are merged into one unified force, consequently simplifying the overall theory and reducing the number of free parameters, of which the SM has nineteen. The concrete motivation is that the three gauge couplings, when evolved to high energy, approach one another near 10^{16} GeV, and that such theories naturally explain the quantization of electric charge [32]. The unification is not exact within the SM alone, and is improved in supersymmetric extensions. A generic prediction of these models is the decay of the proton, which is forbidden in the SM by the accidental conservation of baryon number discussed in Section 1.2, and which is constrained experimentally to a partial lifetime in excess of 10^{34} years [33].

- **The flavor structure:** The SM does not explain why the fermion masses span more than five orders of magnitude, nor why there are exactly three generations. The Yukawa couplings are free parameters fixed only by measurement.

These are just a few of the various issues that arose from the development of the

SM. What is clear is that a new physics theory is needed. The overall collection of theories extending, or partly superseding, the SM is known as physics beyond the Standard Model (BSM). The concepts introduced in this chapter, in particular the suppression mechanisms of Section 1.5 and the open questions listed above, are taken up in Chapter 2, where the specific BSM scenarios underlying the physics analysis presented in this thesis are described in detail.

Chapter 2

Beyond the Standard Model

The open questions left by the SM have motivated a large number of theories that could, at least partially, resolve them. One of the most studied is supersymmetry (SUSY), which offers a solution to the hierarchy problem and, in its R-parity conserving form, a possible dark matter candidate. SUSY extends the SM by adding a new set of particle states. This chapter introduces its theoretical framework and describes how these states would be produced at the LHC, with emphasis on the case in which they are long-lived, which is the signature targeted by the physics analysis presented in Chapter 11.

2.1 The theory of supersymmetry

SUSY [34] theorizes a symmetry between bosons and fermions, under which transformations between fermionic and bosonic fields are permitted. Its fundamental component is the SUSY generator $\hat{Q}$, which turns fermion states into boson states and the other way around:

$$\hat{Q}|\text{fermion}\rangle = |\text{boson}\rangle \qquad \hat{Q}|\text{boson}\rangle = |\text{fermion}\rangle \tag{2.1}$$

The symmetry is a space-time symmetry: $\hat{Q}$ is a fermionic generator carrying spin $1/2$, in contrast to the internal symmetry bosonic generators of the SM. Every SM boson therefore acquires a fermionic superpartner, and every SM fermion a bosonic one. These new states, known as *sparticles*, carry the same gauge quantum numbers as their SM equivalents and differ from them by half a unit of spin. Hence, the partners of the leptons and quarks are the spin-0 *sleptons* and *squarks*, while the partners of the SM gauge bosons are the spin-$1/2$ *gauginos* [34].

Recall Eq. (1.22), which gives the quadratically divergent correction to the Higgs mass produced by a fermion loop. In a supersymmetric theory, each complex scalar S that couples to the Higgs field with coupling λ_S contributes an additional correction of the opposite sign:

$$(\Delta m_h^2)_S = \frac{\lambda_S}{16\pi^2}\left[\Lambda^2 - 2m_S^2 \ln\left(\frac{\Lambda}{m_S}\right) + \dots\right] \tag{2.2}$$

Here Λ is the cutoff scale of the theory and m_S is the mass of the scalar. Each SM fermion is accompanied by two complex scalar partners, one for each chiral state, so that the sum of their contributions cancels the Λ^2 term of Eq. (1.22) exactly, provided that SUSY is an exact symmetry and that the masses and couplings of sparticles are the

same. This means that superpartners degenerate with the SM particles would in any case already have been observed in current experiments. Since they have not, SUSY must be broken. The breaking is introduced through explicit soft terms, which lift the sparticle masses while preserving the cancellation of the quadratic divergence. Many models attempt to explain how this breaking occurs, but the simplest supersymmetric extension of the SM, introducing the smallest number of new particles, is the Minimal Supersymmetric Standard Model (MSSM).

2.1.1 The Minimal Supersymmetric Standard Model

In the MSSM, each field and its superpartner are expressed together as a supermultiplet operating on the SUSY algebra. Chiral supermultiplets pair a spin-1/2 field, that is a fermion, with a complex spin-0 field, while gauge supermultiplets pair a spin-1 field with a spin-1/2 field and thereby describe the vector bosons and their partners.

The Higgs sector requires two chiral supermultiplets rather than one, containing the doublets H_u and H_d together with their spin-1/2 superpartners, the higgsinos. The two doublets carry eight real scalar degrees of freedom. Three of them become the longitudinal polarizations of the $W^\pm$ and Z bosons after electroweak symmetry breaking, by the mechanism described in Section 1.3, and five physical states remain: three neutral scalars h^0, H^0, and A^0, and a charged pair $H^\pm$. In the decoupling limit, in which A^0 is much heavier than the Z boson, the lightest state h^0 acquires SM-like couplings and is identified with the observed Higgs boson.

Once electroweak and supersymmetry breaking contributions are both included, states carrying the same conserved quantum numbers mix. The neutral higgsinos mix with the neutral gauginos to give four neutralinos, denoted $\tilde{\chi}_i^0$ and numbered from lightest to heaviest, while the charged higgsinos mix with the charged winos to give two charginos $\tilde{\chi}_i^\pm$. These states are known collectively as electroweakinos. Their masses and their composition both depend on the mixing, so a given eigenstate can be predominantly bino-like, wino-like or higgsino-like, which in turn governs its production cross section and its decay modes [34]. The same left-right mixing occurs among the squarks and sleptons, though there it is significant only in the third generation, where it is enhanced by the top, bottom and tau Yukawa couplings.

2.1.2 *R*-parity

Supersymmetric theories introduce terms that violate lepton and baryon numbers, expressed as follows:

$$W_{\Delta L=1} = \frac{1}{2}\lambda_{ijk} L_i L_j E_k^c + \lambda'_{ijk} L_i Q_j D_k^c + \mu_i L_i H_u \tag{2.3}$$

$$W_{\Delta B=1} = \frac{1}{2}\lambda''_{ijk} U_i^c D_j^c D_k^c \tag{2.4}$$

Here i, j and k label the different generations, and λ_{ijk}, λ'_{ijk}, λ''_{ijk} and μ_i are the couplings of each interacting term.

In order to cancel the possible baryon and lepton number violating terms in the Lagrangian, a new symmetry is introduced: R-parity. It is a quantum number expressed as

$$P_R = (-1)^{3(B-L)+2s} \tag{2.5}$$

where B and L are the baryon and lepton numbers and s is the spin. SM particles have inherent positive R-parity, while their superpartners have negative R-parity.

The conservation of P_R has fundamental consequences that shape the experimental signature of SUSY at a collider such as the LHC. This conservation makes the lightest supersymmetric particle (LSP) absolutely stable. If the LSP is electrically neutral and interacts only weakly, it traverses the detector without depositing energy, which produces the missing transverse momentum signature on which most SUSY searches at the LHC are based, but also makes it a interesting, potential dark matter candidate.

2.2 Experimental searches for SUSY

Experimental programs searching for SUSY, such as those carried out at the LHC, proceed both directly and indirectly. Direct searches analyze the possible decay products of the new particles produced in pp collisions, while indirect searches use precision measurements of electroweak observables to look for the radiative effects of these particles.

At the LHC, SUSY particles can be produced in pp collisions via electroweak or strong interactions, for instance as shown in Figure 2.1 and Figure 2.2. Production cross sections for a selection of SUSY processes at $\sqrt{s} = 13\,\text{TeV}$ are shown in Figure 2.3.

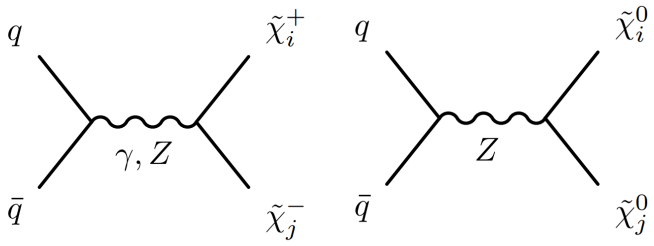

(a) (b)

Figure 2.1: Examples of Feynman diagrams for the electroweak production of sparticles.

Searches at the LHC take advantage of the rare signatures that SUSY particles would leave in the detector in order to construct a selection of events useful for the analysis. Monte Carlo (MC) simulations are generated on grids of different masses, covering the

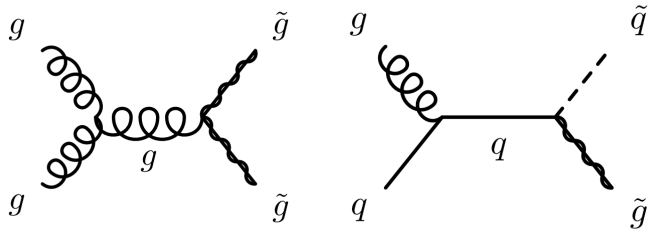

(a) (b)

Figure 2.2: Examples of Feynman diagrams for the strong production of sparticles.

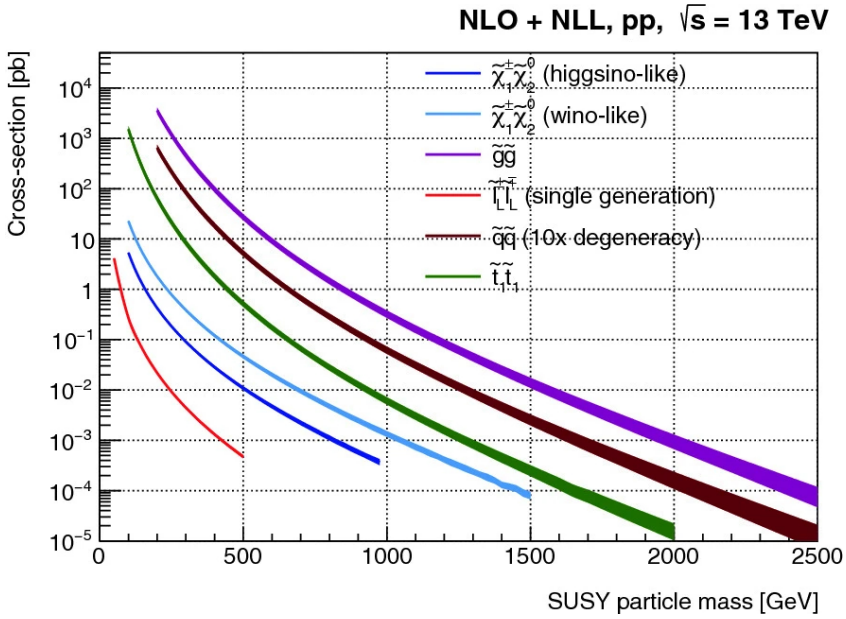

Figure 2.3: Cross sections for the production of SUSY particles in pp collisions at $\sqrt{s} =$ 13 TeV as a function of the particle mass [35].

various scenarios in which the searched-for particles could appear. They provide an efficient way to look for a potential signal and to estimate the likelihood of observing it in the detector. If the observation agrees with the SM prediction, exclusion limits at 95% confidence level on the masses of the SUSY particles in question are set using these same grids, providing a starting point for future searches and a basis for comparison between experiments [9].

2.3 Long-lived particles

Recall Section 1.5, where the decay rates of SM particles and the suppression mechanisms that give them a longer lifetime were described. The same mechanisms can occur in BSM physics, effectively generating long-lived, "exotic" particles (LLPs) [16]. Many BSM models predict long-lived sparticles. Within SUSY, examples include Split SUSY [36, 37], gauge-mediated supersymmetry breaking (GMSB) [38] and, most relevant for this thesis, R-parity-violating (RPV) SUSY [39], which is described below.

A particle of proper lifetime τ produced with velocity β and Lorentz factor γ travels, on average, a laboratory distance

$$d = \beta\gamma c\tau \tag{2.6}$$

before decaying. The signature observed in the detector depends on where this decay occurs, and therefore on the interplay between $c\tau$ and the size of the detector subsystems. The three mechanisms that generate a macroscopic $c\tau$ in BSM models are the same ones identified for the SM in Section 1.5: a small coupling governing the decay, a small mass splitting between the decaying particle and its decay products, and a decay proceeding through a heavy virtual mediator [16].

The resulting signatures, some of which are displayed in Figure 2.4, fall into a small number of classes, each requiring dedicated reconstruction techniques. Decays inside the tracking volume give displaced vertices or displaced lepton pairs; a charged parent whose decay products are too soft to be reconstructed gives a disappearing track; a decay to photons gives non-pointing photons; and a charged parent that traverses the whole detector gives a heavy stable charged particle signature [16]. The case relevant to this thesis is a decay occurring inside the calorimeter volume. The decay products of such a particle reach a given calorimeter cell later than those of a prompt particle originating from the same collision, for two reasons: the parent travels with $\beta < 1$ over the first leg of the trajectory, and the broken path from the interaction point to the decay vertex and then to the cell is longer than the direct line between the interaction point and that cell. The result is a positive time offset with respect to the nominal bunch crossing time, which is the delayed jet signature studied in Chapter 11 and which relies on the timing performance of the calorimeter described in Chapter 5.

In the MSSM, R-parity is taken as an exact symmetry, so the RPV couplings λ, λ' and

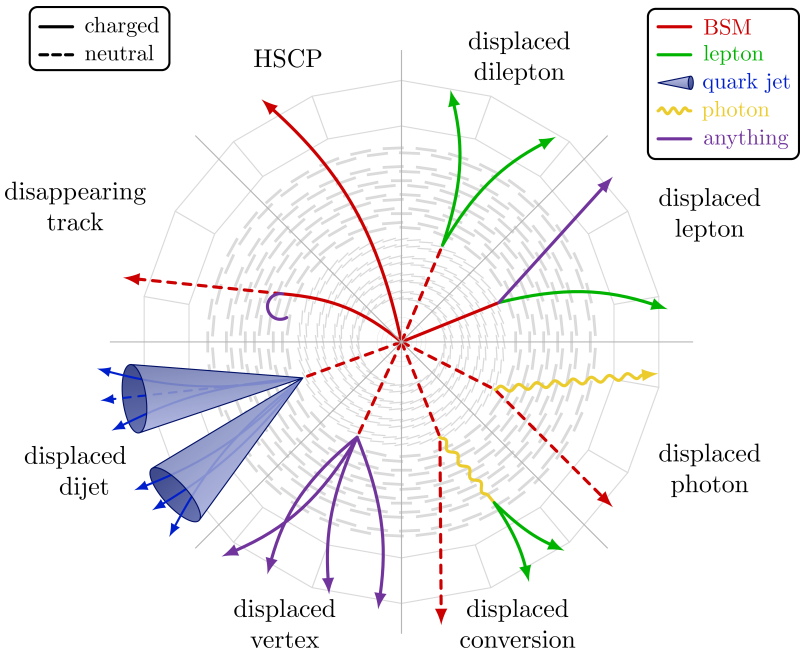

Figure 2.4: A diagram showing the various signatures targeted by LLP searches in ATLAS/CMS. Taken from [40].

λ'' of Eqs. (2.3) and (2.4) are set to zero. In other models, the limits on B- and L-violating processes are relaxed and R-parity is violated. This removes the consequences such as the one mentioned in Section 2.1.2. Here, the LSP is no longer stable but decays to SM particles, so the missing transverse momentum on which conventional SUSY searches rely is strongly reduced, and the final state instead contains the high multiplicity of jets or leptons produced by the LSP decay. The couplings are still assumed to be suppressed, which gives a sparticle decaying via an RPV coupling a lifetime long enough for it to be considered long-lived. This is the case for the RPV electroweak neutralino signal model, which is the subject of the search presented in Chapter 11.

Part II

Experimental Setup

Chapter 3

The Large Hadron Collider

The LHC [41, 42] is the world's largest and most powerful particle accelerator, with a circumference of approximately 26.7 km, housed in a tunnel at a depth ranging from 50 m to 175 m below the surface. It is installed in the tunnel that previously hosted the Large Electron–Positron collider (LEP) at CERN until 2000. The LHC was designed to collide protons at a center-of-mass energy of up to 14 TeV, a target motivated by both theoretical predictions and experimental results from the LEP and Tevatron experiments, which pointed to new physics at the TeV scale. Despite the Standard Model predicting the existence of the Higgs boson, there was no prior indication of its mass; a hadron collider was therefore the optimal choice to discover it, since proton–proton collisions at a center-of-mass energy of 14 TeV allow to explore a continuum of masses due to the composite nature of the proton. This leads to a wide spectrum of final states, maximizing the discovery potential. Moreover, an electron collider at the same energy in the existing 27 km tunnel would not be viable due to the large energy losses from synchrotron radiation, which scale as m^{-4}, where m is the mass of the accelerated particle. Since protons are approximately 2000 times heavier than electrons, these losses are negligible in comparison, making a proton collider far more efficient at the energies required.

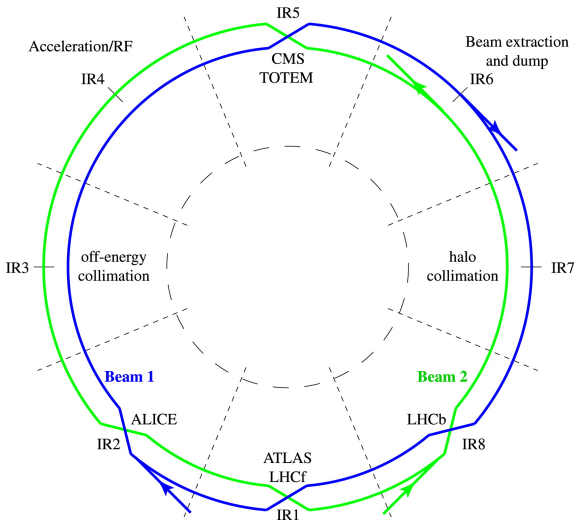

Figure 3.1: The LHC ring with the different interaction regions. Beam 1 circulates clockwise and Beam 2 counter-clockwise [43].

Circular in shape, the ring is divided into eight sectors, each consisting of a 2.4 km arc and a straight section. In the arcs, superconducting dipole magnets bend the beam to

maintain its circular orbit, while the straight sections house the interaction points (IPs) and the superconducting radio-frequency (RF) cavities that provide the acceleration.

There are four main IPs, each dedicated to one of the main experiments at the LHC. The remaining straight sections are used for beam cleaning systems, the superconducting RF cavities, and the beam dumping system, as shown in Figure 3.1. The LHC primarily serves as a proton–proton collider. The two general-purpose experiments, ATLAS (IP1) and CMS (IP5), operate at high luminosity. The other two main experiments operate at lower luminosity: LHCb (IP8) is dedicated to the study of b-physics, while ALICE (IP2) is designed to study heavy-ion collisions and the quark–gluon plasma during the heavy-ion operation of the LHC, in which ions such as lead or oxygen are collided.

3.1 Accelerating protons into collisions

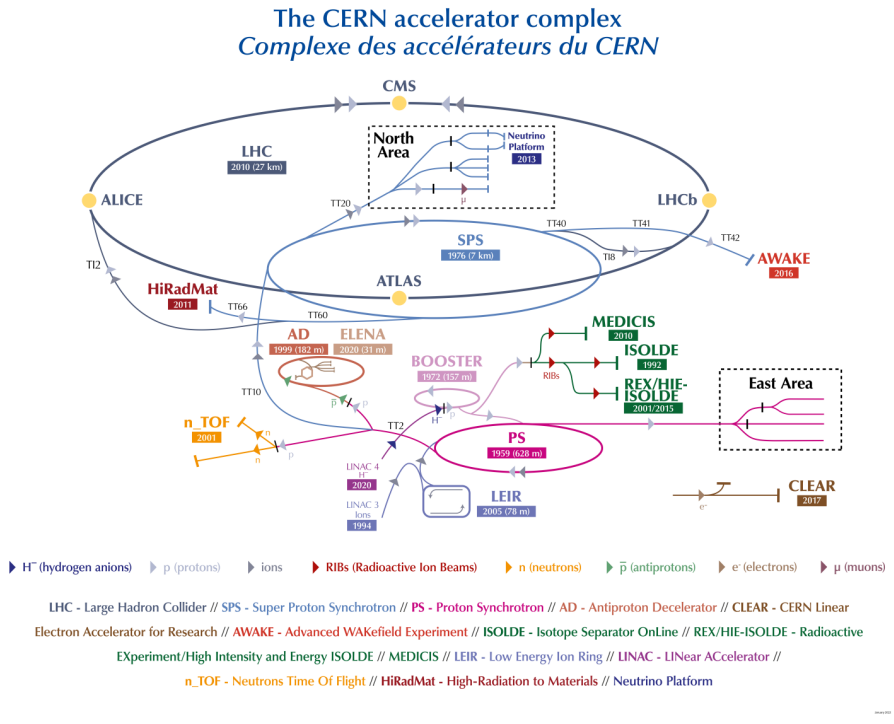

Figure 3.2: The CERN accelerator complex with its pre-accelerators [44].

The acceleration chain begins with a source of negative hydrogen ions (H^-), each consisting of a proton with two electrons. The ions leave the source with an energy of approximately 45 keV, but before they can be injected into the LHC, they must pass through several successive accelerators that extract the protons and accelerate them to nominal energy. The entire accelerator complex used for this is illustrated in Figure 3.2.

First, the ions are accelerated to 3 MeV by a Radio-Frequency Quadrupole (RFQ) and then enter Linac4, an 86-meter-long linear accelerator comprising several types of accelerating structures, which brings them to an energy of 160 MeV. At this point, the H^- ions are injected into the Proton Synchrotron Booster (PSB), where a thin carbon stripping foil removes the two electrons, leaving bare protons. The PSB accelerates these protons to 2 GeV, while also splitting the beam into several consecutive bunches. The protons are then injected into the Proton Synchrotron (PS), where they are accelerated to 26 GeV and further split into more bunches. Next, the beam enters the Super Proton Synchrotron (SPS), a 6.9 km circumference circular accelerator, where it is accelerated to 450 GeV and finally injected into the LHC, close to IP2 for Beam 1 and IP8 for Beam 2. This injection process is repeated several times per ring to fill the LHC according to the nominal filling scheme. During Run 3, the beams are accelerated to a top energy of 6.8 TeV per beam, with up to approximately 2460 colliding bunches.

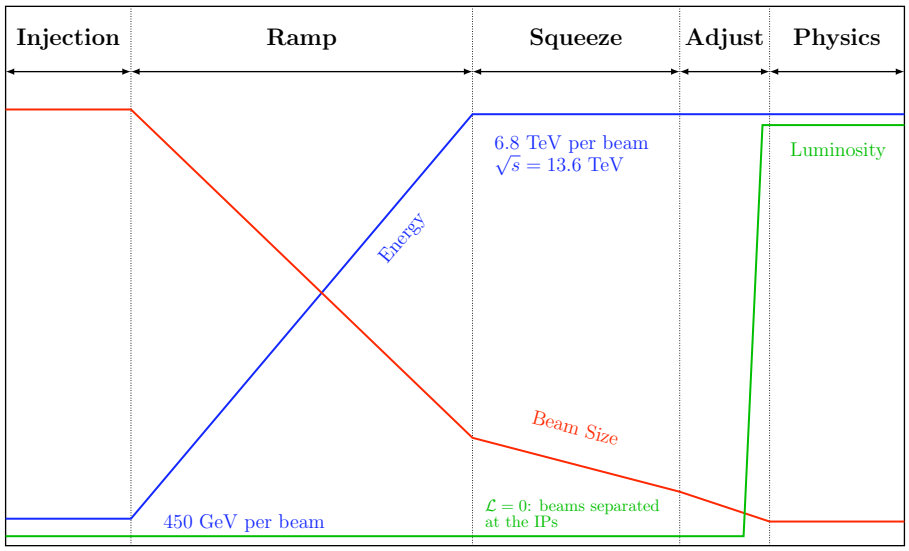

Figure 3.3: Schematic evolution of energy, luminosity and beam size during the steps in the filling scheme.

Once in the LHC for a so-called physics *fill*, the beams go through several steps before reaching the physics stage, at which the proton bunches are at their nominal state and ready to be collided. These steps are schematically shown in Figure 3.3. After *injection*, the beams are *ramped* to high energy, then *squeezed* to a smaller β^*[1], and *adjusted*, resulting in a increase in luminosity, and are finally brought into collision at the interaction point.

[1]The β^* is the value of the betatron function at the interaction point. A smaller β^* corresponds to a more tightly focused beam at the collision point, which increases the probability of proton–proton interactions and thus the luminosity.

3.2 LHC filling scheme

The LHC RF frequency, or how fast the accelerating voltage oscillates, is 400.79 MHz and the revolution frequency of the protons is 11.245 kHz. This yields a total of 35,640 RF buckets around the machine, each of which can in principle be filled with a proton bunch. The nominal bunch crossing frequency is 40 MHz, corresponding to a bunch spacing of 25 ns. There are 3564 possible bunch slots per ring, each identified by a unique bunch crossing identifier (BCID).

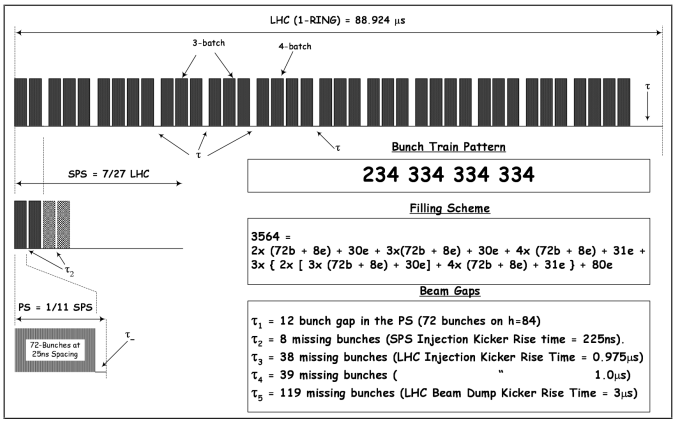

Figure 3.4: Schematic of the nominal 25 ns filling scheme of the LHC. The beam is divided into 39 batches of 72 bunches each, spaced at 25 ns, with a total of 3564 available bunch crossing slots [45].

The injections from the SPS have a bunch train structure, consisting of a certain number of equally-spaced bunches. Between the trains, short gaps must be left for the injection kicker magnets to rise. In addition, a 3 μs abort gap is kept free in order to allow for a safe extraction of the LHC beam by the dump system. The first BCID after the abort gap is by definition numbered as 1. In general, not all bunch slots are filled, and not all filled bunches are paired at every interaction point. The bunches are usually classified into groups of BCIDs[2], including, among others:

- *Paired*: A bunch present in both beams at the same BCID.
- *Unpaired isolated*: A bunch in only one beam, with no bunch in the other beam within 3 BCIDs.
- *Unpaired non-isolated*: A bunch in only one beam, with a nearby bunch in the other beam within 3 BCIDs.
- *Empty*: A BCID with no proton bunch in either beam.

[2]These definitions are taken from [46] and are commonly used within the ATLAS experiment.

3.3 Luminosity, pileup and HL-LHC

One of the key parameters characterizing the performance of the LHC is its ability to produce collisions, quantified by the *instantaneous luminosity*. For two bunches containing n_1 and n_2 protons, colliding at the revolution frequency f_r, the instantaneous luminosity can be written as [47]:

$$\mathcal{L} = \frac{n_1 \, n_2 \, f_r}{4\pi \, \sigma_x \, \sigma_y}, \tag{3.1}$$

where σ_x and σ_y are the transverse beam sizes at the interaction point. The rate of events $\dot{N}$ for a given physics process with cross section σ is directly proportional to the luminosity: $\dot{N} = \mathcal{L} \cdot \sigma$. Precise knowledge of the luminosity is therefore essential for any cross-section measurement, making its determination one of the fundamental tasks at the LHC.

The time-integrated luminosity, $\mathcal{L}_{\text{int}} = \int \mathcal{L} \, dt$, determines the total number of events collected over a given data-taking period. The LHC was designed to achieve a peak instantaneous luminosity of $10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$. This target was first reached in 2016 and subsequently exceeded. During Run 3 (2022–2026), operating at a center-of-mass energy of 13.6 TeV, the peak instantaneous luminosity at ATLAS reached $2.15 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ in 2024, more than twice the original design value [48]. In Run 1 (2010-2012), it was $\sqrt{s} = 7$ and 8 TeV and in Run 2 (2015-2018) 13 TeV. Over the first four years of Run 3, the LHC delivered more than $350\,\text{fb}^{-1}$ of integrated luminosity to ATLAS and CMS in proton–proton collisions, as seen in Figure 3.5.

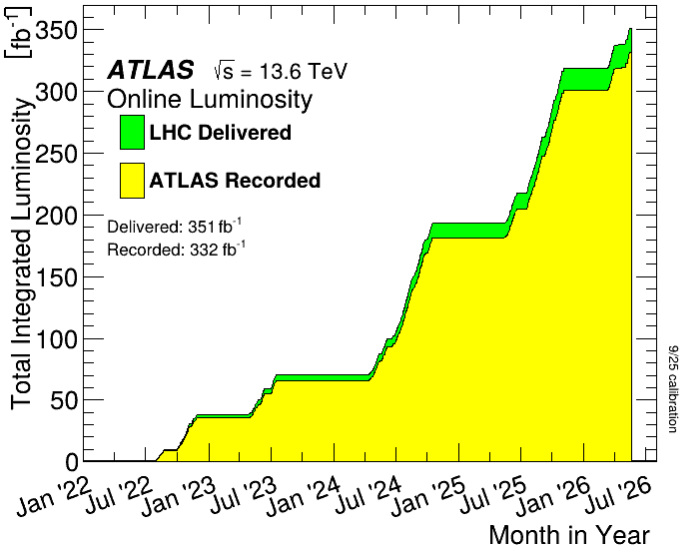

Figure 3.5: The total integrated luminosity during Run 3, for both LHC delivered and recorded by the ATLAS detector [49].

At the high luminosities achieved with the LHC, multiple proton–proton interactions occur simultaneously within each bunch crossing. This phenomenon, known as *pileup*, is characterized by the average number of interactions per bunch crossing, denoted $\langle\mu\rangle$, and poses a significant challenge for the reconstruction and identification of physics objects in the detector. While the LHC was originally designed for an average pileup of about 25, Run 3 operations routinely reached $\langle\mu\rangle \approx 60$, with ATLAS successfully operating at a leveling target of $\langle\mu\rangle = 64$. The evolution of $\langle\mu\rangle$ over the years of Run 3 is shown in Figure 3.6.

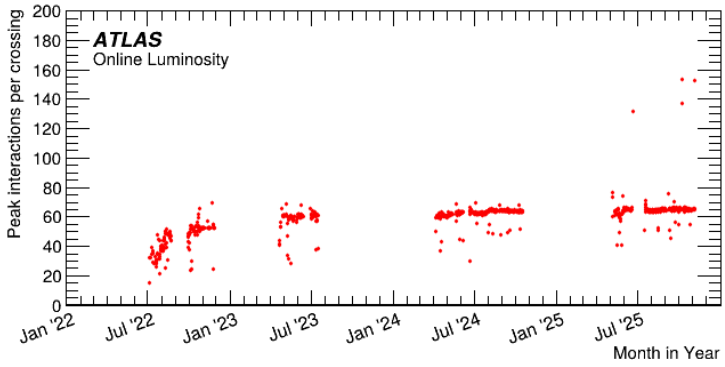

Figure 3.6: Peak interactions per bunch crossings over the years of Run 3 [49].

To maintain manageable pileup conditions while maximizing integrated luminosity, the LHC employs a technique known as *luminosity leveling*, in which the instantaneous luminosity is kept approximately constant at a target value during the physics fill by progressively adjusting the beam separation and β^*, rather than allowing it to decay freely from its peak.

Table 3.1: Summary of key LHC beam parameters: design values and typical Run 3 operational values [43].

Parameter	Design	Run 3 (typical)
Beam energy [TeV]	7.0	6.8
$\sqrt{s}$ [TeV]	14.0	13.6
Number of bunches	2808	∼2460
Protons per bunch	1.15×10^{11}	$\sim 1.6 \times 10^{11}$
Peak $\mathcal{L}$ [$\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$]	10^{34}	2.15×10^{34}
Average pileup $\langle\mu\rangle$	25	∼60

A more comprehensive discussion of the LHC bunch structure, as well as the concepts and determination of luminosity in ATLAS, is provided in Chapter 7.

Chapter 4

The ATLAS experiment

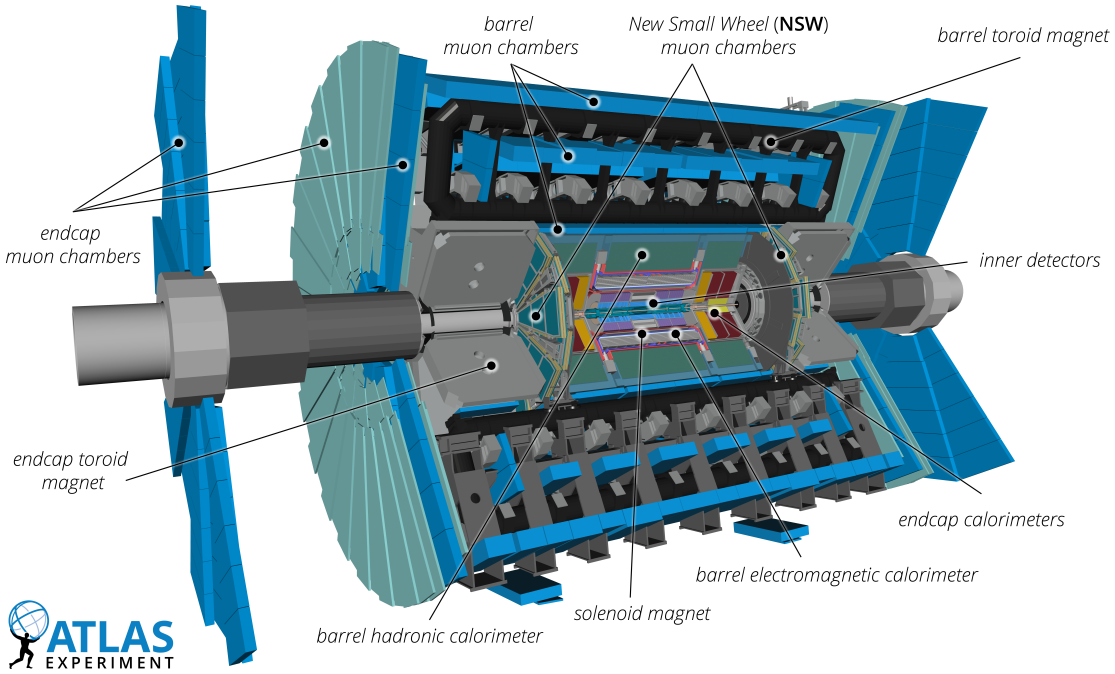

Figure 4.1: Overview of the ATLAS detector and its subdetectors [50].

ATLAS [48] is a general purpose particle detector weighing 7000 t designed with cylindrical shape with 25 m of diameter and 44 m length, as seen in Figure 4.1. It is composed of three main parts: the Inner Detector (ID), surrounded by a superconducting solenoid, the calorimeters, in turn surrounded by a toroid magnet and finally the muon system. Moreover, ATLAS has a two-level trigger system (Level-1 and High-Level), and a data acquisition system. The various subsystems will be described in more detail in the following sections.

4.1 ATLAS coordinate system

Before discussing the various subdetectors, it is important to introduce the coordinate system used in the ATLAS experiment, which is centered at the interaction point of the beams. The system is defined in three spatial coordinates: the z-axis (parallel to the beam direction), the x-axis (pointing toward the center of the LHC ring), and the

y-axis (pointing upwards, away from the center of the Earth), as seen in the sketch in Figure 4.2.

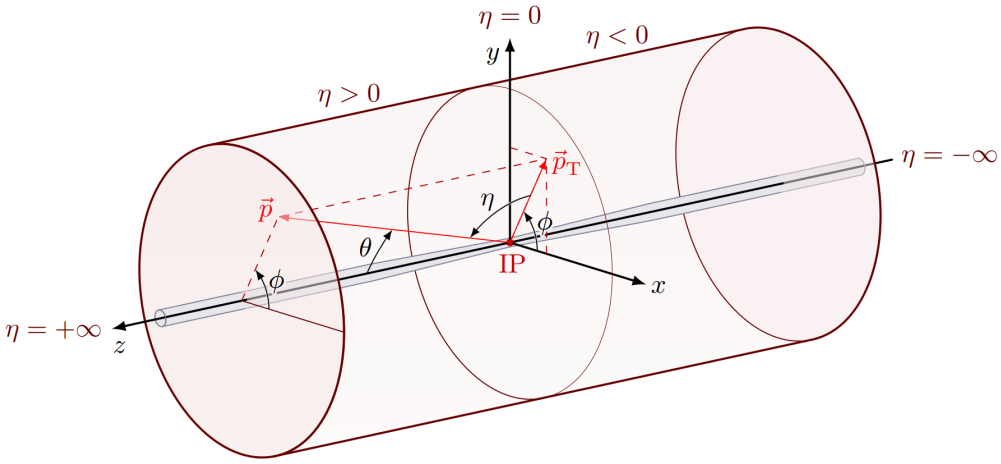

Figure 4.2: The coordinate system of ATLAS relative to the LHC, showing all the coordinates and kinematic variables for particle identification.

In practice, ATLAS analyses more commonly use the angular coordinates η and ϕ, rather than the polar angle θ and azimuthal angle ϕ directly. The azimuthal angle ϕ is defined as

$$\phi = \arctan\left(\frac{y}{x}\right) \tag{4.1}$$

and is measured in the transverse (xy) plane, perpendicular to the beam axis. Since the colliding partons carry negligible initial transverse momentum, momentum conservation can be enforced in the transverse plane, forming the basis for missing transverse momentum ($E_{\mathrm{T}}^{\mathrm{miss}}$) measurements [51]. Furthermore, ϕ is exactly invariant under longitudinal Lorentz boosts, as it is measured in a plane orthogonal to the boost direction.

The polar angle θ, defined as

$$\theta = \arctan\left(\frac{\sqrt{x^2 + y^2}}{z}\right) \tag{4.2}$$

is not used directly. Instead, the quantity of *pseudorapidity* η is preferred, defined as

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right). \tag{4.3}$$

The motivation for this choice is based on the kinematics of hadron collisions. At the LHC, the hard scattering occurs between individual partons (quarks and gluons), each carrying an unknown (usually not equal) fraction of the proton's momentum. Therefore, the center-of-mass frame of each parton–parton collision is boosted along the beam axis

by an unknown, event-by-event varying amount. Differences in the polar angle θ are not preserved under such longitudinal boosts. Instead, differences in *rapidity*,

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \tag{4.4}$$

are exactly invariant, since rapidity transforms additively under longitudinal boosts: $y \to y - y_{\text{boost}}$. *Pseudorapidity* has the advantage of depending only on the polar angle θ, unlike rapidity which requires knowledge of the particle's energy and mass. The connection between the two can be seen by rewriting pseudorapidity as $\eta = \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} \right)$, which is identical in structure to the rapidity with E replaced by $|\mathbf{p}|$. For highly relativistic particles (such as jets, photons, and electrons) where $m \ll |\mathbf{p}|$ and therefore $E \approx |\mathbf{p}|$, pseudorapidity converges to rapidity: $\eta \approx y$.

Moreover, particle production in hadronic collisions is approximately uniform in rapidity in the central region, so a detector segmented uniformly in η sees roughly equal occupancy across its cells. The entire ATLAS detector (calorimeter cells, trigger towers, and tracking layers) is segmented in (η, ϕ) for these reasons.

For particle identification and kinematic reconstruction, the following quantities are defined. The transverse momentum p_{T} is the component of a particle's momentum in the transverse plane:

$$p_{\text{T}} = \sqrt{p_x^2 + p_y^2} = |\mathbf{p}| \sin \theta. \tag{4.5}$$

The transverse energy is defined as

$$E_{\text{T}} = E \sin \theta. \tag{4.6}$$

Finally, the angular distance ΔR between two objects in (η, ϕ) space, commonly used in jet clustering algorithms and isolation criteria, is defined as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \tag{4.7}$$

This metric is approximately Lorentz invariant under longitudinal boosts for relativistic particles, ensuring that jet definitions and isolation requirements have a consistent physical meaning regardless of the unknown boost of each individual collision.

4.2 Inner Detector

The ID, shown in Figure 4.3, is located at the very center of ATLAS and covers a pseudorapidity range of $|\eta| < 2.5$. Its primary purpose is to measure the trajectories of charged particles, or *tracks*, enabling precise reconstruction of their transverse momentum (p_{T}), direction, and interaction vertices (both primary and secondary) [52]. The ID operates inside a 2 T solenoidal magnetic field, which bends the charged particle tracks,

allowing momentum measurements via curvature.

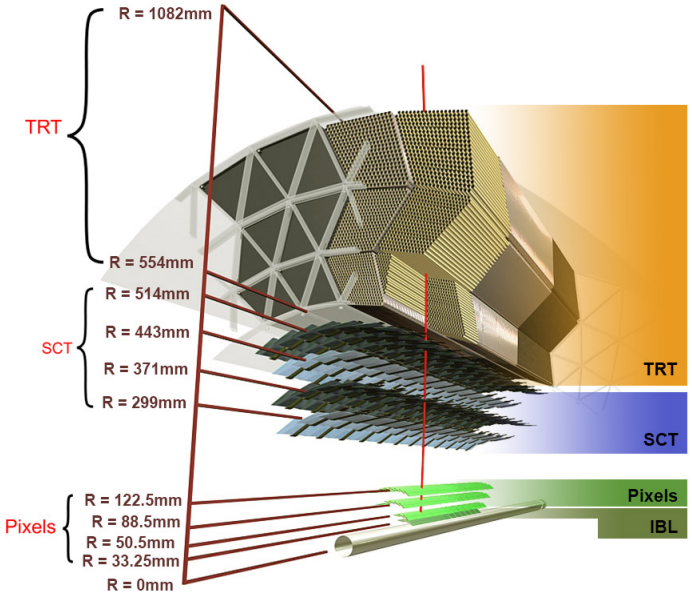

Figure 4.3: The Inner Detector and its subcomponents [53].

The ID consists of four subdetectors: the Insertable B-Layer (IBL) [54, 55], the Pixel Detector [56], the Semiconductor Tracker (SCT) [57], and the Transition Radiation Tracker (TRT) [58]. Each subdetector is designed to provide high spatial resolution and robust tracking performance close to the interaction point.

The IBL is the innermost subdetector. It was installed during the 2013–2014 LHC shutdown as an additional pixel layer between a new, smaller-radius beam pipe and the existing Pixel Detector, at a mean sensor radius of 33.5 mm, compared with 50.5 mm for the B-Layer, the innermost layer of the original Pixel Detector [55]. It consists of 14 carbon composite staves, tilted by 14° in azimuth, each carrying 20 modules equipped with planar or 3D silicon sensors. In total, about 12 million pixels of size 50 µm × 250 µm are read out by FE-I4B front-end chips, covering $|\eta| < 3.0$. Its proximity to the interaction point and its finer pixel pitch improve the impact parameter resolution and vertex reconstruction, while its radiation-hard design preserves the tracking and heavy-flavor identification performance against radiation damage in the B-Layer and the increasing pileup.

The Pixel Detector surrounds the IBL and, before its installation, was the innermost subdetector. It consists of approximately 80 million readout channels distributed over 1744 modules. Each module contains pixels with dimensions of 50 µm × 400 µm, providing excellent spatial resolution near the interaction point. The Pixel Detector is crucial for precise vertex reconstruction and impact parameter measurements, essential

for identifying secondary vertices from heavy-flavor decays.

SCT surrounds the Pixel Detector and is composed of a barrel section and two end-caps. It comprises about 6.3 million readout channels and uses silicon microstrip sensors. Each module consists of two layers of strips with a small angle to allow for precise two-dimensional position measurements. Together with the Pixel Detector, the SCT provides high-precision tracking in the central region and contributes significantly to momentum resolution.

Both the Pixel Detector and SCT are silicon-based detectors. When a charged particle traverses a silicon sensor, it ionizes atoms and creates electron–hole pairs. An applied electric field causes the charge carriers to drift toward the electrodes, inducing a measurable signal. If this signal exceeds a predefined threshold, a hit is registered on that specific sensor element, along with timing information. By combining multiple hits in different layers, a complete track can be reconstructed, and the particle's momentum and direction can be determined.

TRT forms the outermost part of the Inner Detector and covers a pseudorapidity range of $|\eta| < 2.0$. It consists of about 300,000 straw drift tubes, each filled with a gas mixture rich in xenon, which provides high ionization yield. In addition to standard tracking information, the TRT enhances electron identification capability through the detection of transition radiation photons, which are emitted when relativistic charged particles traverse the interfaces of materials with differing dielectric constants. These photons are absorbed in the gas and produce additional ionization, resulting in a higher signal amplitude in the straws.

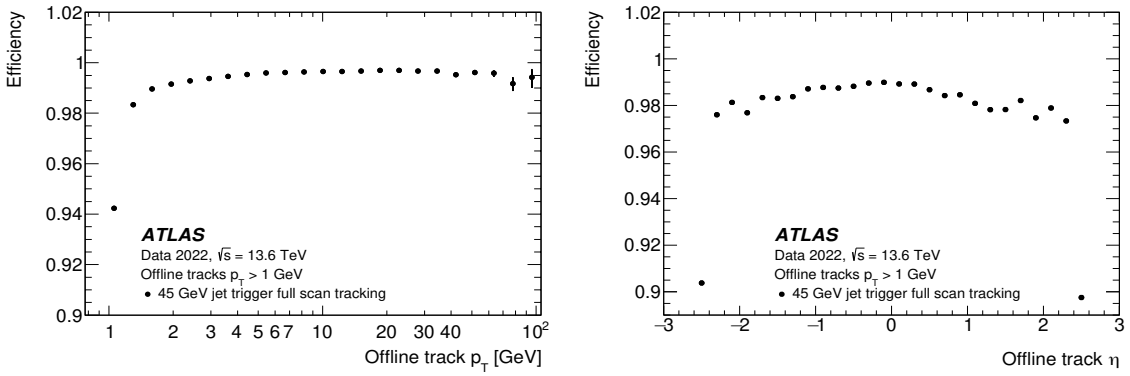

(a) (b)

Figure 4.4: Track finding efficiency with respect to offline tracks versus the offline track p_{T} (a) and η, evaluated for the 45 GeV jet trigger [59].

As evidenced in latest performance studies, the tracking efficiency of the ID is approximately 98.5% relative to the offline reconstruction for events recorded with a 45 GeV jet trigger in 2022 [59]. The efficiency as a function of offline track p_{T} is shown in Figure 4.4a. It is only around 94% at 1 GeV, and reaches a plateau at approximately 5 GeV.

On the other hand, the tracking efficiency against the offline track pseudorapidity performs worse at high $|\jmath|$, while oscillating around 98.5% at central values.

4.3 Calorimeters

ATLAS is equipped with two main calorimeter systems designed to measure the energy and position of particles. They precisely measure electromagnetic particles such as electrons and photons, as well as strongly-interacting particles, including pions and jets. The calorimeters also play a crucial role in the determination of the missing transverse momentum and in particle identification.

The two calorimeter systems, namely the Liquid Argon (LAr) calorimeter and the Tile calorimeter, displayed in Figure 4.5, are both sampling calorimeters composed of several subdetectors arranged symmetrically in ϕ and covering a pseudorapidity range up to $|\eta| < 4.9$ [60].

A more detailed description of the calorimeter systems is provided in Section 5.2.

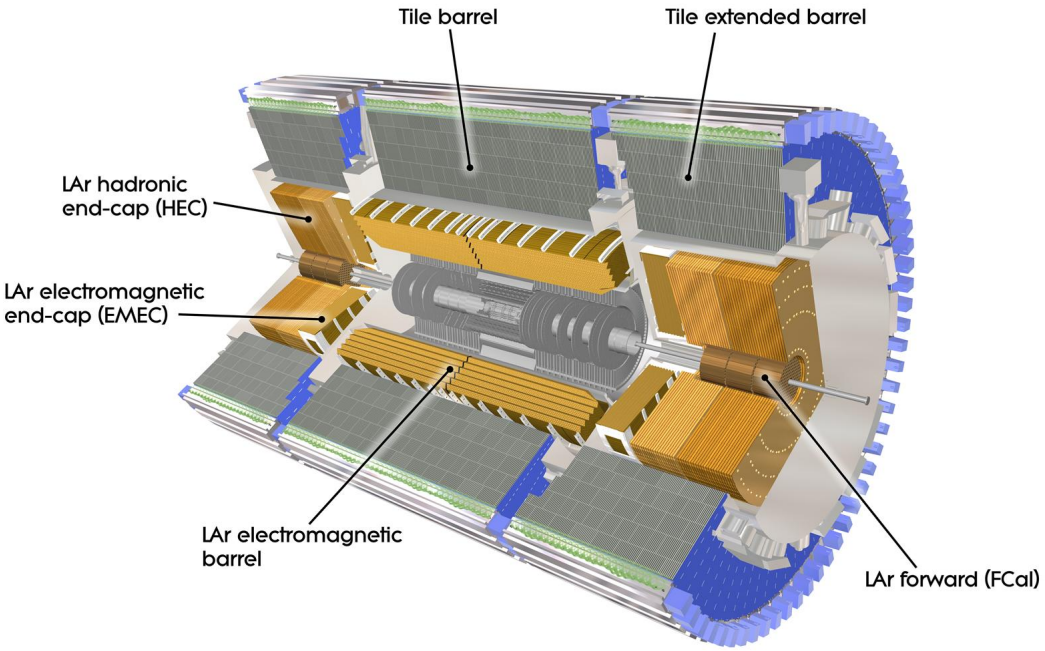

Figure 4.5: The ATLAS calorimeter Tile and LAr, and their respective subsystems [61].

4.4 Muon system

The ATLAS muon spectrometer [62] utilizes the magnetic field generated by a large barrel toroid covering the region $|\eta| < 1.4$, and two smaller end-cap toroids extending the coverage to $1.6 < |\eta| < 2.7$, to detect and measure the trajectories of muons produced in proton–proton collisions. The muon system provides an independent momentum measurement, crucial for high-precision muon reconstruction and identification.

The muon spectrometer consists of several subdetectors, as shown in Figure 4.6: Thin Gap Chambers (TGCs), Resistive Plate Chambers (RPCs), Monitored Drift Tubes (MDTs) [63], and Cathode Strip Chambers (CSCs) [64]. Additionally, the New Small Wheel (NSW) upgrade [65], installed for the Phase-I upgrade of the LHC, during LS2 (2019-2022) for Run 3 (2022-2026), enhances performance in the forward region and improves muon trigger and tracking capabilities.

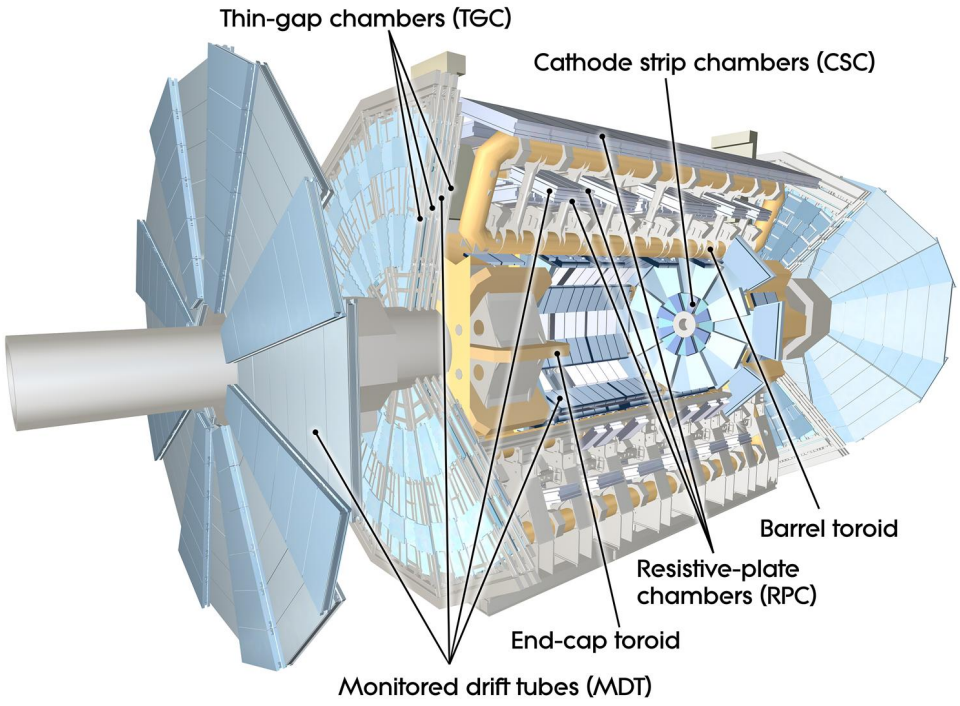

Figure 4.6: The muon system and its subcomponents [66].

TGCs, located in the end-cap ($1.05 < |\eta| < 2.4$), consist of 2.8 mm gas gaps filled with CO_2 and n-pentane (C_5H_{12}). Muons ionize the gas, and charges collected by central wires enable fast timing and coarse position measurements. RPCs, in the barrel ($|\eta| < 1.05$), have three double layers of 2 mm gas gaps between resistive plates with readout strips. The inner layers trigger on low-p_T muons (6–9 GeV), and the outer layer targets muons up to 35 GeV.

Precision tracking is provided by MDTs and CSCs. MDTs cover $|\eta| < 2.7$ with 30 mm aluminum drift tubes filled with argon–CO_2 gas. Drift times yield precise position measurements. CSCs, in the forward region ($2.0 < |\eta| < 2.7$), consist of sixteen multi-wire chambers per end-cap disk, offering high-resolution measurements in η and ϕ.

The NSW replaces innermost end-cap stations ($1.3 < |\eta| < 2.7$), combining small-strip Thin Gap Chambers (sTGCs) and Micromegas detectors. sTGCs provide precise timing and fast triggers, while Micromegas deliver spatial resolution better than 100 µm. Together, they improve fake muon rejection, momentum resolution at high pseudorapidity, and trigger selectivity at high luminosity.

Muon momentum measurements in the spectrometer are independent of those by the Inner Detector, as the toroidal and solenoidal magnets operate separately, enabling cross-checks that enhance overall resolution. Performance of hits and reconstruction efficiency of the muon system are shown in Figure 4.7.

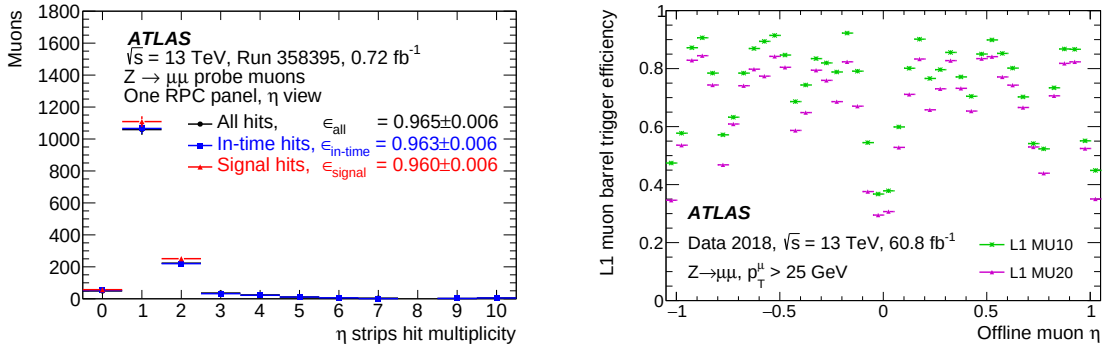

(a) (b)

Figure 4.7: Hit multiplicity and detector efficiency for the η RPC panels (a) and legacy L1 muon barrel trigger efficiency as a function of muon pseudorapidity (b) in Run 2 [67].

4.5 Magnet system

The ATLAS magnet system consists of two main parts: a central solenoid magnet and a set of toroidal magnets. The solenoid provides a magnetic flux density of 2 T to the Inner Detector, enabling precise momentum measurements of charged particles at small radii [68]. The toroidal system, comprising a barrel toroid and two end-cap toroids, supplies magnetic fields of approximately 0.5 T and 1 T, respectively, for the muon spectrometer.

The toroidal system is composed of a barrel toroid and two end-cap toroids, as illustrated in Figure 4.9. Each consists of eight superconducting coils arranged symmetrically around the beam axis. The end-cap coils are rotated by 22.5° relative to the barrel coils, allowing for radial overlap and improved bending power in the transition region.

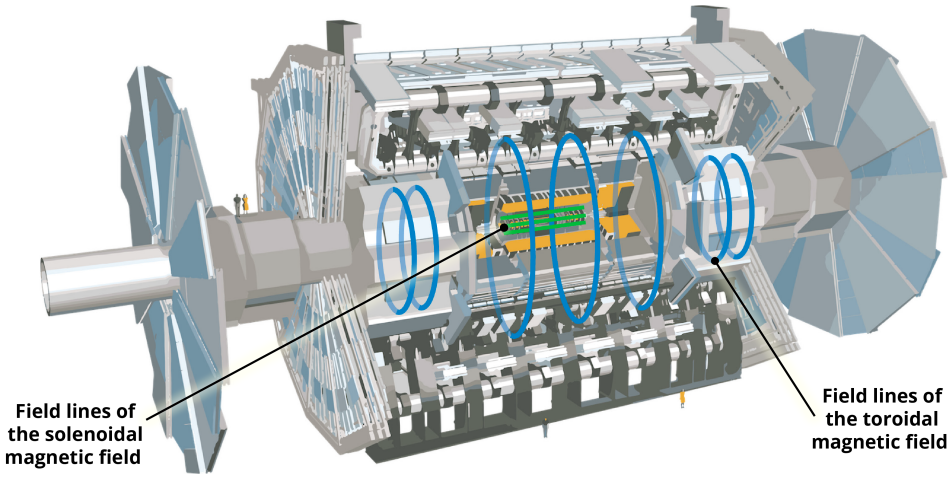

Figure 4.8: The field lines of the magnets within the ATLAS detector [69].

The barrel toroid coils are housed in individual cryostats, while the eight end-cap coils on each side are enclosed within a single large cryostat. The toroidal magnets provide coverage up to $|\eta| < 2.7$.

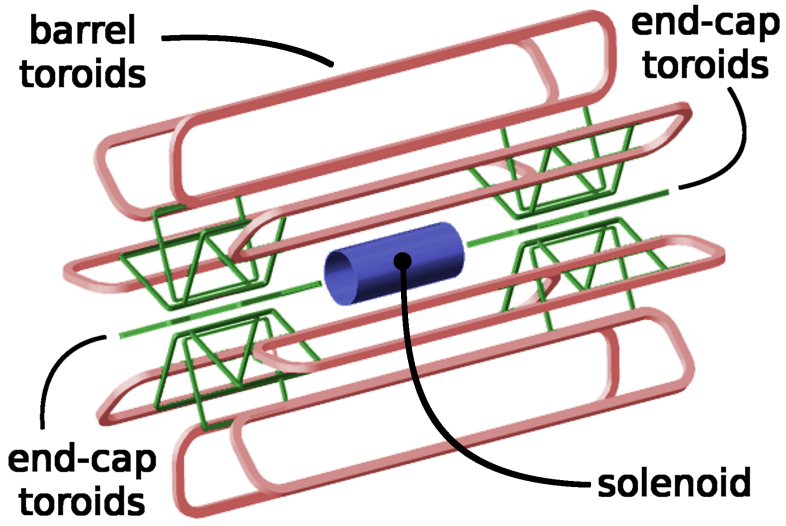

Figure 4.9: A sketch of the ATLAS magnet system with its barrel and end-cap toroids (eight coils each), and the solenoid, located between the Inner Detector and the calorimeters.

The magnetic bending power is characterized by the integral $\int B\,dl$, where B is the magnetic field component perpendicular to the muon trajectory, integrated along the

path between the innermost and outermost muon chambers. The barrel toroid delivers a bending power of 1.5–5.5 T m for $0 < |\eta| < 1.4$, while the end-cap toroids provide approximately 1–7.5 T m for $1.6 < |\eta| < 2.7$. In the transition region ($1.4 < |\eta| < 1.6$), the bending power is reduced due to partial coverage and magnetic field geometry.

4.6 Trigger system

During Run 3 of the LHC, the proton bunch crossing frequency is 40 MHz, corresponding to a collision every 25 ns. This results in an enormous number of events created at each collision, most of which are soft interactions producing low-energy particles and are generally not relevant for precision physics studies nor searches for new physics. To manage this high data rate while retaining events of interest, the ATLAS experiment employs a two-level trigger system designed to rapidly filter and select the most promising events for storage and analysis, as represented schematically in Figure 4.10.

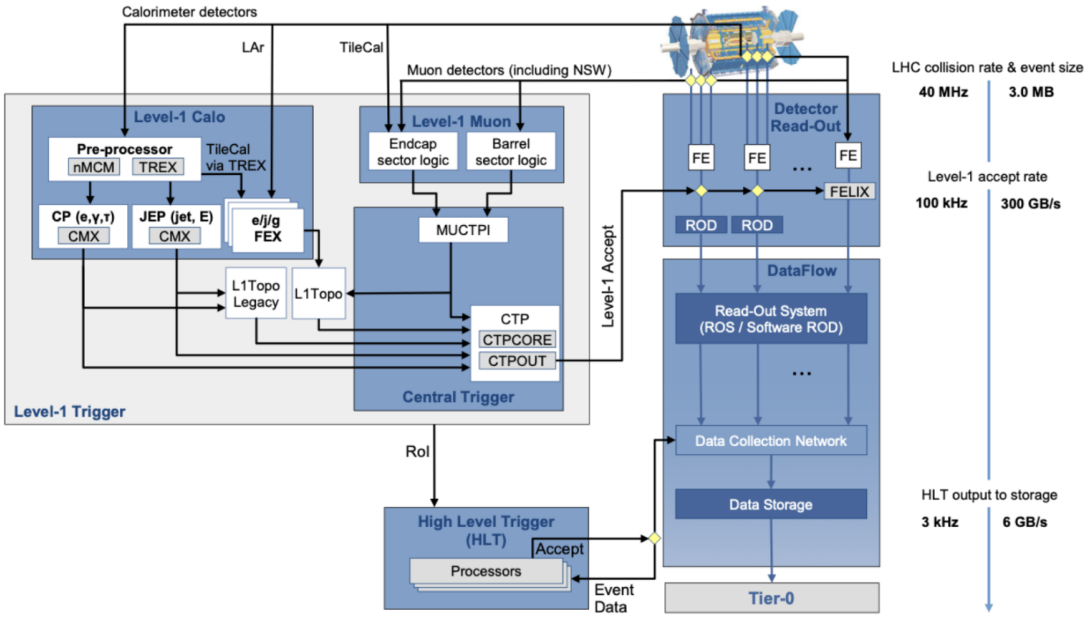

Figure 4.10: The data flow and decision path of the ATLAS trigger system in Run 3 [59].

The Level-1 (L1) trigger [70] is a hardware-based system that provides the first filtering stage, reducing the event rate from 40 MHz to approximately 100 kHz within 2.5 µs, with the following characteristics:

- **Calorimeter triggers:** Identification of electromagnetic and hadronic activity using trigger towers until Run 2 with coarse granularity ($\Delta\eta \times \Delta\phi = 0.1 \times 0.1$). An

example of this can is shown in Figure 4.11(a). Since Run 3, the Digital Trigger has replaced the trigger tower system (more details in Section 6.2).

- **Muon triggers:** Fast signals from RPCs in the barrel and TGCs in the end-caps to identify high-p_{T} muons, as seen in Figure 4.11(b).
- **Missing transverse momentum ($E_{\mathrm{T}}^{\mathrm{miss}}$) and scalar jet sum ($H_{\mathrm{T}}$) triggers:** Real-time calculations from calorimeter data to identify events with significant energy imbalance or high overall activity.
- **Regions of Interest (RoIs):** Identification of small detector regions with interesting activity, used to seed more detailed analysis in the High-Level Trigger.

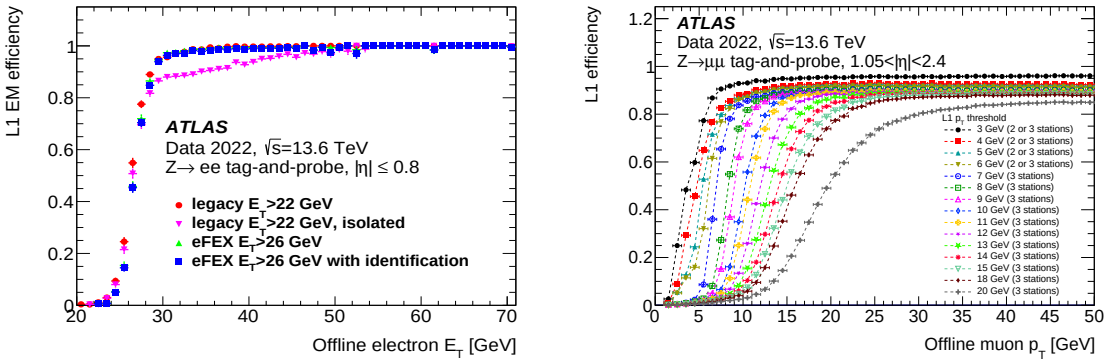

(a) (b)

Figure 4.11: Efficiency of the L1 single-object electromagnetic (EM) trigger (a) and of L1 muon triggers in the end-cap region for various Et and pT thresholds (b) [59].

For Run 3, the L1 system was significantly upgraded with a new Digital Trigger system, laying the groundwork for the High-Luminosity LHC (HL-LHC) Phase-II trigger and improving selectivity without sacrificing efficiency for key physics channels (see Chapter 5 for details).

Events accepted by L1 are passed through the Software Readout Driver (SWROD) [71] to the High-Level Trigger (HLT) [72], a software-based system that performs near-offline-quality reconstruction using the full detector granularity and resolution.

- **Input rate:** Up to 100 kHz from L1.
- **Output rate:** Reduced to a few kHz for offline storage.
- **Object reconstruction:** Full-event reconstruction of electrons, muons, photons, jets, τ-leptons, and $E_{\mathrm{T}}^{\mathrm{miss}}$ using algorithms close to those used offline. The difference in reconstruction efficiency of jets between L1 and HLT is shown in Figure 4.12.

- **Prescales:** Some lower-priority triggers (e.g., low-p_{T} muons or jets) may be prescaled, keeping only a fraction of such events to control bandwidth.
- **RoI-based processing:** Focus on detector regions flagged by L1 to reduce processing time.
- **Inclusive trigger menu:** Includes single-object, multi-object, and topological combinations to maximize efficiency for a wide range of analyses.

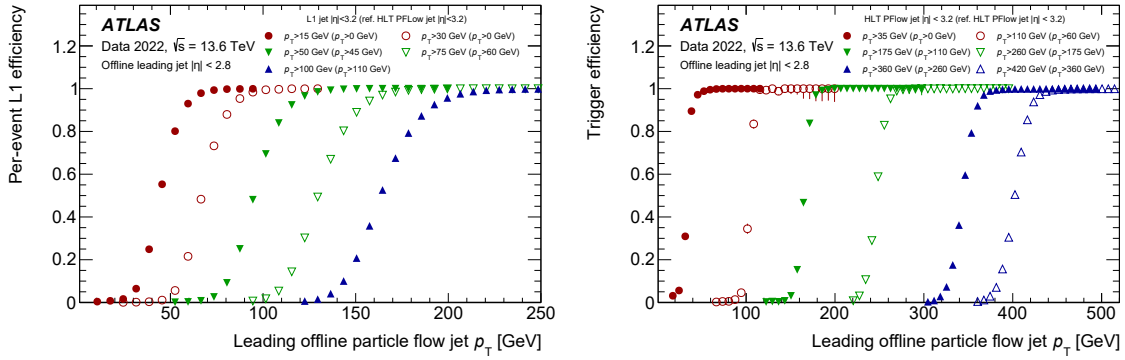

(a) (b)

Figure 4.12: Efficiency of the legacy L1 single jet triggers (a) and efficiency of single jet trigger selection by the HLT (b) in the central region [59].

Once accepted by the HLT, events are divided into streams, with different rates and bandwidths, as illustrated in Figure 4.13:

- **Physics streams:** Events selected by triggers on generic physics objects (electrons, muons, jets, $E_{\mathrm{T}}^{\mathrm{miss}}$, ...) used across the physics program.
- **Calibration and empty bunch streams:** Used for detector monitoring, calibration, and noise studies.
- **Express stream:** A subset of events processed with fast turnaround (within hours) to monitor data quality and detector conditions in near real-time.

All streams are transmitted to Tier-0 at CERN, the central hub of the Worldwide LHC Computing Grid (WLCG) [73], where the initial offline reconstruction is performed and one archival copy of the raw and reconstructed data is stored. The output is then distributed to approximately 13 Tier-1 national and regional computing centers worldwide, which handle long-term storage and scheduled reprocessing campaigns, and in turn supply data to affiliated Tier-2 sites. The roughly 160 Tier-2 sites, hosted at universities and smaller research institutes, are used primarily for Monte Carlo simulation and end-user analysis. Prompt calibration loops process special calibration streams and update

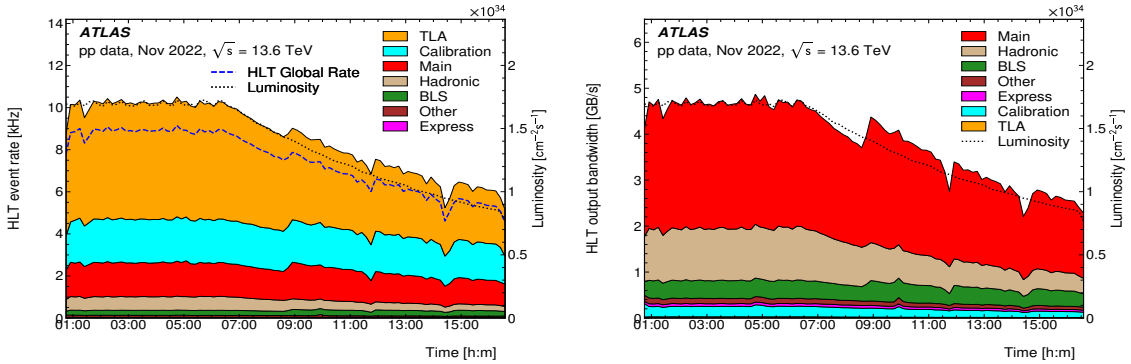

(a) (b)

Figure 4.13: An example of (a) the rate and (b) bandwidth output by the HLT streams in a typical 2022 proton–proton run. The total HLT event rate is lower than the sum of the stream rates, because the same events may be written to multiple streams. The peak bandwidth of the trigger-level analysis (TLA) stream is approximately 25 MB/s, and consequently barely visible in the right hand figure. The event size per stream can vary up to a factor of 300 [59].

detector conditions within 24 hours, improving the accuracy of offline reconstruction. Re-reconstruction campaigns may be launched later to apply updated algorithms and calibrations not available during initial processing.

Chapter 5

The Liquid Argon calorimeter

As introduced earlier, the main work of this thesis is based on a specific subsystem of the ATLAS detector, namely the Liquid Argon calorimeter. For this reason, the following chapter is dedicated only to a more detailed description of this system mentioned briefly in Section 5.2, starting from a general introduction to calorimeters and the concepts that the luminosity methods described in Chapter 8 rely on, through the different components and specification of the LAr calorimeter, to the Phase-I upgrade, central to the work undertaken in this thesis.

5.1 Calorimeters

Calorimeters are types of detectors that determine the energy of an incoming particle by causing it to interact with the detector material and be mostly absorbed by it. When a particle enters a calorimeter, it triggers a chain reaction producing a cascade of secondary particles, commonly referred to as a *shower*. These secondary particles generate a signal whose intensity, for most calorimeters, is directly proportional to the energy of the original particle, allowing the system to measure its properties.

Electromagnetic calorimeters specialize in measuring the energy carried by photons and electrons. In contrast, *hadronic* calorimeters detect the energy of hadrons, such as protons and pions, as well as jets, which are collimated sprays of hadrons produced by the fragmentation of quarks and gluons. In typical particle detectors, the electromagnetic calorimeter is positioned upstream of the hadronic calorimeter. This configuration leverages the fact that photons and electrons are fully absorbed by the electromagnetic section, while hadrons mostly pass through it with minimal energy loss before reaching the hadronic section.

Both electromagnetic and hadronic calorimeters can be classified into two types: *homogeneous* and *sampling*. Homogeneous calorimeters are made entirely of a single material that simultaneously absorbs the particle and produces an electrical or optical signal. In contrast, sampling calorimeters consist of alternating layers of dense absorber material and an active medium. The absorber layers significantly reduce the particle's energy, while the active layers produce a measurable signal whenever charged particles pass through them. This design facilitates the construction of compact calorimeters capable of fully absorbing the energy of incident particles.

5.1.1 Energy measurement

The energy E of an incoming particle in the calorimeter is proportional to the number of particles N generated by the shower caused by the interaction with the detector. This energy has a statistical fluctuation ΔE, that given the proportionality to the number of particles, we can express it as [74]

$$\frac{\Delta E}{E} \propto \frac{\Delta N}{N} \propto \frac{\sqrt{N}}{N} \propto \frac{1}{\sqrt{N}} \propto \frac{1}{\sqrt{E}} \tag{5.1}$$

This shows how, statistically, the energy resolution of a calorimeter improves as the particle energy increases. However, in practical calorimeters, factors such as instrumentation effects and noise must also be considered:

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \tag{5.2}$$

Where the symbol $\oplus$ denotes a quadratic sum, a is the so called *sampling* term, b the *noise* term, and c is a constant.

The sampling term a typically ranges between 8–11% for the LAr calorimeter [75] and is directly influenced by the sampling fraction of the calorimeter. Homogeneous calorimeters display a very small values of a, since the shower is fully absorbed in the active material. Sampling calorimeters, instead, exhibit a larger term due to fluctuations in the energy deposited in the active layers between events. In the ATLAS LAr calorimeter, a presampler was added to correct for energy losses in the upstream dead material and thus avoid a worsening of this term at low energies.

The parameter b represents the combined noise term. Its value depends on the balance between electronics and pileup noise, as these vary inversely with shaping time. Detectors that collect charged particles tend to have higher noise terms due to the preamplifier required in the electronics readout chain. This noise term is particularly relevant for particle energies below 20 GeV.

The constant term c is most significant at high energies and is highly sensitive to the construction quality, calibration, and operation of the detector. It combines effects from material non-uniformities, mechanical imperfections, temperature gradients, radiation damage, and similar factors.

The energy resolution of the LAr calorimeter and the contribution of the different terms can be seen in Figure 5.1.

For a sampling calorimeter such as LAr, it is also useful to define the *sampling fraction* as

$$S_f = \frac{E_{\text{active}}}{E_{\text{active}} + E_{\text{passive}}} \tag{5.3}$$

as only this fraction of the energy is deposited in the active layers of the calorimeter.

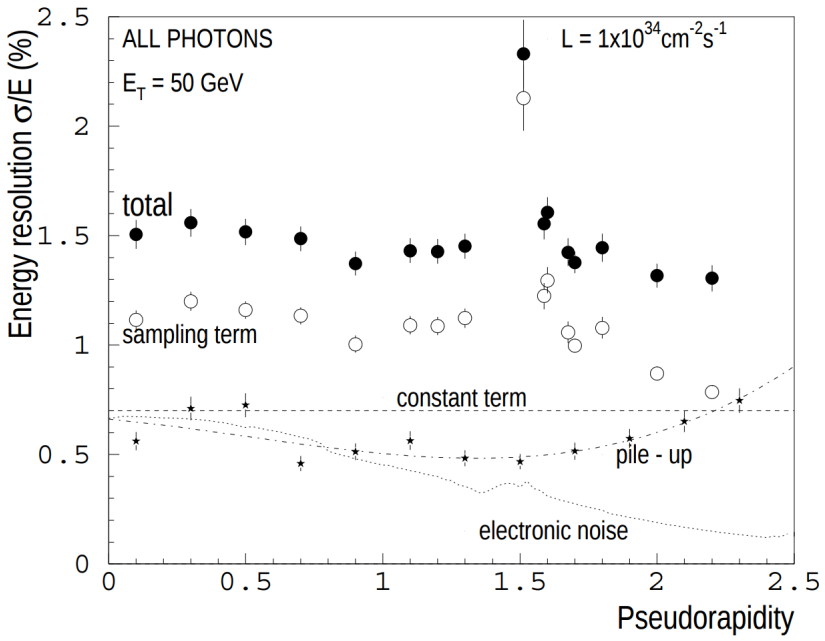

Figure 5.1: Expected energy resolution σ/E for photons with $E_T = 50\,\text{GeV}$ as a function of pseudorapidity, showing the sampling, constant, noise, and pile-up contributions at design luminosity ($\mathcal{L} = 10^{34}\ \text{cm}^{-2}\text{s}^{-1}$), from Monte Carlo simulation [75].

5.1.2 Electromagnetic showers

As an electron or positron traverses matter, it may radiate bremsstrahlung radiation. These emitted photons can in turn initiate electron–positron pair creation. Through this mechanism, a cascade of particles emerges with progressively lower energies, continuing until the individual particles reach an energy threshold too low for pair production. Any remaining energy is absorbed through atomic excitation and ionization processes.

In this context, the radiation length X_0 characterizes each material's response to high-energy electromagnetic particles and provides a practical metric for quantifying energy attenuation.

Expressed in $\text{g}\,\text{cm}^{-2}$, X_0 for monatomic materials may be estimated using:

$$X_0 = \frac{716.4 \cdot A}{Z(Z+1)\ln(Z/287)}\ \text{g}\,\text{cm}^{-2}, \tag{5.4}$$

where Z and A denote the atomic and mass numbers, respectively.

Most incident particles in an electromagnetic shower deposit only a small fraction of their energy within the active detector volume while traversing multiple calorimeter layers. The majority of charged particles inside the shower have limited kinetic energy and travel distances that are typically smaller than the thickness of the calorimeter layers [76]. Containing such showers requires a depth of many radiation lengths, which

grows only logarithmically with the energy of the incident particle. Electromagnetic showers also have lateral expansion stemming from multiple scattering events involving electrons, positrons, and photons. The Molière radius characterizes this transverse extent effectively and is expressed as

$$R_M = X_0 \frac{21\,\text{MeV}}{E_c}, \tag{5.5}$$

where X_0 is the radiation length defined above and E_c denotes the critical energy, approximately given by

$$E_c = \frac{800\,\text{MeV}}{Z + 1.2}, \tag{5.6}$$

For electrons, one radiation length corresponds to the mean path length required for a high-energy electron ($E > 1\,\text{GeV}$) to reduce its energy to $1/e$ of its initial value via bremsstrahlung.

Photons primarily lose energy through pair production, with energy reduced to $1/e$ following a mean free path of $x = 9/7 X_0$.

For accurate electromagnetic energy reconstruction, adequate calorimeter thickness, measured in radiation lengths, must be maintained to capture the full shower. Therefore, materials with high atomic number Z are preferred. The ATLAS electromagnetic calorimeter is more than $22\,X_0$ deep in the barrel and more than $24\,X_0$ deep in the end caps [48].

The following outlines the main mechanisms by which electrons, positrons, and photons interact with matter.

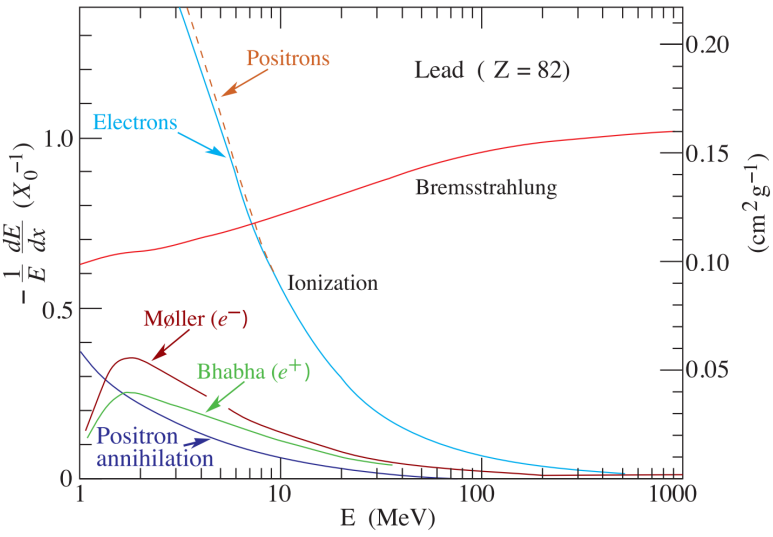

Figure 5.2: Fractional energy loss per radiation length versus electron energy [77].

Electrons and positrons. At low energies electrons and positrons lose energy predominantly through ionization of the medium; above the critical energy E_c (equation 5.6), bremsstrahlung becomes dominant. Figure 5.2 shows the fractional energy loss per radiation length for electrons.

Bremsstrahlung radiation: As an electron or positron with initial energy E_{init} slows down within matter, it loses energy primarily through acceleration-induced radiation in response to the Coulomb interaction with nuclear charges. The emitted photon spectrum obeys a $1/E_{\text{init}}$ distribution, with individual photons typically carrying only a fraction of the parent particle's initial energy.

Photons. Photons interact with matter through pair production, Compton scattering, and photoelectric absorption, with Rayleigh scattering contributing at low energies. The corresponding cross section of each process is shown in Figure 5.3.

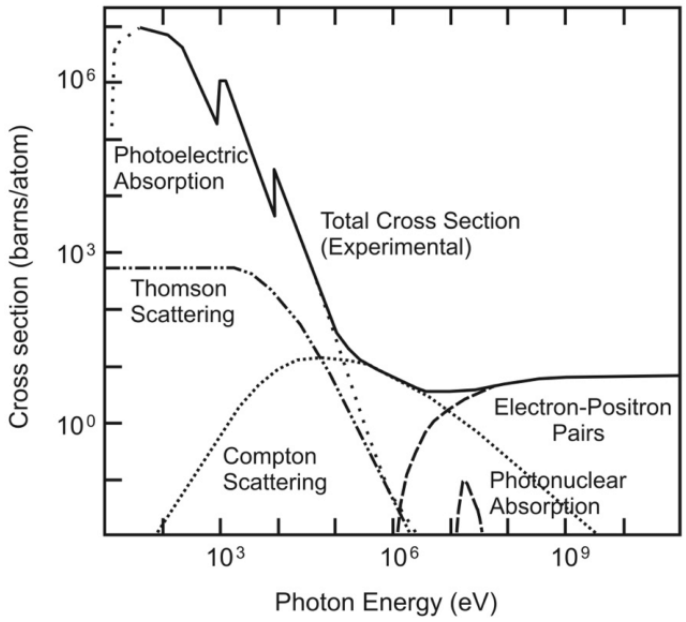

Figure 5.3: Photon cross sections per process versus photon energy [77].

Pair production: When a photon's energy exceeds twice the electron rest mass, it may convert into an electron–positron pair within a nucleus's electromagnetic field. The pair production cross section increases with photon energy, eventually reaching a plateau above 10 GeV. This mechanism becomes the dominant energy-loss channel for particles exceeding a few MeV.

Photoelectric absorption: In this process, an atom absorbs a photon and transitions to an excited state. De-excitation occurs through either characteristic X-ray emission or Auger electron ejection. The interaction cross section scales with the electron density, increasing substantially in high-Z materials where it scales as Z^n with $n \approx 4$–5. Beyond a few MeV, the cross section exhibits energy dependence following E^{-3}.

Rayleigh and Compton scattering: Rayleigh scattering is the coherent elastic scattering of photons from whole atoms; no energy is transferred to the atom, but the photon is deflected. It is significant only below 1 MeV in high-Z materials and plays no role in electromagnetic shower development. Compton scattering results in energy exchange with atomic electrons, frequently ionizing the atom. Between a few hundred keV and several MeV, Compton scattering represents the predominant energy-loss mechanism in typical absorber materials. The interaction cross section approximately scales linearly with Z and inversely with incident energy ($\propto 1/E$).

5.2 The ATLAS calorimeters

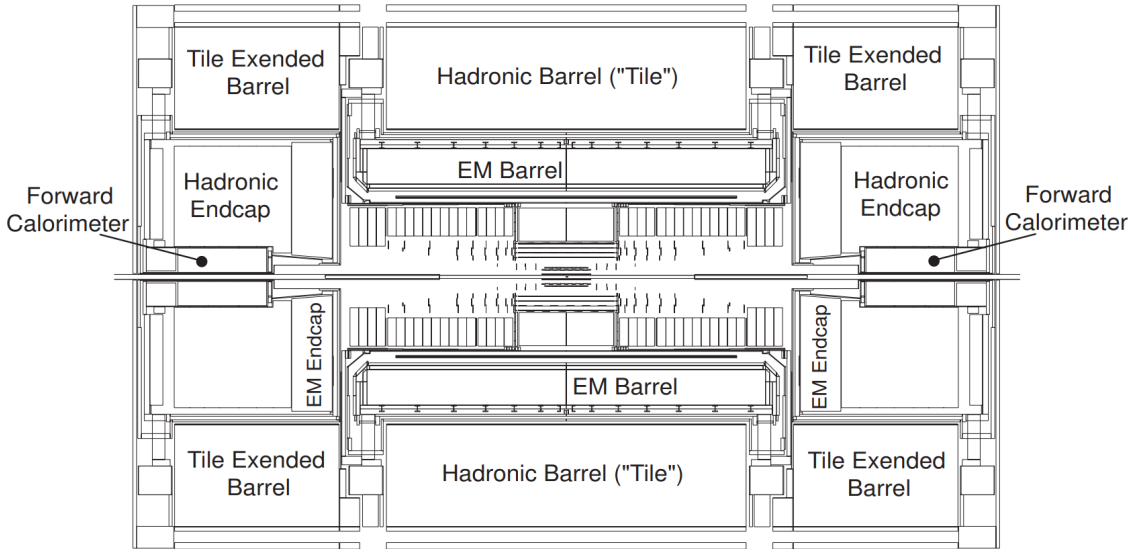

Figure 5.4: Schematic overview of the ATLAS calorimeters.

All calorimeters used in ATLAS are sampling calorimeters. ATLAS is composed of two main calorimeter systems, the Liquid Argon calorimeter [75] and the Tile calorimeter [78]. The LAr system includes the electromagnetic barrel (EMB) and end-cap (EMEC), the hadronic end-cap (HEC), and the forward calorimeter (FCal). The Tile calorimeter comprises a central Long Barrel and two Extended Barrels. The subdetectors are schematically shown in Figure 5.4.

- **Electromagnetic calorimeter:** The EMB covers $|\eta| < 1.475$, while the EMEC extends over $1.375 < |\eta| < 3.2$ on both sides of the detector. Both are built from lead absorbers interleaved with liquid argon, employing an accordion geometry and longitudinally segmented in several layers in the precision region $|\eta| < 2.5$. A detailed description is given in Sections 5.3-5.4.

- **Tile calorimeter:** The Tile calorimeter (TileCal) is the central hadronic calorimeter of ATLAS, surrounding the LAr calorimeter. It is a non-compensating sampling calorimeter that uses steel plates as absorber and plastic scintillating tiles as active material. The scintillation light produced in the tiles is collected by wavelength-shifting fibers and transmitted to approximately 10,000 photomultiplier tubes (PMTs) [79]. TileCal provides measurements of hadrons, jets, hadronic τ-lepton decays, and missing transverse momentum, and also contributes to muon identification and the first-level trigger.

  TileCal is divided into a central Long Barrel, covering $|\eta| < 1.0$, and two Extended Barrels reaching up to $|\eta| < 1.7$. Each section consists of 64 azimuthal modules. The gap between the Long Barrel and the Extended Barrels, needed to route cable services from the Inner Detector and the electromagnetic calorimeter, is instrumented with scintillators to estimate the energy lost in this region.

- **Hadronic end-cap calorimeter:** The HEC uses liquid argon as active material with flat copper plate absorbers. It is housed in the same cryostat as the EMEC, covers $1.5 < |\eta| < 3.2$, and is longitudinally segmented into four sampling layers (two per wheel) with transverse granularity $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for $|\eta| < 2.5$ and 0.2×0.2 for $2.5 < |\eta| < 3.2$.

- **Forward calorimeter:** The FCal provides both electromagnetic and hadronic calorimetry in the region $3.1 < |\eta| < 4.9$ and employs liquid argon with very thin active gaps to cope with the high particle rates. It comprises three longitudinal modules: FCal1 (electromagnetic, copper absorber) closest to the interaction point, followed by the hadronic modules FCal2 and FCal3 (tungsten absorbers).

As this is a thesis focusing on the Liquid Argon calorimeter, a more detailed description of its subdetectors is given in the following sections.

5.3 Electromagnetic barrel calorimeter

The EMB is housed in the barrel cryostat together with the superconducting solenoid, as illustrated in Figure 5.5, integrated to minimize the material in front of the calorimeter (approximately $2.3\,X_0$ at $|\eta| = 0$). It is a lead–liquid-argon sampling calorimeter employing an accordion geometry, which provides full azimuthal coverage

without projective cracks and minimizes the signal-path length for fast readout at the 25 ns bunch-crossing rate.

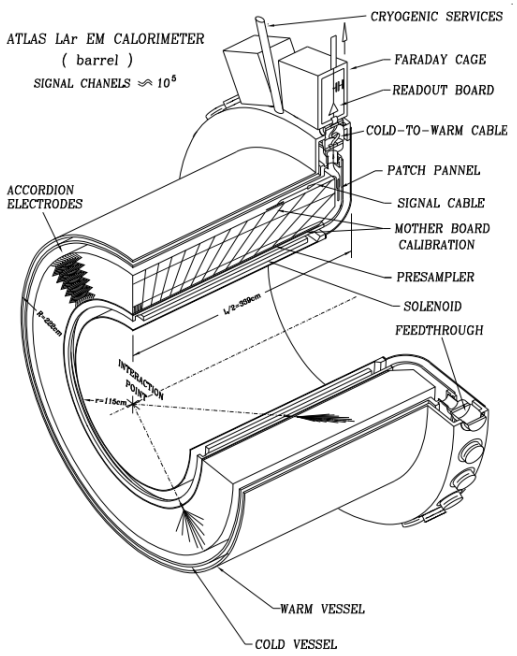

Figure 5.5: A schematic view of the barrel cryostat [75].

It consists of two half-barrels, each with 1024 accordion-shaped lead absorbers interleaved with readout electrodes centered by honeycomb spacers, maintaining a constant liquid argon gap of 2.1 mm at a nominal voltage of 2000 V. An example of such an absorber is seen in Figure 5.6. The total depth exceeds 22 X_0, around $\eta = 0$ and increases in $|\eta|$ to about 33 X_0 at $|\eta| = 1.3$. Each module is longitudinally segmented into four samplings whose granularity and depth are functional to the reconstruction of electromagnetic showers.

- **Presampler.** A thin (11 mm) active LAr layer with no absorber, covering $|\eta| < 1.52$, with granularity $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$. Its purpose is to correct, on a cell-by-cell basis, for the energy lost by electrons and photons in the upstream material (Inner Detector, services, cryostat walls and solenoid).

- **Front layer (strips).** The first accordion sampling, segmented into very fine strips along η with $\Delta\eta \times \Delta\phi = 0.0031 \times 0.1$ (a pitch of $0.025/8$ in η) over $|\eta| < 1.4$. The fine η granularity enables rejection of $\pi^0 \to \gamma\gamma$ from prompt photons up to a few tens of GeV of transverse energy, and also provides photon pointing capability. The layer is about 4.3 X_0 deep; counting the roughly 1.7 X_0 of inactive material in front of the active accordion volume, it ends at about 6 X_0.

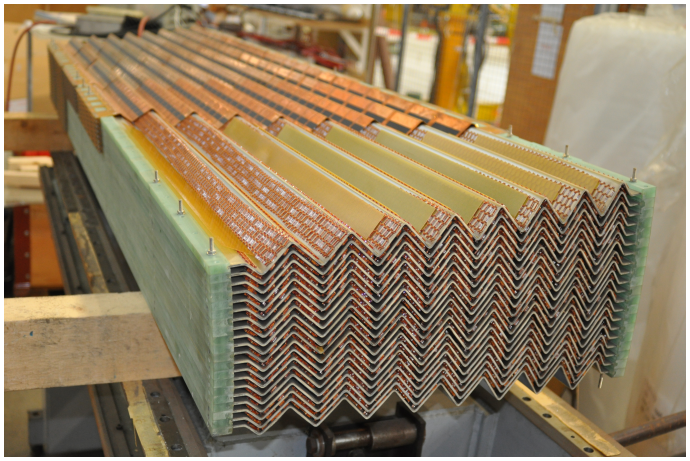

Figure 5.6: One of the accordion-shaped absorber modules during construction [80].

- **Middle layer.** The second and thickest accordion sampling, with square towers of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. It absorbs the bulk of the electromagnetic shower energy and drives the energy resolution, with a thickness of $\sim 16\,X_0$ at $\eta = 0$ (from $\sim 4\,X_0$ to $\sim 20\,X_0$ measured from the front of the active accordion). The transverse cell size is matched to the Molière radius of lead/LAr, so that an EM shower deposits the largest fraction of its energy in a 3×7 cluster of middle cells in $\eta \times \phi$.
- **Back layer.** The third accordion sampling, with coarser granularity $\Delta\eta \times \Delta\phi = 0.05 \times 0.025$. It catches the longitudinal tail of high-energy showers and estimates energy leakage beyond the calorimeter. Near $\eta = 0$, the middle layer is truncated at $\sim 22\,X_0$ of total depth, including the inactive material in front, leaving a minimum of about $2\,X_0$ for the back layer.

The granularities and depths of all four samplings are summarized in Table 5.1 and illustrated in Figure 5.7.

Table 5.1: Granularity and depth of the EMB samplings at $\eta \approx 0$. Depths vary with η and are thicknesses of the active layers; the back layer is a minimum. The inactive material in front of the accordion is listed separately and is not part of the accordion depth.

Sampling	$\Delta\eta$	$\Delta\phi$	Depth [X_0]
Presampler	0.025	0.1	~ 0.1
Front	0.025/8	0.1	~ 4.3
Middle	0.025	0.025	~ 16
Back	0.05	0.025	$\gtrsim 2$

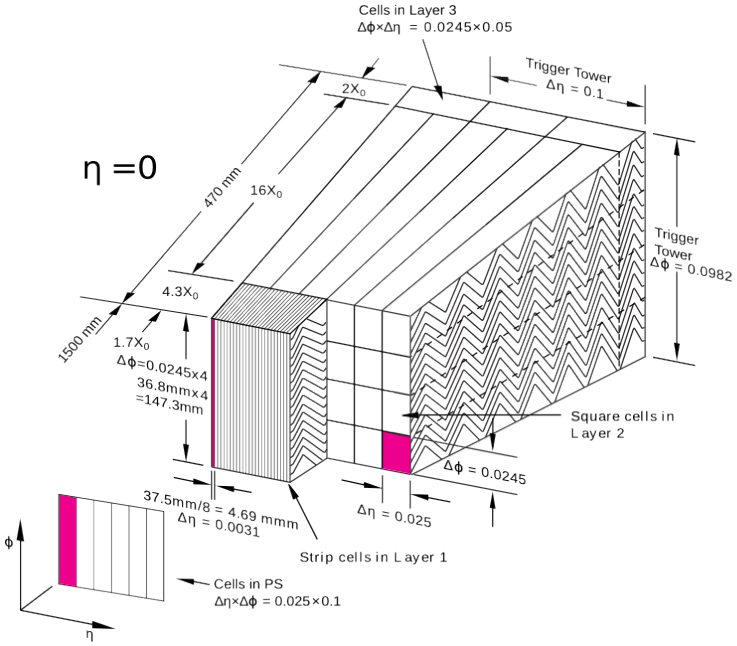

Figure 5.7: Structure of a module of the LAr EMB. The four samplings (presampler, front, middle and back) are shown with their respective granularity in ϕ and η [81].

5.4 End-cap calorimeters

Each end-cap cryostat contains the Electromagnetic End-cap Calorimeter, the Hadronic End-cap Calorimeter, and the Forward Calorimeter, as shown schematically in Figure 5.8.

The EMEC shares the same lead–liquid-argon accordion technology as the EMB but differs in several aspects. The accordion waves run radially rather than axially, and the liquid argon gap is not constant but increases with radius, varying from 0.9 to 2.8 mm in the outer wheel ($1.375 < |\eta| < 2.5$) and from 1.8 to 3.1 mm in the inner wheel ($2.5 < |\eta| < 3.2$). To compensate, the high voltage is adjusted in discrete steps from 1000 V to 2500 V as a function of $|\eta|$. The total depth exceeds 24 X_0. In the precision region ($1.5 < |\eta| < 2.5$), the EMEC reproduces the EMB longitudinal segmentation. An end-cap presampler, thinner than the barrel one and covering only $1.5 < |\eta| < 1.8$, is followed by front, middle, and back accordion samplings of comparable granularity. In the inner wheel, the accordion is segmented into only two longitudinal samplings with coarser transverse granularity.

The HEC departs significantly from the accordion design, adopting instead a conventional parallel-plate geometry with flat copper absorbers and liquid argon gaps of 8.5 mm, shown in Figure 5.9. Each end-cap consists of two wheels of 32 modules: the front wheel (HEC1) with 25 mm thick plates and the rear wheel (HEC2) with 50 mm thick plates. Each wheel is itself segmented into two longitudinal samplings. The trans-

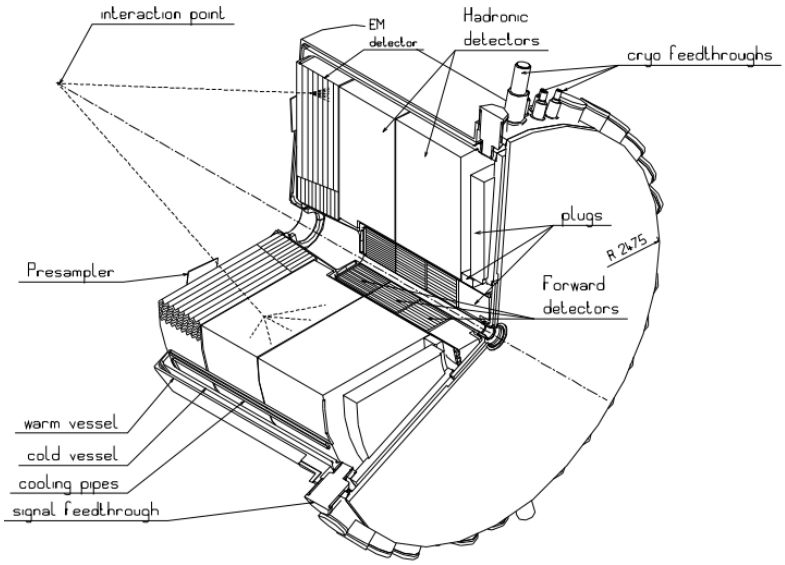

Figure 5.8: A schematic view of the end-cap cryostat. It contains three subdetectors, the EMEC, HEC and FCal.

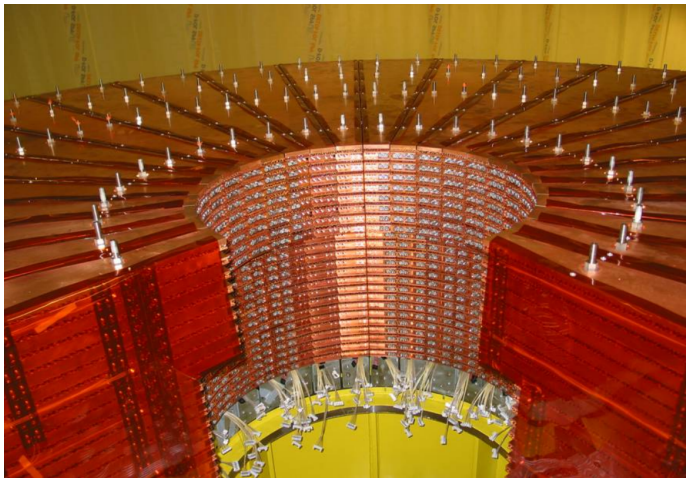

Figure 5.9: A slice of typical HEC copper plate absorbers [82].

verse granularity is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for $|\eta| < 2.5$ and 0.2×0.2 for $2.5 < |\eta| < 3.2$. Within each gap, three electrodes form an electrostatic transformer structure, creating an effective double gap of 3.6 mm that reduces the required voltage to 1800 V and limits ion-buildup effects. Cryogenic GaAs preamplifiers mounted directly on the HEC wheels inside the cryostat further optimize the signal-to-noise ratio.

5.5 Forward calorimeter

The FCal covers the highest pseudorapidity region ($3.1 < |\eta| < 4.9$) and is integrated in the end-cap cryostat, within the inner bore of the HEC, approximately 4.7 m from the interaction point. It comprises three longitudinal modules: FCal1, the electromagnetic module closest to the interaction point, uses copper absorbers; FCal2 and FCal3 are hadronic modules using tungsten absorbers. Rather than an accordion or parallel-plate geometry, all three modules employ a tubular electrode structure in which metal rods (anodes) are centered inside metal tubes (cathodes), arranged in a hexagonal pattern of approximately 10,000 tubes per module. An example of this structure is illustrated in Figure 5.10.

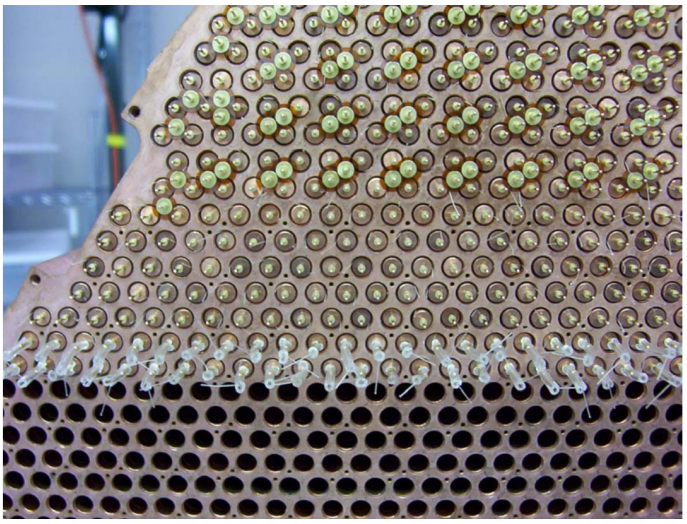

Figure 5.10: A detail of the structure of the FCal metal rods [83].

This design produces extremely narrow liquid argon gaps, essential to prevent positive-ion buildup at the very high particle fluxes in the forward region. FCal1 has gaps of only 269 µm with a drift time of about 60 ns; FCal2 and FCal3 have gaps 50% and 100% larger than FCal1, respectively, permitted by the lower ionization density deeper in the shower. In all three modules, the rod material matches the absorber matrix (copper or tungsten), and the radiation-hard materials throughout ensure stable long-term performance.

5.6 Readout electronics and trigger system

5.6.1 Readout electrodes

As mentioned in the earlier section, the calorimeter is composed of liquid argon gaps alternating with the absorber material. Inside of these gaps, a high voltage (HV) is applied around the signal electrodes, which are composed of three layers of copper separated by Kapton (see Figure 5.11). The two outer layers are held at the HV potential, which is distributed through resistors made of resistive paint silk-screened on the outer copper, while the central layer collects the signal through capacitive coupling [84, 85].

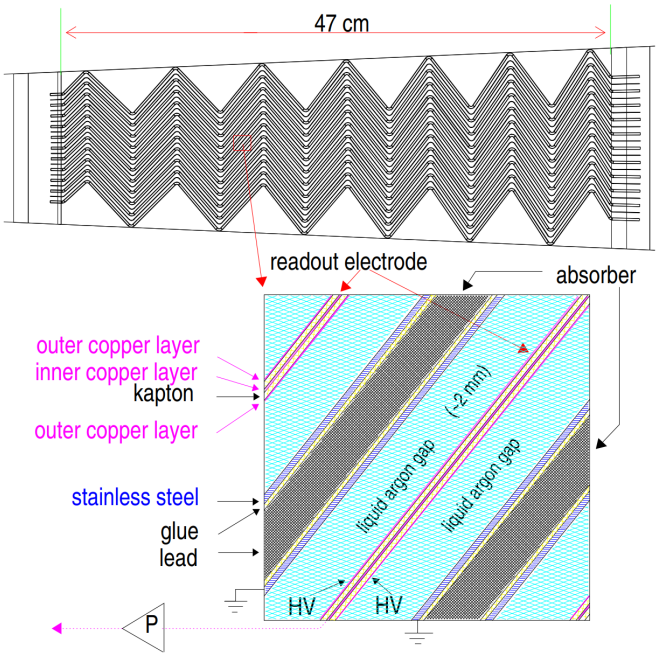

Figure 5.11: The structure of the electrodes within the liquid argon gaps in an EMB module [86].

The HV creates an electric field, where the electrons produced by the ionization in the liquid argon are drifted toward the electrodes in the central layer, inducing a drift current that results from the sum of both gaps surrounding the electrode. This induced current is the main building block of the readout signal in the calorimeter. Its duration is set by the drift time $t_d = w_{gap}/v_{drift}$, where the drift velocity v_{drift} grows with the electric field, so that t_d is fixed locally by the gap width w_{gap} and by the applied HV [86]. In the EMB, the 2.1 mm gap on each side of the electrode corresponds to a drift time of about 450 ns at 2000 V [85]. Each subdetector requires a distinct electrode geometry suited to its absorber structure and radiation environment, as illustrated in Figure 5.12. In the EMB, accordion-folded copper–Kapton electrodes conform to the lead absorber shape, maintaining a uniform 2.1 mm gap at a nominal voltage of 2000 V [75]. The EMEC

uses the same accordion technology but with radially oriented waves and a gap that varies with radius; the HV is adjusted in steps from 1000–2500 V across $|\eta|$ to maintain adequate charge collection despite the non-uniform gap (see Section 5.4). The HEC employs flat annular electrodes in the electrostatic transformer arrangement, and the FCal uses a tubular rod-in-tube geometry with gaps as narrow as 269 µm, requiring correspondingly lower voltages.

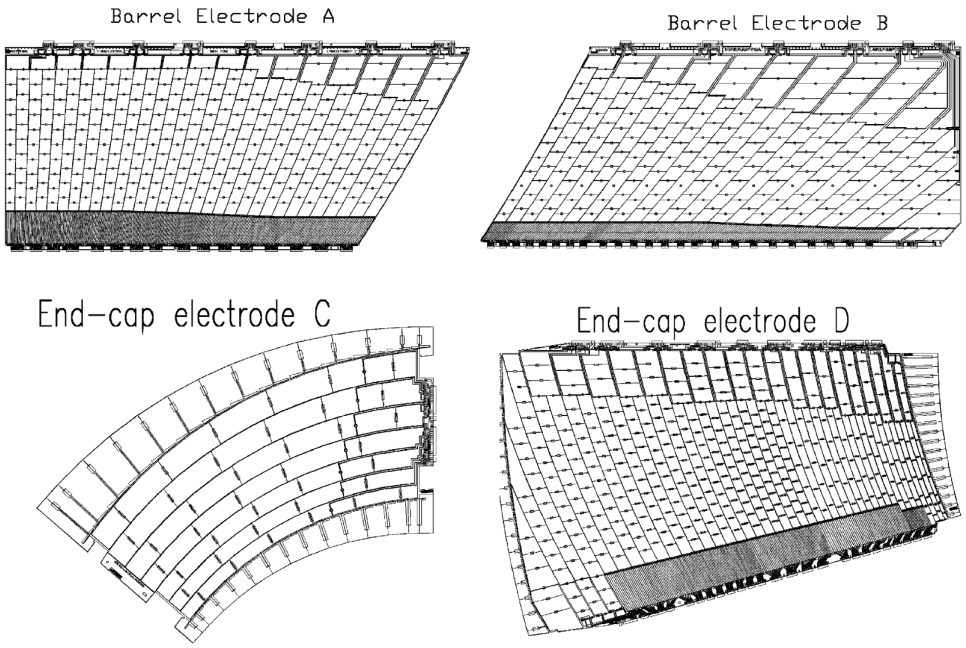

Figure 5.12: Layout of the different electrode types in the subdetectors [84].

The HV is distributed to the calorimeter in sectors. In the EMB, each side of an electrode is powered independently, in sectors of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ corresponding to 32 electrodes, so that a faulty supply affects only one of the two gaps contributing to the signal of a cell. In the EMEC, the outer (inner) wheel is divided into seven (two) HV sectors in $|\eta|$, each subdivided in ϕ into groups of 24 (4) electrodes, again powered independently on each side [85]. The nominal HV is therefore a function of $|\eta|$ only, and takes the same value in all the ϕ-sectors at a given $|\eta|$. The electrodes also define the cell geometry. The segmentation in η and in depth is etched into the copper of the electrodes, and in the EMB two electrode types are used, covering $|\eta| < 0.8$ (type A) and $0.8 < |\eta| < 1.475$ (type B), matching absorbers with lead thicknesses of 1.53 mm and 1.13 mm, respectively. The segmentation in ϕ is obtained instead by summing adjacent electrodes: 16 for a cell of the front layer and 4 for a cell of the middle and back layers [85].

These construction choices determine the shape of the signal. The ionization pulse of a cell is fixed by its drift time, by the inductance L and capacitance C of the cell and of

its connection to the summing boards, which depend on the cell size, and by the shaping of the readout electronics described in the following. These are the parameters through which the physics pulse is predicted from the calibration pulse [85]. Since the accordion structure, the electrode layout and the HV settings repeat identically in ϕ, cells located at the same η in a given layer share the same drift time and nearly the same LC. Hence, to a good approximation, the pulse shape is a function of the layer and of η alone.

The LAr calorimeter readout electronics are designed to measure energies from about 50 MeV to 3 TeV, by combining three overlapping gain scales with ratios of approximately ten between consecutive gains [60]. The triangular ionization currents induced in the drift gaps are shaped and continuously sampled at 40 MHz, synchronously with the LHC bunch-crossing clock, as illustrated in Figure 5.13. While awaiting a Level-1 (L1) trigger decision, the samples are held in analog pipeline memories located on the detector. Upon an L1 accept, a fixed number of samples around the pulse maximum are digitized and transmitted to the back-end electronics, where the signal amplitude and timing are reconstructed. The L1 accept rate was 75 kHz during Run 1 and has been 100 kHz since Run 2 [59].

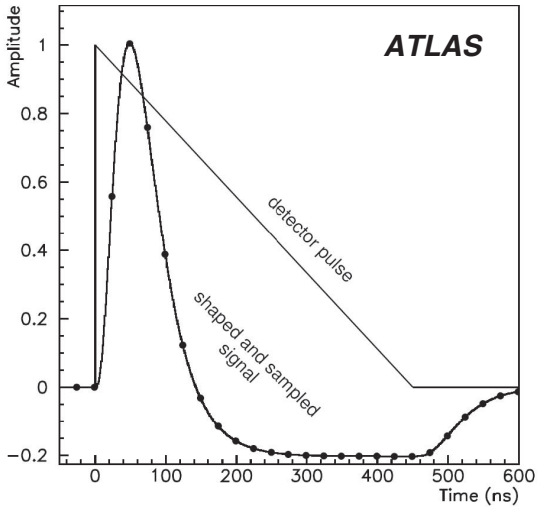

Figure 5.13: Triangular ionization current in a LAr cell and the shaped and amplified signal. The black points on the shaped pulse indicate the ADC samples taken every 25 ns, synchronously with the LHC bunch-crossing clock [60].

Two parallel readout paths coexist in the LAr system. The *precision (or main) readout path* digitizes the shaped pulses on the front-end boards only after an L1 accept and feeds the reconstructed cell energies to the ATLAS High-Level Trigger and offline reconstruction. The *trigger readout path* delivers calorimeter information at every bunch crossing to the L1 calorimeter trigger system. A schematic overview of the readout architecture as operated during Run 1 and Run 2 is shown in Figure 5.14. The architecture

is split between *front-end (FE) electronics*, located directly on the cryostats, and *back-end (BE) electronics*, housed in the USA-15 underground service cavern.

5.6.2 Front-end electronics

The 58 FE crates are mounted directly on the LAr cryostats, in the gap between the barrel and the end-cap cryostats and on the outer faces of the end-cap cryostats. The high voltages required to drift the ionization electrons across the LAr gaps, in the range of roughly 250 to 2500 V depending on the subdetector, are brought into the cryostat through warm feedthroughs located on its upper surface.

Each FE crate hosts five types of boards [87]:

- **Front-End Boards (FEBs):** the FEBs read out and digitize the signals from the LAr cells. Each FEB is a ten-layer printed circuit board processing 128 channels. The raw ionization signals are amplified by preamplifiers, then split and shaped into bipolar pulses by shaper chips with a time constant $\tau_{sh} = RC \approx 15$ ns, producing three overlapping gain scales. Gain selector chips choose the appropriate gain for each channel on the basis of the peak sample value. The shaped signals are sampled at 40 MHz and held in Switched Capacitor Array analog pipelines with up to 144 cells, corresponding to a maximum L1 latency of 2.5 µs. Upon an L1 accept, typically two to six samples per channel are digitized by 12-bit analog-to-digital converters (ADCs), multiplexed, and transmitted to the back-end over 1.6 Gbps optical links;

- **Layer Sum Boards (LSBs):** the LSBs perform the first stage of the analog sums feeding the L1 trigger path. Shaped signals from groups of cells are summed on the LSBs and then routed, via the FE-crate baseplane, either to the Tower Builder Boards (legacy, up until the middle of Run 3) or to the LAr Trigger Digitizer Boards (LTDBs; Phase-I configuration, see Chapter 6);

- **Tower Builder Boards (TBBs):** in the legacy trigger configuration, the partial sums produced by the LSBs were further combined on the TBBs to form Trigger Towers (TTs) of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, without longitudinal segmentation. The resulting analog signals were transmitted over to the L1 calorimeter trigger in USA-15 [88]. This path has been fully retired after the commissioning of the Phase-I Digital Trigger.

- **Calibration Boards:** the Calibration Boards inject precisely known voltage pulses into dedicated resistors mounted on the mother-boards inside the cryostats, reproducing as closely as possible the physics signal. The injected pulse has a fast rising edge matching that of the triangular physics signal but an exponential decay, due to the inductively injected current;

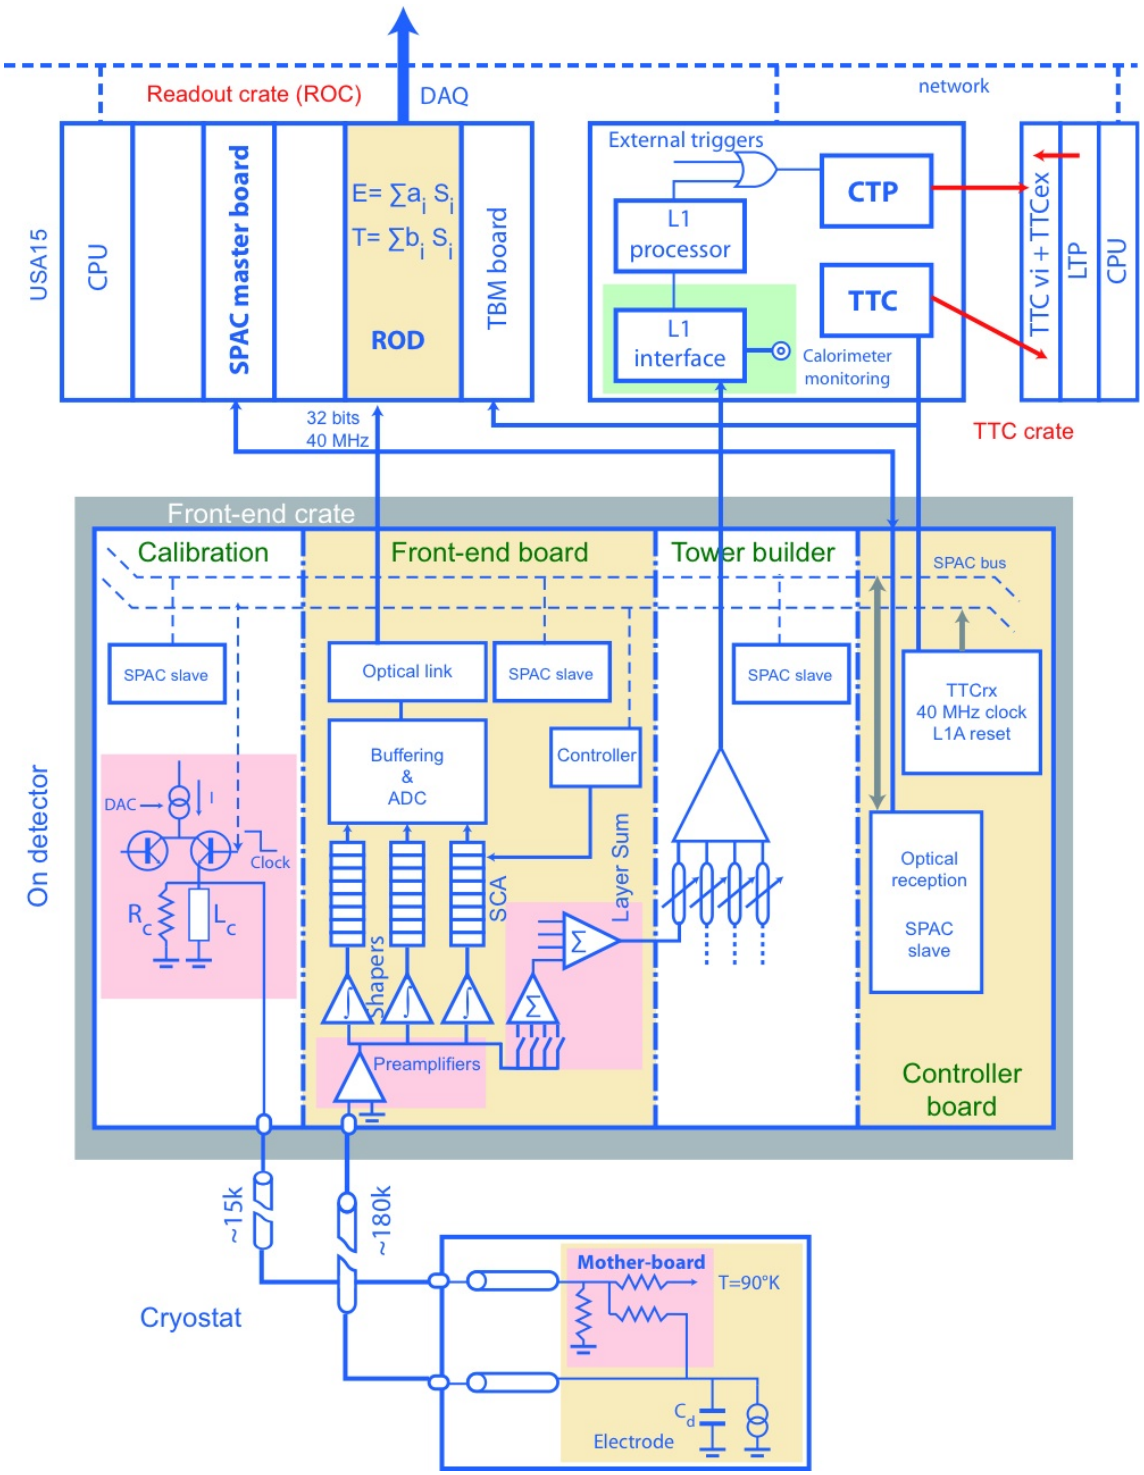

Figure 5.14: Schematic diagram of the legacy LAr readout electronics system, in the configuration operated during Run 1 and Run 2. The front-end electronics (bottom) are mounted on the cryostat, while the back-end electronics (top) are located in the USA-15 service cavern [81].

- **Controller Boards:** the Controller Boards receive the 40 MHz LHC clock and the Trigger, Timing and Control (TTC) signals and distribute them to the other boards of the FE crate, together with the configuration and control messages.

5.6.3 Back-end electronics

The back-end electronics filter, format, and monitor the digitized LAr signals, and reconstruct for each channel the energy E, the signal timing, and a pulse-quality factor. The data arrive at the back-end through 1.6 Gbps optical links from the FEBs and are processed inside VME-based Readout Driver (ROD) crates installed in USA-15 [89]. Each ROD crate contains:

- **ROD boards:** the RODs receive the digitized samples from the FEBs at the L1 accept rate and compute, for each channel, the deposited energy, the signal time, and a pulse-quality factor. The energy and time are reconstructed using an *optimal filtering* algorithm [90],

$$E = \sum_{i=1}^{n} a_i \, (S_i - p), \qquad E \cdot t = \sum_{i=1}^{n} b_i \, (S_i - p), \tag{5.7}$$

  where S_i are the ADC samples, p is the pedestal, and the coefficients a_i and b_i are precomputed from the expected pulse shape and the noise autocorrelation so as to minimize the variance on E and t. More on the optimization process of the optimal filtering coefficients (OFCs) is described in Section A.2. The quality factor measures the agreement of the reconstructed pulse with the expected shape and is used to flag mismeasured signals. Each ROD Digital Signal Processor (DSP) is capable of approximately 5.7×10^9 instructions per second, which corresponds to a processing budget of about 10 µs per event at a 100 kHz L1 rate for the full reconstruction of the FEB samples [89];

- **CPU Board:** a VME-based processor that controls and monitors the ROD crate and provides the interface to the ATLAS data acquisition and run-control infrastructure;

- **SPAC Master Board:** configures the FE and BE boards and reads back their registers through the Serial Protocol for ATLAS Calorimeters (SPAC);

- **Trigger Busy Module (TBM):** receives the trigger, timing and control signals, including the LHC clock and the synchronous commands, and propagates a busy signal back to the Central Trigger Processor whenever the crate is unable to accept further events.

5.7 Noise and pileup

At the design LHC luminosity of $\mathcal{L} = 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the vast majority of the interactions occurring are soft and produce a diffuse, roughly η-uniform background of low-momentum particles, the pileup introduced in Section 3.3, which is overlaid on every recorded event. From minimum-bias simulations used for the LAr design, the predicted charged-particle multiplicity is about $\langle dN_{\mathrm{ch}}/d\eta \rangle \approx 7.9$ with an average transverse momentum of $\langle p_{\mathrm{T}} \rangle \approx 460\,\mathrm{MeV}$ per particle [75], as illustrated in Figure 5.15.

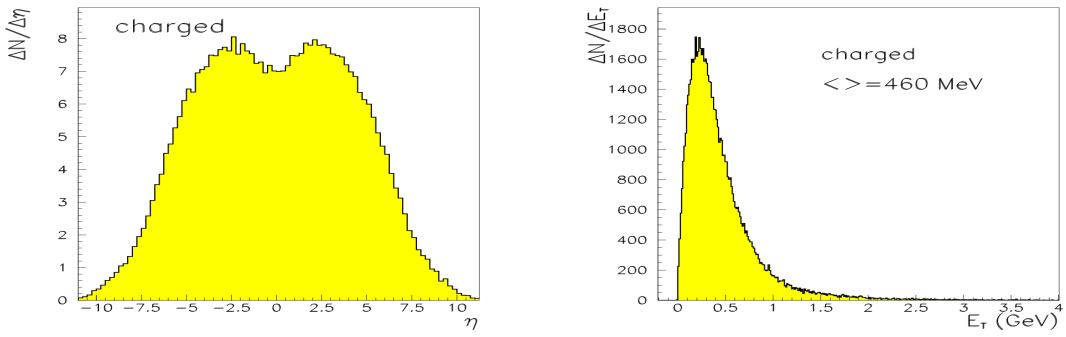

(a) (b)

Figure 5.15: Predicted charged-track multiplicity as a function of pseudorapidity (a) and transverse-momentum spectrum (b) for minimum-bias events at $\sqrt{s} = 14\,\mathrm{TeV}$, used as input to the LAr noise and pileup studies [75].

It is convenient to distinguish two contributions to pileup. *In-time* pileup is the energy deposited in the same bunch crossing as the hard-scatter event of interest, by the additional ∼22 soft interactions expected at design luminosity [75][1]. *Out-of-time* pileup is the residual energy from interactions that took place in adjacent bunch crossings. Because the LAr ionization signal has a slow (∼450 ns) triangular tail (see Figure 5.13), a single cell is sensitive to soft deposits from many tens of bunch crossings before and after the triggered one.

To suppress the long ionization tail and limit the integration window in time, the LAr front-end electronics applies a bipolar CR–RC^2 shaping to the signal which, by construction, integrates to zero in time. As a consequence, when many minimum-bias events are overlaid uniformly across the bunch train, the positive lobe of the in-time signal is exactly canceled, on average, by the negative lobes of the out-of-time signals from neighboring bunch crossings. The mean energy contributed by pileup to a given calorimeter cell is therefore zero. In a typical electromagnetic cluster the residual mean is only a few tens of MeV [75].

[1] In Run 3, with leveling at $\langle \mu \rangle = 64$, they are about 60, while they are expected to reach 200 at the HL-LHC.

While the mean of the pileup contribution vanishes, its *fluctuations* do not. Pileup acts as an additional noise source, also known as *pileup noise*, whose root mean square (RMS) in a given cell is proportional to the quantity

$$I = \left(\sum_i \frac{f_i^2}{f_{\max}^2} \right)^{1/2}, \tag{5.8}$$

where f_i is the amplitude of the shaped pulse at the i-th bunch crossing relative to the same hard interaction, normalized to its peak value $f_{\max}$ [75]. The sum runs over all bunch crossings to which a given cell is sensitive. Equation (5.8) makes it explicit that pileup noise depends on both the shape of the pulse and the LHC bunch structure. A longer, slower pulse spans more bunch crossings, includes more f_i terms in the sum, and therefore gives a larger pileup noise. For a typical electromagnetic cluster, the resulting pileup-noise RMS at design luminosity is of order $\sigma_{\text{pu}} \sim 300\,\text{MeV}$ in E_{T} [75], which is comparable to the electronic noise contribution and therefore non-negligible for the calorimeter resolution at low energies.

A natural way to reduce I, and hence the pileup noise, is to make the shaped pulse shorter by reducing its peaking time t_p. Doing so, however, comes at a cost: shortening t_p widens the bandwidth of the front-end amplifier and chain, which increases the *electronic* noise approximately as $1/\sqrt{t_p}$, while the pileup noise scales approximately as $\sqrt{t_p}$ [75]. The total cell noise, obtained by summing the two contributions in quadrature, therefore exhibits a minimum at an intermediate value of t_p, as shown in Figure 5.16. This minumum sets the LAr shaper design: the chosen peaking time of $t_p \approx 45\,\text{ns}$ minimizes the total noise at the nominal high-luminosity operating point of the LHC [75].

This unavoidable pileup noise is, from the point of view of energy reconstruction, a nuisance that the LAr shaper design was optimized to suppress. From a different point of view, however, it is also a probe of the instantaneous luminosity. As Eq. (5.8) makes explicit, the pileup contribution to a cell is built from the soft interactions taking place in the bunch crossings to which that cell is sensitive, and its magnitude therefore grows with the mean number of inelastic interactions per bunch crossing μ, itself proportional to the per-bunch luminosity through Eq. (7.5). In other words, the very fluctuations that limit the calorimeter resolution at low energies carry, encoded in their amplitude, an instantaneous measurement of how many soft collisions are taking place in the LHC. The dependence of this dispersion on the luminosity, and its use as a luminosity estimator, are developed in Chapter 8.

This chapter has described the LAr calorimeter, from its physical subdetectors through the readout chain and its noise characteristics. The Phase-I upgrade of the trigger readout, which is the central technical subject of this thesis, is described in Chapter 6. The luminosity measurement method that exploits the pileup-driven RMS of the digitized samples discussed above is developed in Chapter 8; thus not using the reconstructed cell energy, which has already been filtered, but the RMS of the ADC

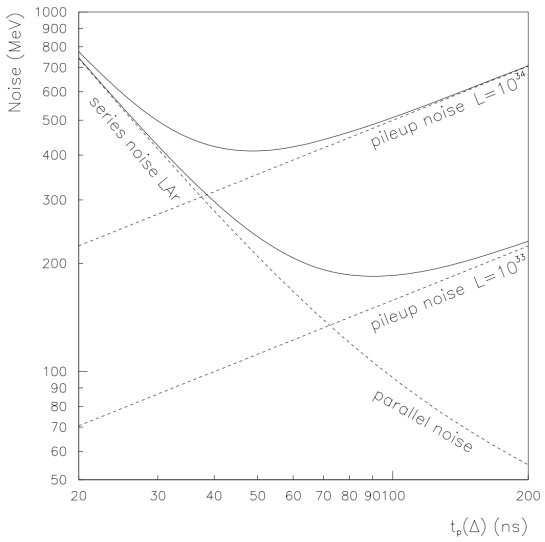

Figure 5.16: Optimization of the LAr shaper peaking time t_p: electronic noise decreases as $\sim 1/\sqrt{t_p}$ while pileup noise increases as $\sim \sqrt{t_p}$, so that the sum in quadrature has a minimum at intermediate t_p. The chosen value of $t_p \approx 45\,\text{ns}$ minimizes the total noise at the nominal LHC luminosity [75].

samples, which preserves the full pileup-driven dispersion before any optimal filter is applied.

Chapter 6

The Digital Trigger upgrade

The ATLAS detector upgrade program is structured in two phases aligned with the long shutdowns of the LHC. The Phase-I upgrade was installed during Long Shutdown 2 (LS2, 2019–2022) and commissioned during the first years of Run 3, targeting the instantaneous luminosities of $\mathcal{L} \sim 2 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and average pileup $\langle\mu\rangle \sim 60$ foreseen for this running period [91]. The main upgrade implemented in the LAr calorimeter is the introduction of the new Digital Trigger (DT) readout path for the LAr calorimeters, which is the subject of this chapter, and the new Feature EXtractor (FEX) processors (eFEX, jFEX, gFEX) of the upgraded Level-1 calorimeter trigger (L1Calo), which uses the DT information to build electromagnetic objects, jets, and global quantities such as missing transverse momentum for the L1 trigger decision [92]. In parallel, the muon spectrometer received a new inner forward detector, the New Small Wheel [65], to maintain the L1 muon trigger rate and resolution at Run 3 luminosities, and the trigger and data acquisition (TDAQ) system was upgraded with a new readout architecture based on the Front-End Link Interface eXchange (FELIX) [92].

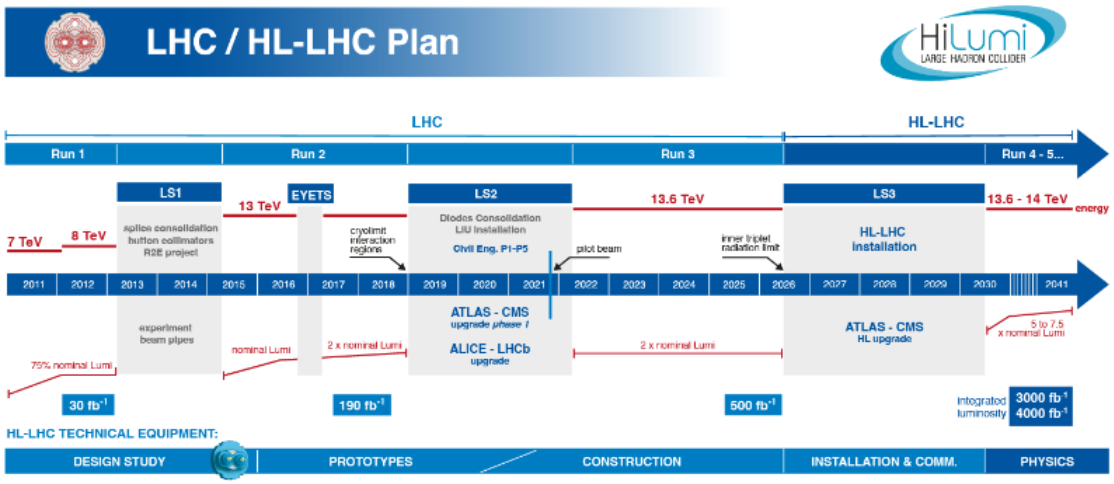

Figure 6.1: Overview of the LHC and HL-LHC long-term plan, showing the center-of-mass energy of the collisions and the expected instantaneous luminosity. The cumulative integrated luminosity at the end of each data-taking period is also indicated [93].

The Phase-I upgrade is designed to be forward-compatible with the Phase-II upgrade, which will be installed during Long Shutdown 3 (LS3) to prepare the detector for

the HL-LHC [93]. At the HL-LHC, ATLAS will operate at instantaneous luminosities of up to $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ with an average pileup $\langle\mu\rangle$ of about 200, and is expected to collect a total integrated luminosity of $3000\,\mathrm{fb}^{-1}$ over roughly a decade of running [94]. The LHC and HL-LHC plan is shown in Figure 6.1.

Meeting this goal requires a deep restructuring of the ATLAS readout and trigger systems: a completely new all-silicon Inner Tracker (ITk) replaces the current Inner Detector [95, 96], the LAr and Tile calorimeter front-end electronics are replaced so that every bunch crossing is digitized and streamed off detector for trigger and readout [97, 98], and a new High-Granularity Timing Detector (HGTD) improves pileup rejection in the forward region [99]. The trigger keeps the two-stage structure of the current Level-1 / HLT system, at higher rates and with a longer latency: a single hardware level, called Level-0 (L0) in the TDAQ design report, with an accept rate of 1 MHz and a latency of up to 10 µs, followed by a software Event Filter (EF) that reduces the output rate to 10 kHz [100]. The plans for tracking in the EF were later revised in an amendment to that report [101]. The LAr Phase-I components (LTDBs, baseplanes and the LAr Digital Processing System, LDPS) are kept during the HL-LHC era, so the Digital Trigger developed for Run 3 directly becomes part of the Phase-II Level-0 calorimeter trigger [97].

Within this broader program, the LAr Phase-I upgrade introduced a new Digital Trigger readout path for the Liquid Argon calorimeters, which became the sole operational L1 calorimeter trigger input path from the 2024 data-taking period onward, after several years of parallel commissioning alongside the legacy analog system [102]. The DT replaces the coarse analog Trigger Tower sums of the legacy system with longitudinally segmented *Super Cells* (SCs), digitizes them on detector at 40 MHz, and performs all trigger-level signal processing in the off-detector LDPS. The remainder of this chapter describes the physics requirements that motivated this upgrade (Section 6.1), the SC scheme and the associated front-end and back-end electronics, the commissioning of the Digital Trigger, and its measured performance in Run 3.

6.1 Physics requirements for Phase-I and HL-LHC

With the discovery of the 125 GeV Higgs boson in Run 1, the ATLAS physics program has shifted toward precision measurements of Higgs boson couplings, searches for rare and forbidden decays, and searches for physics beyond the Standard Model [91, 103]. Many of these measurements rely on final states at transverse momenta close to the electroweak scale, where the trigger system is the dominant source of acceptance loss. The Phase-I and Phase-II upgrades are therefore driven not by the desire to trigger on higher-p_{T} objects, which scale naturally with luminosity, but by the need to *preserve* low p_{T} thresholds on electrons, photons, taus, jets, and $E_{\mathrm{T}}^{\mathrm{miss}}$ while keeping the overall L1 output rate within its allowed range of 100 kHz in Run 3 and 1 MHz at the HL-LHC [92,

100].

The main difficulty is the scaling of pileup with luminosity. Both the occupancy of the calorimeter and the event-by-event fluctuations of the underlying event grow with $\langle\mu\rangle$, and running with a fixed L1 trigger menu at higher $\langle\mu\rangle$ produces higher rates dominated by QCD background and by pileup-induced objects. The standard way to keep the rate constant is to raise the L1 p_{T} thresholds, which directly reduces the signal acceptance of analyses whose final states peak near the threshold. Some examples [91]:

- For vector-boson-fusion $H \to \tau\tau$ in the $e\tau_{\mathrm{had}}$ channel, raising the L1 electron threshold from 20 to 25 GeV and the τ threshold from 40 to 50 GeV reduces the signal acceptance by about 37%.
- For $H(\to b\bar{b}) + Z(\to \nu\nu)$, the L1 $E_{\mathrm{T}}^{\mathrm{miss}}$ threshold drives the offline selection; keeping the L1 threshold at a level that allows an offline requirement of $E_{\mathrm{T}}^{\mathrm{miss}} > 160$ GeV rather than 200 GeV preserves a factor of about two in signal acceptance.
- For $H \to WW^*$ with dilepton final states, higher single-lepton thresholds would cost about 24% of the acceptance in the $ee\nu\nu$ channel and 12% in $e\nu\mu\nu$.

Meeting these requirements from the LAr and L1Calo side is achieved through improvements provided by the Phase-I upgrade [91]:

1. **Finer transverse granularity and longitudinal segmentation.** The legacy Trigger Towers are replaced by the Super Cell scheme, giving the L1 calorimeter trigger access to the same shower-shape information used offline and at the HLT for electron, photon, and tau identification. More details are given in Section 6.2.1.
2. **Shower-shape discrimination at L1.** Variables such as the lateral shower width in the middle layer ($w_{\eta,2}$), the ratio of energies in 3×2 versus 7×2 SC groups (R_η), and the fraction of energy in the back layer (f_3) can now be computed in the Feature EXtractor firmware. These variables, which closely mirror the variables used in the offline electron and photon identification and significantly improve background rejection against QCD jets at fixed signal efficiency [91], are designed to reduce the L1 electron trigger rate by up to a factor of four at $E_{\mathrm{T}} \approx 20$ GeV for the same electron efficiency, relative to the legacy TT-based selection [91].
3. **ROD-style optimal filtering at L1.** The LAr Trigger prOcessing MEzzanines (LATOMEs; see Section 6.2.3) reconstruct each SC transverse energy every 25 ns by applying OFCs to the stream of ADC samples, in direct analogy with the algorithm running in the precision-readout RODs (see Eq. (5.7)). This significantly improves both the L1 energy resolution and the BCID assignment compared to the legacy analog peak-finder, which sharpens the turn-on curves of electromagnetic, jet, and $E_{\mathrm{T}}^{\mathrm{miss}}$ L1 triggers and reduces the probability of mis-assignment of out-of-time signals to the wrong bunch crossing.

4. **Extended dynamic range.** The 12-bit ADCs on the LTDBs cover transverse energies up to roughly 100 GeV in the front layer and 400 GeV in the middle layer without saturation, with a dynamic range per SC exceeding that of the legacy trigger path across η. This prevents saturation of the L1 inputs for high-E_{T} objects relevant to searches for heavy resonances and high-mass BSM signatures.

These improvements allow the L1 calorimeter trigger to maintain, or even lower, the Run 2 p_{T} thresholds on electrons, photons, taus, and jets while keeping the total L1 rate within the limits at Run 3 luminosities, and they provide the basis of the Level-0 calorimeter trigger that will operate at the HL-LHC [92, 100]. The precision measurements and searches listed above therefore retain acceptance where, without the Phase-I upgrade, they would have been forced to either raise their offline thresholds or accept significant signal loss.

6.2 Phase-I upgrade and the Digital Trigger

Until the first part of Run 3, the LAr contribution to the L1 calorimeter trigger was derived from coarse analog sums of cell, or Trigger Towers, of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ and without longitudinal segmentation [75, 92]. At the higher instantaneous luminosities and pileup levels of Run 3, such coarse granularity would have forced significant increases of the L1 transverse momentum thresholds to stay within the allowed L1 output rate. The Phase-I upgrade addresses this by replacing the analog TT sums with longitudinally segmented Super Cells, illustrated in Figure 6.2, by digitizing SC signals on the detector, and by performing all trigger-level signal processing in the off-detector digital electronics. This allows to use shower-shape information at L1 and applying the same optimal filtering algorithms used in the precision readout [91, 102].

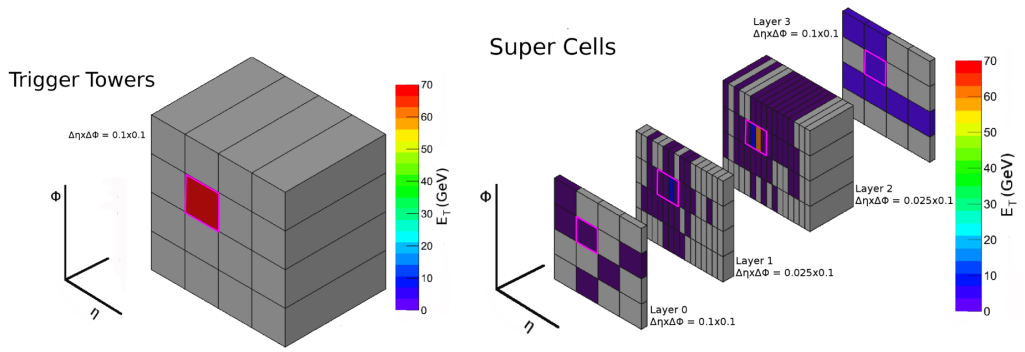

Figure 6.2: Comparison of the legacy Trigger Tower scheme (left) and the Super Cell scheme (right) for an electron energy deposit of $E_{\mathrm{T}} = 70$ GeV. The SCs preserve the longitudinal layer information and provide finer transverse granularity in the EM front and middle layers [91].

6.2.1 The Super Cell scheme

The SC scheme follows the natural longitudinal segmentation of the calorimeter. In the front and middle electromagnetic layers the transverse granularity is refined from 0.1×0.1 to $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$, while the presampler and the back electromagnetic layer retain $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. In the barrel, this corresponds to ten SCs per legacy trigger tower, with fewer SCs per TT in the end-cap regions. Overall, the more than 182,000 LAr readout cells are summed into approximately 34,000 SCs, providing up to ten times more trigger signals than the legacy system and preserving layer information that is otherwise lost in the analog TT sum. A direct comparison of the two schemes is shown in Table 6.1 and Figure 6.3.

Table 6.1: Transverse granularity of the electromagnetic calorimeter in the four longitudinal layers, expressed in elementary readout cells, in the legacy Trigger Tower scheme, and in the Phase-I Super Cell scheme. For each layer, $n_\eta \times n_\phi$ gives the number of elementary cells summed into one Trigger Tower or Super Cell, and $\Delta\eta \times \Delta\phi$ is the resulting transverse size. Values correspond to the EMB and EMEC outer wheel ($|\eta| < 2.5$) [91].

	Layer	Elementary Cell	Trigger Tower		Super Cell	
		$\Delta\eta \times \Delta\phi$	$n_\eta \times n_\phi$	$\Delta\eta \times \Delta\phi$	$n_\eta \times n_\phi$	$\Delta\eta \times \Delta\phi$
0	Presampler	0.025×0.1	4×1	0.1×0.1	4×1	0.1×0.1
1	Front	0.003125×0.1	32×1		8×1	0.025×0.1
2	Middle	0.025×0.025	4×4		1×4	0.025×0.1
3	Back	0.05×0.025	2×4		2×4	0.1×0.1

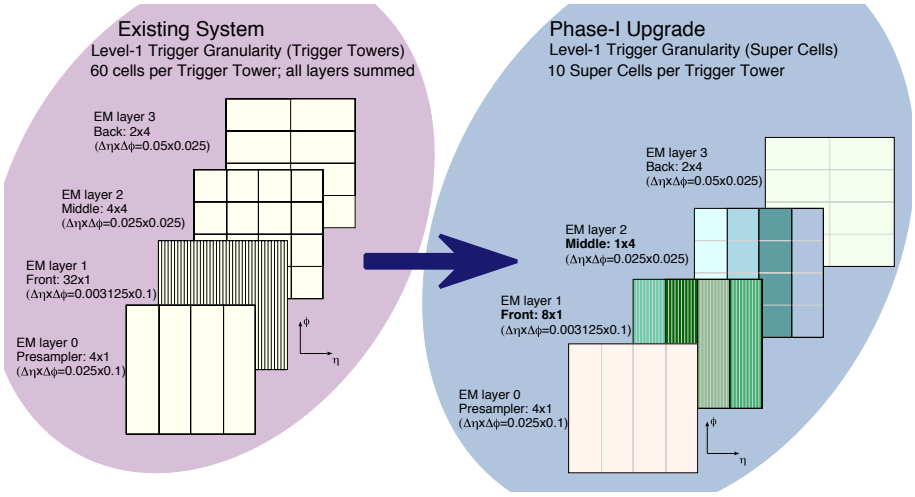

Figure 6.3: The comparison between the two schemes in terms of L1-trigger granularity for each layer of the EM part of the detector [91].

The finer granularity, combined with per-layer information, and the shower-shape

discriminants, translate into sharper turn-on curves and reduced trigger rates for the same L1 p_{T} thresholds, as discussed in Section 6.3.

6.2.2 Front-end electronics

The signal path to the new DT branches off from the main readout at the linear mixer stage of the FEBs. A schematic view of the full trigger readout electronics chain is given in Figure 6.4. The following front-end components were modified or added during LS2 [91, 102]:

- **New Layer Sum Boards.** The legacy LSBs, plugged onto the FEBs, produced analog sums at the TT granularity (0.1×0.1). The new LSBs instead form four SCs per TT in the front and middle EM layers, and a single SC in the presampler and back layer.
- **New baseplanes.** A total of 114 new baseplanes route the much larger number of SC analog signals from the FEBs to the new digitizer boards.
- **LAr Trigger Digitizer Boards.** 124 LTDBs are installed, one per baseplane in each EMB, HEC and EMEC Standard crate, and two per baseplane in each EMEC and FCal crate. Each LTDB digitizes up to 320 SC signals at 40 MHz with 12-bit ADCs, and multiplexes groups of eight channels into 5.12 Gbps serial streams that are transmitted off-detector on VCSEL-driven optical fibers via the MiniTransmitter (MTx) modules.

The analog section of the LTDB provides gain, shaping, and offset such that the 12-bit ADC dynamic range is exploited fully. The digital section, organized around the LOCx2 serializer, the GBTx link and the GBT-SCA slow-control chip, makes sure that BCID counters are reset synchronously with the LHC clock cycle, so that each ADC sample is associated with a given bunch crossing [102].

6.2.3 Back-end electronics: the LAr Digital Processing System

The off-detector side of the DT is implemented in the LDPS, which receives the LTDB ADC streams, reconstructs the SC transverse energy and timing at 40 MHz, and distributes the results to the L1 calorimeter trigger and to the ATLAS data acquisition system. The LDPS comprises 31 LAr Digital Processing Blades (LDPBs) based on the Advanced Telecommunications Computing Architecture (ATCA), located in three ATCA shelves. Each LDPB consists of a carrier board, the LAr Carrier (LArC), equipped with up to four LATOMEs. Each LATOME is built around a high-bandwidth field-programmable gate array (FPGA) that processes data from up to 320 SCs. Across the full system the LDPS receives approximately 25 Tbps of ADC data from the LTDBs

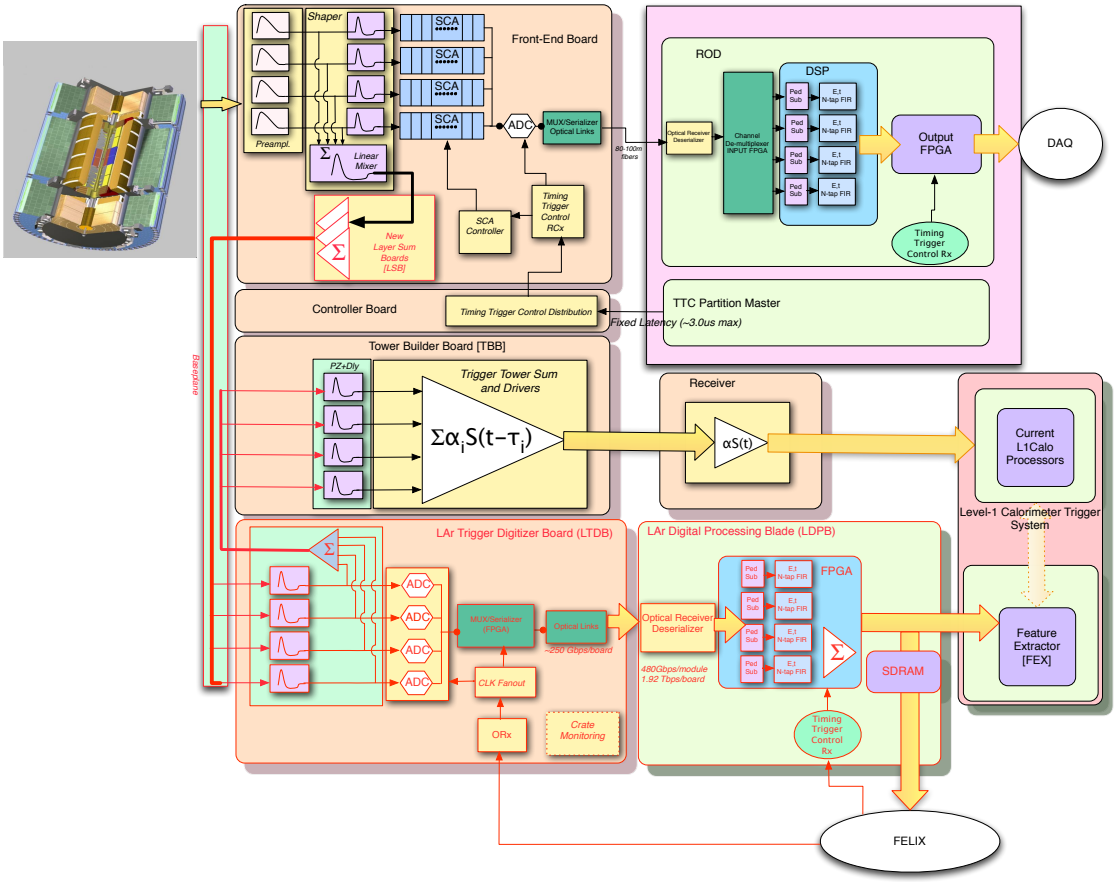

Figure 6.4: The overview of the LAr electronics architecture, with the upgraded components outlined in red [91].

and transmits approximately 41 Tbps of reconstructed SC transverse energies to the L1 calorimeter FEXs [91, 102].

Every 25 ns, the LATOME firmware applies a digital optimal filter, similarly to the main readout (see Eq. (5.7)) to the stream of ADC samples from each SC to compute the transverse energy E_{T} and the pulse peaking time t,

$$E_{\mathrm{T}} = \sum_{i=1}^{n} a_i \left(S_i - p - \delta p_{\mathrm{BCID}}\right), \qquad E_{\mathrm{T}}\, t = \sum_{i=1}^{n} b_i \left(S_i - p - \delta p_{\mathrm{BCID}}\right), \tag{6.1}$$

where S_i are the digitized ADC samples, p is the electronic pedestal, δp_{BCID} is a bunch-crossing-dependent baseline correction, and a_i, b_i are the precomputed OFCs for each SC. The baseline term δp_{BCID} is indexed by the BCID and is updated regularly to compensate for the residual offset induced by out-of-time pileup. The bipolar shaping used in the LAr front-end keeps the ideal bunch-integrated response at zero, but at Run 3 pileup levels and with a pulse extending well beyond the 25 ns bunch spacing, the average baseline of a given BCID can drift away from the electronic pedestal, as illustrated in Figure 6.5.

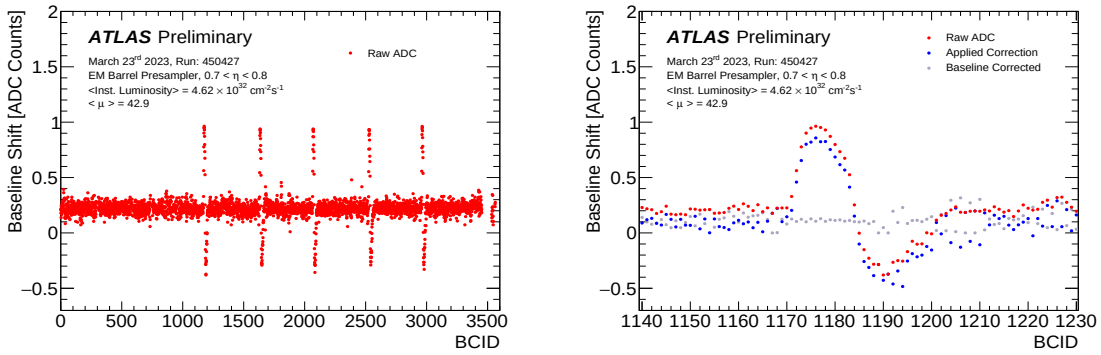

(a) (b)

Figure 6.5: Measurement of the baseline shift as a function of BCID of presampler SCs in the Electromagnetic Barrel at for the full orbit (a) and the beginning of a bunch train (b). The red markers show the values of the baseline shift, while the gray data-points represent correction applied. Finally, the blue markers represent the baseline-corrected values [104].

After OFC reconstruction, the E_{T} values are encoded on ten bits and transmitted over optical links to the three FEXs: the eFEX, which builds electromagnetic objects; the jFEX, which builds hadronic jets; and the gFEX, which computes global quantities such as missing transverse momentum for the L1 trigger decision [92].

It is worth noting that while the OFC-based E_{T} is the quantity used by the L1 calorimeter trigger, the raw ADC samples S_i used as input to Eq. (6.1) are themselves recorded for a subset of events in a special monitoring stream. These raw samples,

before optimal filtering, carry the bunch-by-bunch fluctuations of the in-time and out-of-time pileup and were the basis of the preliminary luminosity measurement method developed in this thesis (see Chapter 8) but also in the performance studies described in the Appendix A.

6.3 Performance of the Digital Trigger in Run 3

The Phase-I hardware was installed during LS2 and commissioned in parallel with the legacy analog trigger during the early Run 3 data-taking period in 2022 and 2023. In this parallel configuration, both the TT-based and the SC-based L1 trigger decisions were computed, but only the legacy path was used for physics data-taking, while the DT path was validated against the precision readout and against the legacy TT trigger on collision data [102]. After this validation period, and following the demonstration of near-full SC coverage and accurate energy and timing reconstruction, the DT became the sole operational L1 calorimeter trigger input path for ATLAS from the start of the 2024 data-taking period. The legacy analog chain has since been retired for physics use.

The DT path has been validated on proton–proton collision data collected throughout Run 3. During the high-luminosity running period of 2024, 99.6% of SCs were operational, with the remaining channels disabled or masked, demonstrating the robustness of the installed front-end and back-end electronics. Furthermore, the average pulse shapes reconstructed from SC signals in physics events agree closely with the reference pulses derived from dedicated calibration runs. These results are illustrated in Figure 6.6

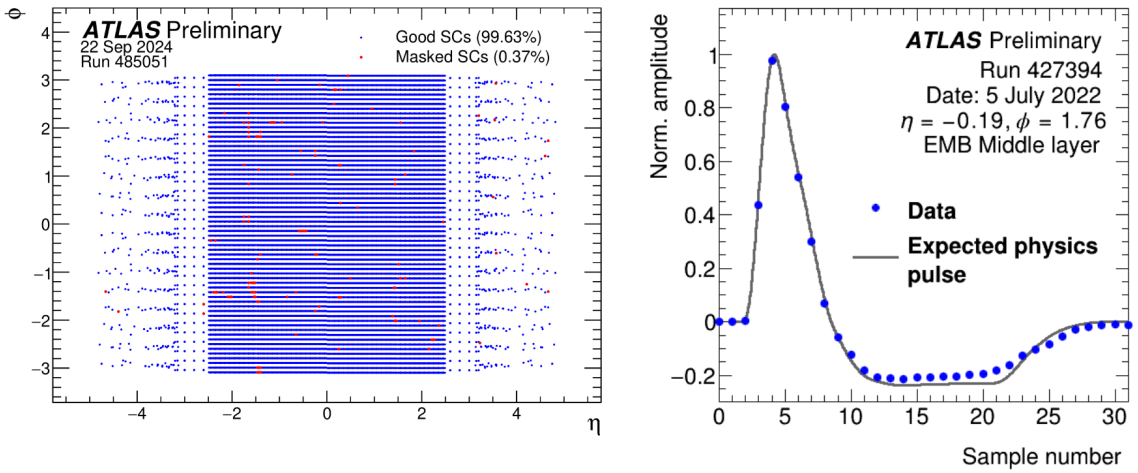

Figure 6.6: Super Cell performance in Run 3: (left) coverage map showing the fraction of operational SCs during high-luminosity running in 2024; (right) average pulse shapes reconstructed from SC signals compared with the expected pulses derived from calibration data [105].

The energy and timing reconstruction from the DT have been cross-checked against the precision (main) LAr readout on a SC-by-SC basis. The SC transverse energy shows an excellent linear correlation with the sum of the corresponding precision-cell energies, as illustrated in Figure 6.7. The reconstructed SC timing is narrowly distributed around the nominal bunch crossing, and the timing-based BCID-selection logic implemented in the LATOME firmware successfully suppresses early and late pulses, reducing the probability of assigning an SC pulse to the wrong bunch crossing.

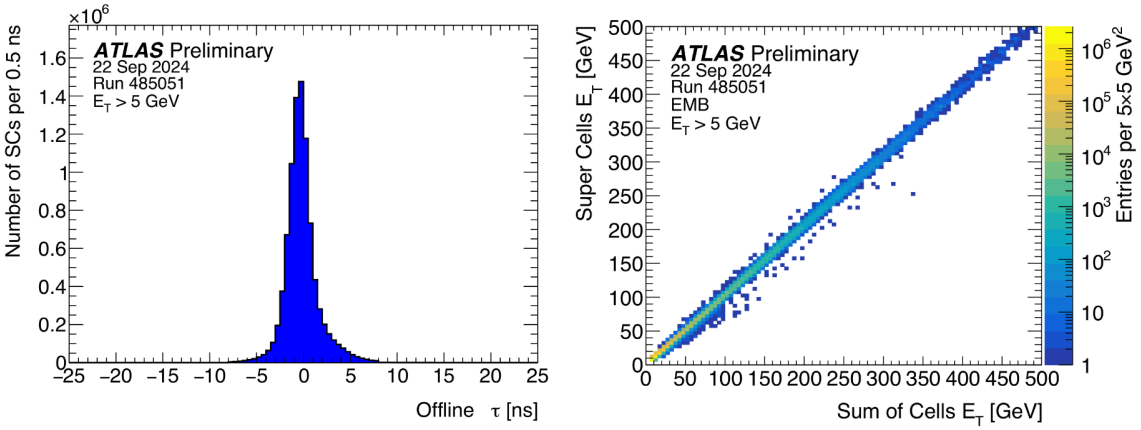

Figure 6.7: Validation of the SC energy and timing reconstruction: (left) narrow SC timing distribution centered around zero; (right) linear correlation between the SC transverse energy and the sum of the corresponding precision-cell energies from the main LAr readout [105].

When the SC information is used as input to the upgraded L1 calorimeter trigger, the improved transverse granularity and the access to layer information translate into sharper trigger turn-on curves for electrons and jets as a function of the offline transverse momentum, at fixed nominal L1 thresholds. The SC-based L1 path reduces the electron and jet trigger rates by roughly 10% relative to the legacy TT-based system, and equivalently allows the L1 thresholds to be lowered while keeping the overall L1 rate under control [106]. This improvement is visible in the electron and jet turn-on curves shown in Figure 6.8.

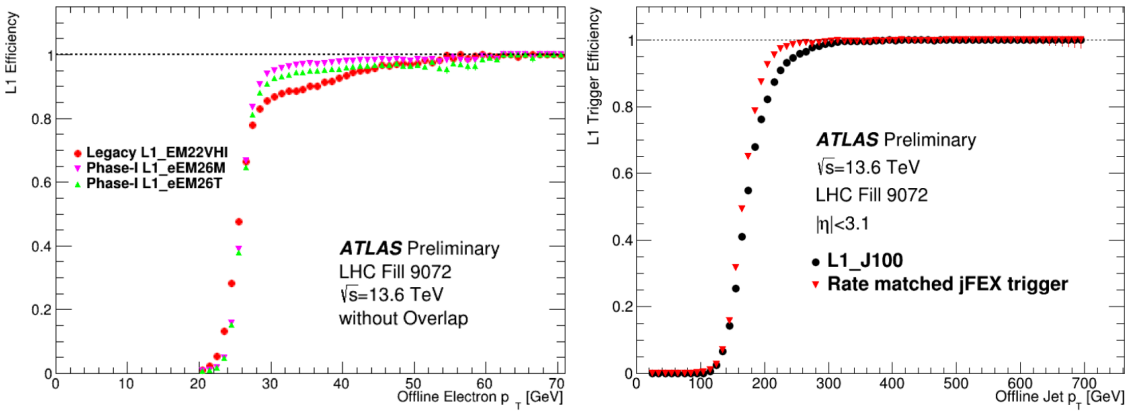

Figure 6.8: L1 calorimeter trigger turn-on curves for electrons (left) and jets (right), comparing the SC-based DT path with the legacy TT path, as a function of offline transverse momentum [106].

Part III

Luminosity Measurement

Chapter 7

Luminosity in high-energy physics experiments

Luminosity is a core concept in high-energy physics that quantifies the potential of a particle collider to produce events of interest, such as the creation of new particles or the observation of rare processes. It directly affects the rate of collisions and the likelihood of detecting significant phenomena, and is therefore paramount for every physics analysis at high-energy experiments. This chapter provides a comprehensive overview of luminosity, its importance, the factors affecting it, and its implications for high-energy physics experiments.

7.1 Definition and physical meaning

The instantaneous luminosity $\mathcal{L}$ can be expressed in terms of the LHC beam parameters (see Eq. (3.1)), but also in terms of the event rate of a given physical process and its cross section as

$$\sigma_{pp\to X} = \frac{R_{pp\to X}}{\mathcal{L}} = \frac{1}{\mathcal{L}} \cdot \frac{dN_{pp\to X}}{dt} = \frac{N_{pp\to X}}{\mathcal{L}_{\text{int}}}, \tag{7.1}$$

where the last equality applies the integrated luminosity over the relevant data-taking interval, $\sigma_{pp\to X}$ is the process cross section, and $N_{pp\to X}$ is the number of observed events, which depends on the knowledge of detector acceptances and biases, reconstruction efficiencies, and background estimation [47, 107]. The standard unit of *instantaneous* luminosity is $\text{cm}^{-2}\,\text{s}^{-1}$. The total, or *integrated*, luminosity accumulated over time is

$$\mathcal{L}_{\text{int}} = \int \mathcal{L}\, dt, \tag{7.2}$$

and is measured in units of cm^{-2}, conventionally expressed in inverse barns ($1\,\text{b}^{-1} = 10^{24}\,\text{cm}^{-2}$).

Essentially, luminosity serves as the fundamental scaling factor connecting theory and experiment. The expected event yield for a given process is

$$N_{\text{events}} = \mathcal{L}_{\text{int}} \times \sigma_{\text{process}}, \tag{7.3}$$

so that a higher integrated luminosity increases the probability of observing rare phenomena, improves the statistical precision of measurements (with statistical uncertain-

ties scaling as $1/\sqrt{N}$ for Poisson-distributed event counts), and enables stringent tests of the Standard Model. An accurate luminosity measurement is therefore paramount for the extraction of cross sections and for meaningful comparisons between measured and predicted yields.

This is especially true when the experimental techniques to measure signal rates are advanced enough, and the understanding of acceptances, detector biases, and reconstruction efficiencies reaches the sub-percent level, so that the final precision of physics measurements, i.e., of the cross section, is dominated by the luminosity uncertainty. The theoretical determination of the cross section depends on the partonic cross section and the PDFs, as described in Section 1.6; the calculation of the partonic cross section is an entirely theoretical exercise based on higher-order perturbative QCD [107]. The importance of luminosity precision is even more relevant when searching for "new physics" beyond the Standard Model, which requires a precise measurement of luminosity to determine background estimates and signal sensitivity.

A good example of a highly precise measurement is the production rate of W and Z bosons measured by ATLAS using 2011 data at $\sqrt{s} = 7\,\text{TeV}$ [108]. The measurement reaches a total experimental uncertainty of about 0.5% (not counting the 1.8% luminosity uncertainty), meaning that the main limitation comes from systematic effects rather than statistics. On the theory side, predictions are also very precise, with the dominant source of uncertainty being the knowledge of the proton's internal structure (the PDFs), at the few-percent level [107].

A more accurate luminosity measurement would allow tighter tests of the internal consistency of the Standard Model and provide stronger constraints on the PDFs. This is precisely why achieving a luminosity measurement with an uncertainty below a few percent is such an important experimental goal.

7.2 Luminosity at proton collider experiments

So far, the difference between *instantaneous* and *integrated* luminosity has been explained. Another important distinction is between *absolute* and *relative* luminosity, both essential to convert a measured signal rate for a physics process into a cross section determination. The absolute luminosity is determined on an absolute scale that is calibrated, at one point in time, via dedicated procedures such as the van der Meer (vdM) scan, described in detail in Section 7.4, and directly affects the reported cross section. The relative luminosity, measured through so-called *luminosity monitors*, validates the stability and long-term consistency of the instantaneous luminosity, which is essential to evaluate the time integral correctly [107].

7.2.1 Relative luminosity measurements

The *bunch* luminosity in a proton collider can be defined as

$$\mathcal{L}_b = \frac{R_{\text{ref}}}{\sigma_{\text{ref}}}, \tag{7.4}$$

where R_{ref} is the rate of a reference collision process and σ_{ref} is the corresponding cross section. If one chooses inelastic pp collisions as the reference, the equation becomes

$$\mathcal{L}_b = \frac{\mu f_r}{\sigma_{\text{inel}}}, \tag{7.5}$$

where μ is the mean number of inelastic interactions per bunch crossing, also known as the *pileup parameter*, f_r is the revolution frequency, and σ_{inel} is the inelastic pp cross section. Using this, the total instantaneous luminosity is [109]

$$\mathcal{L}_{\text{inst}} = \sum_{b=1}^{n_b} \mathcal{L}_b = n_b \langle \mathcal{L}_b \rangle = n_b \frac{\langle \mu \rangle f_r}{\sigma_{\text{inel}}}. \tag{7.6}$$

Thus, the total instantaneous luminosity can be determined by any method that measures μ. The cross section σ_{inel} is a constant that merely sets the scale, so it cancels out in relative luminosity measurements. These methods, or *luminometers*, can be classified as [107]:

- *Event counting:* uses the fraction of bunch crossings in which a given detector registers an "event" satisfying a given selection requirement. For instance, a bunch crossing can be said to contain an "event" if at least one pp interaction in that crossing induces at least one observed hit in the detector being considered.

- *Hit counting:* the number of hits per bunch crossing (for example, the number of electronic channels or energy clusters above a specified threshold) in a given detector is used as the proxy.

- *Particle counting:* uses the distribution of the number of particles per bunch crossing (or its mean), obtained from reconstructed quantities such as calorimeter energy, or from other observables proportional to the instantaneous particle flux traversing the detector (e.g. the ionization current drawn by a LAr calorimeter HV line, or the signal-charge measurement performed by LUCID-2 (LUminosity measurement using Cherenkov Integrating Detector), as described in Section 7.3).

These luminometers measure a fraction of the inelastic interactions per bunch crossing, so Eq. (7.5) can be rewritten as

$$\mathcal{L}_b = \frac{\epsilon \mu f_r}{\sigma_{\text{inel}}} = \frac{\mu_{\text{vis}} f_r}{\sigma_{\text{vis}}}, \tag{7.7}$$

where ϵ is the efficiency of detecting one inelastic collision in the luminometer, and $\mu_{\text{vis}} = \epsilon\mu$ is the *visible interaction rate*, i.e. the average number of collisions per bunch crossing detected by the luminosity monitor. The *visible cross section* σ_{vis} is the calibration constant, specific to each detector and algorithm, that relates the measured interaction rate to the absolute bunch luminosity $\mathcal{L}_b$. Determining μ_{vis} for a given detector constitutes the relative luminosity measurement; it is complementary to the absolute calibration of σ_{vis}, as described in the next section.

7.2.2 Absolute luminosity measurement

One further categorization is the absolute *specific luminosity* [107], expressed as the luminosity *per bunch* and per unit bunch intensity, which characterizes the fundamental performance of the colliding bunches in the LHC, as described briefly in Section 3.2. With bunches colliding every 25 ns at nominal Run 3 pileup of $\langle\mu\rangle \approx 60$, each bunch crossing produces on average several tens of inelastic interactions.

The algorithm-specific visible cross section σ_{vis} introduced in Eq. (7.7) is a calibration constant that sets the absolute luminosity scale for a given algorithm. It is determined using the van der Meer beam-separation scan [109] calibration method.

Once σ_{vis} is known, the pileup parameter μ can be calculated as

$$\mu = \frac{\mu_{\text{vis}}\,\sigma_{\text{vis}}}{\sigma_{\text{inel}}}. \tag{7.8}$$

Because σ_{inel} is not precisely known, ATLAS uses a standard reference value of $\sigma_{\text{inel}} = 80\,\text{mb}$ for pp collisions at $\sqrt{s} = 13.6\,\text{TeV}$.

The parameter μ and its average value $\langle\mu\rangle$ are important for characterizing the performance of the LHC and its detectors. However, the luminosity calibration based on σ_{vis} does not depend on the chosen reference value of σ_{inel}.

If σ_{vis} remains stable over time and does not depend on LHC conditions (such as the number of bunches or the value of $\langle\mu\rangle$), then measuring μ_{vis} at any moment is sufficient to determine the instantaneous luminosity per bunch.

In reality, no luminosity algorithm is perfectly stable or fully independent of LHC conditions. Comparing multiple algorithms and detectors is therefore necessary to obtain a precise luminosity measurement and a reliable estimate of its uncertainty [109].

vdM programs are organized during LHC operation and are the main technique to set the absolute luminosity scale: from the measurable parameters of the colliding bunches at the peak of the scan, the visible cross section of inelastic pp collisions can be calculated using Eq. (7.5). After determining the absolute scale from these observables in low-μ vdM running conditions, a *calibration transfer* procedure is used to extrapolate the calibration to nominal, high-μ conditions.

7.3 Luminosity measurements in ATLAS

ATLAS employs several subdetectors and methods to accurately measure luminosity throughout its data-taking campaigns. This section describes the various approaches, from absolute luminosity measurements using vdM scans and the primary ATLAS luminometers LUCID-2 [110] and the Beam Conditions Monitor (BCM) [111], which provide the per-bunch online and offline measurements, to the other subdetectors used as relative luminosity monitors, namely the Inner Detector track-counting algorithms (Pixel and SCT) and the LAr and Tile calorimeter systems [109]. These complementary measurements are essential for validation cross-checks, for monitoring the long-term stability of the luminosity, and for the determination of the uncertainty on the total integrated luminosity. The most recent ATLAS luminosity results are summarized at the end of the chapter.

7.3.1 LUCID-2 and BCM

The LUCID-2 Cherenkov detector provides the official ATLAS online luminosity. It is composed of 16 PMTs located on both forward arms of the ATLAS detector (sides A and C), around the beam pipe. Its main mechanism uses Cherenkov light produced in the quartz windows of the PMTs. The PMT windows are coated with ^{207}Bi sources, used to provide calibration signals, as illustrated in Figure 7.1. Calibrations are performed between LHC fills, allowing the high voltage applied to the PMTs to be adjusted over time to compensate for aging effects.

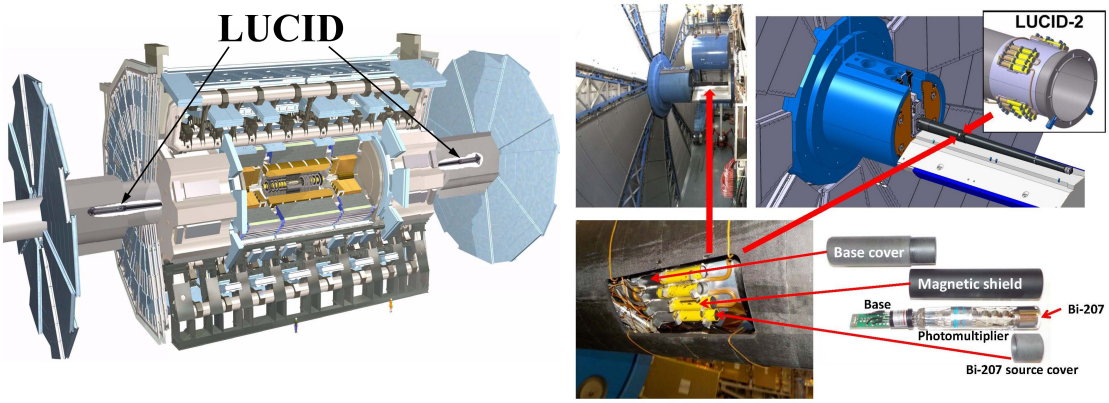

Figure 7.1: LUCID-2 [110] in the ATLAS detector (left) and a detail of the system's PMTs and ^{207}Bi sources used for calibration (right) [112].

The BCM detector [111] consists of two groups of four diamond sensors, arranged in a cross pattern and placed symmetrically on both sides of the interaction point, at $z = \pm 184\,\text{cm}$ and $R = 5.5\,\text{cm}$. Its main role is to monitor particle fluxes close to the beam pipe, protecting the innermost ATLAS components and triggering a beam abort

if radiation levels become too high. Thanks to its fast readout electronics, the BCM can also serve as a luminosity detector.

The simplest LUCID-2 algorithm uses a single PMT and counts an "event" whenever the PMT signal exceeds a given threshold (a "hit"), corresponding to one or more inelastic pp interactions detected in a given bunch crossing (an "event-counting" luminometer, see Section 7.2.1). Assuming the number of inelastic pp interactions per bunch crossing follows a Poisson distribution, the probability of recording no event, $P_0 = e^{-\mu_{\text{vis}}}$, is related to the single-PMT visible interaction rate $\mu_{\text{vis}} = \epsilon\mu$ by [109]

$$P_{\text{evt}} = \frac{N_{\text{evt}}}{N_{\text{BC}}} = 1 - e^{-\mu_{\text{vis}}}, \tag{7.9}$$

where N_{evt} is the number of recorded events and N_{BC} is the total number of bunch crossings. Inverting this gives

$$\mu_{\text{vis}} = -\ln\left(1 - P_{\text{evt}}\right) = -\ln\left(1 - \frac{N_{\text{evt}}}{N_{\text{BC}}}\right). \tag{7.10}$$

In a higher-pileup environment, the probability of no signal being recorded decreases drastically, reaching *saturation* when $N_{\text{evt}}/N_{\text{BC}} \approx 1$. To overcome this limit, hit-counting algorithms can be used, where the total number of hits across all channels is counted. Equation (7.10) can then be rewritten as

$$\mu_{\text{vis}} = -\ln\left(1 - \frac{N_{\text{hit}}}{N_{\text{BC}}\, N_{\text{ch}}}\right), \tag{7.11}$$

where N_{hit} is the number of hits per bunch crossing and N_{ch} is the number of detector readout channels [113].

LUCID and BCM both support event-counting and hit-counting luminosity measurements through several algorithms, each defined by a different counting condition. In the *EventOR* algorithm, a bunch crossing is recorded as an event if at least one hit is detected on either the A side or the C side of the detector. In the *EventAND* algorithm, a bunch crossing is counted only if at least one hit is seen on both sides simultaneously. The coincidence condition can be satisfied either by a single pp interaction or by separate hits on each side coming from different interactions within the same bunch crossing, which makes the expression for μ_{vis} more involved.

For each data-taking year, the most stable algorithm is selected as the nominal one, while the remaining algorithms serve as cross-checks and help evaluate systematic uncertainties. As shown in Table 7.1, the nominal LUCID algorithms varied across Run 2: in 2015 and 2018 a single PMT on the C side (PMT C9 and PMT C12, respectively) was used with event counting, while in 2016 and 2017 the BiHitOR algorithm relied on hit counting from bismuth-calibrated PMTs on both the A and C sides of LUCID.

Equations (7.10) and (7.11) are inherently non-linear in N_{evt} and N_{hit}, so it is not

possible to simply sum the rate measurements over different BCIDs and then obtain μ_{vis}, unless μ_{vis} is the same for all BCIDs. At the LHC the bunch intensity can vary by 10–20% from bunch to bunch [107], making it paramount for ATLAS to maintain bunch-by-bunch measurements, provided first by LUCID and BCM, and potentially by new methods such as the one developed in this thesis.

Table 7.1: Nominal LUCID algorithms used for the luminosity measurement in Run 2, the corresponding visible cross section σ_{vis} measured by the vdM calibration, the peak $\langle\mu\rangle$, and the fraction of bunch crossings $f_{\text{no-hit}}$ where no hit was recorded at the respective μ [109].

Year	Algorithm	σ_{vis} [mb]	Peak $\langle\mu\rangle$	$f_{\text{no-hit}}$ [%]
2015	PMT C9	6.540	16	27.0
2016	BiHitOR	6.525	41	3.5
2017	BiHitOR	6.706	60	0.6
2018	PMT C12	6.860	55	0.9

7.3.2 Luminometers in ATLAS

The main ATLAS relative luminosity monitors are based on track-counting algorithms, using data from the silicon detectors (the SCT and Pixel detectors), and on calorimeter-based algorithms, exploiting the gap currents in the LAr calorimeter and the PMT currents in the Tile calorimeter.

The track-counting algorithms measure the visible interaction rate μ_{vis} per bunch from the mean number of reconstructed tracks per bunch crossing over a luminosity block, using a dedicated reconstruction stream with tracks satisfying the "TightPrimary" quality requirements outlined in Ref. [114]. The precision of this measurement is limited at the very low pileup of vdM scans, making it difficult to calibrate via the vdM scans themselves. Instead, the calibration against LUCID is performed in the same LHC fills as the vdM scans, using periods of stable head-on collisions in ATLAS (typically when the vdM scans are being executed in CMS).

The calorimeter-based luminosity measurements use data from the ATLAS calorimeters. In the case of LAr, the measurement is based on the EMEC and FCal HV system: the current drawn through the HV power supplies, produced by traversing charged particles ionizing the liquid argon, is proportional to the instantaneous luminosity. Due to its dependence on slow ionization currents, this method cannot resolve individual bunches and is therefore limited to bunch-integrated measurements. Furthermore, its sensitivity is not sufficient at the low luminosities of vdM running for direct calibration, so track-counting measurements are used to cross-calibrate it in high-luminosity fills [109]. The methods developed in this thesis, also based on the LAr calorimeter, aim to overcome these limitations: the sample-based methods of Chapters 8 and 9 can

resolve individual bunches, and the pulse-rate method of Chapter 10 retains sensitivity down to vdM luminosities.

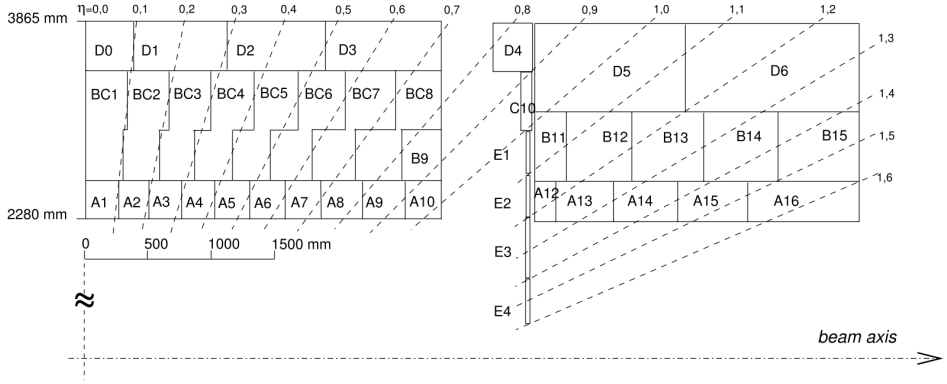

Figure 7.2: A schematic diagram of the cells in the Tile calorimeter. The D and E cells are the ones used for luminosity measurements [115].

The Tile calorimeter, using steel as absorber and plastic scintillators as active medium, reads out signals with wavelength-shifting fibers connected to PMTs. The current induced in the PMTs is proportional to the number of particles passing through the corresponding TileCal cell, and therefore to the instantaneous luminosity. The luminosity measurement is performed with the Tile D-cells (in particular D6, the closest to the interaction point, used for the Run 2 luminosity measurement [109]), which however have insufficient sensitivity to be calibrated directly with vdM scans, similarly to the LAr EMEC and FCal gap-current method. The Tile E-cells (in particular E3 and E4, the closest to the beamline, as shown in Figure 7.2) are instead sensitive enough at the vdM scale, making them well suited to constrain calibration-transfer uncertainties, but they are not stable enough for long-term stability studies.

7.4 Van der Meer scans and calibration transfer

The calibration of the absolute luminosity in ATLAS is performed with vdM beam-separation scans [107, 109] in special LHC fills, normally at very low pileup. The scan is performed in discrete steps in the x and y directions, and for each step the value of μ_{vis} is measured. The visible cross section σ_{vis} is then determined from the resulting vdM scan curve. An example of such a scan curve from 2017 is shown in Figure 7.3.

At zero crossing angle, the bunch instantaneous luminosity can be expressed as

$$\mathcal{L}_b = f_r \, n_1 n_2 \int \rho_1(x,y) \, \rho_2(x,y) \, dx \, dy, \tag{7.12}$$

where f_r is the revolution frequency, n_i is the population of bunch i, and $\rho_i(x,y)$ is the normalized particle density in the transverse plane. The proton densities can be

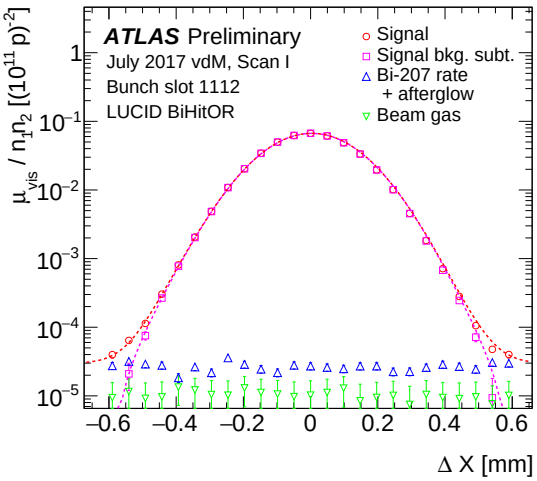

Figure 7.3: The visible interaction rate μ_{vis} per unit bunch population product $n_1 n_2$ is shown as a function of the horizontal beam separation Δx, as measured by the LUCID BiHitOR algorithm during the July 2017 vdM session. Each point corresponds to a discrete step; at minimal separation μ_{vis} reaches its maximum because the beams are completely overlapping [109].

factorized into horizontal and vertical parts,

$$\rho(x, y) = \rho_x(x)\,\rho_y(y), \tag{7.13}$$

so that Eq. (7.12) reduces to

$$\mathcal{L}_b = \frac{f_r\, n_1 n_2}{2\pi\, \Sigma_x\, \Sigma_y}, \tag{7.14}$$

where Σ_x and Σ_y are the effective widths of the beam overlap region, measured from the scan curve.

The calibration of a given detector or algorithm is then performed by equating Eqs. (7.7) and (7.14), leading to

$$\sigma_{\text{vis}} = \mu_{\text{vis}}^{\text{max}}\, \frac{2\pi\, \Sigma_x\, \Sigma_y}{n_1 n_2}, \tag{7.15}$$

where $\mu_{\text{vis}}^{\text{max}}$ is the visible interaction rate corresponding to the maximum of the scan curve.

The determination of σ_{vis} via the vdM absolute calibration requires bunch-by-bunch capability and sufficient sensitivity at low pileup; in ATLAS this restricts the direct vdM calibration to LUCID and BCM. The track-counting algorithm, although bunch-resolved, is cross-calibrated to LUCID in the same fills (Section 7.3.2), and all other luminometers are calibrated by normalization to the LUCID measurement.

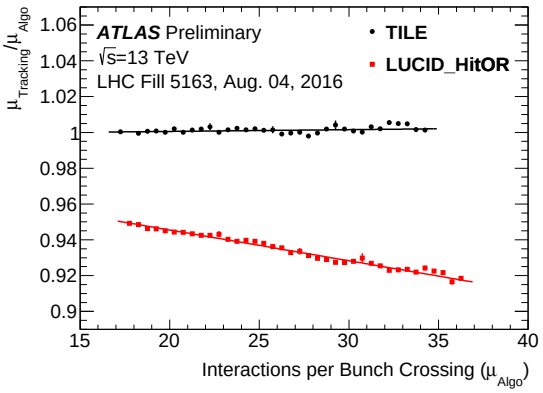

Figure 7.4: Ratio of the instantaneous luminosities from track counting and LUCID (in red), and from track counting and TileCal luminosity (in black), as a function of $\langle\mu\rangle$ [109].

Once this is done, the LUCID calibration determined at the low pileup of vdM scans must be *transferred* to high-μ data-taking conditions. Due to the non-linearity of LUCID with μ, a correction of $\mathcal{O}(10\%)$ is required at nominal high-μ running conditions [107]. The non-linearity is determined by comparison with the track-counting algorithm in a long nominal physics run with a gradual decrease in $\langle\mu\rangle$ over time. This is illustrated in Figure 7.4, where the non-linearity is evident in the ratio between track-counting and LUCID measurements, in contrast to the same ratio for TileCal-based luminosity. The dependence is parameterized as a polynomial in $\langle\mu\rangle$,

$$\langle\mu_{\text{corr}}\rangle = p_0 \langle\mu_{\text{uncorr}}\rangle + p_1 \langle\mu_{\text{uncorr}}\rangle^2, \tag{7.16}$$

and the resulting parameters are used to correct the LUCID luminosity in subsequent runs.

7.5 ATLAS Run 2 luminosity results

Once the absolute LUCID luminosity is determined by the vdM scan and converted to high-μ conditions through the calibration transfer, it is essential to quantify possible divergences over time. To do so, μ-corrected LUCID integrated luminosity measurements are compared run by run with the measurements from other luminometers, normalized to LUCID in a reference run. An example for the year 2018 is shown in Figure 7.5, where various detectors are normalized to LUCID and the ratio per run is plotted as a function of time.

The overall goal is to compute the final luminosity value for each data-taking period, which is then used as a systematic uncertainty input in physics analyses. Most analyses use the full Run 2 dataset, so a precise determination of the integrated luminosity and

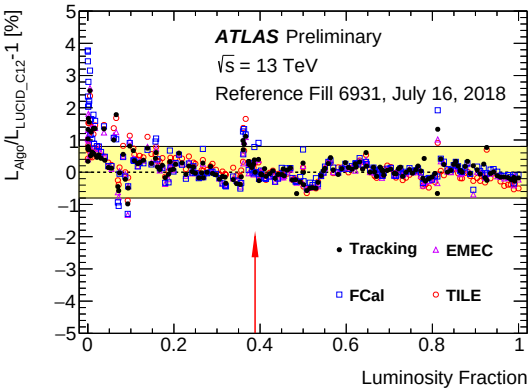

Figure 7.5: Fractional differences in run-integrated luminosity between LUCID and the track-counting, Tile, EMEC and FCal measurements, as a function of the cumulative delivered luminosity normalized to the 2018 total. The luminosity measurements from the other detectors have been normalized to that of LUCID in the reference run indicated by the arrow [109].

its uncertainty for the four years of Run 2 is needed. The total integrated luminosity measured for the full Run 2 amounts to $140.1(12)\,\text{fb}^{-1}$, corresponding to a relative uncertainty of 0.83% [109]. This represents the most precise luminosity measurement at a hadron collider to date, improving on the earlier preliminary result by a factor of two. A more detailed breakdown is shown in Table 7.2.

Table 7.2: Extract of integrated luminosity and systematic uncertainty results for the Run 2 data-taking periods. Full table can be found in [109]. Contributions marked * are fully correlated between years.

Data sample	**2015**	**2016**	**2017**	**2018**	**Comb.**
Integrated luminosity (fb^{-1})	3.24	33.40	44.63	58.79	140.07
Total uncertainty (fb^{-1})	0.04	0.30	0.50	0.64	1.17
Uncertainty contributions (%):					
Subtotal vdM calibration	0.96	0.70	0.99	0.93	0.65
Calibration transfer*	0.50	0.50	0.50	0.50	0.50
Calibration anchoring	0.22	0.18	0.14	0.26	0.13
Long-term stability	0.23	0.12	0.16	0.12	0.08
Total uncertainty (%)	1.13	0.89	1.13	1.10	0.83

Chapter 8

Luminosity measurement with the Digital Trigger

In the following chapters, the main work of this thesis is presented. First, a feasibility study of the novel technique to measure the luminosity with the Digital Trigger is described, in which its preliminary performance and systematic effects are analyzed, and its performance is then compared with the current ATLAS luminosity methods. Then, a study with official ATLAS physics data is presented, and approved results are shown. Finally, the possibility of bunch-by-bunch measurements and the use of the LATOME Pulse Rate Monitor as a luminometer are explored.

8.1 Introduction to the methods and datasets

The first approach to use the Digital Trigger as a luminometer was performed using the RMS of the ADC samples recorded at every bunch crossing, each 25 ns, by the Digital Trigger. This work is described in detail in Section 8.2, where a private stream of data from LAr called the MON data (the processing and analysis of these data was part of the work performed for the LAr Data Quality, see Section A.3) was used. These data were chosen for a feasibility study due to the access to both raw and baseline-corrected samples, and between 6 and 32 samples per channel and per event[1]. This is far more information than is available in the standard physics data streams. However, due to its high prescale factor (only 1 out of every 60,000 events in the stream is saved), the dataset could not be used for more precise and official estimates. The low statistics did, however, keep the dataset compact and easily storable and accessible, unlike the official ATLAS physics data, where a single run can occupy terabytes of disk space.

When the project was mature enough, the study was moved to datasets from official ATLAS data streams. In this thesis, these are often called "SWROD data", as they are processed through the SWROD, as described in Section 4.6. There, only targeted runs could be studied, due to the substantial data volume per run, but with the high statistics available per run, the method could be probed at its full potential. It is worth noting that, unlike the MON stream, the standard ATLAS physics data streams carry only baseline-corrected samples rather than raw ADC samples. This distinction is inconsequential for the RMS method, which relies on the dispersion of those samples, but becomes critical

[1]The MON configuration used throughout this analysis stores 11 samples per channel and per event; the per-LB processing assumes this value.

for the bunch-by-bunch study described in Chapter 9. Four streams were used in this part of the project:

- `physics_Main`: the primary general-purpose physics stream, containing events accepted across the bulk of the ATLAS trigger menu. It carries the highest recorded event rate and the largest per-run statistics.
- `physics_L1Calo`: a stream collecting events selected by L1 calorimeter-based triggers, used primarily for calorimeter trigger performance monitoring and commissioning studies.
- `physics_ZeroBias`: a stream of events recorded one full LHC orbit after a high-energy electron or photon has been accepted by the L1 trigger, without any physics requirement imposed on the recorded event itself.
- `calibration_LArPEBDigitalTrigger`: a LAr-specific calibration stream using Partial Event Building (PEB), which records only the LAr Digital Trigger data fragment for selected events, keeping the event size small. This is also the SWROD monitoring stream referenced later in this section, and it is referred to as `SWROD_PEB` in the following.

The `physics_Main` stream was found to yield the best resolution to luminosity, and for this reason it was used to produce ATLAS public plots [116], published in a peer-reviewed proceeding [117]. An important effect of the choice of a stream is the related trigger bias, which is discussed in Section 8.4.1.

Apart from the RMS, another observable was used to measure the luminosity bunch by bunch, namely the mean pedestal ADC value of samples as a function of BCID, as outlined in Chapter 9. Here, the MON data proved to be particularly valuable, as the only information on the effect of the bunch structure is stored in the raw samples (the goal of the baseline correction is to cancel this effect, making the baseline-corrected samples ineffective in this context). In May 2026, the readout recipe of LAr was changed and set to store raw data in the SWROD monitoring stream (the `SWROD_PEB` stream). This resulted in raw samples being available in SWROD data for one nominal physics pp run (run 520264, fill 11712, May 8 2026), providing highly useful additional data for the bunch-by-bunch luminosity study.

The final part of the project explores the possibility of the LATOME Pulse Rate Monitor (PRM) as a luminosity detector, outlined in Chapter 10. This is a function implemented in the LATOME firmware that computes the rate of pulses with energy exceeding a configurable threshold value per SC. This observable is mainly used to monitor the activity and performance of the SCs in real time. The data are archived in the TDAQ monitoring data archive by the Persistent Back-End for the ATLAS Information System of TDAQ (P-BEAST) [118]. The same system was used to retrieve the archived data for processing.

8.1.1 Integration with ATLAS luminosity framework

As described in Section 7.3, ATLAS uses LUCID as its primary luminosity detector, alongside other tracking-based and calorimeter-based luminosity methods, such as the LAr gap currents drawn by the EMEC and FCal and the photomultiplier currents of TileCal. These calorimeter-based methods provide only bunch-integrated luminosity and cannot resolve individual bunch crossings. New relative luminosity methods are increasingly in demand, both to complement LUCID and to provide per-bunch luminosity measurements and algorithms suitable for calibration transfer, given the challenging pileup conditions expected during HL-LHC operation. As described in the following sections, the LAr Digital Trigger has the potential to fulfill all of these roles.

A relative luminosity measurement typically relies on fast counting devices that provide a proportional signal to luminosity. The key requirements for such a monitor are [47]:

- Must cover a large dynamic range: $10^{27}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ to $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ [47]. This range already needs to be extended in practice: ATLAS Run 3 instantaneous luminosities exceed $2 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, and the HL-LHC is expected to reach up to $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, with μ up to 200 [93].
- Must operate at high speed, ideally resolving individual bunches.
- Must be reproducible from run to run.
- Should operate under varying machine conditions (e.g. with and without crossing angle, different optics (β^*), etc.).
- Must be compatible with different beam species (protons, ions).
- Can be integrated into a feedback system to optimize luminosity.

When these requirements are met, the instrument provides a reliable signal proportional to the instantaneous luminosity. A central part of the feasibility study presented in Section 8.2 is to demonstrate that the DT RMS method satisfies these criteria as a relative luminosity monitor, as detailed in the following sections.

In this context, the DT is an ideal candidate: the RMS of its samples is straightforward to compute, exhibits good stability over time and linearity with luminosity. Furthermore, because the DT provides data at every bunch crossing, and BCID information is stored every 25 ns, this method has the potential to deliver bunch-by-bunch luminosity measurements, enabling a direct measurement of $\mathcal{L}_b$ for each colliding bunch, as will be described in Chapter 9.

In addition, the LATOME PRM may achieve the sensitivity required at vdM scales, making the DT a valuable asset for luminosity calibration transfers. Together, these capabilities position the LAr Digital Trigger as one of the most versatile systems in the

ATLAS luminosity infrastructure, with the potential to serve as a relative luminometer with bunch-by-bunch capability.

8.2 The RMS of Digital Trigger samples as a luminosity measurement method

Since the mean number of interactions per bunch crossing is directly proportional to the per-bunch luminosity, any observable that scales with μ can serve as a luminosity estimator, as described in Section 7.2.1. The dispersion of Digital Trigger ADC samples is such an observable. This dispersion is evaluated by taking the RMS of the samples available per SC for each triggered event, and averaging them over the whole detector for each luminosity block. As pileup increases, additional energy deposits from out-of-time and in-time collisions broaden the sample distribution, leading to a larger RMS that grows with μ and hence with the luminosity. This is evident in Figure 8.1, where simply plotting this observable as a function of luminosity block reveals a clear similarity to the ATLAS online luminosity profile.

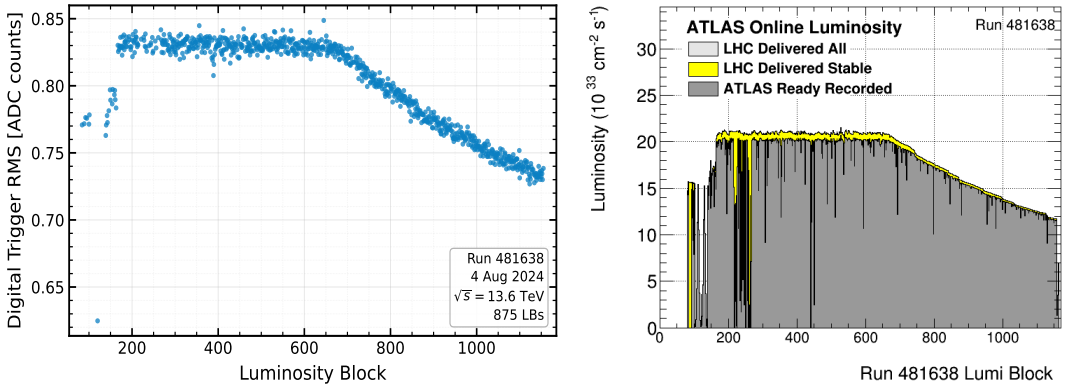

(a) (b)

Figure 8.1: (a) The RMS of the Digital Trigger samples of run 481638 as a function of luminosity block, and (b) the corresponding ATLAS online luminosity profile per LB. A strong similarity between the two can be seen directly.

8.2.1 Theoretical foundation

The correlation observed in Figure 8.1 can be traced back to the pileup noise introduced in Section 5.7. The bipolar CR–RC^2 shaping applied by the LAr front-end electronics integrates to approximately zero in time, so when soft interactions are overlaid uniformly along a bunch train, the positive lobe of each in-time deposit is canceled, on average, by the negative lobes of out-of-time deposits from neighboring crossings.

The mean pileup contribution to a sample vanishes [75], its fluctuations, however, do not. The pileup contribution to a given SC is the superposition of the deposits of the soft interactions taking place in the bunch crossings to which that SC is sensitive, weighted by the shaped pulse amplitude at each crossing, as expressed by Eq. (5.8) [75]. The number of such interactions grows with the mean number of inelastic interactions per bunch crossing μ, which is itself proportional to the bunch luminosity through Eq. (7.5). The dispersion of the samples therefore increases with the luminosity.

The observable exploited in this chapter is the RMS of the raw ADC samples, which preserves this pileup-driven dispersion before any optimal filter or baseline correction is applied. Two contributions determine its value: the electronic noise of the readout chain, which is independent of the beam conditions and sets a floor on the measured dispersion, and the pileup term, which grows with μ and therefore with the luminosity. At luminosities low enough that the pileup contribution is small compared with this electronic noise floor, the measured RMS is dominated by the electronic noise and carries little information on the beam conditions. As the luminosity increases, the pileup term progressively dominates and the RMS becomes an increasingly sensitive probe of the instantaneous luminosity.

The relation between the sample RMS and the luminosity is established empirically, run by run, by fitting the measured RMS against the reference ATLAS online luminosity. Over the luminosity range spanned within a single fill, the correlation is found to be well described by a simple linear relation,

$$\mathrm{RMS} \;=\; k\,\mathcal{L} + \mathrm{RMS}_0, \tag{8.1}$$

where the slope k quantifies the sensitivity of the observable to the luminosity and the intercept RMS_0 absorbs the luminosity-independent contribution, and is therefore expected to sit close to the electronic noise floor; this expectation is verified directly in Section 8.3.5 using empty bunch crossings. The quality of this linear description is demonstrated for the full detector in Section 8.3, and per layer and per "scieta"[2] in Section 8.3.4. Inverting Eq. (8.1) gives the luminosity estimate associated with a measured sample RMS,

$$\hat{\mathcal{L}} \;=\; \frac{\mathrm{RMS} - \mathrm{RMS}_0}{k}, \tag{8.2}$$

which is the estimator used throughout this chapter, and which takes the same form per channel i in Eq. (8.3) when the per-layer and per-scieta estimates are combined in Section 8.3.4.

Towards low luminosity, where the electronic noise dominates the dispersion, the response flattens and the sensitivity of the method is progressively lost; this is observed both as a degraded performance at vdM luminosity scales (Section 8.3.10). Conversely,

[2]Scieta, or Super Cell index η, is the integer index that labels a Super Cell's position along the pseudorapidity (η) direction within a given calorimeter layer and region.

at high μ the pileup term dominates and the method performs best, which makes it naturally suited to the conditions foreseen for the HL-LHC.

8.3 Feasibility study

As established in Section 8.2.1, the RMS of the samples is a proxy for the pileup, and therefore for the luminosity. In order to demonstrate its capabilities as a relative luminometer, it is compared against the ATLAS online luminosity.

The procedure is performed as following:

1. Compute the RMS of the samples and average it per luminosity block (LB), as shown previously in Figure 8.1.

2. Plot it against the ATLAS online luminosity to establish a correlation, fit the resulting distribution, and extract the residuals with respect to the fit, as done in Figure 8.2.

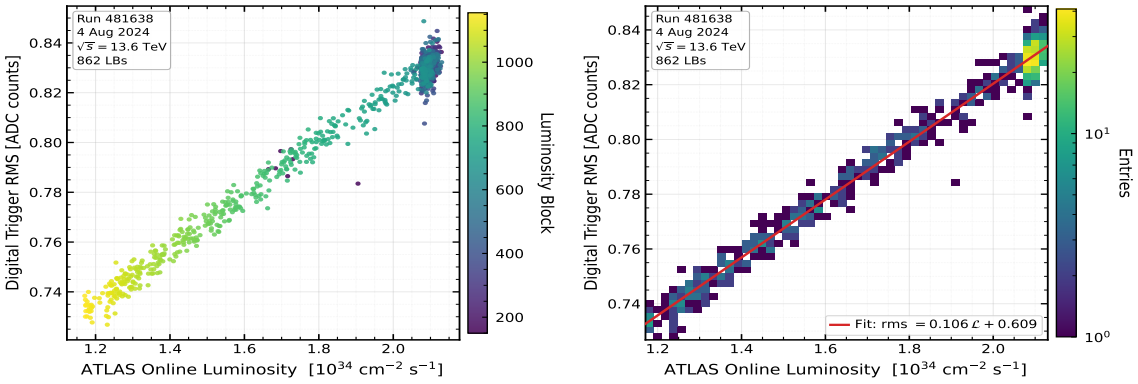

(a) (b)

Figure 8.2: (a) The ATLAS online luminosity per LB as a function of the Digital Trigger sample RMS for run 481638, with the color scale indicating the luminosity block. (b) The same correlation shown as a two-dimensional histogram with the linear fit overlaid. The first 100 LB are excluded from the fit (see text).

3. Fit the residual distribution with a Gaussian and extract its parameters, as seen in Figure 8.3.

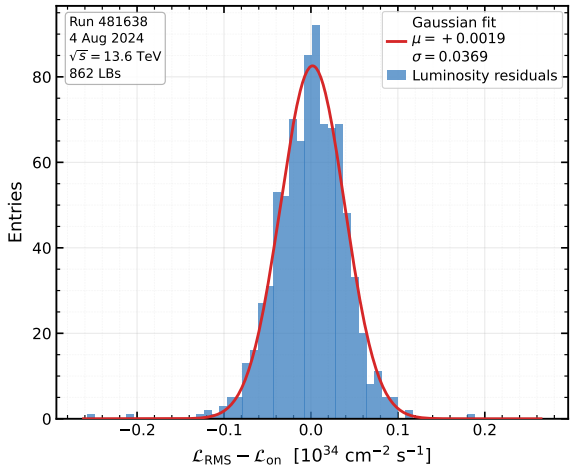

Figure 8.3: The distribution of the absolute residuals of the RMS-to-online-luminosity fit for run 481638, with a Gaussian fit overlaid. The fitted mean and width in luminosity units ($10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) are given in the legend.

4. Use the width of the Gaussian fit as the luminosity resolution estimator, normalized by the online luminosity to obtain a relative estimate of the precision. This is done in Figure 8.4.

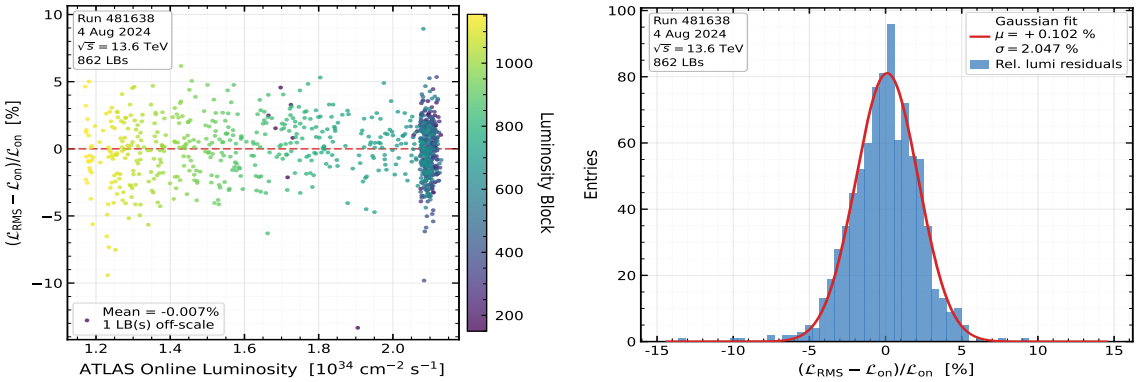

(a) (b)

Figure 8.4: (a) The relative residuals of run 481638 as a function of the Digital Trigger RMS, and (b) their distribution with a Gaussian fit overlaid. The fitted width corresponds to a resolution on the luminosity of approximately 2%.

The relative resolution obtained with this method is a quantity mainly driven by the long-term stability and linearity of the luminosity measurement in ATLAS, which is one of the contributions to the uncertainty described in Section 7.5, and Table 7.2.

It can be seen from the previous figures that a preliminary 2% resolution is already achieved in the relative measurement with respect to the online luminosity, as evidenced

in Figure 8.4b. As noted in Figure 8.2, the first 100 LBs of the run are excluded from the fit. The reason is understood by looking at the official ATLAS online luminosity profile in Figure 8.1: those LBs include zero luminosity and the beginning of a short emittance scan (see Section 8.3.9). Repeating the same exercise without excluding those LBs degrades the result. This suggests that there is a range of systematic effects that need to be analyzed, in order to establish the range over which the RMS remains a reliable proxy for the luminosity. These effects can be studied in terms of layer or subdetector, bunch structure, accelerator conditions, and luminosity scale. The next sections explore the extent of these effects in order to constrain the capabilities of the DT RMS as a luminometer.

8.3.1 Data quality cuts

The MON data used for this analysis are very low-level data at Super Cell scale, used to spot channels with defects, dead readout and general issues in the performance of the Digital Trigger. By design, therefore, no data-quality selection is applied to the output.

In the context of this luminosity analysis, data-quality cuts are necessary to measure the luminosity as efficiently as possible. These cuts are applied on several levels:

- **Super Cell level**: masked and dead Super Cells are filtered out.
- **Event level**: entries with invalid samples are rejected. These are flagged with an ADC count of -1 in the raw samples, which corresponds to -0.125 in the baseline-corrected scale, the two values being related by the factor of 8 between the two scales.
- **Luminosity block level**: LBs without the *Stable Beams* flag, LBs containing emittance scans, and LBs with zero instantaneous luminosity are removed from the main analysis, and LBs shorter than 10 s.

These cuts are applied for the main luminosity analysis, but are dropped for specific studies such as the effects of emittance scans on the estimation, and the analysis of vdM scans (see Section 8.3.9).

8.3.2 Raw or baseline-corrected samples?

The main advantage of the MON data is the inclusion of both "raw" and baseline-corrected samples. This distinction raises the question of whether the choice affects the luminosity estimation. The baseline correction accounts for the effects of the bunch structure, in which a baseline shift is caused at the beginning of each bunch train. This correction has severe consequences for the per-bunch luminosity estimation, but for the preliminary per-LB analysis described in this section, the baseline correction appears

to have limited effect. This is shown in Figure 8.5, where the probability densities of the fit residuals are superimposed for the raw and the baseline-corrected samples of the same run, using identical selection cuts. A slight offset in the mean of the distribution is visible, but the relative resolution is comparable.

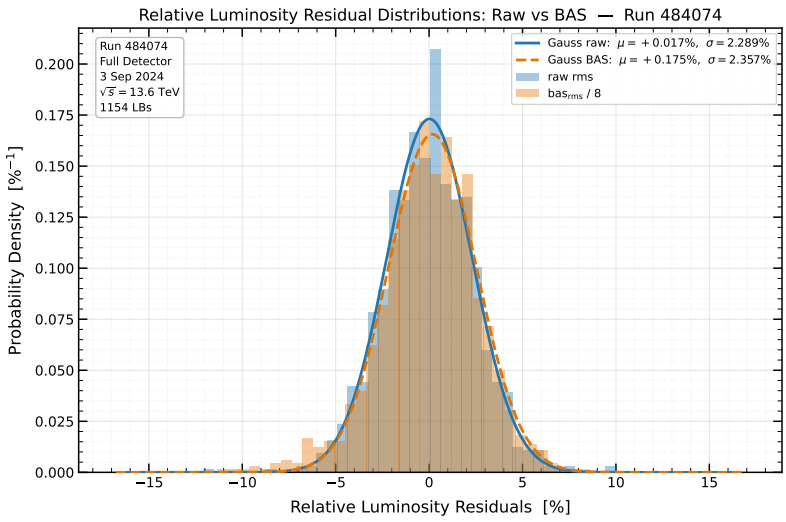

Figure 8.5: The distribution of the relative residuals of the fit of the RMS per LB to the online luminosity, for raw (blue) and baseline-corrected (orange) samples of run 484074. A slight offset in the mean is visible, but the resolution is comparable.

Given this similarity, all plots in the remainder of the preliminary study of the RMS method are produced using raw samples, so as to keep consistency with the rest of this thesis, where raw samples are used for the per-bunch luminosity measurement in Chapter 9.

8.3.3 Estimators per layer and per η coordinate

One of the main systematic effects is the contribution of each subdetector and layer to the estimation. This can be further subdivided per scieta, where the contribution of each slice of the η coordinate is analyzed. Due to the differences in the distribution of pileup between the central and the forward regions, the subdetector dependence of the luminosity estimation appears to be substantial. An example is shown in Figure 8.6, where the relative residual distributions obtained from the EMB presampler and the FCal1 layer are compared, and yielding a quite different resolution. An overview of the resulting relative resolution obtained using each respective layer of LAr is shown in Figure 8.7. Here, it is evident that each subdetector has a different response and contribution to the luminosity measurement that can be easily used for optimization with simple weighting procedures.

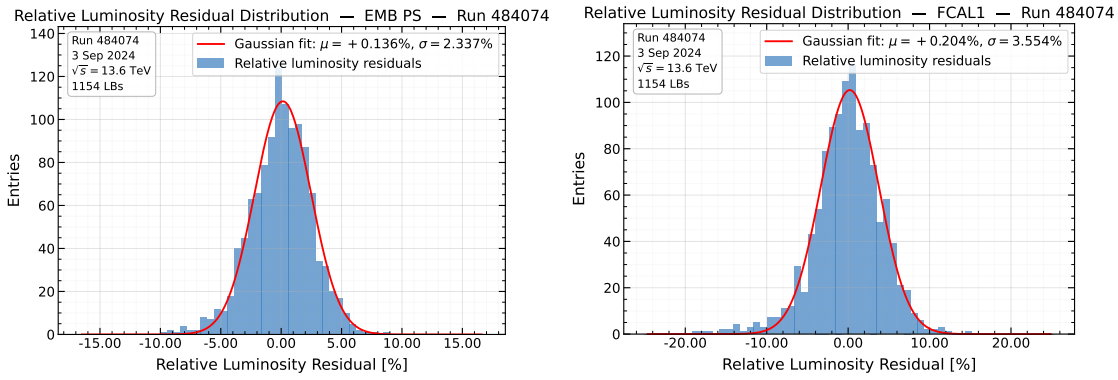

(a) (b)

Figure 8.6: The relative residual distribution obtained using samples from (a) the EMB presampler and (b) the FCal1 layer, for run 484074, each with a Gaussian fit overlaid.

Moreover, the layer dependence can be further factorized in scieta coordinates, indicating a strong η dependence of the sensitivity of the LAr calorimeter to the luminosity. Figure 8.8 shows the sensitivity and the resolution of the estimators per scieta coordinate for four representative detector/layer combinations.

8.3.4 Weighting per layer and per η coordinate

Once the effects of the dependence per layer and per scieta are known, they can be used to optimize the luminosity estimation. A layer and scieta weighting can be performed, in which each individual layer (or scieta bin) i provides, per luminosity block, its own calibrated luminosity estimate $\hat{\mathcal{L}}_i$, obtained from a linear calibration against the reference ATLAS online luminosity,

$$\hat{\mathcal{L}}_i = \frac{\mathrm{RMS}_i - \mathrm{RMS}_{0,i}}{k_i}, \tag{8.3}$$

of the same form as Eq. (8.2). Each channel, that is, each scieta-layer combination, is then characterized by a resolution σ_i. The channels are then combined into a single estimator

$$\hat{\mathcal{L}}_{\mathrm{weighted}} = \frac{\sum_i w_i \hat{\mathcal{L}}_i}{\sum_i w_i}, \tag{8.4}$$

where the weight w_i sets how strongly a given channel contributes to the combination. A $1/\sigma$-*weighted* combination is computed, with $w_i = 1/\sigma_i$, which favors channels with a smaller residual spread. The results for run 484074 are shown in Figure 8.9 for the per-layer weighting and in Figure 8.10 for the per-scieta weighting, both exhibiting an improvement in the relative resolution of the measurement of 10% and 17%, respectively.

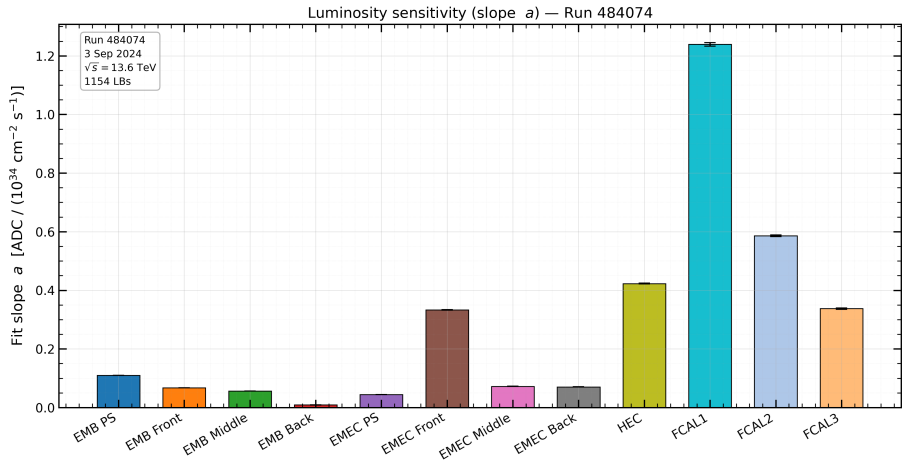

(a)

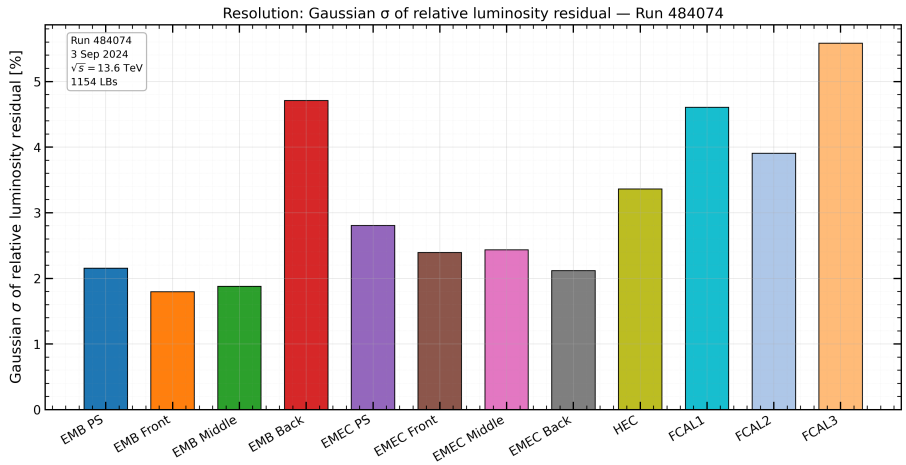

(b)

Figure 8.7: Overview of the resulting (a) sensitivity, i.e. the fit slope k, and (b) relative resolution obtained using samples from each different layer of the LAr calorimeter, for run 484074. The highly variable values indicate a strong layer dependence of the sensitivity of the luminosity measurement.

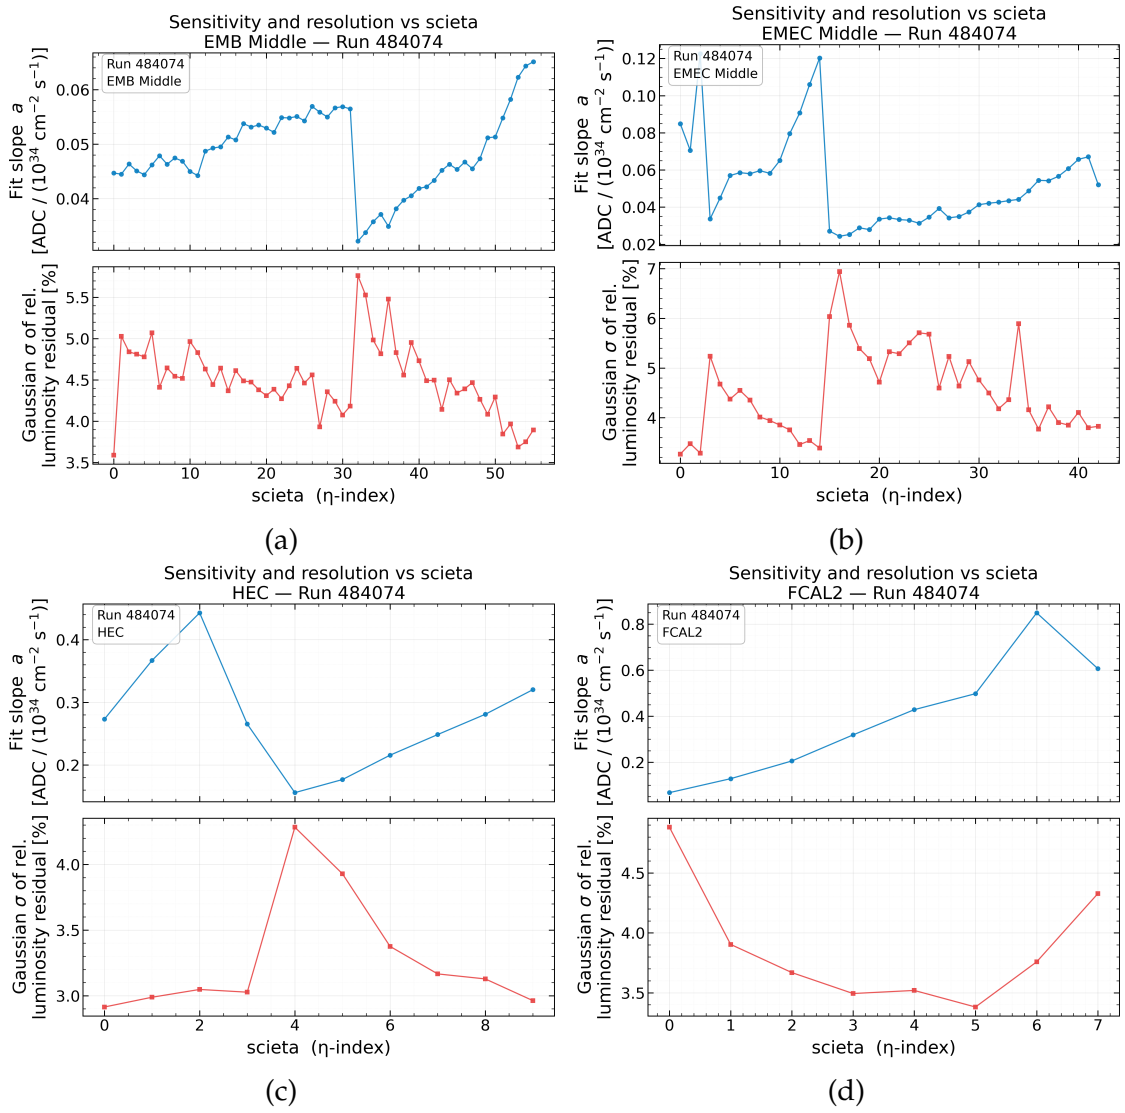

Figure 8.8: The fit slope (sensitivity) and the Gaussian resolution of relative residuals, as a function of scieta, for four exemplary detector/layer combinations from run 484074: (a) EMB middle, (b) EMEC middle, (c) HEC, and (d) FCal2. The sensitivity shows a strong dependence on η, reflecting the underlying calorimeter geometry. In the EMB middle layer, the step at scieta 32 corresponds to the boundary between the A and B barrel electrodes (see Figure 5.12), where the lead absorber thickness changes and the sampling fraction changes accordingly.

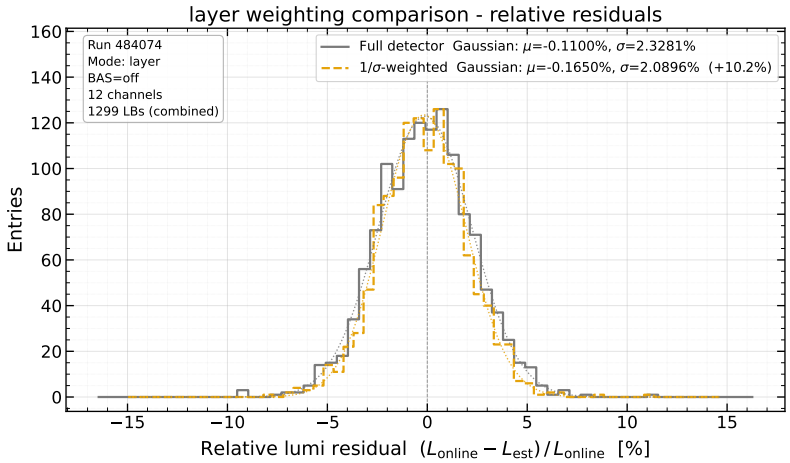

Figure 8.9: Residual distributions of the combined per-layer estimators for run 484074: $1/\sigma$-weighted, compared against the full-detector (unweighted) baseline. Gaussian fits are overlaid and the fitted widths are given in the legend.

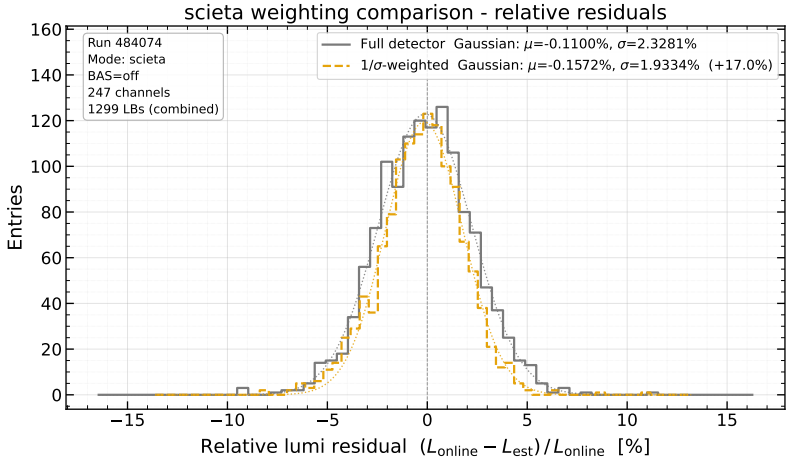

Figure 8.10: Residual distributions of the combined per-scieta estimators for run 484074: $1/\sigma$-weighted distribution, compared against the full-detector (unweighted) baseline. Gaussian fits are overlaid and the fitted widths are given in the legend.

8.3.5 Filled bunches versus empty bunches

A further systematic effect to be studied is the contribution of empty bunch crossings to the measurement. As mentioned in Section 3.2, the filling scheme of the LHC includes empty bunch crossings, that is, BCIDs where there is no proton bunch in either beam.

To distinguish between these different types of bunch crossings, ATLAS classifies each BCID into a set of predefined *bunch groups*. Each bunch group is a mask over the 3564 BCID slots in the LHC orbit, classifying BCIDs according to properties such as whether they contain colliding bunches, whether they fall in the machine's abort gap, or whether they are reserved for detector calibration.

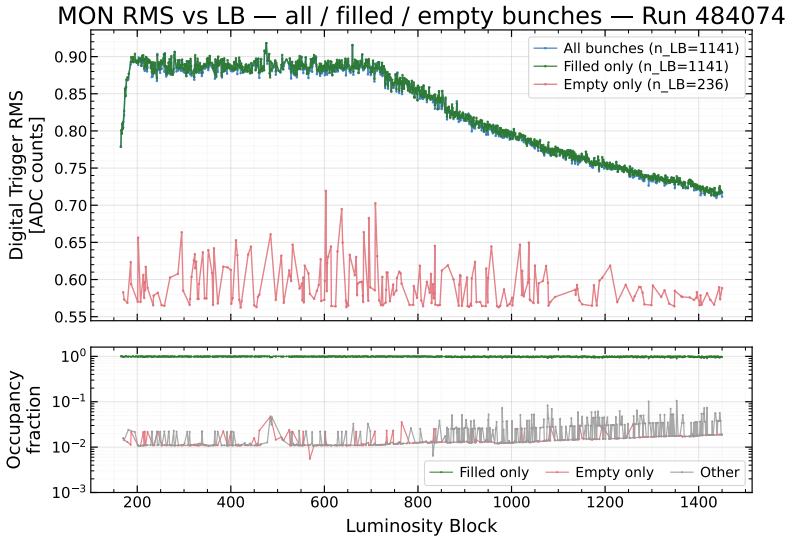

Figure 8.11: The top panel shows the RMS of the DT samples as a function of LB for all bunches (blue), filled bunches (green) and empty bunches (red), for run 484074. The bottom panel shows the fraction of events per LB belonging to each category of bunch crossings.

For the purpose of this study, the following combination of BCID cuts was used to filter out empty bunch crossings, using data averaged over the whole detector:

- **BCRVeto** excludes the small window of BCIDs surrounding the Bunch Counter Reset (BCR) signal, which is issued once per orbit in the abort gap to resynchronize the bunch counter. Crossings in this window are unreliable for BCID-level bookkeeping and must be removed regardless of whether they are filled or empty.
- **AllWithoutCalreq** removes BCIDs reserved for calibration requests (CalReq), such as TileCal laser or charge-injection pulses fired during the abort gap. These calibration pulses can register detector activity in bunch crossings that contain no proton bunch, so excluding CalReq slots prevents calibration signals from being mistaken for genuine empty-bunch background.

- **Empty** selects the BCIDs where neither beam contains a proton bunch, directly implementing the filled-versus-empty distinction of interest for this study.

As seen in Figure 8.11, the luminosity information is carried overwhelmingly by the filled bunches, and the events stored in the MON data originate almost entirely from the filled bunch group. The empty-bunch RMS sits close to the fitted intercept RMS_0 of the full-detector calibration, consistent with RMS_0 representing the electronic noise floor discussed in Section 8.2.1.

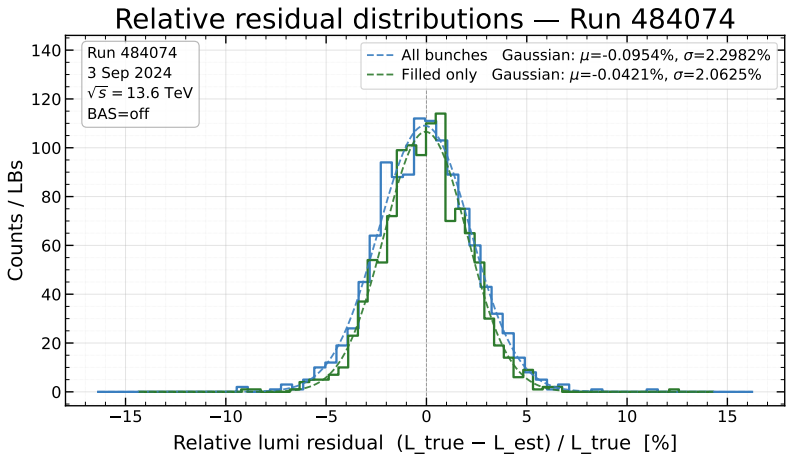

Figure 8.12: The relative residual distributions of run 484074 when using samples from all bunches (blue), and from only filled bunches (green). Both distributions are centered on zero, with an improvement in resolution when using only filled bunches.

The effect of removing events from empty bunches is shown in Figure 8.12, where the relative residual distribution is plotted for all bunches and for filled bunches only. Both distributions are centered around zero, and the resolution improves from 2.3% to 2.1% when only filled bunches are used.

8.3.6 Effects of bunch structure

In order to study the effects of bunch structure, several runs with diverse LHC conditions were chosen. Suitable runs are found during an LHC ramp-up period, where the number of colliding bunches and the train structure change from run to run. The runs used are listed in Table 8.1.

Figure 8.13 shows how the luminosity measurement depends on the number of colliding bunches, and hence on the luminosity scale. The fits parameters vary considerably with the number of colliding bunches, but are fairly aligned once the measurement is normalized by the number of colliding bunches. The RMS residual dependence on the number of colliding bunches likely comes from larger fluctuations at the start of a train, where negative shaping from earlier BCIDs is missing.

Table 8.1: The runs used for the bunch-structure study, from the 2025 LHC ramp-up period.

Run	Date	Coll. bunches
502536	17 Jul 2025	434
502557	18 Jul 2025	1154
502578	19 Jul 2025	2162
502684	21 Jul 2025	2448

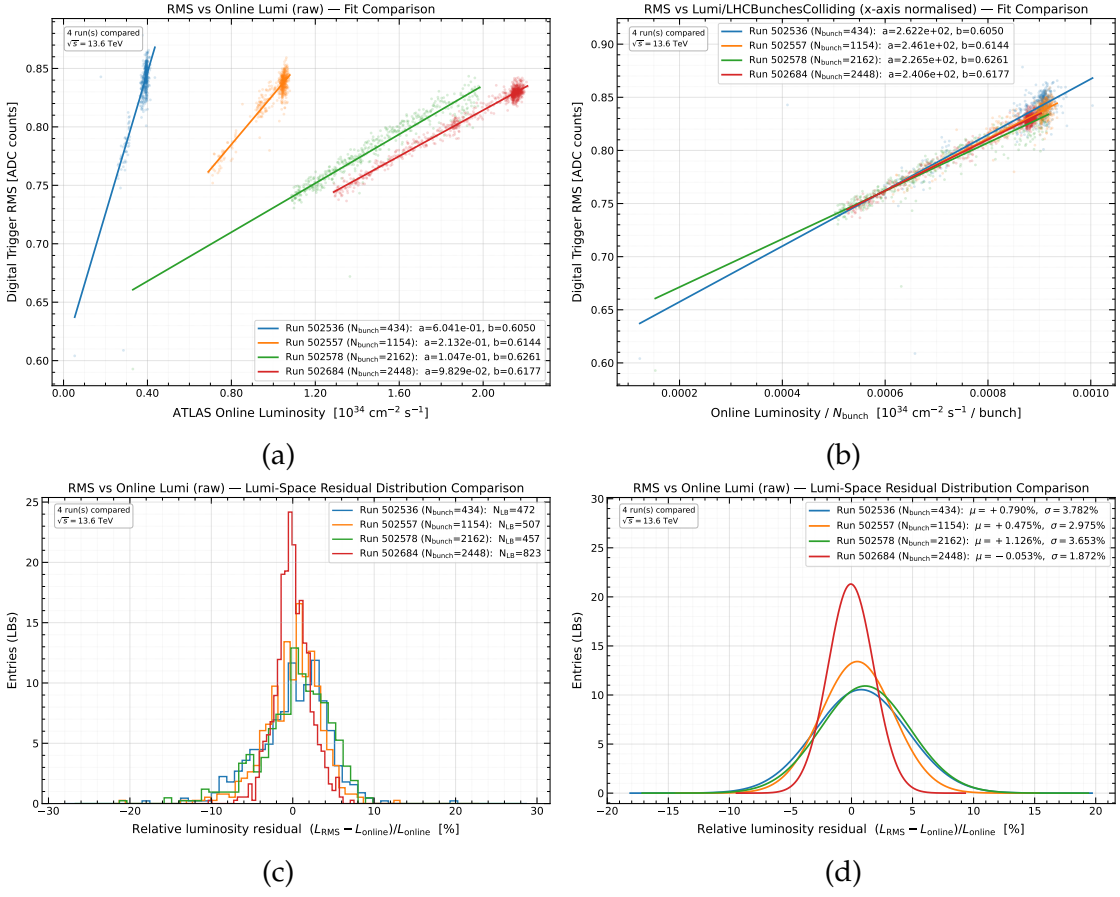

(a) (b) (c) (d)

Figure 8.13: The DT samples RMS as a function of the ATLAS online luminosity for runs with different bunch structure (a), and the corresponding plot after normalizing the measurement by the number of colliding bunches (b). All fits align when normalizing per bunch, except for the run with the lowest number of colliding bunches, run 502536. The corresponding relative residual distributions (c) and their respective Gaussian fits (d) are also plotted.

The RMS method performs well for nominal runs and high luminosities, but its sensitivity degrades at lower luminosity. This could be due to the low statistics available in these types of run (fewer LBs are available), which is made worse also by the high prescale factory of the MON stream, as already mentioned in Section 8.1.

8.3.7 Cross-year calibration: 2024 versus 2025 data

Two nominal runs from 2024 and 2025 are chosen to study the reproducibility of the measurement across different years: 486834 (15 Oct 2024, 2340 colliding bunches) and 509235 (26 Oct 2025, 2448 colliding bunches).

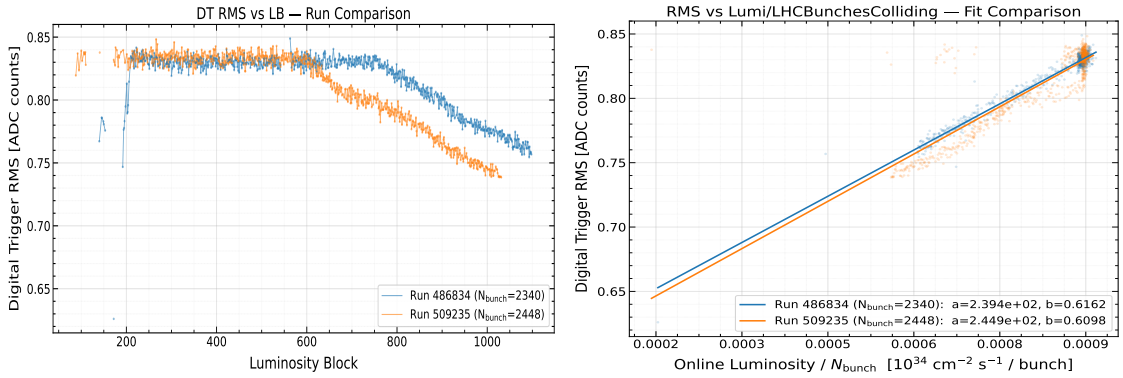

(a) (b)

Figure 8.14: The RMS of the samples from the Digital Trigger as a function of LB (a) and the corresponding fits, normalized by the number of colliding bunches (b) for nominal physics runs from 2024 (blue) and 2025 (orange).

The results seen in Figure 8.14 show correlation fits with very similar values of slope and intercept. Given the different number of colliding bunches in the nominal configuration for 2024 and 2025, the two sets of luminosity values can be normalized per bunch, which brings the two fitted slopes slightly closer, albeit with a slight difference. This indicates good reproducibility and validity of the RMS method across different data-taking years, but the difference might be large enough that a per-year, or per-period, recalibration is required to achieve full reproducibility of the results.

8.3.8 Luminosity measurement in heavy ion runs

One of the main characteristics of a functional luminosity monitor is its sensitivity to different types of collision data. Part of the 2025 data-taking campaign was reserved for heavy-ion (HI) collisions, specifically PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.36\,\mathrm{TeV}$. Run 510639, from 17 November 2025, was chosen to study this, as a representative nominal HI run with 1032 colliding bunches and a bunch spacing of 50 ns.

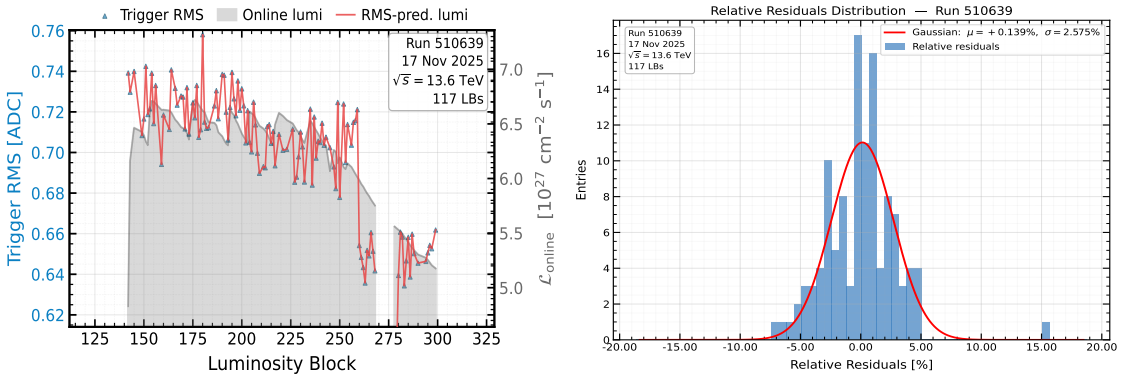

(a) (b)

Figure 8.15: The RMS of the DT samples (blue), the ATLAS online luminosity (gray), and the RMS-predicted luminosity (red), as a function of LB for the heavy-ion run 510639 (a), and the corresponding distribution of the relative residuals of run 510639, with a Gaussian fit overlaid. The distribution is centered close to zero, but its width is considerably larger than for nominal pp physics, and the tails are visibly non-Gaussian (b).

As shown in Figure 8.15(a), the RMS method retains some sensitivity to HI data, but the measurement is considerably noisier than for nominal pp physics data. The residual distribution, shown in Figure 8.15(b), is centered close to zero, but its width is considerably larger than for data from proton–proton collisions.

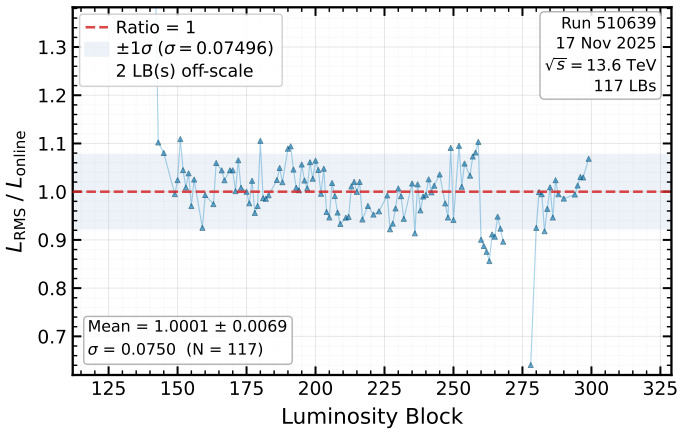

Figure 8.16: The ratio between the luminosity estimated with the RMS method and the online luminosity, as a function of LB, for run 510639. The mean ratio is consistent with unity, with a small spread.

Figure 8.16 shows the ratio between the RMS-derived luminosity and the ATLAS online luminosity. The ratio is consistent with unity, with a small spread, indicating that the method retains sensitivity at heavy-ion luminosity scales while being statistically limited.

8.3.9 Effect of emittance scans on the luminosity estimation

One of the first issues encountered in the measurement of the luminosity with the RMS method is the lack of sensitivity to *emittance scans*. An emittance scan is a dedicated van der Meer-style beam scan (see Section 7.4) in which the two colliding beams are separated and swept transversely across one another in small steps, while the resulting interaction rate is recorded at each separation. Emittance scans are performed during normal physics data-taking within an LHC fill, typically at the start or at the end of the fill, as visible in Figure 8.17 for run 484074.

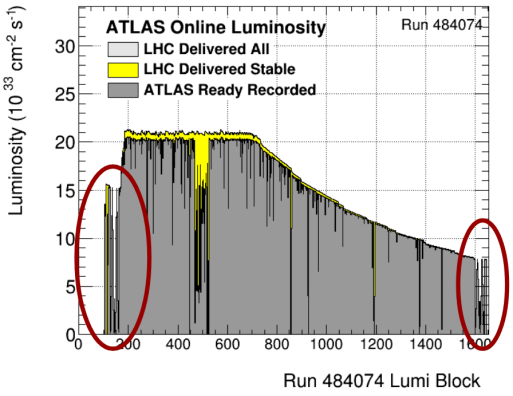

Figure 8.17: The ATLAS online luminosity profile of run 484074, with the emittance scan periods highlighted in red.

Emittance scans intentionally move the beams away from head-on collisions, so they produce a distinctive, transient signature, in which the instantaneous luminosity drops sharply as the beams are separated and recovers once they are back to the nominal collision point.

The RMS method is largely insensitive to the emittance scans, as shown in Figure 8.18, where the correlation between the RMS and the online luminosity per LB is plotted for run 484074 with and without emittance scans. All the outliers that do not follow the linear correlation originate from the emittance scans. The RMS is therefore not a reliable proxy for the luminosity during the scans, and the corresponding LBs must be filtered out to optimize the luminosity measurement.

The origin of this behavior could be a decorrelation induced by per-LB averaging, as during the scan the beams are moved in steps lasting a few seconds, so the luminosity varies appreciably *within* a single LB. On the other hand, highly prescaled events recorded by the MON server are not distributed uniformly in time within the LB, but rather are assigned with a much coarser granularity. This means that the two averages are effectively taken over different time profiles.

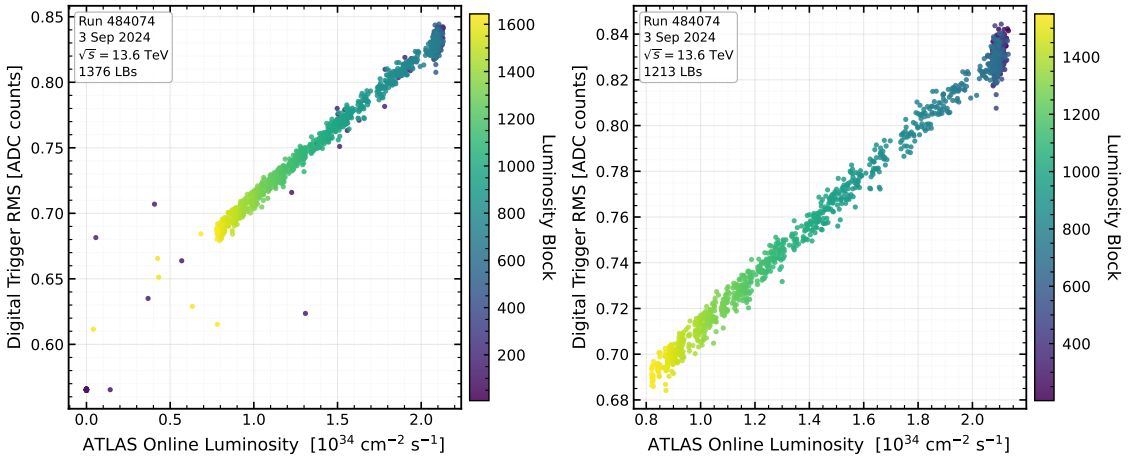

(a) (b)

Figure 8.18: The correlation between the RMS and the online luminosity of run 484074 (a) with the emittance scans included, and (b) with the scans filtered out. All the outliers in (a) originate from the emittance scans at the start and at the end of the run, indicating weak sensitivity of the RMS method during the scans.

8.3.10 Van der Meer scans

As described in Section 7.4, van der Meer scans are among the most important procedures in luminosity determination. It is useful for a luminosity measurement technique to be sensitive to these scans, and consequently to contribute, for example, to the luminosity calibration transfer. The 2025 ATLAS vdM program was performed in July, and run 502235, shown in Figure 8.19(a), was chosen for this study.

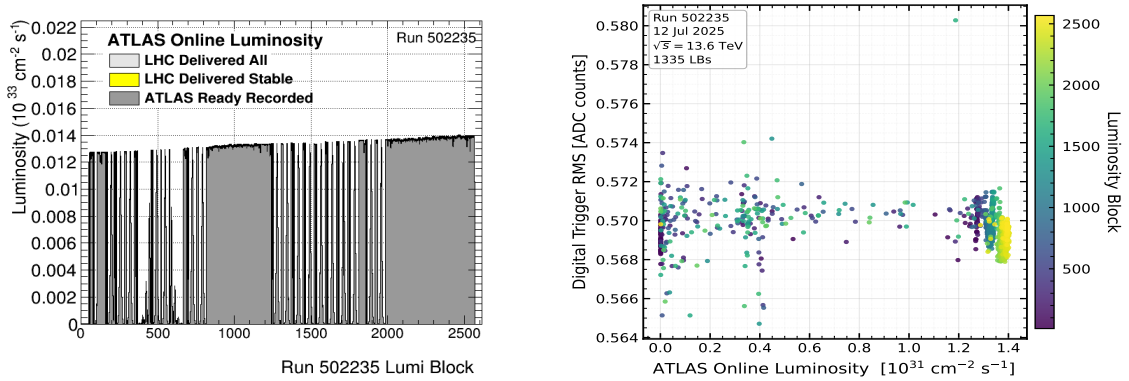

(a) (b)

Figure 8.19: The ATLAS online luminosity as a function of LB in the main run of the 2025 ATLAS vdM program, run 502235 (a), and the corresponding correlation of the RMS of the DT samples per LB as with the ATLAS online luminosity (b). No linear correlation between the two is observed.

When plotting the RMS of the DT samples per LB as a function of the ATLAS online luminosity, as seen in Figure 8.19(b), the result shows no correlation at all. At vdM luminosity scales, the RMS method seems to lose sensitivity almost completely. This is also apparent in Figure 8.20, where the RMS of the samples is superimposed on the online luminosity as a function of LB.

The origin of the low sensitivity are most likely related to the dispersion of the samples falling below the electronic noise represented by the intercept RMS_0 of Eq. (8.1) at such extreme LHC conditions, where $\langle\mu\rangle \approx 0.5$ and the instantaneous luminosity roughly three orders of magnitude below nominal pp physics levels.

The loss of sensitivity is not, in principle, irreducible. A dedicated unprescaled stream, fed by a random L1 item gated on colliding BCIDs, would remove the contribution of empty crossings identified in Section 8.3.5 and increase the available statistics by orders of magnitude, limited in practice by the readout bandwidth that could be allocated to it. Such a stream could not be commissioned in time within the Run 3 data-taking program, so whether it would provide the improvement required to reach a useful resolution at vdM luminosities could not be tested. The next opportunity to do so will be the first vdM program of Run 4.

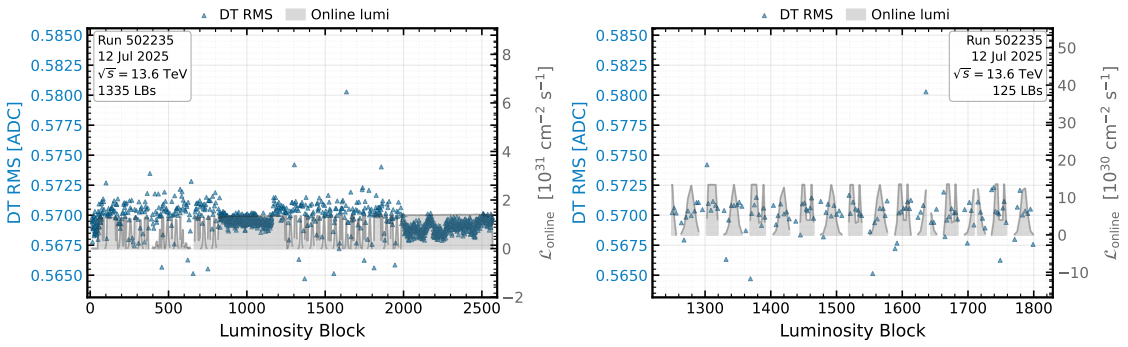

(a) (b)

Figure 8.20: The RMS of the DT samples and the ATLAS online luminosity, superimposed as a function of LB, (a) for the whole run 502235, and (b) for the selected LBs during the scan. Both show weak sensitivity of the RMS to the online luminosity at vdM scan conditions.

8.3.11 Long-term stability

While it is important for a luminosity monitor to have sensitivity to different luminosities and LHC conditions, and to have good resolution, it is also critical for a relative luminometer to be *stable* over time. As explained in Section 7.2, the main task of a relative luminometer is to monitor changes in the instantaneous luminosity over a time interval, so stability over this time interval is key.

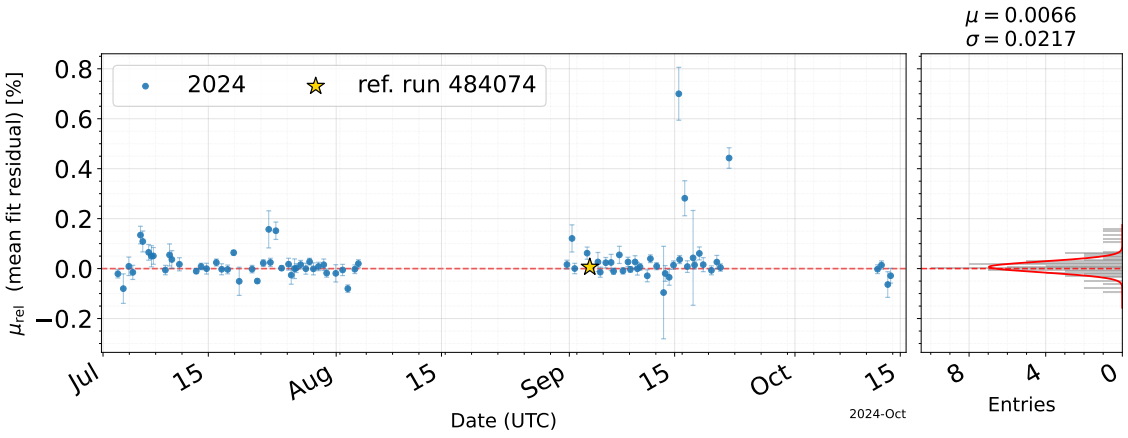

(a)

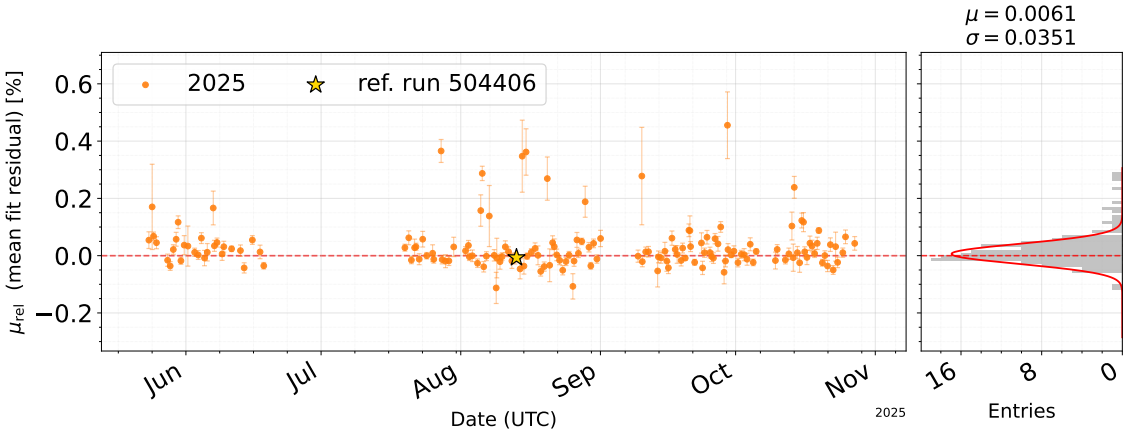

(b)

Figure 8.21: The mean of the relative residual distribution of the DT sample RMS per LB with respect to the ATLAS online luminosity, per run, for the 2024 pp data taking (a), using the fit parameters of run 484074 as the reference calibration, and the corresponding plot for 2025 pp data taking (b), using the fit parameters of run 504406 as the reference calibration.

In order to study this, all available MON data from the 2024 and 2025 Good Run Lists (GRLs) were gathered and analyzed using a reference run for each year. The RMS per LB for these runs was computed and correlated with the ATLAS online luminosity, and the resulting fit parameters were applied to all other available runs to estimate the luminosity. The mean of the corresponding residual distribution was extracted as a proxy for the stability over time.

The results are shown in Figure 8.21(a) for the 2024 pp data taking (several runs were not available in MON data in 2024 due to issues encountered during the commissioning of the DT and with the server machines used for the MON data), and in Figure 8.21(b) for the 2025 pp data taking. The RMS method is stable over time, showing a relative residual mean centered on zero for most runs in the two years, and forming an overall distribution centered on zero with a comparable width in both years.

8.3.12 Preliminary comparison with other luminosity methods

The luminosity measurement with the Digital Trigger sample RMS using the MON data has so far shown promising results as a standalone technique. The next step is to compare these results with the existing luminosity methods available in ATLAS, which were introduced in Section 7.3. For this study, the following methods are included, using the "preferred" algorithm (the standard algorithm chosen by ATLAS for that specific data-taking period), or another available algorithm where the preferred one is not available:

- **LUCID**: as the main luminosity detector of ATLAS, it is chosen as the reference, and the algorithm used is `LUCID_Preferred`, or `LUCID_HITOR_BI2` when not available.
- **Track counting**: another main luminosity monitor, using `TightModLumi`, or `TrackLumi` when not available.
- **EMEC and FCal**: the LAr gap currents drawn by the EMEC and FCal, which provide a complementary luminosity measurement from the LAr calorimeter itself, using the `EMEC_Preferred` and `FCal_Preferred` algorithms, or `EMEC_C_Onl` and `FCal_C_Onl` when not available.

A first look at this comparison is shown in Figure 8.22, where the luminosity estimated by each method is superimposed and plotted as a function of LB. The DT RMS method using the MON data shows good sensitivity to the online luminosity, but the measurement is noisier than the other methods at LB-scale resolution.

The same behavior is seen in Figure 8.23, where the ratio of each estimated luminosity to LUCID is plotted as a function of LB. The RMS method has a large spread but is well centered on unity, indicating good stability but a statistical limitation compared

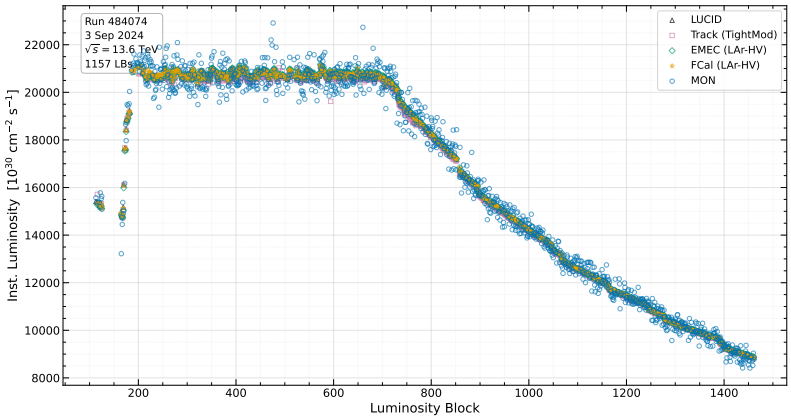

Figure 8.22: The RMS-estimated instantaneous luminosity and the luminosity measured with LUCID, track counting, and the EMEC and FCal LAr gap currents, as a function of LB, for run 484074.

with the other methods. Figure 8.24 shows the same ratio as a function of the LUCID instantaneous luminosity, where the same behavior is observed.

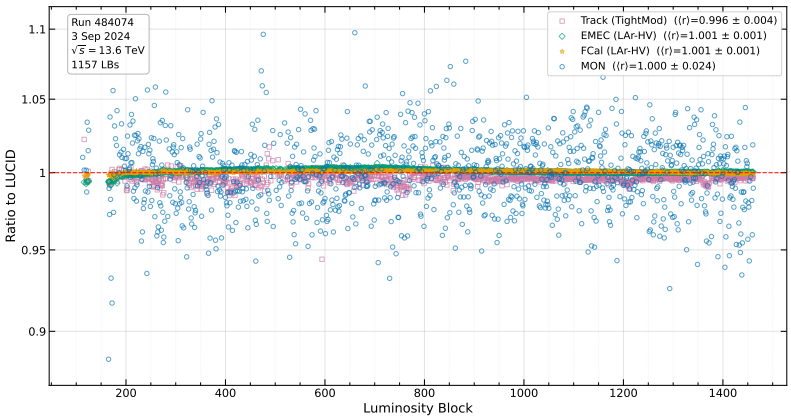

Figure 8.23: The ratio to LUCID of the DT RMS and of the other ATLAS luminosity methods, as a function of LB, for run 484074.

As done in Section 8.3.11, it is useful to examine the long-term stability of the DT RMS technique compared with the other ATLAS luminosity measurement methods. The mean ratio of the luminosity estimated by each method to LUCID is plotted as a function of time, each data point representing one run. The results are shown in Figure 8.25 for the GRL runs available from 2024, and in Figure 8.26 for 2025.

All four methods are centered on unity across both years, suggesting that the DT RMS carries no significant long-term drift with respect to LUCID. The run-to-run spread of the DT RMS ($\pm$0.010 in 2024 and $\pm$0.013 in 2025) is larger than that of the track-counting and LAr gap-current methods ($\pm$0.003–0.007). The DT RMS is thus the least

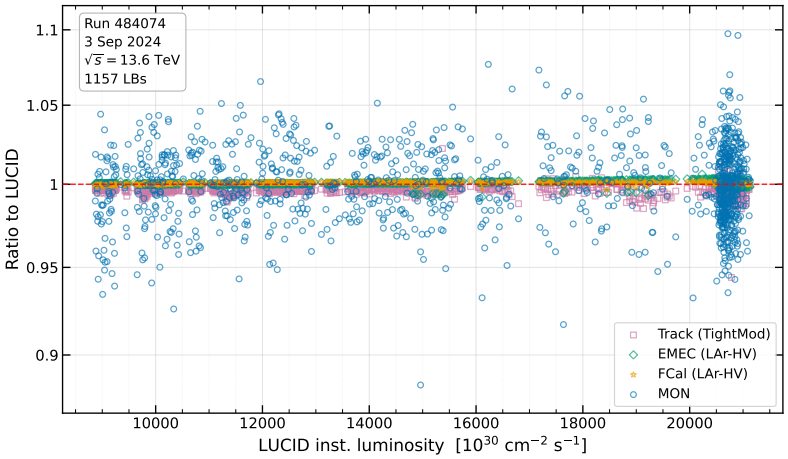

Figure 8.24: The ratio to LUCID of the DT RMS and of the other ATLAS luminosity methods, as a function of the LUCID instantaneous luminosity, for run 484074.

precise of the four on a per-run basis. This is a direct consequence of the limited statistics of the heavily prescaled MON stream and it is substantially improved when the study is moved to the high-statistics SWROD data. A repetition of this study with `physics_Main` data is presented in Section 8.4.3.

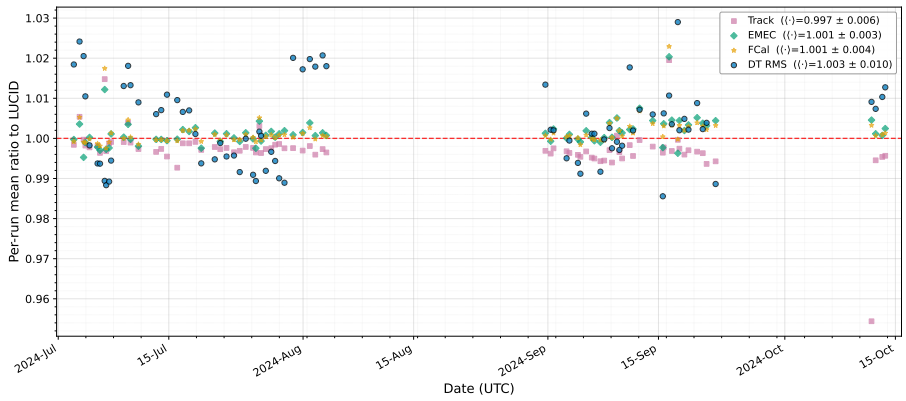

Figure 8.25: The per-run mean ratio to LUCID of the DT RMS and of the other ATLAS luminosity methods, as a function of time, for the available GRL runs in 2024. Each data point represents one run, with the mean and standard deviation over all runs given in the legend.

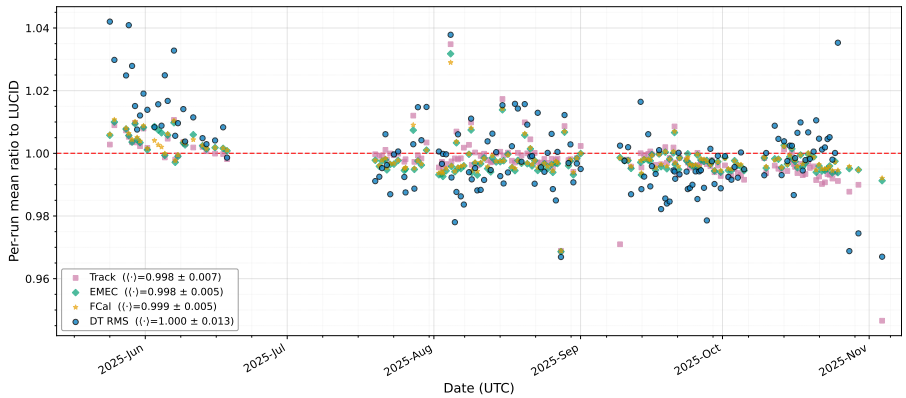

Figure 8.26: The per-run mean ratio to LUCID of the DT RMS and of the other ATLAS luminosity methods, as a function of time, for the available GRL runs in 2025. Each data point represents one run, with the mean and standard deviation over all runs given in the legend.

8.4 Study of the SWROD data

As seen in the previous section, the MON data proved to be a useful dataset to perform a preliminary study on the luminosity measurement with the RMS of Digital Trigger samples, especially in the context of long-term effects and stability. The shortcomings of this dataset were, however, apparent when studying low-luminosity runs, runs with different bunch structure and, most importantly, vdM scans.

The SWROD data, instead, are a good candidate for targeted studies to examine the limits that could not be investigated with the MON data and to probe the precision that can be achieved by the system when using the most effective and statistically largest dataset. Following on the analysis of the systematics performed during the preliminary study, new data quality cuts were applied to the data, namely the filtering of events from empty bunch groups (see Section 8.3.5).

First, the nominal physics run 509803 from 2025 and the high-μ run 509891 are analyzed, producing a set of plots approved by the ATLAS Collaboration, as presented in Section 8.4.2, to demonstrate the capabilities of the RMS method when using the dataset with most statistics, namely the ATLAS `physics_Main` stream. Then, a variety of runs are analyzed using the other available streams, described earlier in Section 8.1, and the effects of trigger biases when using these different streams are described. Finally, a targeted comparison with the current ATLAS luminosity methods is performed, this time using the `physics_Main` dataset.

8.4.1 Comparison of streams and trigger biases

While the access to different streams of data can be very useful to compare different measurements and obtain the maximal resolution for the study, it is necessary to study

the effects of the trigger on the measurement, and whether it introduces biases.

The recorded samples enter the RMS computation only for events accepted by the trigger, so each stream samples the LHC orbit and the collision population differently, through two distinct mechanisms. The first is the BCID sampling profile. Physics triggers fire almost exclusively on colliding bunch crossings, so streams such as `physics_Main` populate the filled bunch groups, where the samples carry the pileup-driven dispersion. The `calibration_LArPEBDigitalTrigger` stream, in contrast, also accepts events in empty and inter-train BCIDs, whose samples contain only electronic noise. These events dilute the per-LB RMS and degrade its correlation with the luminosity, as already established in Section 8.3.5.

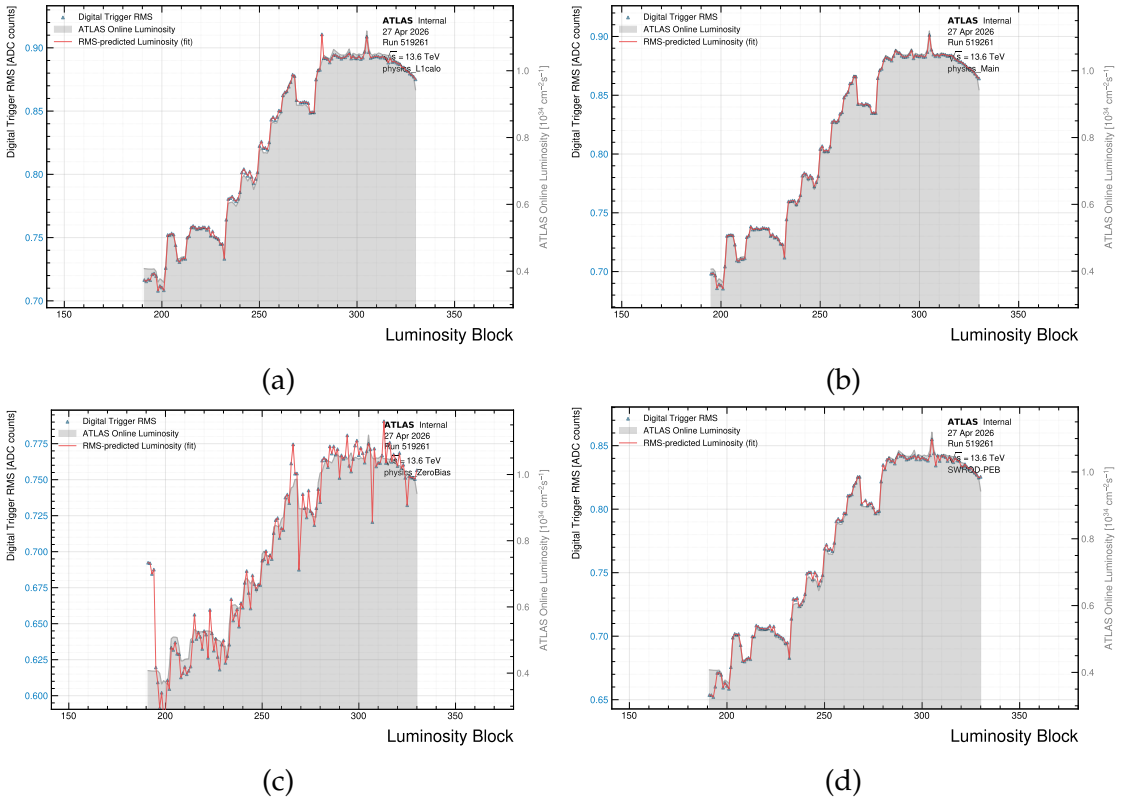

Figure 8.27: The RMS of the Digital Trigger samples superimposed on the ATLAS online luminosity, with the RMS-predicted luminosity from the linear fit overlaid, as a function of LB for run 519261, using the `physics_L1Calo` (a), `physics_Main` (b), `physics_ZeroBias` (c), and `SWROD_PEB` (d) streams.

The second mechanism is an event-content bias. Physics triggers preferentially select crossings with large energy deposits, so the accepted events are systematically more active than an unbiased sample at the same μ, and the size of this bias depends on the trigger menu composition of each stream. The `physics_ZeroBias` stream, whose accept rate is proportional to the luminosity but whose recorded events carry no selection on

their content, provides a reference that is free of content bias by construction, although its statistics are very limited.

One way to study this is to first compare the measurement of the same run with different streams, obtain the resolution and see whether the different results are due to higher or lower statistics (as is the case of `physics_Main` and `physics_ZeroBias`).

The datasets of run 519261 (27 April 2026, 576 colliding bunches), were processed for each different stream (`physics_Main`, `physics_L1Calo`, `SWROD_PEB`, and `physics_ZeroBias`), and Figure 8.27 shows the results: for each stream, the RMS seems to align to the online luminosity, with the exception of `physics_ZeroBias` being fairly noisy.

The more interesting result is shown in Table 8.2, where the relative resolution computed from the residuals for the different streams, showing `physics_Main` as the most precise data stream. One thing worth of notice is to look at the number of total events processed: `SWROD_PEB` is the stream with most events, without being the most precise. At the same time, `physics_ZeroBias` has a remarkably low number of events processed compared to the others, but nonetheless achieving a 1% resolution toward the online luminosity.

Table 8.2: The different streams used in this study of run 519261, the number of events process and the respective relative resolutions.

Stream	Nr of Events	Rel. Resolution [%]
`SWROD_PEB`	2,216,516	0.39
`physics_Main`	1,975,046	0.19
`physics_L1Calo`	257,837	0.40
`physics_ZeroBias`	40,765	1.0

This result suggests that, as expected, the number of events affects the resolution of the measurement; however, it is not the only factor at play. The trigger selection also affects the result: `physics_Main` triggers on filled bunches and on events with physics content, while `SWROD_PEB` collects a sizable contribution from empty bunches. This is evident from Figure 8.28, showing that the fraction of events from filled bunches is much higher for `physics_Main` than for the `SWROD_PEB` stream, which has a large number of events coming from inter-train BCIDs.

However, the resolution by itself cannot show whether a stream measures the luminosity with a *bias*. A constant offset in the RMS, such as the one expected from the event-content selection of the physics triggers, is absorbed into the intercept RMS_0 of the per-run calibration of Eq. (8.1). A trigger bias becomes a systematic effect only if it changes with the luminosity or with time within the fill, for example through prescale changes as the luminosity decays. This is checked by comparing the residual of each stream.

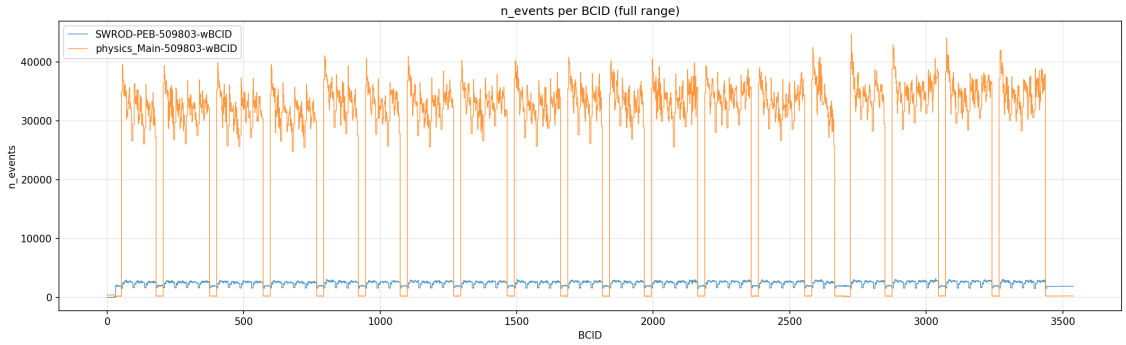

(a)

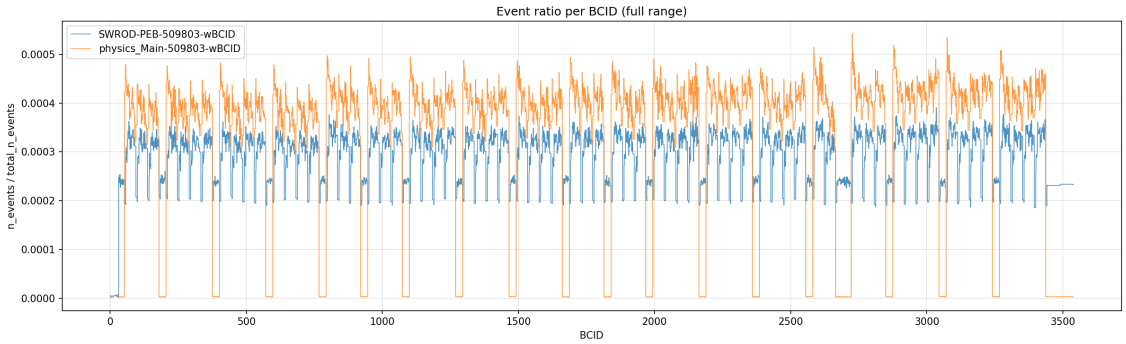

(b)

Figure 8.28: (a) The number of events per BCID for the `SWROD_PEB` (blue) and `physics_Main` (orange) streams of run 509803, and (b) the same distributions, each normalized to the total number of events of its own stream. The `physics_Main` events are concentrated in the filled BCIDs, while the `SWROD_PEB` stream collects a sizable fraction of its events from inter-train BCIDs.

Figure 8.29 shows the difference between the residual of each stream and that of `physics_Main`. In run 519268, `physics_L1Calo` and `SWROD_PEB` agree with `physics_Main` within about 0.3% above $\mathcal{L} \approx 0.9 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, while `physics_L1Calo` deviates by up to $(1.9 \pm 0.6)\%$ towards lower luminosities. In run 519261 the differences reach $(1.0 \pm 0.3)\%$ for `physics_L1Calo` and about 0.6% for `SWROD_PEB`. `physics_ZeroBias` is compatible with `physics_Main`, within uncertainties of about 0.4–1% over most of the range.

These checks do not exclude a trigger bias. However, within the two runs analyzed, they limit its luminosity-dependent part to about 1%, with larger deviations still possible at the lowest luminosities of a fill.

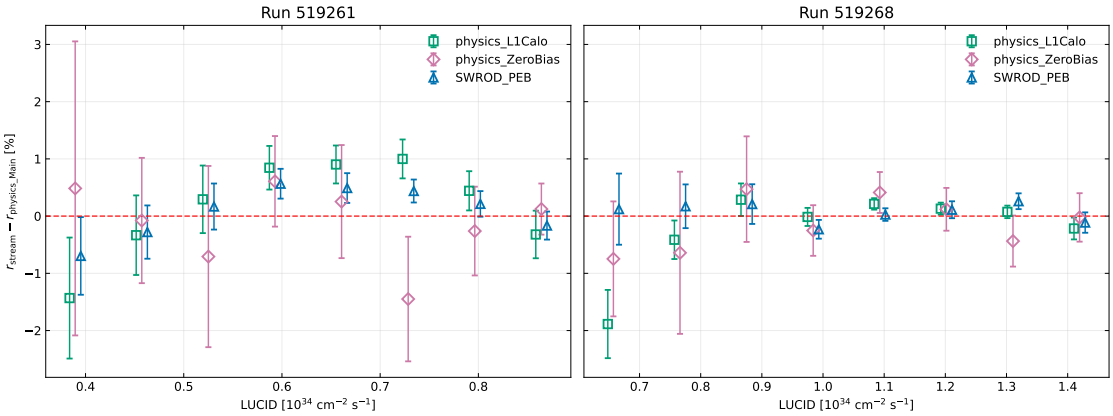

Figure 8.29: Difference between the relative residual to LUCID for each stream analyzed and `physics_Main`, in bins of LUCID luminosity, for runs 519261 (left) and 519268 (right). Each stream is calibrated with its own linear fit. The uncertainties combine the statistical error of the mean in each bin.

Given that `physics_Main` is the data stream that provides the most precise determination of the luminosity, we can use it as a baseline to probe the limits of the measurement at nominal and HL-LHC conditions, as described in the next section, and use it to compare with other ATLAS methods, as it will be shown in Section 8.4.3.

8.4.2 Approved results with ATLAS physics data

In this section, the results of the study using `physics_Main` data for the analysis of a nominal 2025 physics run and the analysis of a high-μ run are presented.

The data are from stable proton–proton collision events from run 509803, recorded on November 3, 2025, with an average $\langle\mu\rangle = 58.6$. In this run, 2448 colliding bunches were present in stable beam conditions and typical Run 3 conditions of the LHC and ATLAS. Figure 8.30 shows a strong correlation between the RMS and the ATLAS online luminosity profile as a function of LB, and Figure 8.31a shows a linear correlation

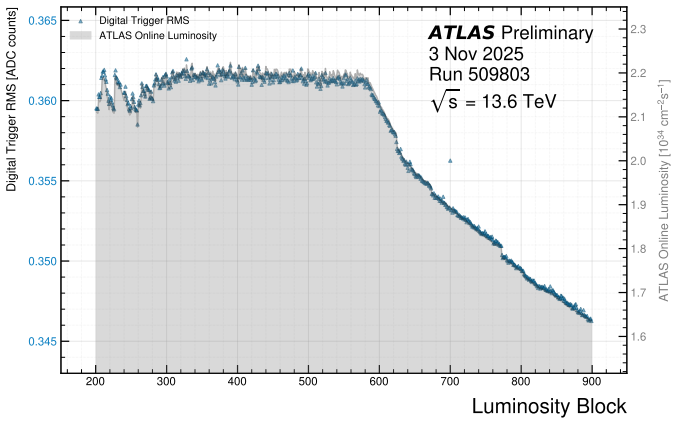

Figure 8.30: The RMS of the Digital Trigger samples (blue) superimposed on the ATLAS online luminosity (gray) as a function of LB, for the nominal physics run 509803, using the `physics_Main` stream [117].

between the RMS of the samples and the ATLAS online luminosity. Most strikingly, Figure 8.31b is evidence of the optimal resolution that the RMS method can yield, reaching a 0.04% relative resolution with respect to the ATLAS online luminosity.

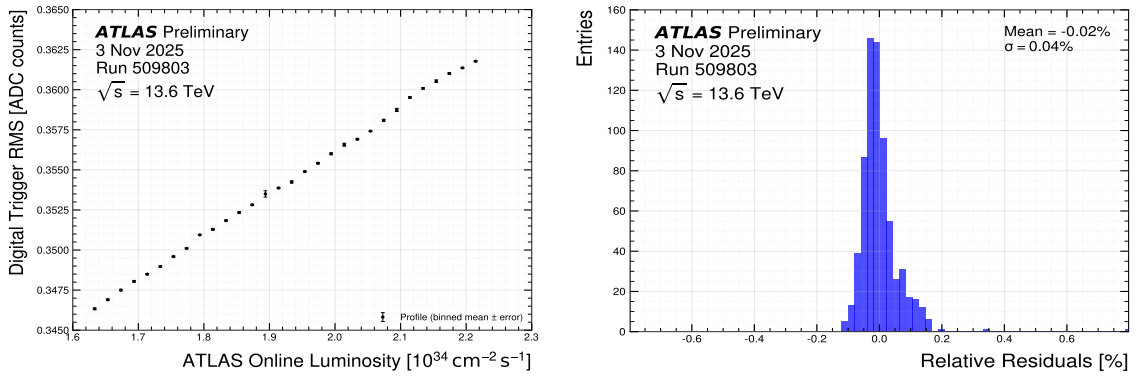

(a) (b)

Figure 8.31: (a) Profile of the Digital Trigger sample RMS as a function of the ATLAS online luminosity for run 509803, showing a linear correlation, and (b) the distribution of the relative residuals of the linear fit, with the fitted mean and width given in the legend [117].

Similarly, the same analysis using the `physics_Main` stream was performed on stable proton–proton collision events from run 509891 on November 5, 2025, with an average $\langle\mu\rangle = 120.6$. The RMS of the samples superimposed on the ATLAS online luminosity is shown in Figure 8.32, exhibiting again a strong correlation.

Figure 8.33 shows a strongly linear correlation between the RMS and the ATLAS online luminosity, and the distribution of the residuals, reaching a relative resolution

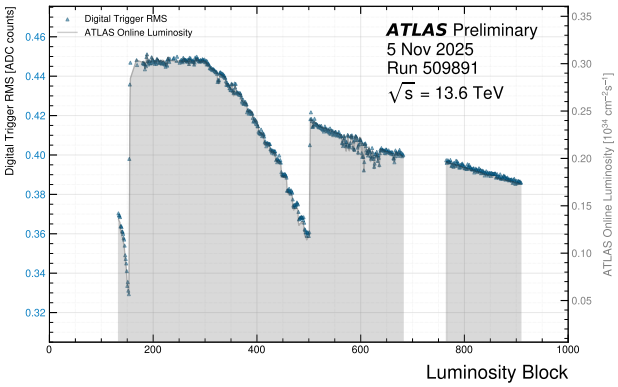

Figure 8.32: The RMS of the Digital Trigger samples (blue) superimposed on the ATLAS online luminosity (gray) as a function of LB, for the high-μ run 509891, using the `physics_Main` stream [117]. The RMS shows strong correlation to the online luminosity and sensitivity to small changes in it, as evidenced in the dips in luminosity around LB 600, closely followed by the RMS.

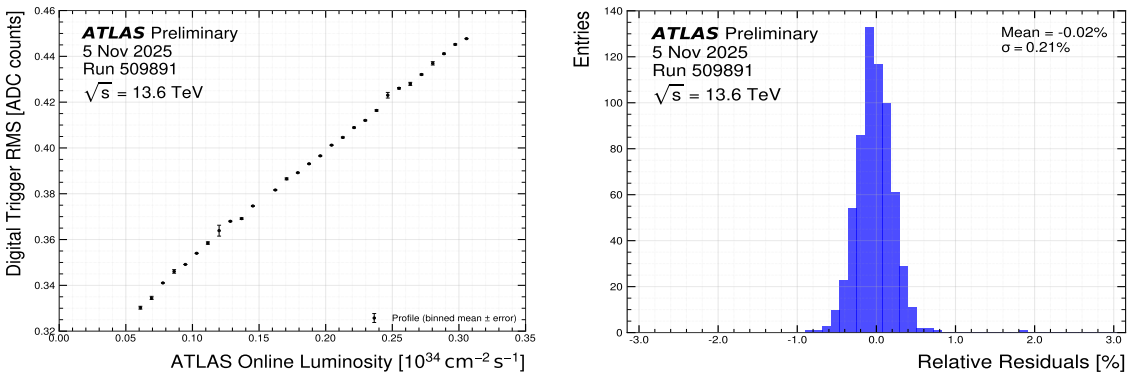

(a) (b)

Figure 8.33: (a) Profile of the Digital Trigger sample RMS as a function of the ATLAS online luminosity for the high-μ run 509891, and (b) the distribution of the relative residuals of the linear fit, with the fitted mean and width given in the legend [117].

of 0.21%. As visible from the vertical-axis ranges of Figures 8.30 and 8.32, the instantaneous luminosity of run 509891 is almost an order of magnitude lower than that of run 509803, despite its higher $\langle\mu\rangle$; this is consistent with the fewer colliding bunches typical of a dedicated high-μ configuration. As discussed in Section 8.2.1, a lower instantaneous luminosity reduces the pileup contribution to the sample dispersion relative to the fixed electronic noise floor, degrading the resolution of the RMS-based estimate; the smaller number of usable LBs at these conditions further limits the statistics entering the fit. These results suggest that the measurement of the luminosity with the RMS method can be used in nominal physics conditions as well as in conditions similar to the ones foreseen during the HL-LHC, making it a strong candidate as an ATLAS relative luminosity monitor during Run 4.

8.4.3 Comparison with ATLAS luminosity methods

Once it was established that `physics_Main` is the stream that yields the most precise measurement of the luminosity, it was relevant to repeat the comparison to other ATLAS luminosity methods performed in Section 8.3.12.

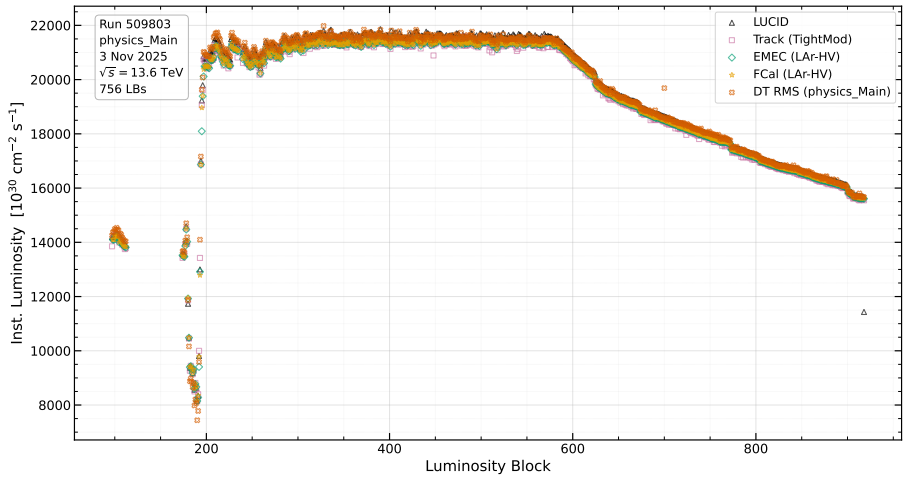

Figure 8.34: The instantaneous luminosity estimated with the DT RMS using the `physics_Main` stream, and the luminosity measured with LUCID, track counting, and the EMEC and FCal LAr gap currents, as a function of LB, for run 509803.

Figure 8.34 shows a much more stable measurement compared to the one with the MON data from Figure 8.22, giving instead a result more in line with the other luminosity methods in ATLAS. The result is seen clearly also in Figure 8.35, where the ratio to LUCID per LB is plotted for each method. Quite strikingly, the RMS method with the `physics_Main` dataset proves to be the closest to LUCID, boasting a mean ratio of 1.000 ± 0.016, compared with, for instance, the mean ratio of track counting to LUCID, 0.990 ± 0.014, and of the LAr gap-current methods, 0.992 ± 0.014.

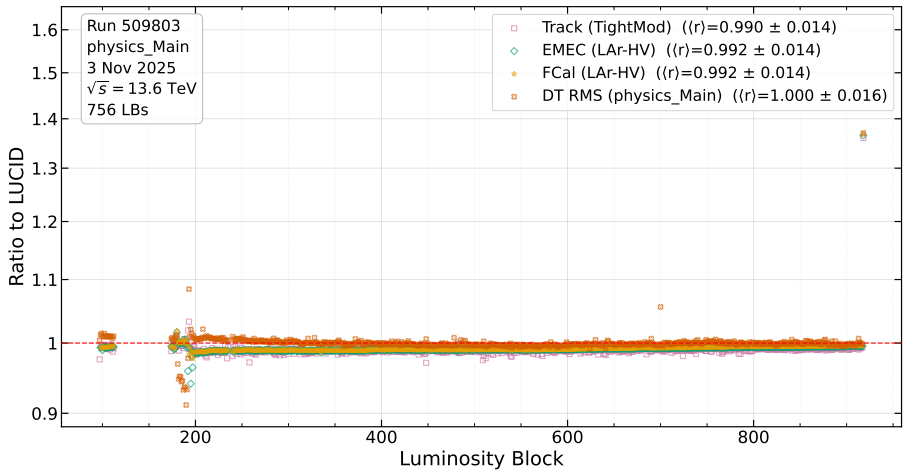

Figure 8.35: The ratio to LUCID of the DT RMS (`physics_Main`) and of the other ATLAS luminosity methods, as a function of LB, for run 509803. The mean ratio and its standard deviation for each method are given in the legend.

However, in Figure 8.36, the main weakness of the RMS method is evident. The plot shows the ratio to LUCID for each method, as a function of the LUCID instantaneous luminosity: while having a comparable ratio to other methods at higher luminosity, the RMS method seems to diverge slightly toward lower luminosities, indicating an insensitivity to low luminosities already seen with the preliminary study using the MON data, and confirmed here also with the `physics_Main` stream dataset.

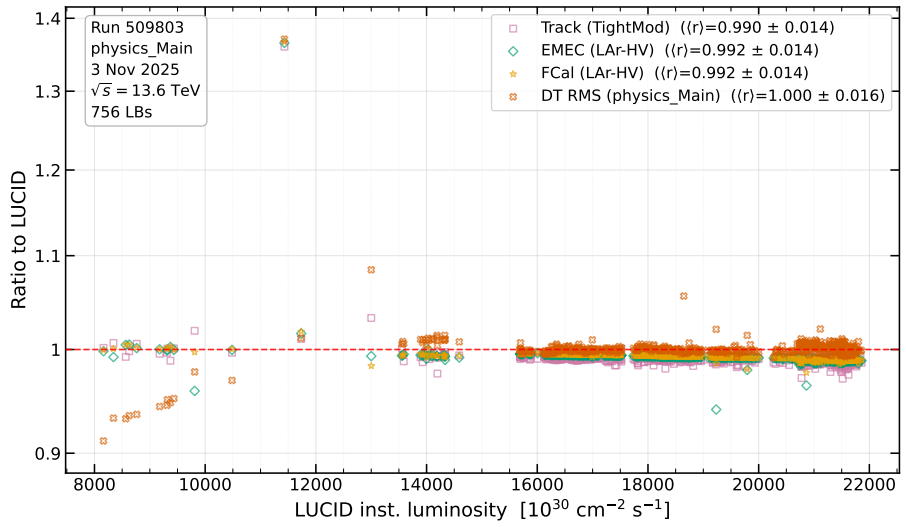

Figure 8.36: The ratio to LUCID of the DT RMS (`physics_Main`) and of the other ATLAS luminosity methods, as a function of the LUCID instantaneous luminosity, for run 509803.

As mentioned earlier in Section 8.1, the `physics_Main` dataset, while being the optimal stream for the measurement, is not as accessible, given the large size of the data for each run. For this reason, a long-term stability study using the `physics_Main` dataset could not be performed. This is a limitation of offline analyses such as the one performed in the thesis. However, if this type of luminosity measurement were to be implemented online, or in the firmware for Run 4, these limitations would be easily overcome.

Chapter 9

Bunch-by-bunch luminosity measurement

Measuring the luminosity bunch by bunch is an important but simultaneously difficult task to perform. As described in Section 7.3, only a few systems in ATLAS are capable of measuring luminosity per bunch. One of the main advantages of the Digital Trigger, apart from its stability and linearity, is its storage of BCID information for each sample. This chapter describes a novel method to measure the luminosity *per bunch* using exactly that BCID information in the samples, i.e., the mean ADC pedestal measured per BCID. This is effectively different from the RMS method described earlier, as it does not use the dispersion of the samples, but rather the dependence of the mean ADC pedestal on the bunch trains structure.

At first, data from the MON stream were used, given the access to raw samples without baseline correction, which therefore carry the full information of each bunch. The only exception is run 520264, for which, as mentioned earlier in Section 8.1, the monitoring data from the SWROD (the `SWROD_PEB` stream) was recorded with no baseline correction due to a change in readout configuration. Therefore, the higher statistics from the SWROD of this run were used for this study.

The chapter first sets out the convolution model and some possible methods that can be used for the inversion, validates them on a toy Monte Carlo, and then applies them to the measured mean pedestal of run 520264. A comparison with the luminosity per bunch from LUCID and track counting is performed to assess the performance of each method, and a final section extracts a data-driven kernel from the LUCID per-bunch luminosity in order to separate the limitations of the inversion algorithms from those of the assumed pulse shape.

9.1 Idea and theoretical background

The goal of this study is to reconstruct an independent estimate of the instantaneous luminosity separately for each of the $N_{\mathrm{BCID}} = 3564$ possible BCIDs that make up one LHC orbit, using DT data. The starting point is to look at the mean sample pedestal per BCID, shown in Figure 9.1 for a specific scieta coordinate. The bunch structure is evident from the plot, and the natural next step is to understand what this signal consists of.

At every bunch crossing, depending on the pileup, LAr pulses are recorded, such as the one shown in Figure 9.2. The designed bipolar shape offers a way of self-correcting

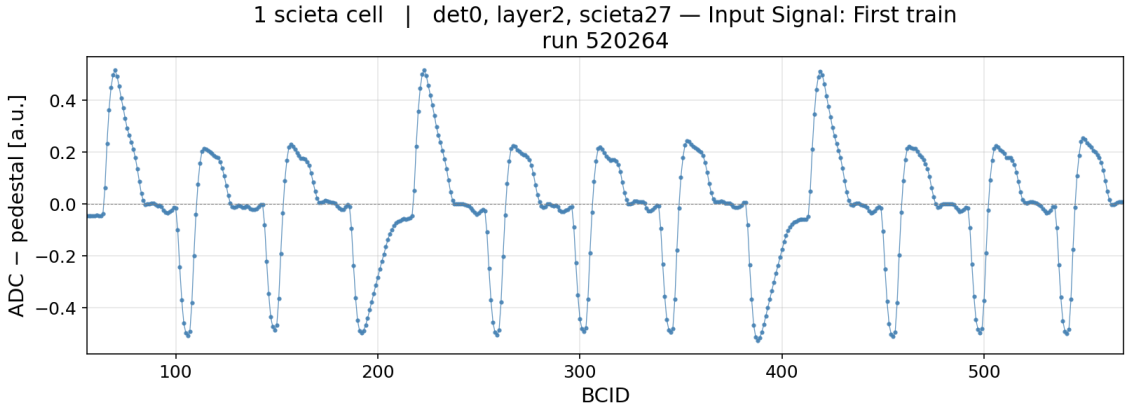

Figure 9.1: The mean of the samples with the pedestal subtracted as a function of BCID, for the first bunch train in the orbit, for an example scieta coordinate in the EMB middle layer. The bunch and train structure is clearly visible in the plot.

the baseline shifts caused by the structure of the bunch trains, effectively create a signal such as the one shown in Figure 9.1, with negative undershoots in the inter-train and inter-bunch regions.

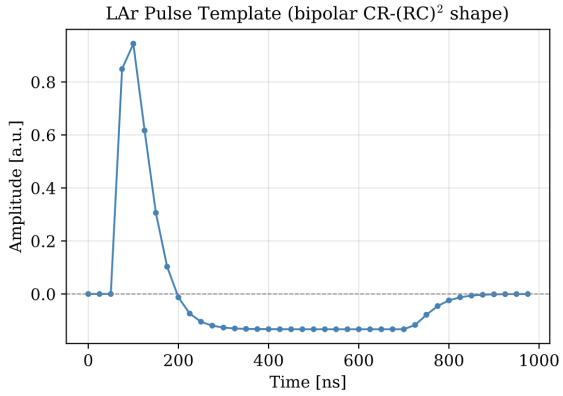

Figure 9.2: Simulation of an ideal template of a LAr pulse. Being a bipolar pulse lasting for more than 800 ns, the contribution of these pulses will shape the signal by creating a pedestal shift.

The pulses extend more than 800 ns, but the bunches collide every 25 ns, so multiple pulses overlap at any instant, contributing to the signal, as illustrated in the example in Figure 9.3, where 9 pulses with varying amplitude are simulated and plotted to show the contribution to the signal. Therefore, this signal can be seen as a *convolution* of the luminosity with the pulse shape as its kernel. Denoting by $L[i]$ the true instantaneous per-bunch luminosity $\mathcal{L}_b$ (Eq. (7.5)) in bunch crossing i, and by $h[j]$ the LAr pulse giving the fraction of pulse amplitude observed j bunch crossings after the first one, the

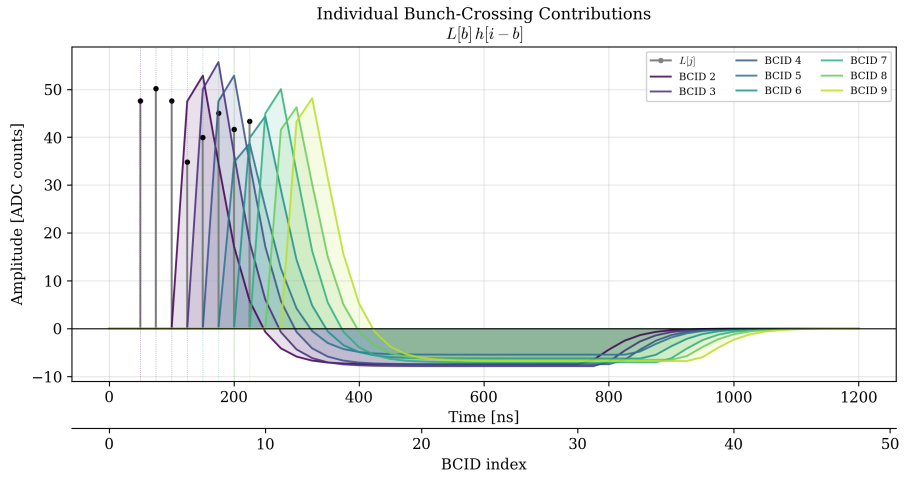

Figure 9.3: Individual contributions to the observed signal from nine consecutive bunch crossings. Each bunch crossing b produces the same pulse shape h, scaled by its own luminosity $L[b]$ and shifted to its arrival time; the long bipolar tail of each pulse extends well beyond the 25 ns bunch spacing, so many pulses contribute to any given sample.

expected signal $S[i]$ in BCID i is defined as

$$S[i] \;=\; \sum_{j} L[i-j]\,h[j] \;\equiv\; (L * h)[i], \tag{9.1}$$

i.e. a discrete convolution of the true per-bunch luminosity with the LAr pulse shape kernel. This is clearly illustrated in Figure 9.4, where the final "detected" simulated signal created by the convolution of superimposed pulses shown previously in Figure 9.3 is plotted. This is precisely the structure already seen in the data in Figure 9.1, where the measured mean pedestal per BCID is the sum of the contributions of all the superimposed pulses recorded. Essentially, it is the signal in which the luminosity is hidden, or *convolved*. Consequently, the luminosity per bunch can be retrieved by *deconvolving* the signal with the known kernel.

Writing Eq. (9.1) in matrix form, we let $\mathbf{L}$ collect the per-bunch luminosities and $\mathbf{S}$ the observed signal over one orbit. The convolution with kernel h is a linear operator, so there exists a matrix such that

$$\mathbf{S} \;=\; A\,\mathbf{L} \tag{9.2}$$

Deconvolution is the inverse problem of estimating $\mathbf{L}$ from the observed $\mathbf{S}$ and the known (or assumed) kernel h.

So, a direct estimator is simply

$$\hat{\mathbf{L}} \;=\; A^{-1}\mathbf{S}. \tag{9.3}$$

For any length-N_{BCID} sequence $x[n]$ defined on one orbit, the discrete Fourier trans-

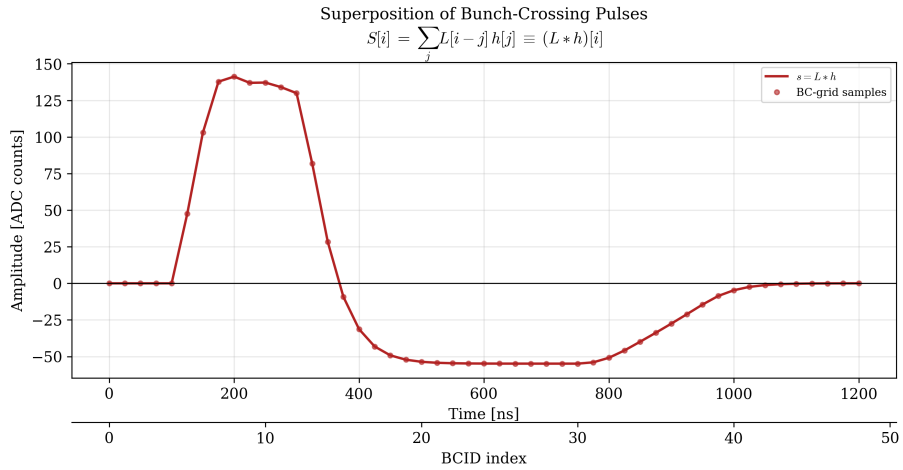

Figure 9.4: Sum of the individual bunch-crossing contributions of overlapping pulses. The observed signal is not a sequence of resolved single-bunch pulses but their coherent sum, in which the undershoot of earlier pulses partially cancels the peak of later ones.

form (DFT) is

$$\tilde{x}[k] \;=\; \sum_{n=0}^{N_{\mathrm{BCID}}-1} x[n]\, e^{-2\pi i\, nk/N_{\mathrm{BCID}}}, \qquad k = 0, 1, \ldots, N_{\mathrm{BCID}} - 1, \tag{9.4}$$

where the index k is the frequency. This, in this notation $\tilde{S}[k]$, $\tilde{h}[k]$ and $\tilde{L}[k]$ are the DFTs of the observed signal, of the pulse kernel and of the true luminosity respectively.

The convolution of Eq. (9.1), which couples every BCID to its neighbors, becomes

$$\tilde{S}[k] \;=\; \tilde{L}[k]\, \tilde{h}[k], \tag{9.5}$$

that is, N_{BCID} independent scalar equations, one per frequency bin. Each of them can be solved for $\tilde{L}[k]$ by a single division, so that we can compute, bin by bin,

$$\tilde{L}[k] \;=\; \frac{\tilde{S}[k]}{\tilde{h}[k]}. \tag{9.6}$$

The luminosity is then recovered by transforming back with an inverse DFT. This is equivalent to Eq. (9.3), and it is faster to evaluate through the *fast Fourier transform* (FFT). Hence, the direct candidate for the deconvolution is a classical FFT division followed by an inverse FFT.

In this study the same template LAr kernel is given to different methods, so that any difference in performance can be ascribed to the inversion algorithm alone and not to the quality of the kernel. The three options chosen are the following:

- **Direct algebraic inversion**, Eq. (9.3), and its frequency-domain form Eq. (9.6).

Exact when A is invertible, i.e., the **FFT method**.

- **Inverse matrix method**, which replaces the exact inverse by the minimizer of the squared residual via the normal equations

$$\hat{\mathbf{L}} \;=\; (A^{\mathsf{T}}A)^{-1}A^{\mathsf{T}}\,\mathbf{S}. \tag{9.7}$$

- **Iterative statistical inversion**, of which Richardson–Lucy (RL) [119, 120] is the choice: a multiplicative expectation-maximization scheme for Poisson-distributed data, derived under the explicit assumption that both the kernel and the solution are non-negative, defined as

$$\hat{L}^{(n+1)}[i] \;=\; \hat{L}^{(n)}[i] \cdot \frac{1}{\sum_j h[j]} \sum_j h[j]\, \frac{S[i+j]}{\big(h * \hat{L}^{(n)}\big)[i+j]}, \tag{9.8}$$

9.2 Toy Monte Carlo simulation

A toy Monte Carlo simulation is used to validate and compare the deconvolution methods before applying them to real data. A synthetic ground-truth luminosity profile $L[i]$ is generated over a toy LHC orbit, with the luminosity in each filled BCID drawn from a Poisson distribution. This ground truth is convolved with a LAr pulse template, shown in Figure 9.2, following Eq. (9.1), to produce a toy signal $S[i]$, to which Gaussian noise is added to emulate the effect of electronic noise on the measured pedestal.

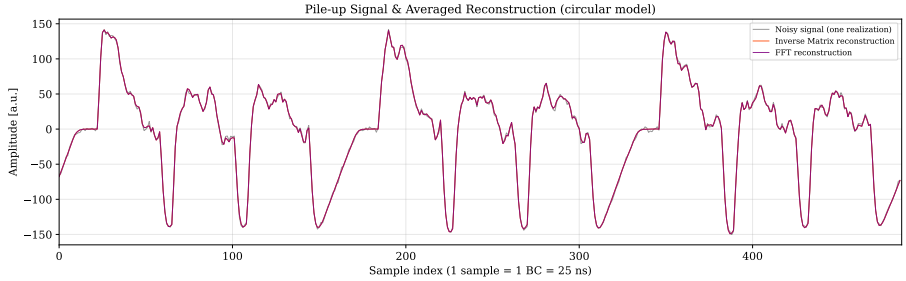

Figure 9.5: Noisy toy signal for one noise realization, together with the inverse matrix and FFT reconstructions, averaged over noise realizations. It is evident that both methods reconstruct the simulated luminosity almost perfectly.

The inverse matrix and FFT estimators of Eqs. (9.7) and (9.6) are applied to $S[i]$ and compared against the known $L[i]$, as done in Figure 9.5, repeating the procedure over many noise realizations to quantify the reconstruction root-mean-square error (RMSE) and the relative error of each method.

Richardson–Lucy deconvolution [119, 120] could not be included as it is in this comparison, as a fully positive kernel is required. A variant of the kernel that is non-negative

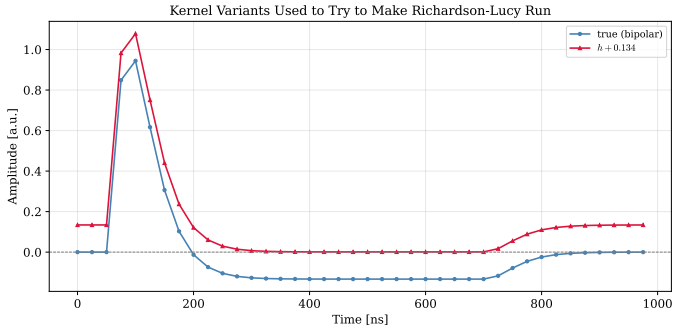

Figure 9.6: The true bipolar LAr pulse template and the non-negative variant tested with Richardson–Lucy, $|h|$ and $h + c$.

by construction, $h[j] + c$ for a constant offset c large enough to make every sample non-negative (an example is seen in Figure 9.6), was tested as a possible workaround.

Table 9.1 summarizes the results at noise $\sigma = 2.0$ with 200 noise realizations. Both FFT and inverse matrix methods methods recover the input luminosity to sub-percent accuracy, with a Pearson correlation above 0.9999 in either case. The inverse matrix and the FFT are found to agree perfectly, as mathematically they are reduced to exactly the same bin-by-bin division of Eq. (9.6). The modified kernel for the Richardson–Lucy method, on the other hand, gives a reconstruction an order of magnitude worse, with the best RMSE 3.13 compared to 0.26 of the other two.

Table 9.1: Toy MC reconstruction accuracy for the inverse matrix, FFT and variant RL methods, averaged over 200 noise realizations at $\sigma = 2.0$.

Method	RMSE	Correlation	Rel. error
Inverse Matrix	0.260	0.99994	0.44%
FFT	0.260	0.99994	0.44%
RL+offset	3.131	0.99148	6.25%

This is visualized in terms of residuals for all methods as in Figures 9.8-9.9. Replacing h by $h + c$ makes the method usable, as it retains some sensitivity as shown in Figure 9.7, but not comparable to the other two in terms of performance.

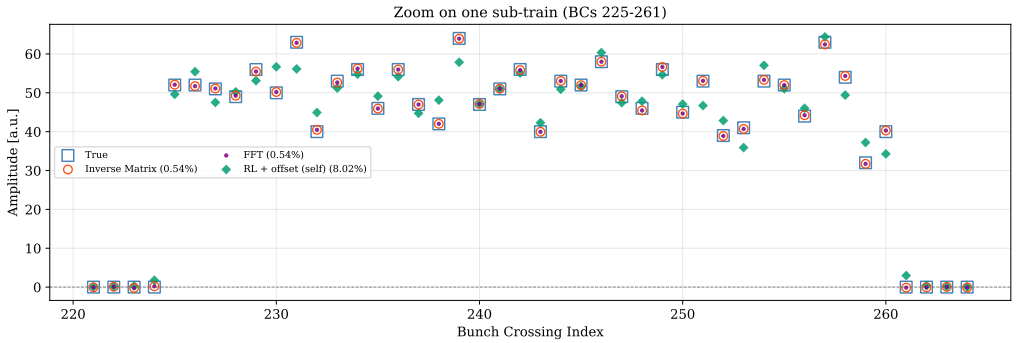

Figure 9.7: True vs. reconstructed luminosity for a representative bunch train, for the inverse matrix, FFT and the variant RL methods.

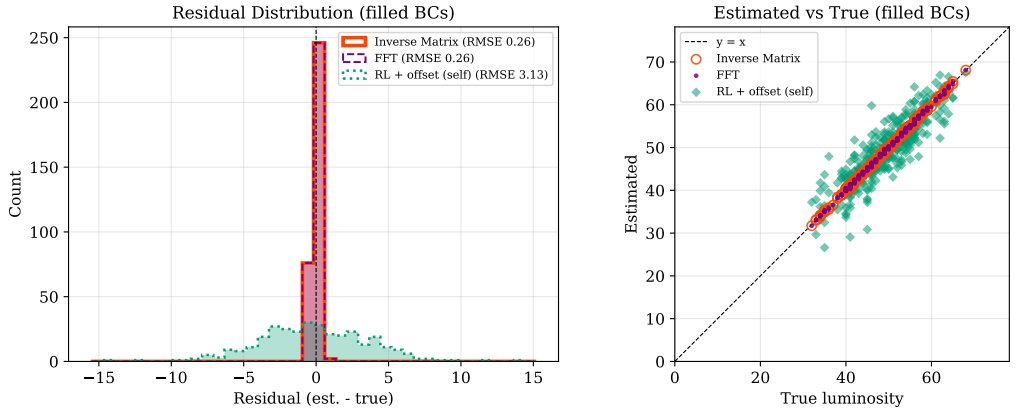

Figure 9.8: The distribution of the residuals between the estimated and true luminosity in the filled BCIDs (left), and the estimated versus true luminosity in the filled BCIDs, with the diagonal $y = x$ overlaid (right), for the inverse matrix (red) and FFT (purple), and variant RL (green) methods.

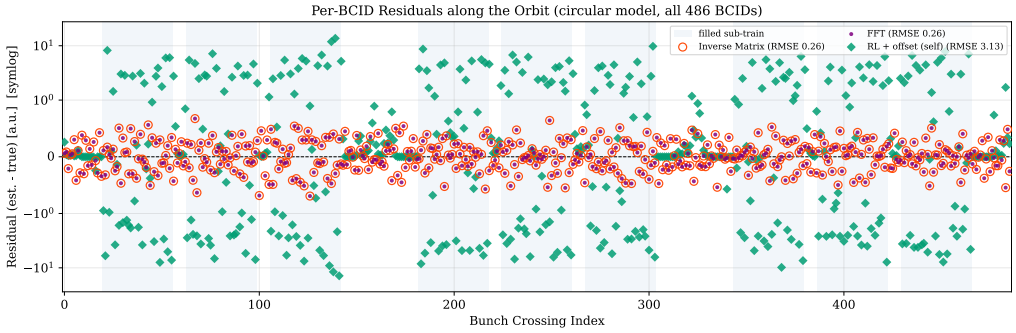

Figure 9.9: The residuals between the estimated and true luminosity as a function of BCID, for the inverse matrix (red) and FFT (purple) and variant RL (green) methods.

9.3 Deconvolution of the mean pedestal

The toy Monte Carlo of the previous subsection established that the FFT, the inverse matrix are possible candidates for the deconvolution to retrieve the luminosity by bunch. This section applies them to the measured mean pedestal of run 520264, where no ground truth exists and the reconstruction must instead be validated against the official ATLAS luminometers.

We start with the input signal, the mean ADC value measured per BCID, and with the kernel. The kernel is chosen from a set of pulse shapes measured in dedicated delay runs during the LHC ramp-up, collected per scieta, given the common electronics shared by cells (and therefore Super Cells) located at the same scieta (more details were given in Section 5.6.1). In these dedicated runs, delays were applied in steps of roughly 3 ns and average pulse shapes were reconstructed for all delay values, then normalized and aligned in time to yield the final pulse shape used in this study. An example of such pulse shape measured from LHC data is illustrated in Figure 9.10.

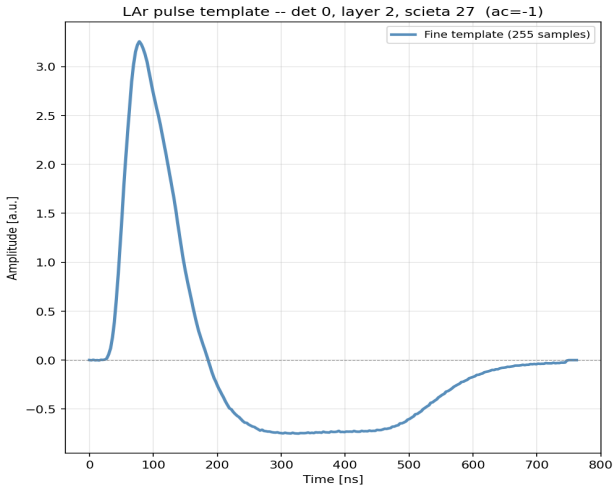

Figure 9.10: The pulse shape for scieta 27 in the EMB middle layer measured using LHC data from a special delay run.

The idea is to deconvolve the luminosity on a per scieta basis using the corresponding pulse shape, and finally to average the results over all available scieta coordinates to retrieve the final, full-detector luminosity per bunch.

An example of this is shown in Figure 9.11, where the mean signal per BCID is plotted together with the luminosities per bunch retrieved with the inverse matrix and with the FFT, for a train in the middle of the orbit, using only the signal and kernel from one scieta coordinate. The two reconstructions are indistinguishable, which is the expected result already seen in the toy Monte Carlo.

It is evident that the measurement from only one scieta slice is not optimal, as seen in the noisy inter bunch intervals, but it already shows promising results: the bunch-

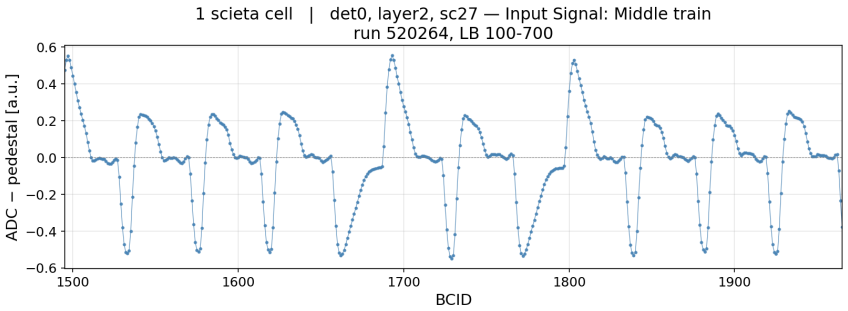

(a)

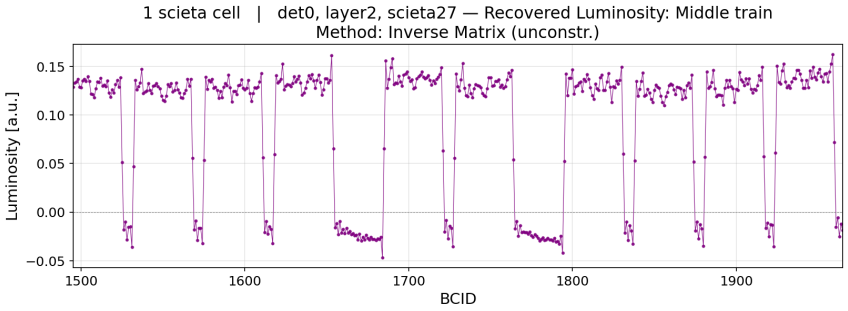

(b)

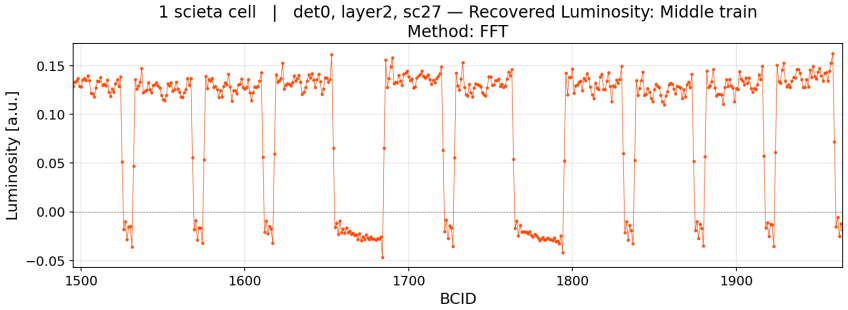

(c)

Figure 9.11: The mean pedestal per BCID used as input signal (a), and the corresponding reconstructed luminosities using the inverse matrix (b) and the FFT (c) for the middle train in the orbit in run 520264, using only the pulse shape and the data from scieta 27 on EMB middle layer.

to-bunch structure within a train is followed closely, at the price of a negative baseline between trains and of artifacts producing a dip toward the end of the orbit.

Because the inverse matrix and the FFT return the same solution, showing both in every subsequent figure would be redundant. From here on the unconstrained result is therefore represented by the FFT alone, which evaluates that same solution far more cheaply through the fast Fourier transform.

The next step is to compute this for each scieta available, and average the various measurements for the whole detector. The results of this are shown in Figure 9.12, where the retrieved luminosity is plotted over each BCID in the full orbit. The same artifacts are present also here, as seen in the overshoot at the start of the orbit and the recurring negative baseline.

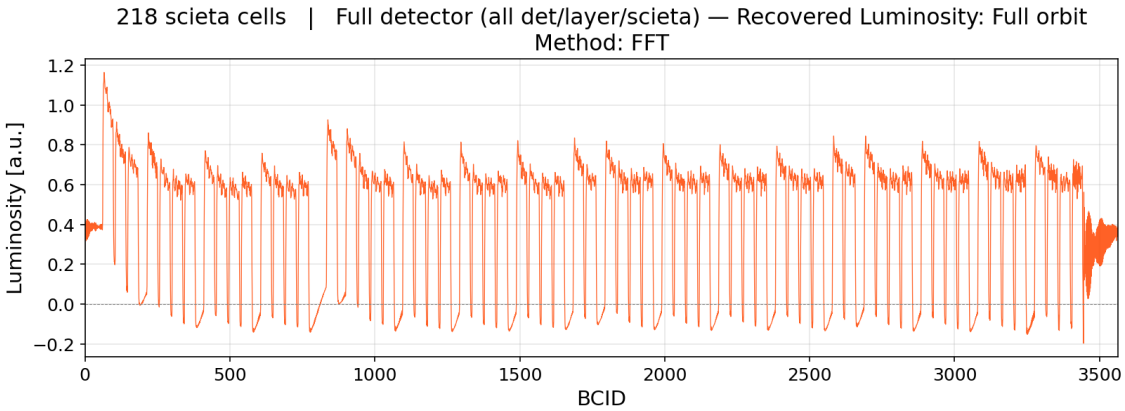

Figure 9.12: The luminosity retrieved via FFT deconvolution and averaged over all the scieta coordinates for the full orbit in run 520264.

Figure 9.13 shows example trains at the start, middle and end of the orbit. These plots show that the method has great sensitivity for the BCIDs within the filled bunches, and that the inter-train intervals are better reconstructed by averaging over all the SCs in the detector, compared to the results obtained when using only one scieta coordinate in Figure 9.11.

Overall, it is clear that the bunch-by-bunch luminosity can be retrieved from the mean ADC pedestal of the Digital Trigger samples using an FFT deconvolution. However, the measurement is approximate at best as several artifacts remain.

The most probable cause of these artifacts is the kernel. The pulse shape used for this study was obtained in dedicated delay runs in 2025, while the data investigated are from 2026, so the template may differ from the pulse shape that actually produced the recorded data. Another contribution is also possible: there is the possibility that the real pulses have a tail extending beyond the 800 ns spanned by the 32 samples that can be measured experimentally by the Digital Trigger, so that the tail is truncated in the template by construction. Both point to the same conclusion, namely that an alternative,

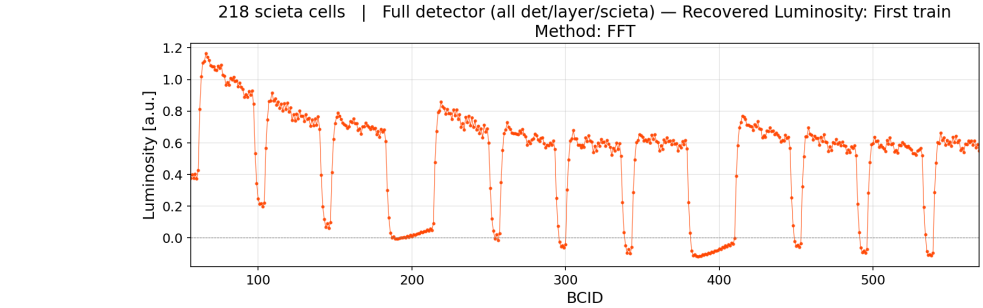

(a)

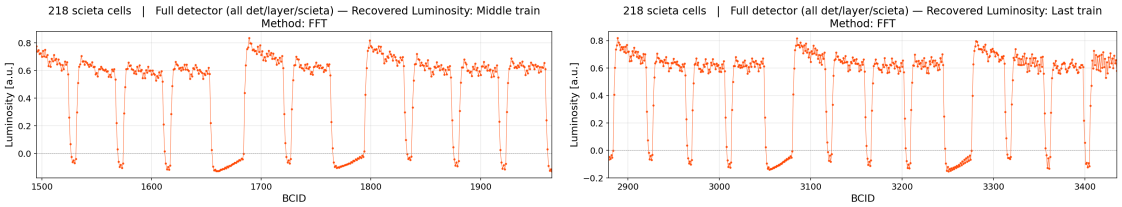

(b) (c)

Figure 9.13: The luminosity retrieved via FFT deconvolution and averaged over all the scieta coordinates in run 520264. The first train (a), a train in the middle of the orbit (b), and the last train (c) are shown.

dedicated pulse shape is needed to optimize the measurement of the luminosity per bunch. This is tested directly in Section 9.5, where a kernel extracted from the data is substituted for the template and the artifacts are found to follow the kernel rather than the inversion algorithm.

The next step is to use this result and compare it to ATLAS luminosity measurements per bunch with LUCID, as it is presented in the next section.

9.4 Comparison with LUCID luminosity per bunch

The results obtained in the previous section using the data from run 520264 are now compared to the LUCID and track-counting measurements per bunch. For this study, the values obtained by averaging the luminosity measured per scieta are used. The deconvolved values from LAr are averaged over scieta and scaled to LUCID using a constant α. The scaled values are then plotted with the corresponding per-BCID measurements of the LUCID and track-counting luminosity determinations.

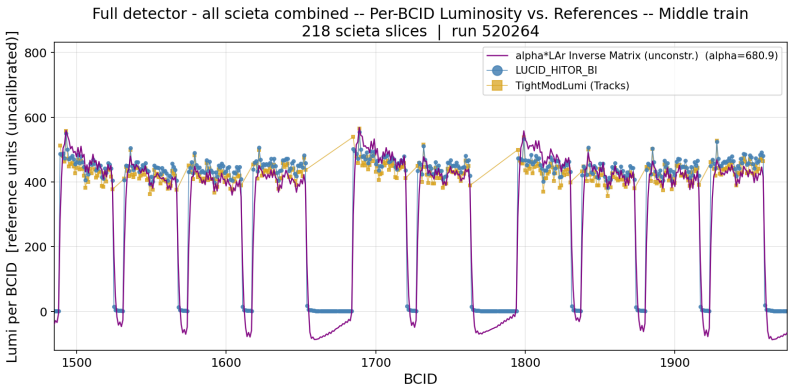

(a)

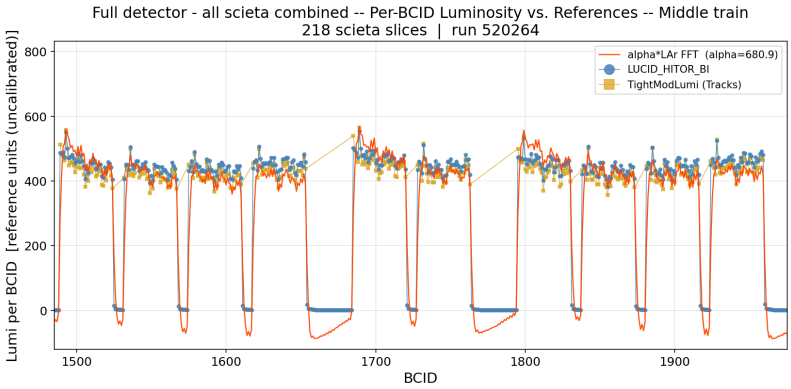

(b)

Figure 9.14: The estimated luminosity obtained by deconvolving the LAr sample pedestal, as a function of BCID, scaled by α and compared to the corresponding measured values of LUCID (blue) and track counting (yellow), for a train in the middle of the orbit: inverse matrix (a) and FFT (b) reconstructions.

Figure 9.14 shows the comparison with LUCID for the inverse matrix and the FFT. The methods once again yield the same performance, both reconstructing the bunches fairly well, but both exhibit small overshoots at the beginning of bunch trains and negative luminosity at the inter-bunch intervals.

The artifacts described in Section 9.3 become very evident once compared to refer-

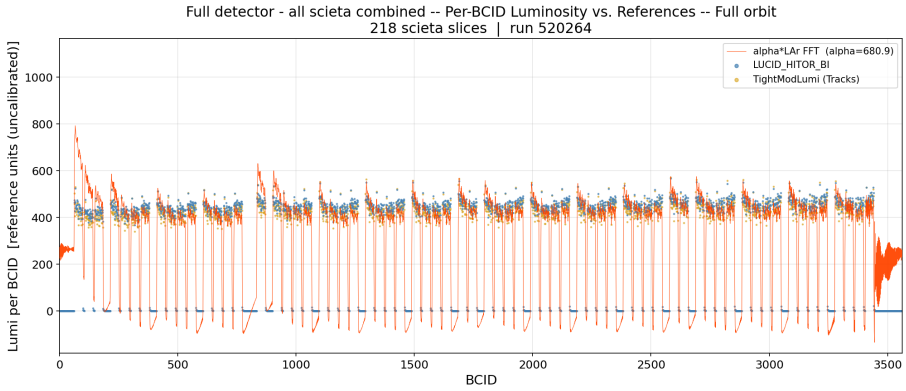

Figure 9.15: The estimated luminosity using the uncoinstrained method as a function of BCID scaled and compared to the corresponding measured values of LUCID (blue) and track counting (yellow) for the full orbit in run 520264.

ence luminosity measurements over the full orbit, as shown in Figure 9.15: the overshoot is predominant in the first train, and a residual undershoot persists in between bunches throughout the orbit. For trains in the middle and at the end of the orbit, however, the measurement is fairly stable and comparable to LUCID and track counting.

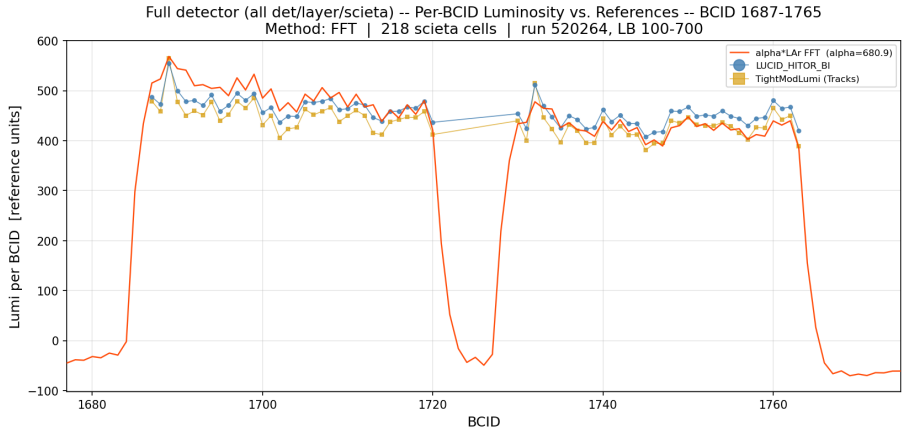

Figure 9.16: A detailed comparison between the estimated luminosity using the deconvolved LAr samples pedestal and the corresponding LUCID and track-counting luminosities per bunch for a BCID range in the middle of the orbit.

Nevertheless, looking at a detailed plot of the estimated luminosity for a short BCID range in the middle of the orbit, as shown in Figure 9.16, one can confirm the efficacy of the method: the luminosity is reconstructed well when compared to LUCID and track counting, exhibiting strong sensitivity per bunch.

The artifacts shown earlier most likely originate from the choice of kernel, but another cause could be the formulation itself, which offers a free parameter at every BCID

of the orbit and therefore reconstructs noise into crossings that carry no bunch at all.

This could be solved by removing those free parameters, which is possible because the LHC filling scheme is known: which BCIDs can carry a bunch, and which BCIDs were not read out at all, are both known offline. Two pieces of prior information can therefore be added to the inversion, effectively applying a *constrained* version of the method. First, a *column restriction*: the luminosity is zero at every BCID that carries no bunch. Then, a *row deletion*: the BCIDs for which no sample was recorded carry no measurement, and the values standing in for them should not be allowed to constrain the fit. A pedestal term is fitted alongside, as a single explicit scalar. This is what is meant by *constraining* the deconvolution in the following, and it can be done for either of the two formulations, though not in the same way.

For the inverse matrix the constraints are exact and immediate, because both are operations on the matrix of Eq. (9.2): this means dropping columns, and that the rows of unmeasured BCIDs are deleted from the system. The luminosity is then zero by construction at the empty BCIDs. The normal equations of the reduced system are formed and solved directly, exactly as in Eq. (9.7), so this is the same estimator evaluated on the constrained operator. It is referred to as the *constrained inverse matrix* in the following.

For the FFT the same two constraints cannot be written down in closed form, as the discrete Fourier transform of Eq. (9.4) is defined on the complete periodic sequence. The procedure consequently becomes *iterative*: the frequency-domain division of Eq. (9.6) and the constraints in the BCID domain are applied in alternation, and the cycle is repeated until the solution stabilizes. This estimator is referred to as the *constrained* FFT in the following.

The effect of the reconstruction with the constrained FFT method is shown in Figure 9.18: the sensitivity is retained, and the negative luminosities in the inter-bunch intervals are eliminated, while the small overshoots at the start of trains remain. Constraining the inversion is therefore what removes the dominant artifact of the unconstrained result, though the overshoot at the start of each train remains unsolved, most likely originating from the choice of kernel. And while the results look good, this is only possible by setting constraints offline knowing the BCID structure.

With the measured template pulse shape as kernel, however, the constrained inverse matrix does not share this improvement, as shown in Figure 9.17: it is zero by construction at the empty BCIDs, but develops an oscillation from bunch to bunch *within* the trains, together with negative drops at the end of the orbit, and its residual against LUCID is the worst of any estimator considered here (Table 9.2).

Overall, deconvolving the luminosity from the mean pedestal of the ADC samples of the Digital Trigger therefore shows remarkable promise, though with the residual artifacts still remaining. Since those artifacts can be attributed to the pollution from empty BCIDs, and the assumed pulse shape, the natural next step is to replace the

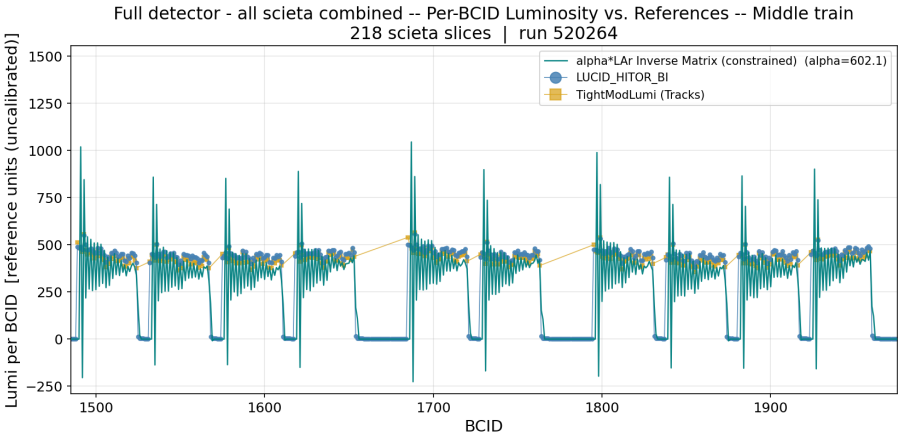

Figure 9.17: The estimated luminosity using the constrained inverse matrix with the 2025 template pulse shape, for a train in the middle of the orbit. The constraints remove the inter-bunch undershoot, but with this kernel the reconstruction develops a marked bunch-to-bunch oscillation within the train.

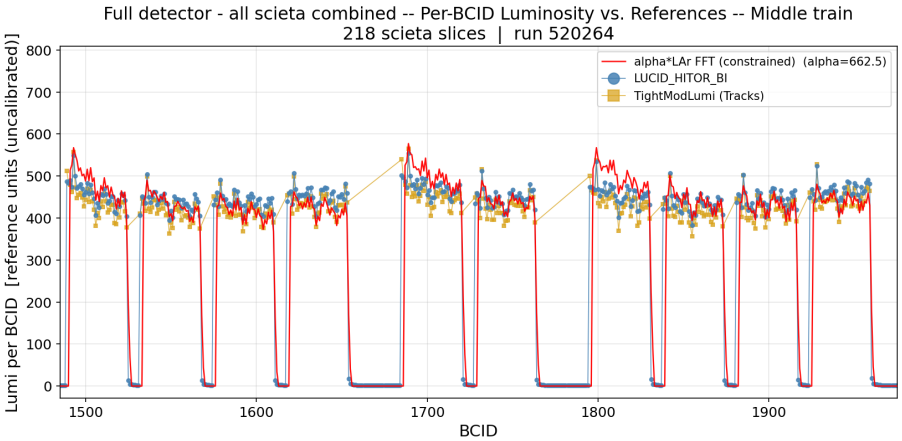

Figure 9.18: The estimated luminosity using the constrained FFT method for a BCID range in the middle of the orbit. The constraints seem to resolve the undershoots in the inter-bunch BCIDs while retaining the small overshoots at the beginning of the trains but also the strong sensitivity.

template by a kernel obtained from the data themselves, while utilizing the knowledge of the BCID scheme to constrain the methods and finally probe the maximal capabilities of the methods. This is done in Section 9.5.

9.5 Kernel extraction with LUCID

In order to solve the issues with the choice of pulse shape highlighted in the previous section, and to establish whether the problem actually lies with the choice of kernel or with the deconvolution method itself, a new, data-driven kernel was extracted from the data using the per-BCID luminosity measured by LUCID.

The idea to extract an "ideal" kernel is to reverse the roles of the two factors in the convolution of Eq. (9.1). In the deconvolution, the kernel h is assumed known and the per-bunch luminosity $\mathbf{L}$ is the unknown. Here, the luminosity is instead taken as known and set to the LUCID per-BCID orbit L^{ref}, averaged over the luminosity blocks of the run, while the kernel become the unknown, and retrieved for each scieta-bin. An example of such extracted kernel is shown in Figure 9.19.

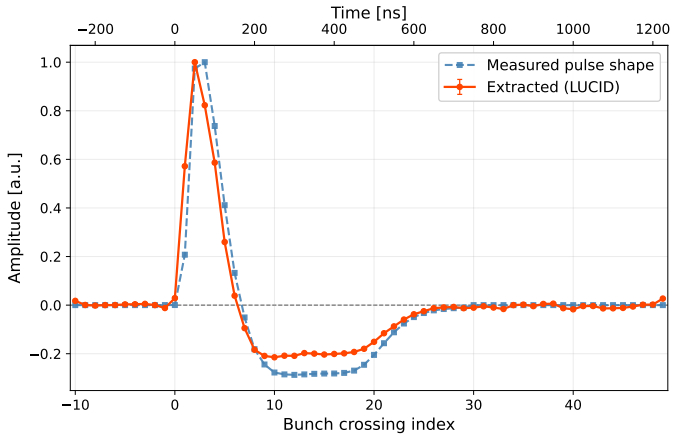

Figure 9.19: The pulse shape extracted with LUCID superimposed over the "template" measured pulse shape.

The same convention to compare with LUCID and track counting was set up and the results for a train in the middle of the orbit are shown in Figure 9.20. With the extracted kernel both constrained estimators track the references closely and reconstruct the inter-train BCIDs with no negative undershoot, in contrast with the template result of Figure 9.14. The constrained inverse matrix is the more uniform of the two across the train, while the constrained FFT retains a small negative offset with respect to the references.

The stability of the constrained inverse matrix is evidenced in Figure 9.21, where the luminosity is almost perfectly reconstructed throughout the full orbit. The constrained

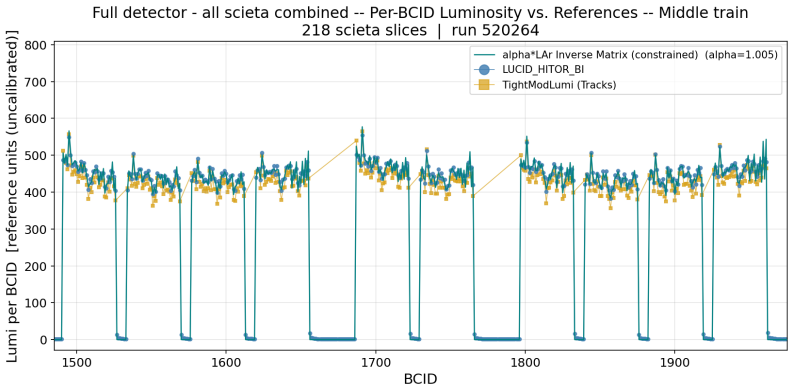

(a)

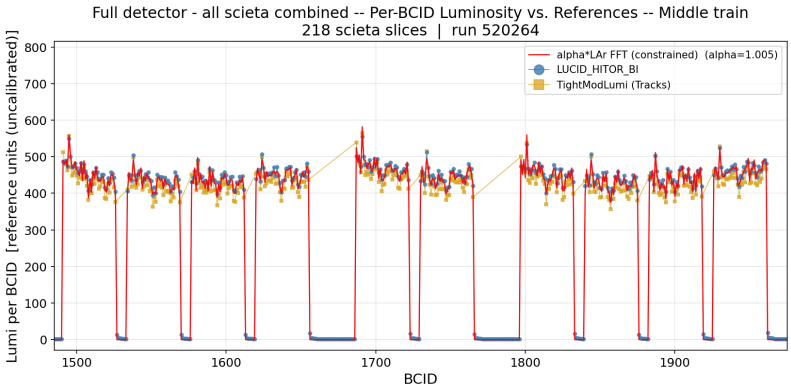

(b)

Figure 9.20: The estimated luminosity using the kernel extracted from LUCID as a function of BCID scaled and compared to the corresponding measured values of LUCID (blue) and track counting (yellow) for a train in the middle of the orbit, for the constrained inverse matrix (a) and the constrained FFT (b). Both constrained estimators now follow the references closely.

FFT reproduces the same trains well, but retains a residual overshoot in the first trains, of the kind already seen with the measured template pulse shape, and exhibits a slow negative trend toward the end of the orbit.

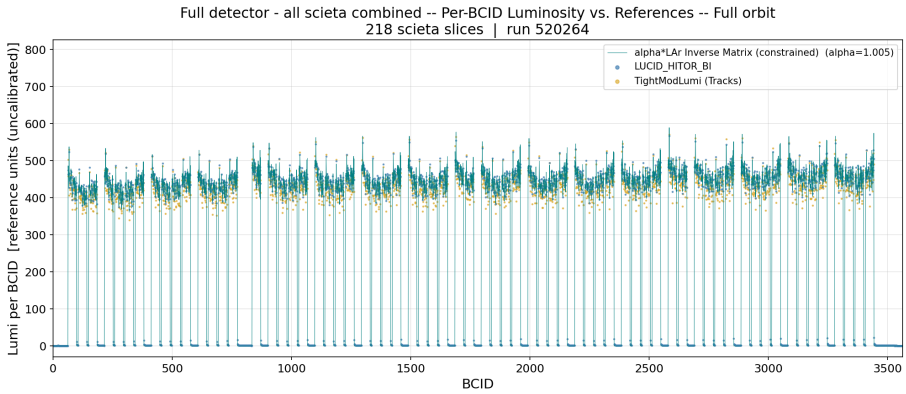

(a)

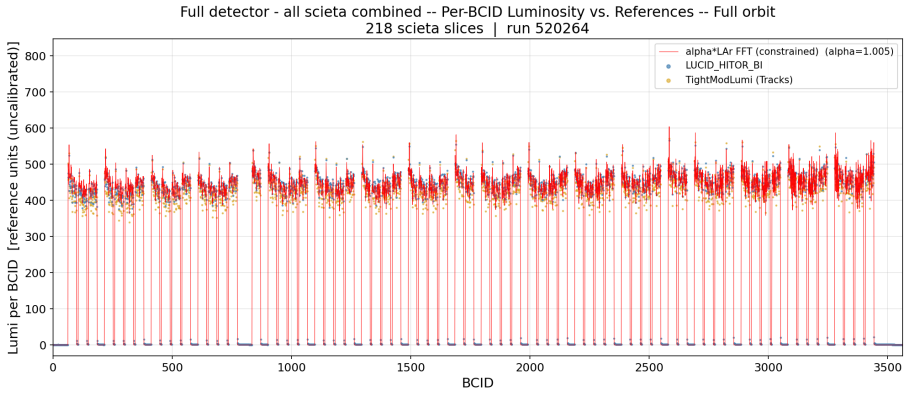

(b)

Figure 9.21: The estimated luminosity using the LUCID-extracted kernel for the constrained inverse matrix (a) and the constrained FFT (b) as a function of BCID scaled and compared to the corresponding measured values of LUCID (blue) and track counting (yellow) for the full orbit in run 520264.

The promising results of this method are further shown in Figure 9.22, where the residuals versus LUCID are plotted for all BCIDs in run 520264, giving a mean of -0.10% and an RMS of 3.4% for the constrained inverse matrix, and a mean of -0.099% and an RMS of 3.1% for the constrained FFT. The two are equivalent to within a few per mille in RMS, but they are not equivalent in the *shape* of the residual: that of the constrained inverse matrix is flat across the whole orbit, whereas that of the constrained FFT is visibly wider at the two ends of the orbit and tighter in the middle.

Table 9.2 shows the residuals against LUCID for the two kernels and the estimators considered. The unconstrained entry is a single row, since the inverse matrix and the FFT are essentially the same estimator. Imposing the constraints on the FFT improves the

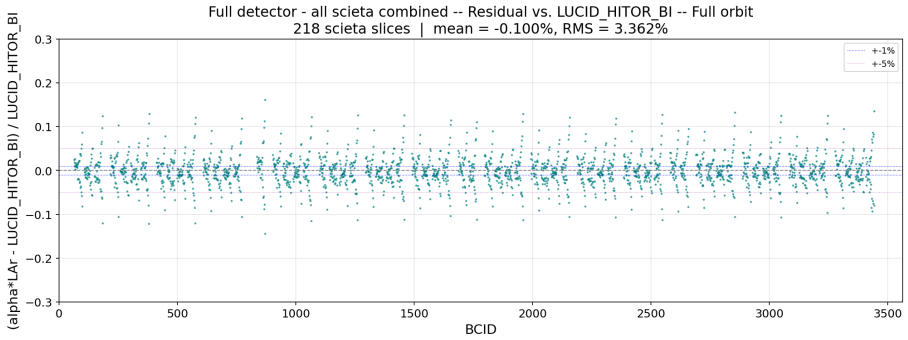

(a)

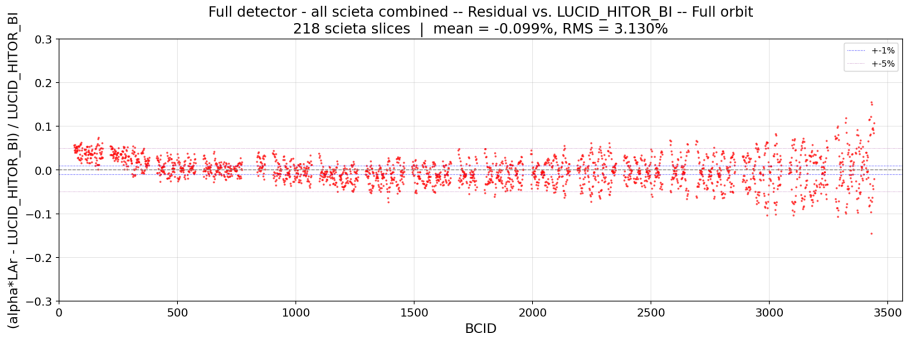

(b)

Figure 9.22: The relative residual $(\alpha\,\mathcal{L}_{\mathrm{LAr}} - \mathcal{L}_{\mathrm{LUCID}})/\mathcal{L}_{\mathrm{LUCID}}$ of the constrained inverse matrix (a) and the constrained FFT (b) with the LUCID-extracted kernel, as a function of BCID, for all filled BCIDs of the full orbit in run 520264. The full detector is used, averaging 218 scieta slices. The dotted lines mark the $\pm 1\%$ and $\pm 5\%$ levels; the mean and RMS of the distribution are given in the title.

template result by a third, from an RMS of 11% to 7.5%, whereas replacing the template by the extracted kernel brings both constrained estimators close to 3%, the constrained inverse matrix in particular from 34% down to 3.4%. The unconstrained inversion is the only entry that does not benefit from the better kernel, its dominant error coming from the treatment of the unmeasured BCIDs rather than from the assumed pulse shape. The kernel, and not the inversion algorithm, is therefore what limits the measurement.

It is with the LUCID-extracted kernel that the constraints improve the inverse matrix method. Once the kernel is extracted from the data its oscillations and negative dips are removed, and the residual falls from 19% to 3.4%.

Table 9.2: Relative residual $(\alpha\,\mathcal{L}_{\mathrm{LAr}} - \mathcal{L}_{\mathrm{LUCID}})/\mathcal{L}_{\mathrm{LUCID}}$ with respect to the LUCID per-bunch luminosity, using the 2025 template pulse shape and the kernel extracted per scieta bin from LUCID. All the filled BCIDs of the full orbit of run 520264 are used, averaging 218 scieta slices.

	2025 template		LUCID-extracted	
Method	Mean [%]	RMS [%]	Mean [%]	RMS [%]
Unconstrained[1]	−1.102	11.073	−3.030	19.476
FFT (constrained)	−0.663	7.452	−0.099	3.130
Inverse matrix (constrained)	−15.322	34.407	−0.100	3.362

This latest study demonstrated the importance of the choice of kernel. Here, a "perfect" kernel is extracted from an already known target luminosity, which shows that both constrained methods perform well when given an ideal kernel: the constrained FFT reaches an RMS of 3.1% against LUCID and the constrained inverse matrix 3.4%, against 19% for the unconstrained inversion on the same kernel. Taking the two together, the constrained inverse matrix is adopted as the best-performing reconstruction of this chapter. The two methods are separated by only a few per mille of RMS, but the constrained inverse matrix achieves it with a residual that is uniform over the whole orbit, whereas the residual of the constrained FFT starts degrading at the orbit ends where its constraint is imposed only approximately. More work is expected to refine the methods to remove the artifacts seen, and most of all a system to obtain optimal pulse shapes that can be used for an effective reconstruction of the luminosity per bunch.

A promising direction to achieve this is to *learn* the inverse of the convolution instead of inverting it analytically. A deep neural network can be trained to map the measured mean-pedestal orbit directly onto the per-bunch luminosity, using simulated orbits of the kind described in the toy Monte Carlo study earlier in this chapter as training data, where the true per-bunch luminosity is known by construction, and then validated on

[1] A single row, since the unconstrained inverse matrix and the unconstrained FFT are the same estimator on the forward operator and agree to numerical precision.

data against the per-BCID luminosity measured by LUCID. The advantage of this approach is that the training set need not be generated with a single fixed pulse shape, sinces by sampling the training orbits across a distribution of pulse-shape variations, the network learns a reconstruction that is robust to the mismatch between the measured template and the pulse shape actually produced under the running conditions of the data.

Chapter 10

LATOME Pulse Rate Monitor as a luminometer

The two techniques presented so far both exploit the digitized ADC samples that the Digital Trigger records for a subset of events: the dispersion of the samples in Chapter 8, and their mean pedestal per BCID in Chapter 9. Both were shown to achieve promising results, the latter reproducing the per-bunch reference luminosity once a kernel appropriate to the data is used, but both are constrained by ATLAS trigger decisions: the samples are available only for events accepted by a trigger and written to a specific stream, so the statistics and the BCID sampling profile of the measurement are inherited from the trigger menu and from the prescale of that stream, as explained in Section 8.1. On top of that, the RMS method was observed to degrade where the pileup contribution to the samples becomes small compared with the electronic noise, which is, together with the dependence on the trigger, the origin of the likely loss of sensitivity at low luminosity described in Section 8.2.1 and of the near-complete loss of correlation at vdM scales reported in Section 8.3.10. The per-bunch method carries a limitation of its own, namely its dependence on the accuracy of the assumed pulse shape, as established in Section 9.5.

This chapter explores a different observable computed by the DT, namely the rate at which the LATOME firmware reconstructs pulses above a configurable energy threshold in each Super Cell. This quantity is computed continuously, for every bunch crossing of every orbit, independently of any trigger decision, and it is published by a dedicated monitoring function of the Digital Trigger back-end, the PRM. The chapter first describes what the PRM is and where it sits in the Phase-I readout architecture, then establishes the principle by which its pulse rate can be related to the instantaneous luminosity, and finally presents the study of its performance as a luminometer.

10.1 Monitoring event rates with the PRM

One of the main systems used to monitor and investigate issues in the Super Cells introduced by the Phase-I upgrade of LAr (see Section 6.2) is the Pulse Rate Monitor. The PRM is an operational tool monitoring the rate at which a given SC is producing pulses, highlighting whether a channel is dead, masked, mistimed, or affected by the pedestal issues discussed in Section A.5, usually exhibiting an anomalous rate.

The PRM is implemented in the LATOME firmware in the back-end of the LAr

Phase-I system, described in Section 6.2.3. It monitors the rate of digitized pulses per SC whose reconstructed energy exceeds a threshold defined per SC and per run. A graphical interface for the online monitoring of these rates was also implemented, the LArWebApp (also known as *LArSoup*), visualizing them per SC for each layer of the LAr calorimeter. An example of the rates monitored during a Stable Beams proton–proton run can be seen in Figure 10.1.

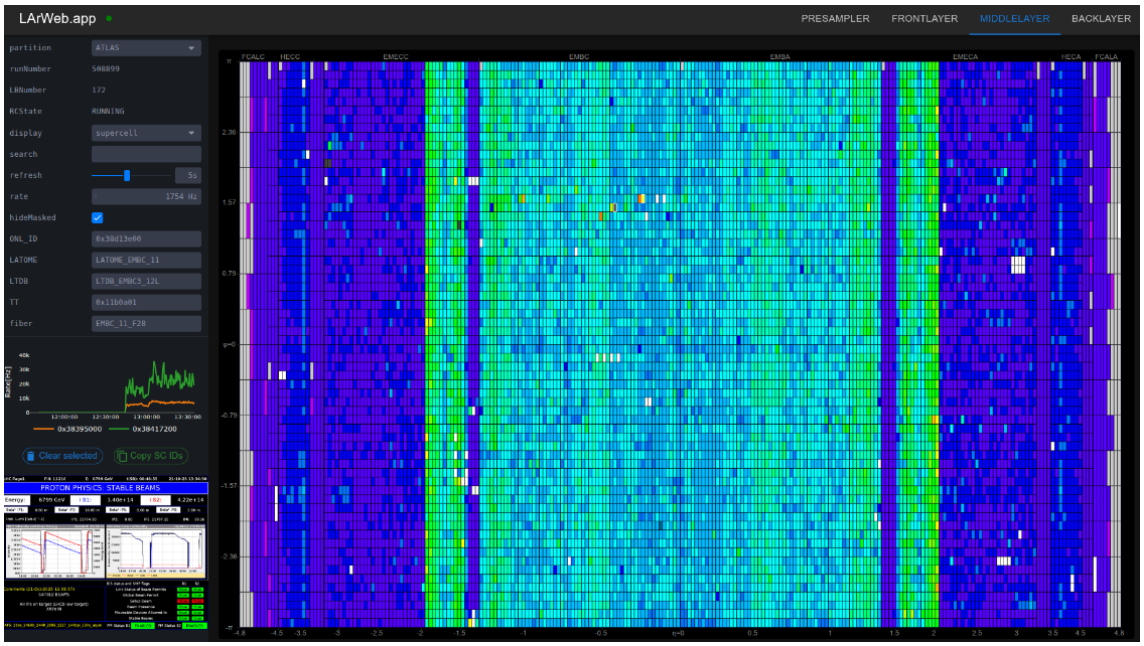

Figure 10.1: A screenshot of the LArWebApp for the monitoring of SC rates [121].

10.1.1 Derivation of the pulse rate in the LATOME firmware

As described in Section 6.2.3, every 25 ns the LATOME firmware applies the optimal filter of Eq. (6.1) to the stream of ADC samples of each SC, obtaining the transverse energy E_{T} and the signal timing t from the samples S_i after subtraction of the electronic pedestal p and of the BCID-dependent baseline correction δp_{BCID}. For each SC j, a counter in the firmware is incremented whenever the reconstructed transverse energy assigned to a bunch crossing satisfies

$$E_{\mathrm{T},j} \; > \; E_{\mathrm{T},j}^{\mathrm{thr}}, \tag{10.1}$$

where the threshold $E_{\mathrm{T},j}^{\mathrm{thr}}$ is configured per SC and per run, as twice the electronic noise of the SC, $E_{\mathrm{T},j}^{\mathrm{thr}} = 2\,\sigma_j$, with σ_j derived for each run from minimum bias data. The energy is computed at 40 MHz, but due to coverage limits in the current firmware, the SCs are split in 6 groups and 54 SCs are processed at a time. The counter in the LATOME is reset every 2 s, and given the iteration through the 6 SCs groups, it leads to an measurement

every 12 s for each SC. The counter contents are then published to the ATLAS online information service, from which they are archived by P-BEAST, the Persistent Back-End for the ATLAS Information System of TDAQ [118]. The pulse rate of SC j is then

$$r_j = \frac{N_j}{\Delta t}, \tag{10.2}$$

with N_j the number of counts accumulated over the readout interval Δt. The same archived values retrieved through P-BEAST are the input to the study presented in the remainder of this chapter, as anticipated in Section 8.1.

The LATOME processes every bunch crossing of every orbit for every SC, so the PRM rate is not subject to the trigger selection, to the prescales, or to the BCID sampling profile that shape the sample-based datasets of Section 8.1. Furthermore, it does not feed the E_{T} stream transmitted to the eFEX, jFEX and gFEX, it plays no role in the L1 trigger decision, and it is not part of the event data recorded through the TDAQ readout chain. Hence, the PRM carries no per-event information whatsoever.

10.2 PRM as a luminometer

Recall Section 7.2.1: since the bunch luminosity is proportional to the mean number of inelastic interactions per bunch crossing through Eq. (7.5), any observable that scales with μ can serve as a luminosity estimator. The pulse rates in the Digital Trigger are such an observable. The classification of Section 7.2.1 places the PRM in the *hit-counting* category, as it counts the number of electronic channels whose reconstructed energy exceeds a specified threshold. Hence, the PRM can function as a hit-counting relative luminometer built out of the Digital Trigger.

As performed for the RMS method, the relation between the measured rate and the luminosity in the context of this study is established empirically by fitting the rate against the reference ATLAS online luminosity run by run,

$$r = k\,\mathcal{L} + r_0, \tag{10.3}$$

of the same form as Eq. (8.1), and inverted as in Eq. (8.2) to obtain the luminosity estimate.

As done in the previous chapters for the RMS method, a study was performed to probe the capabilities, and the weaknesses, of the PRM as a luminosity measurement method. First, the method is demonstrated in nominal proton–proton conditions in Section 10.3 and its stability over time at such conditions, a key requirement for a relative luminometer discussed in Section 7.2, is assessed in Section 10.7. Then, its dependence to layer and the scieta coordinate is studied and used to optimize the measurement, in Section 10.4. Finally, the sensitivity of the PRM at vdM luminosity scales, where the RMS method clearly failed, is explored in Section 10.6. It will be seen how the layer-scieta

optimization proves to be critical at the vdM regime.

10.3 Nominal pp physics run

To first establish if the PRM works as a luminometer, it is necessary to demonstrate the linearity with data in controlled, nominal conditions. Run 499844, a nominal run from the 2025 proton–proton campaign, is chosen for this analysis.

The easiest, and coarser, way to demonstrate this is to compute the mean SC rate per LB, and plot it against the ATLAS online luminosity. Figure 10.2 shows exactly this, exhibiting strong correlation and a rate-predicted luminosity very close to the online luminosity. Figure 10.3(a) confirms the correlation, showing a profile of the strongly linear correlation between the Super Cell rate and the online luminosity, while in (b) the corresponding residuals are plotted, resulting in a 0.83% relative resolution to the luminosity measurement.

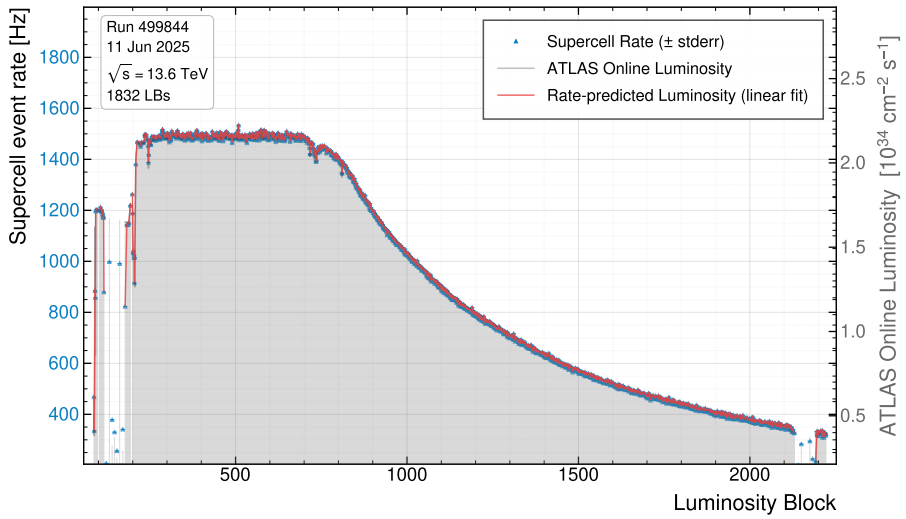

Figure 10.2: The full detector-averaged SC rate (blue dots) and the ATLAS online luminosity (gray area) plotted as a function of LB for run 499844. The luminosity estimated by fitting the SC rate is outlined in red.

One thing worth of notice is the first and last LBs of the run: here, emittance scans are present, but, contrary to what is done for the RMS-based method, they are considered for the measurement. The PRM shows certain sensitivity toward emittance scans and very low luminosities, though still not optimally, as seen in Figure 10.4. The main reason can be found in the layer and η dependence of the measurement: as it is soon shown in Section 10.4, the PRM exhibits different scales of SC rate in each layer of the detector, which renders the averaging between SC rates detrimental. This leads therefore to different, localized, sensitivities that cannot be resolved together. Studying this dependence is critical for the measurement of vdM scan, as explained in Section 10.6.

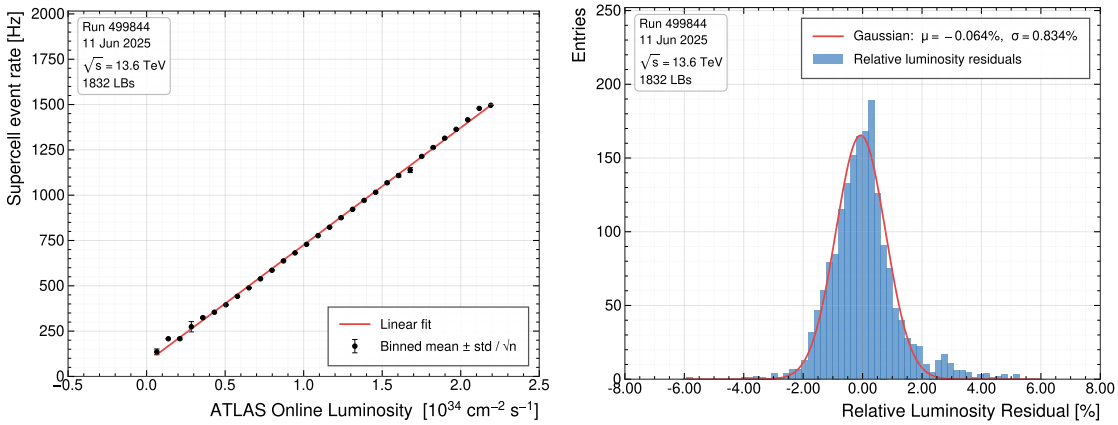

(a) (b)

Figure 10.3: Profile plot of the correlation between the SC rate and the online luminosity, with the respective fit overlayed on top in red (a) and the resulting relative residuals distribution plotted in (b). The Gaussian mean and resolution are shown in the plot.

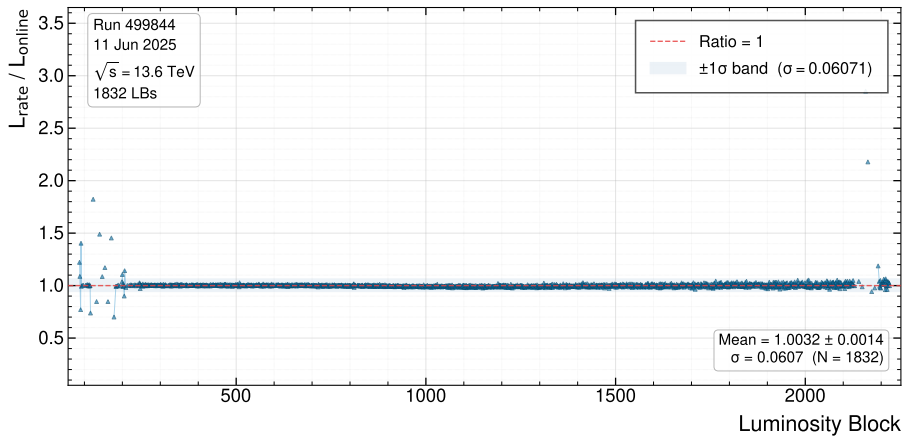

Figure 10.4: The ratio between the predicted luminosity using the PRM and the online luminosity as a function of LB.

10.4 Layer and η-coordinate optimization

In order to study the difference in sensitivity of each layer to the luminosity, the SC rate is correlated to the online luminosity, and the resulting relative resolution of the residuals is computed per layer. This is shown in Figure 10.5. EMB middle proves to be the best-performing layer, while EMEC back gives the worst measurement resolution. This can also be seen in the detailed plots in Figure 10.6, where the rate is superimposed over the online luminosity as a function of LB. Here, one can clearly notice how the measurement from the EMEC back layer is substantially more noisy than the one from EMB middle.

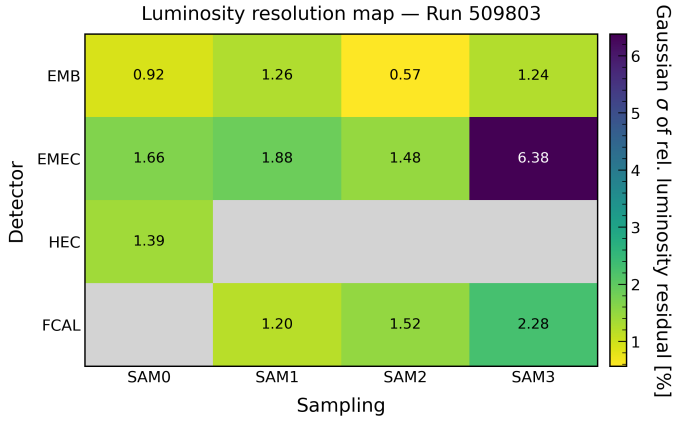

Figure 10.5: A heatmap showing the resulting relative resolution of the measurement for each subetector/layer combination. Slightly better results are seen for EMB, while EMEC back layer shows the worst resolution.

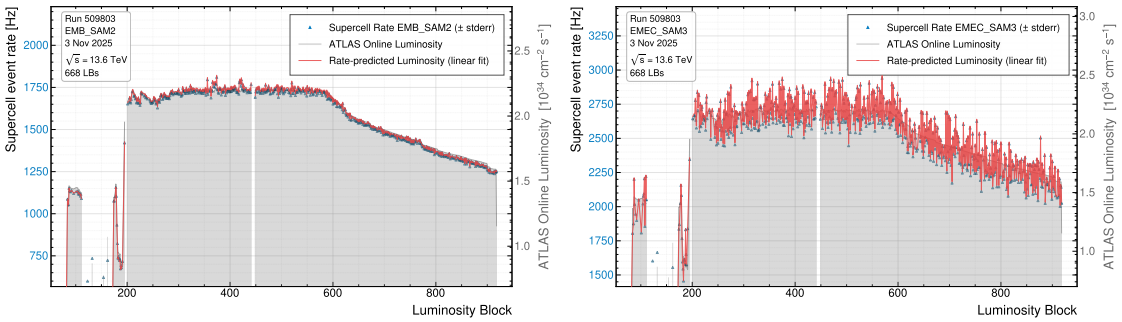

(a) (b)

Figure 10.6: The SC rate (blue dots) and the ATLAS online luminosity (gray area) plotted as a function of LB for run 509803 for the EMB middle (a) and EMEC back layer (b). The luminosity estimated by fitting the SC rate is outlined in red.

More than the noisy measurement per se, the real difference is the scales of the rates. The SCs of each layer appear to have extremely diverse scales when it comes to pulse

rates, as evidenced in detail in Figure 10.7. Here, the mean SC rate of each layer is plotted as a function of online luminosity in run 509803. The scales go from the order of 10 Hz for the presampler layers and EMB front, up to 10^3 Hz for other layers, which clearly demonstrates that averaging the measurements over the whole detector can in many cases be detrimental.

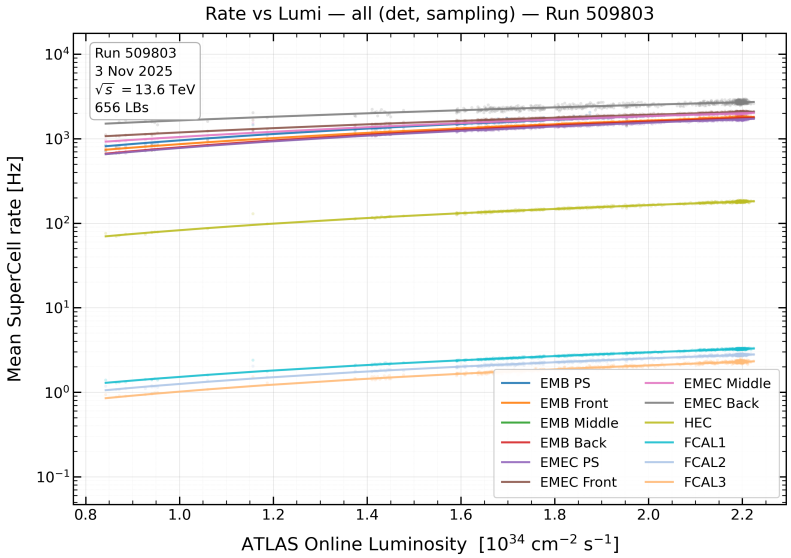

Figure 10.7: The mean Super Cell rate for each layer as a function of the online luminosity for run 509803. It is clear that the vastly different scales going from the order of 10 Hz for the presampler layers and EMB front, to 10^2 Hz for HEC, up to 10^3 Hz the remaining layers, for can make the averaging detrimental for the measurement.

10.5 High-μ run

As described earlier, the PRM is in the LATOME firware. As such, it will be inherited by the LAr system in Run 4, and during the ongoing LS3, new functionalities or improvements can be added to optimize this method for Run 4 luminosity measurements. For this reason, it is paramount to investigate the performance of the PRM in conditions similar to those of the HL-LHC. Run 509891, already studied in Section 8.4.2, is a run recorded in 2025, with conditions resembling the HL-LHC, namely reaching an average $\langle\mu\rangle = 120.6$.

When performing the measurement averaged over the whole detector, the results in Figure 10.8 shows how, while there is some correlation, the measurement is quite noisy. The degraded result originates from the averaging process, due to highly diverse scales in the rates, and low-performing layers. This is confirmed in Figure 10.9, where the resulting relative resolution of the residual distribution of each layer is plotted, showing clearly the EMEC back layer degrading the overall measurement.

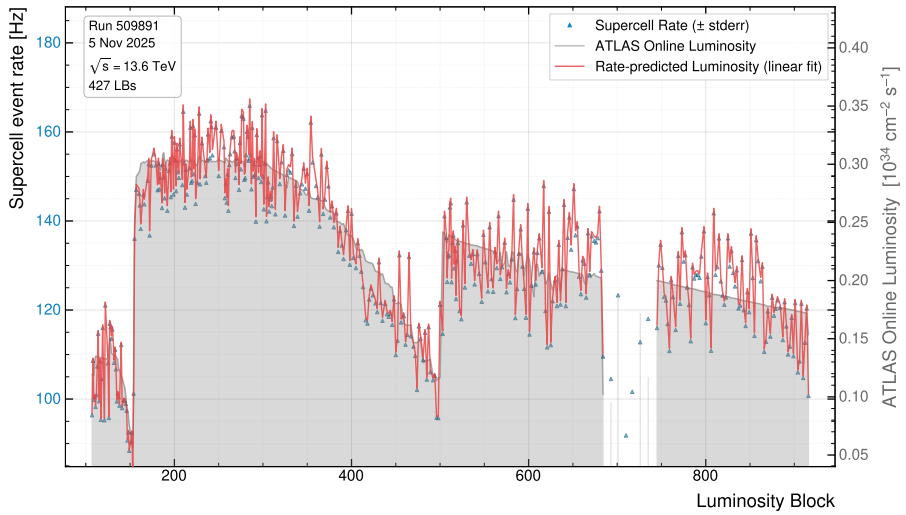

Figure 10.8: The SC rate (blue dots) and the ATLAS online luminosity (gray area) plotted as a function of LB for run 509891. The luminosity estimated by fitting the SC rate is outlined in red.

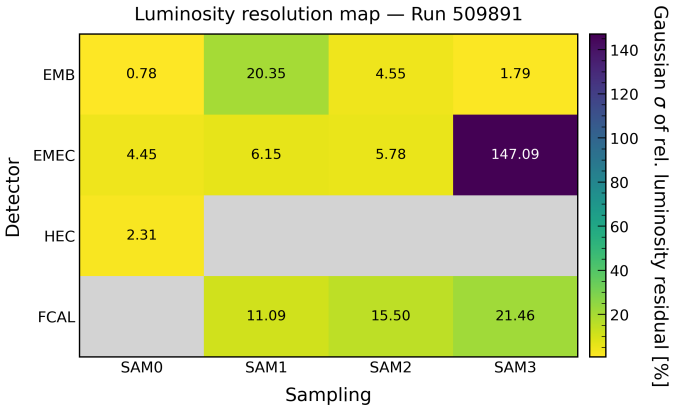

Figure 10.9: A heatmap showing the resulting relative resolution of the measurement for each subetector/layer combination. Slightly better results are seen for EMB, while EMEC back layer shows the worst resolution.

Instead, if the measurement of the luminosity is performed only with the best performing layer, in this case the EMB presampler, the results are massively improved, as seen in Figure 10.10. Here, the correlation is clearly improved, and resulting in a relative resolution of 0.79%, hence demonstrating the capabilities of the PRM as a luminometer also at HL-LHC conditions. However, for next developments, it will be paramount to carefully optimize the energy thresholds of the PRM to reach the best possible sensitivity.

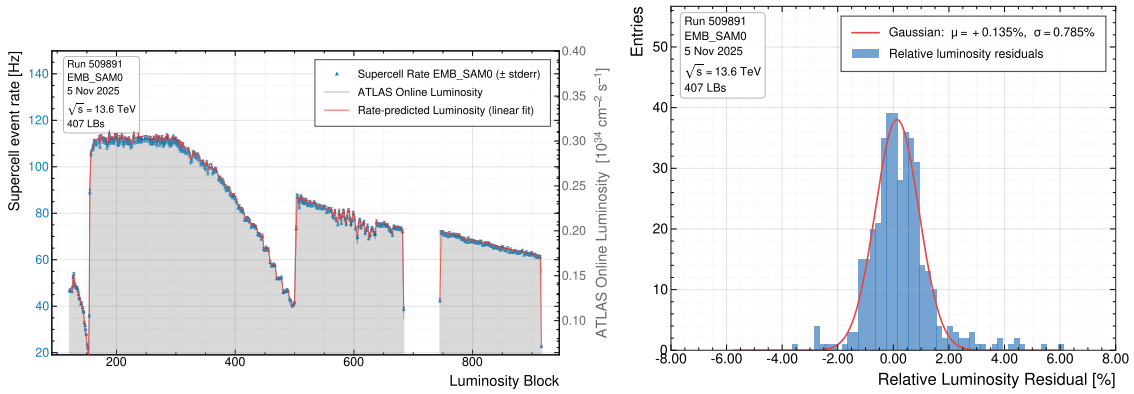

(a) (b)

Figure 10.10: The SC rate (blue dots) and the ATLAS online luminosity (gray area) plotted as a function of LB for run 509891 (a), and the resulting distribution of the residuals, showing a relative resolution of 0.79% (b). Both plots are produced using only pulse rates from the EMB presampler.

10.6 Sensitivity to vdM scans

The most promising use of the PRM is to probe what was not possible to reach with the RMS-metod, namely vdM scans. As described in Section 8.3.10, the RMS of the DT samples is completely insensitive to vdM scans luminosities.

To optimally probe vdM scans with the PRM, it is first necessary to study the layer and η-coordinate dependence of the measurement. It was seen in Section 10.4 that the scales of the rates vary widely between different parts of the detector, which would make a measurement of vdM averaged over the full detector useless. This is once again confirmed in Figure 10.11, where, also at vdM luminosities, the scale of the SC rates varies strongly among layers.

In order to further optimize the measurement, the measurement is repeated for each scieta-coordinate in each layer, and the corresponding relative resolution is computed. The best performing scieta-coordinates are plotted in Figure 10.12. This only gives a qualitative, coarse idea of which scieta-coordinates to target to achieve a better mea-

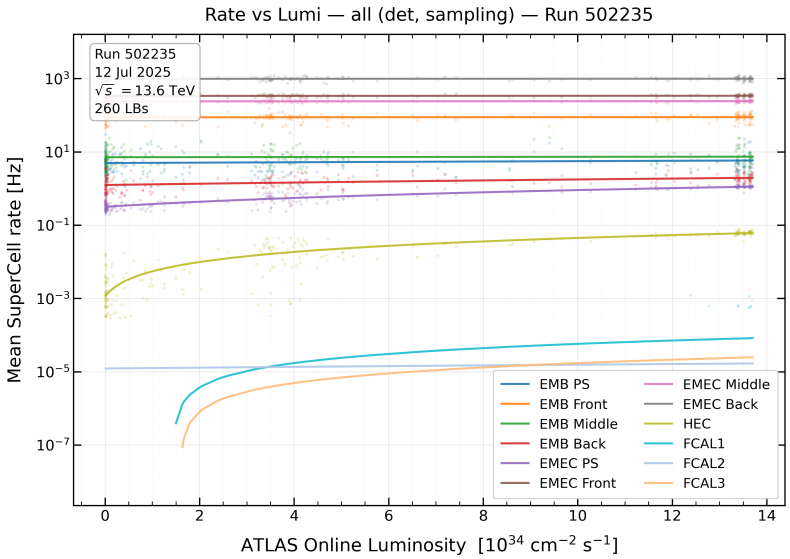

Figure 10.11: The mean Super Cell rate for each layer as a function of the online luminosity for the 2025 vdM scan in run 502235. Also for this run, the scales are vastly different in each layer.

surement, showing that scieta 29 in the EMEC middle layer gives the most optimal measurement.

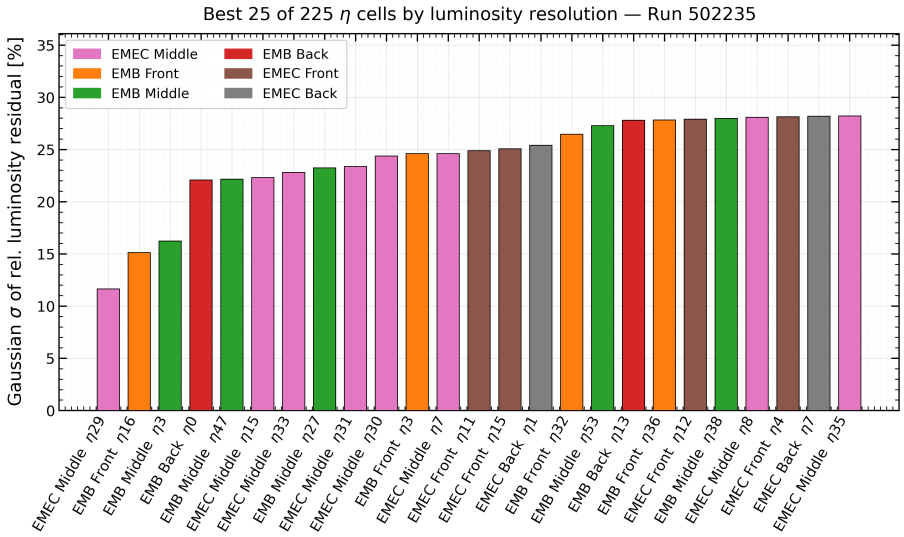

Figure 10.12: The resulting relative resolution of residuals for the best-performing scieta-coordinates. Scieta 29 in the EMEC middle layer appears to be the scieta-coordinate with the best measurement.

Once scieta-29 in the EMEC middle layer is chosen as the reference coordinate for the vdM scan measurement, the SC rates are correlated and fit against the online luminosity,

and plotted in Figure 10.13. The plot shows strong sensitivity to the luminosity in the vdM, something that could not be achieved with the RMS method. The linearity between the two is further confirmed in Figure 10.14, where a profile plot of the correlation between the SC rate and the online luminosity, with the corresponding linear fit, is shown.

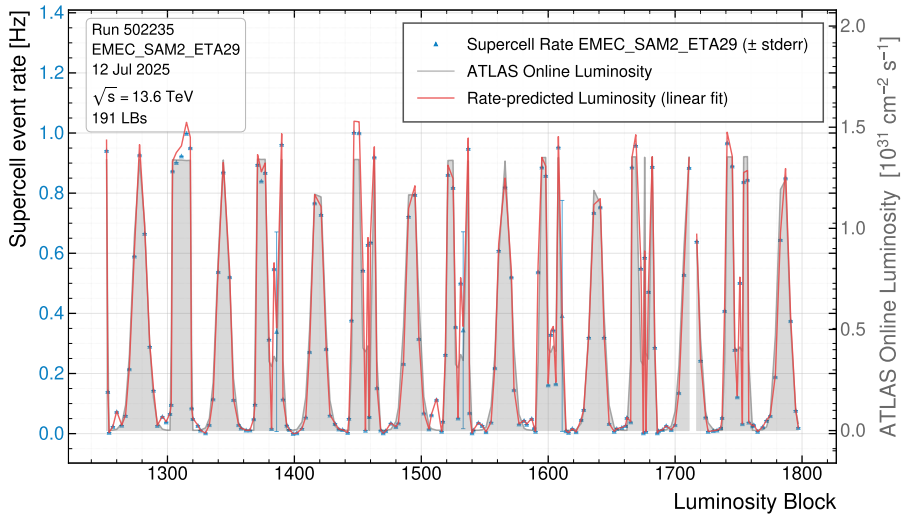

Figure 10.13: The SC rate (blue dots) and the ATLAS online luminosity (gray area) plotted as a function of LB for a selected LB range of the 2025 vdM scan (run 502235). The luminosity estimated by fitting the SC rate is outlined in red, and appears to be strongly correlated to the online luminosity.

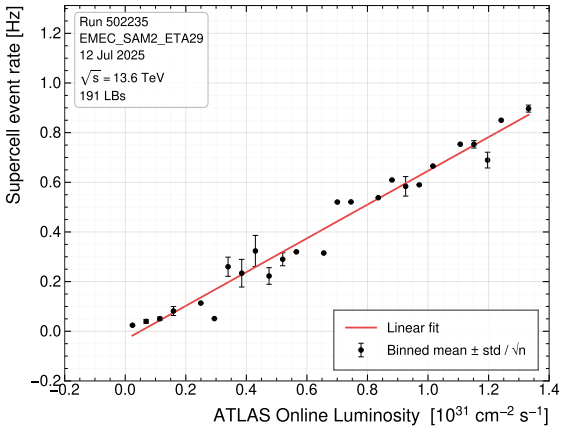

Figure 10.14: Profile plot of the correlation between the SC rate and the online luminosity for the 2025 vdM scan in run 502235. The corresponding linear fit is overlayed in red.

10.7 Long-term stability

While it has been demonstrated so far that the PRM shows strong correlation to the ATLAS online luminosity for the most diverse LHC conditions, it is necessary to analyze its stability over time. The same measurement, using only data from the EMB middle layer, as it was shown to be the best performing layer in nominal pp conditions (see Figure 10.5), is performed on all available runs from the 2025 GRL, using the linear fit obtained from 499912 as calibration and applied then to all other runs.

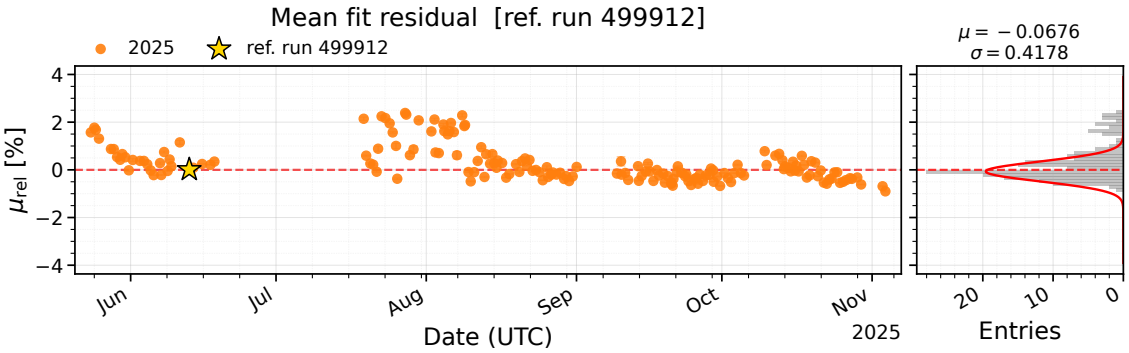

Figure 10.15: The relative mean of the residual distribution for all available GRL runs from 2025, computed using the coefficients from linear fit of run 499912 as calibration, each dot representing one run. The resulting distribution displayed on the right panel is centered around zero, but clearly showing different populations due to LUCID-2 issues in the beginning of 2025 data-taking, and due to LAr issues in August.

Figure 10.15 shows the relative mean of the residual distribution using this method for all analyzed runs, showing a fair stability over time, as the mean of most runs is mostly concentrated within a 1% range. What is striking is also the presence of a smaller populations at around 2%, when analyzing the resulting distribution. While the deviations in the beginning of the data-taking were related to LUCID-2 issues, the rest could likely originate from changes in the LATOME firmware occurring during the 2025 *pp* data-taking period, but also the continuous change of rate threshold happening run by run.

Part IV

Search for Long-Lived Particles

Chapter 11

Search for long-lived particles with a delayed-jet signature

As described in Chapter 2, SUSY is one of the most common BSM theories searched at the LHC. The relevant models presented in this thesis are long-lived supersymmetric particles introduced in Section 2.3.

Standard physics analyses in ATLAS rely on reconstruction techniques intended to reconstruct particles produced directly from the pp interaction point. For searches targeting LLPs, where particles travel for considerable distances in the detector with decay products not originating from the interaction point, or even particles that traverse the whole detector without decaying, non-standard reconstruction techniques are needed.

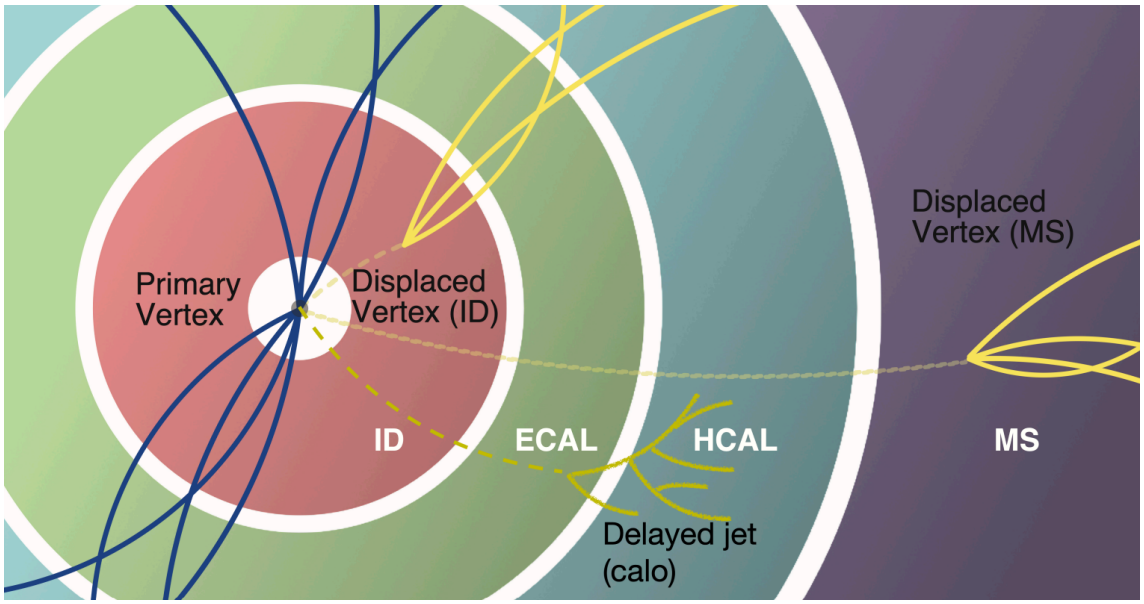

Figure 11.1: A schematic illustration of an LLP decay with a displaced signatures in a typical detector. Adaptation from [16].

Many LLP analyses are carried out at LHC experiments, targeting a diverse array of detector signatures, as shown in Figure 11.1. Examples include searches for charged LLPs decaying in tracking detectors and resulting in disappearing tracks [122], or searches for LLPs producing displaced vertices or tracks away from the interaction point [123]. Another targeted signature is delayed calorimeter deposits, such as those left by displaced photons [124, 125] or by jets, as illustrated schematically in Figure 11.1.

This chapter presents the search for long-lived particles decaying with a *delayed-jet* signature in the calorimeters, providing an introduction to the analysis and describing the data and simulated event samples used. At the time of writing the analysis is still ongoing, so the overall status of the search is presented. The work I carried out on the truth studies and on the definition of the signal regions is then described in more detail in Sections 11.6 and 11.7.

11.1 Analysis overview

As stated previously, many LLP searches exploit the spatial displacement of the particle's decay. These searches therefore rely on charged-particle tracking, which leads to a considerable loss of sensitivity for decays occurring beyond the Inner Detector, where a displaced vertex can be reconstructed, at $R \approx 300\,\text{mm}$ [123]. The limits of the Inner Detector extend to $R \approx 1.1\,\text{m}$ and is the boundary drawn in Figure 11.5. When investigating hadronically decaying LLPs, this proves to be a issue, which can be solved by targeting, instead, the jets produced by decays occurring in the calorimeters.

Hence, the search presented in this thesis targets this specific region of the ATLAS detector by exploiting a different observable, namely the arrival time of energy deposits. As introduced in Section 2.3, a massive LLP travels with $\beta < 1$, leading to a delayed arrival time (see Eq. (2.6)). This delay originates from the reduced velocity of the LLP compared to prompt particles and from the increased path length, as the LLP travels from the interaction point to a displaced vertex and subsequently decaying and leaving a deposit in a calorimeter cell. This clearly leads to a offset in time with respect to the nominal bunch crossing time where the particle was produced. Hence, the targeted signature is a jet built from ordinary calorimeter jet constituents arriving with a delay in the order of a few ns. This proves to be a largely unexploited observable, while at the same time complementary to tracking-based searches.

A set of dedicated HLT trigger chains selecting on jet time were implemented with this task, and operated from July 2024 until the end of Run 3. More details are presented in Section 11.4.

The signature targeted by the analysis is hadronically decaying LLPs, heavy and slow enough to produce a measurable delay. The targeted masses can be of the order of $\mathcal{O}(100\,\text{GeV})$ to a few TeV with lifetimes of $\tau \sim 1\,\text{ns}$ upwards. However, the dedicated triggers of Section 11.4 accept jets delayed by up to 15 ns. On the other hand, the samples simulated for this analysis cover a parameter space of $m_{\tilde{\chi}_1^0} = 1000$–$2000\,\text{GeV}$ and $\tau = 1$–$30\,\text{ns}$, chosen to target the existing Run 2 displace-vertex exclusion limits described in Section 11.2.2, that only reach about 1.6 TeV near $\tau = 0.1$ ns, and weaken at longer lifetimes [123].

The main signal model consists of direct electroweak production of long-lived neutralinos decaying to three jets via RPV coupling, a model that was introduced in Sec-

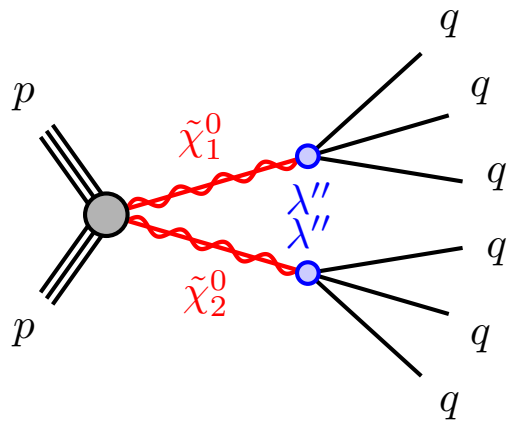

Figure 11.2: The Feynman diagram of electroweakly produced RPV neutralinos decaying to jets. In this analysis, the samples are normalized to the summed higgsino cross sections of Table 11.1.

tion 2.3. The corresponding Feynman diagram is shown in Figure 11.2. Two further signal models are targeted, although their sample production is not yet mature at the time of writing: strong production of the same long-lived neutralinos (mainly via gluon–gluon fusion channel), and heavy higgs decaying to a pair of long-lived scalars.

As it will be described later, the signal yields for this model are expected to be small. Hence, the analysis is designed to reject background down to the level of a few events, as already done in various other LLP searches. Four background sources are considered, all estimated from data. These backgrounds can be divided in two groups:

- **Non-collision background (NCB):** includes cosmic muons and beam-induced background (BIB).
- **Collision background:** includes jets from satellite collisions and mistimed SM jets.

This search uses the estimation of these background for a two-stage analysis: a preselection is applied to centrally reconstructed data using the background estimation, and the surviving events are then re-reconstructed with additional cell-level information, which are normally not kept for jets. Eventually, the final selection is done on these events (the procedure is also known as *event-picking*), using timing quantities not available in the standard data format.

The analysis is blinded, though one run was unblinded for background studies, and the overall background estimations are still ongoing, hence a only brief description of their development is presented in Section 11.5. Studies of the truth record and the definition of the signal regions (SRs) are described in Sections 11.6 and 11.7.

11.2 Data and simulated event samples

11.2.1 Collision data

The analysis uses proton–proton collision data recorded with the ATLAS detector in Run 3 at $\sqrt{s} = 13.6\,\text{TeV}$. The dedicated delayed-jet HLT triggers have been in the physics menu since July 2024, collecting data until the end of LHC operations in July 2026, amounting to an approximate integrated luminosity of $210\,\text{fb}^{-1}$. The final integrated luminosity and its related uncertainty (and Good Run List) is not yet finalized at the time of writing. The dataset is blinded, and all selection and background-estimation studies described here are performed on signal-free samples or on a single unblinded run, run 482028, recorded in 2024.

11.2.2 Simulated signal samples

The signal samples used in this analysis were generated with MadGraph5_aMC@NLO [126] using the `RPVMSSM_UFO` model [127], and showered with Pythia8 [128] using the A14 tune [129] and the NNPDF2.3LO parton distribution functions [130].

As mentioned previously, the benchmark process studied in the analysis is the electroweak production of a neutralino pair, $pp \to \tilde{\chi}_1^0\tilde{\chi}_2^0$, with each one decaying to three quarks via the R-parity-violating coupling λ'' of Eq. (2.4). A single mass parameter $m_{\tilde{\chi}_1^0}$ labels each mass point, and since the coupling is small, both neutralinos are long-lived. The two LLPs are produced back to back, each decaying to a three-jet system.

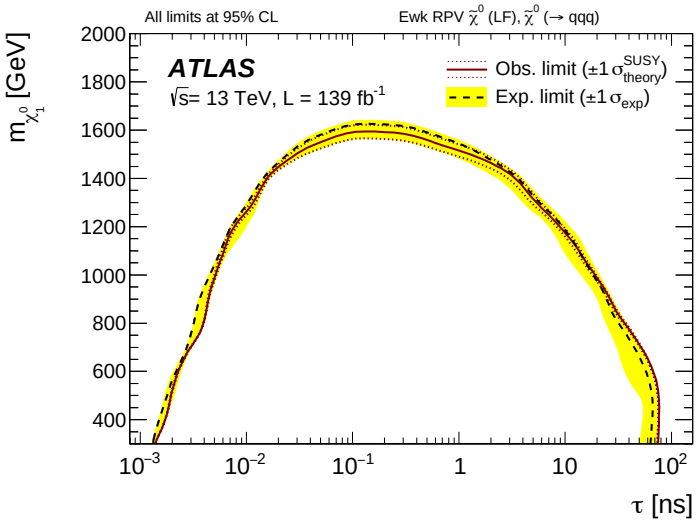

Figure 11.3: Observed and expected 95% confidence-level exclusion limits on the lightest neutralino mass as a function of its proper lifetime, for the electroweak RPV model. Taken from [123].

An extended grid of mass and lifetime points was produced, and the derived datasets cover eleven mass points from 1000 to 2000 GeV in steps of 100 GeV, each with lifetimes of 1, 3, 10 and 30 ns.

The grid targets the exclusion limits set in Figure 11.3, that reach their maximum of approximately 1.6 TeV at around $\tau = 0.1$ ns and then falls with increasing lifetime. The extension of the grid to $\tau = 30$ ns therefore has the potential to explore the space that Run 2 vertex-based searches did not reach [123]. Furthermore, masses below approximately 1 TeV are not included in the grid, as already within the exclusion limits set.

The signal samples are normalized to the total production cross section of the degenerate higgsino states, summed over the four channels $\tilde{\chi}_1^0\tilde{\chi}_1^+$, $\tilde{\chi}_1^+\tilde{\chi}_1^-$, $\tilde{\chi}_1^0\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_1^-$. With the chargino and the second neutralino only about 1 GeV above the lightest neutralino, every channel decays to it with soft, largely unobserved products and therefore yields the same experimental signature of two displaced three-jet decays, which is why the four are summed rather than only the $\tilde{\chi}_1^0\tilde{\chi}_2^0$ channel used. The values are calculated by the LHC SUSY Cross Section Working Group [131] for degenerate higgsino production at $\sqrt{s} = 13.6$ TeV, computed at $\text{NNLO}_{\text{approx}}$ plus next-to-next-to-leading-logarithmic (NNLL) resummation with Resummino 3.1.1 and the `PDF4LHC21_40` set, and are listed in Table 11.1. Given the small values of the cross section, shown in Figure 11.4, signal yields are expected to reach up to about 86 events at the most favorable grid point for the targeted luminosity in this search. This will be shown in Section 11.7, which further confirms that the suppression of the background is paramount for the sensitivity of the analysis.

Unless stated otherwise, the distributions and plots shown in the following sections are for a single mass–lifetime point chosen as a reference benchmark $m_{\tilde{\chi}_1^0} = 1000$ GeV, $\tau = 3$ ns, chosen to validate the truth studies and the signal regions.

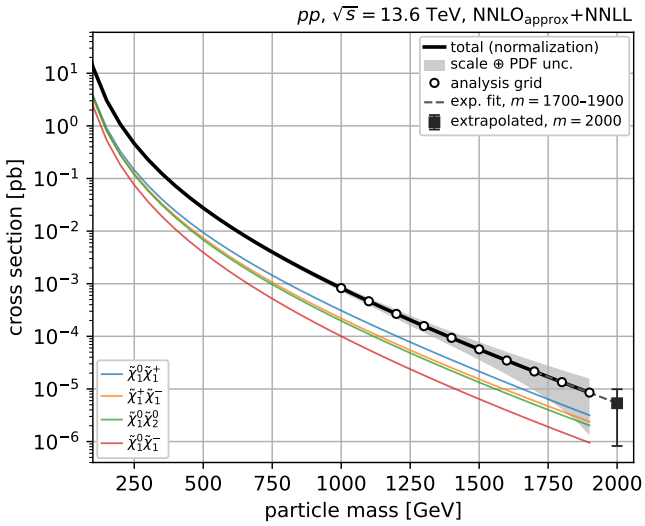

Figure 11.4: Production cross sections of the degenerate higgsino states at $\sqrt{s} = 13.6$ TeV, for the four production channels and for their sum, computed at NNLO$_{\text{approx}}$+NNLL by the LHC SUSY Cross Section Working Group [131]. Calculated values are available up to 1900 GeV; the dashed line is an exponential fit to the total over 1700–1900 GeV, used to extrapolate to the 2000 GeV grid point (filled marker), the only one for which no calculation exists. The signal samples are normalized to the total.

Table 11.1: Summed production cross sections of the four degenerate higgsino channels ($\tilde{\chi}_1^0\tilde{\chi}_1^+$, $\tilde{\chi}_1^+\tilde{\chi}_1^-$, $\tilde{\chi}_1^0\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_1^-$) at $\sqrt{s} = 13.6$ TeV for the eleven mass points simulated in this analysi [131].

$m_{\tilde{\chi}_1^0}$ [GeV]	σ [pb]
1000	8.25×10^{-4}
1100	4.64×10^{-4}
1200	2.67×10^{-4}
1300	1.57×10^{-4}
1400	9.36×10^{-5}
1500	5.66×10^{-5}
1600	3.47×10^{-5}
1700	2.15×10^{-5}
1800	1.35×10^{-5}
1900	8.52×10^{-6}
2000	5.35×10^{-6}

11.3 Event reconstruction

Jets are reconstructed with the anti-k_t algorithm [132, 133] with radius parameter $R = 0.4$ from topological cell clusters calibrated at the electromagnetic scale (EMTopo jets). No displaced-specific reconstruction is used. Topological clusters are built from the cell energy significance and do not require tracks, so jets from LLP decays beyond the tracker are reconstructed like any other jet. The jet direction is defined from the primary vertex to the calorimeter deposit, which differs from the direction of the decay products as seen from the decay vertex (Section 11.6.3). The jet time, instead, is built up from cell-level information. In ATLAS, the time of a topological cluster is normally computed using the energy-squared-weighted average of its constituent cell times [134]:

$$\bar{t}_{\text{cluster}} = \frac{\sum_i \left(w^{\text{geo}}_{\text{cell},i} E_i\right)^2 t_{\text{cell},i}}{\sum_i \left(w^{\text{geo}}_{\text{cell},i} E_i\right)^2}, \tag{11.1}$$

where E_i and $t_{\text{cell},i}$ are the energy and the time-of-flight-subtracted time of cell i and $w^{\text{geo}}_{\text{cell},i}$ is its geometrical weight.

An alternative, resolution-weighted definition,

$$\bar{t}^{\text{res}}_{\text{cluster}} = \frac{\sum_i t_{\text{cell},i}/\sigma_i(E)^2}{\sum_i 1/\sigma_i(E)^2}, \tag{11.2}$$

gives a modest improvement in resolution and a mean cluster time better centered on the expected value of zero. Both definitions are carried through the analysis, and the choice of jet-time variable is currently open as the work on the backgrounds estimation is still ongoing.

The quantity that drives the jet-time resolution is the energy deposited in the individual cells: cell time and energy are reconstructed together by the optimal filtering algorithm of Eq. (5.7), whose variance on t scales inversely with the signal amplitude [90]. At jet level this cell-energy dependence translates into a resolution that improves with the jet transverse momentum, since more energetic jets deposit more energy per cell and in more cells. An absolute delay of a given size is therefore far more significant for a high-p_{T} jet than for a soft one, which makes a resolution-scaled variable more appropriate than the time in absolute terms. The jet-time resolution is thus parameterized as a function of p_{T} [135]

$$\sigma_t^{\text{jet}}(p_{\text{T}}) = e^{A+Bp_{\text{T}}} + C, \qquad A \approx -0.8, \quad B \approx -0.014\,\text{GeV}^{-1}, \quad C \approx 0.68\,\text{ns}, \tag{11.3}$$

and the delay significance is defined as

$$t/\sigma_t = \frac{t_{\text{jet}}}{\sigma_t^{\text{jet}}(p_{\text{T}})}. \tag{11.4}$$

11.4 Dedicated HLT triggers for delayed jets

The search relies on a set of HLT chains developed within this analysis effort and running since early July 2024. Two timing hypotheses are implemented on top of standard jet selections:

- `XtimingY` requires a jet with an absolute delay between X and Y nanoseconds;
- `XtimeSigY` requires a jet whose delay is at least X standard deviations, evaluated with Eq. (11.3), with Y the absolute maximum delay in nanoseconds.

The jet-time computation is inexpensive, so the chains add negligible CPU cost and, being based on the signal feature itself, achieve low rates. Representative rates measured on enhanced-bias data in 2024 are given in Table 11.2.

Table 11.2: Rates of the dedicated delayed-jet HLT chains, measured on enhanced-bias data in 2024.

Chain	Rate [Hz]
`HLT_2j45_2j55_3timeSig_L14jJ40`	2.7 ± 0.6
`HLT_2j100_2timeSig15_L1jJ90`	2.7 ± 0.6
`HLT_j220_j150_2timing15_L1jJ160`	0.6 ± 0.3
`HLT_j200_2timing15_L1jJ160`	2.1 ± 0.5
`HLT_j200_3timeSig15_L1jJ160`	1.5 ± 0.5
`HLT_j250_3timeSig15_L1jJ160`	0.3 ± 0.2
`HLT_j250_2timing15_L1jJ160`	0.6 ± 0.3

11.5 Backgrounds estimation

The estimation of the backgrounds is still ongoing at the writing of this thesis. For this reason, only a small introduction of the background sources is presented.

In the absence of a mechanism producing genuinely delayed jets, the Standard Model generally cannot populate the signal regions. Backgrounds therefore arise either from mismeasurements or from processes that are not pp collisions at the nominal bunch-crossing time. Four sources of background are identified:

Mistimed SM jets: Prompt jets generally peak at $t = 0$ with an intrinsic width set by the calorimeter timing resolution and by the spread of the collision time and vertex position along the beam. Reconstruction imperfections and residual out-of-time pileup generate non-Gaussian tails that extend into the targeted signal region.

Satellite collisions: Collisions between protons occupying non-nominal RF buckets, also known as *satellite* collisions, produce jets that are genuinely early or late. The dominant satellite population is estimated to arrive approximately 5 ns from the nominal crossing time. This background is irreducible but localized in time.

Cosmic muons: A cosmic muon undergoing a large energy loss in the calorimeter can seed a jet-like deposit. Its arrival time is entirely uncorrelated with the LHC clock, so it is uniformly distributed in time, and its rate is independent of beam conditions.

Beam-induced background: Particles arriving along the beam line, from beam halo interacting with collimators or from beam–gas interactions, shower in the calorimeters and can fake jets with a large apparent delay.

11.6 Truth studies

Before the signal samples can be used to define a selection, the truth record must be understood: which particles the generator actually writes out, where they decay, and how the resulting partons map onto reconstructed jets. This is critical step for rare signatures, such as the one investigated in this search, because the standard truth-jet collections are built assuming production at the interaction point, as it is described in Section 11.6.3. The truth studies start from the LLP itself to the reconstructed jets, so the decay kinematics of the neutralinos are examined first, followed by their decay products and the truth jets built from them, and finally the matching of these truth jets to reconstructed jets.

11.6.1 LLP decay kinematics

The starting point is the position of the LLP decay, which determines both which detector systems can observe it and the size of the expected delay. Figure 11.5(a) and (b) show where the neutralinos decay at $m_{\tilde{\lambda}_1^0} = 1000\,\text{GeV}$ for the four available lifetimes, in the transverse radius R_{decay} and along the beam axis in $|z_{\text{decay}}|$.

The limits of the Inner Detector, the electromagnetic and the hadronic calorimeter are overlaid on both plots. In both coordinates the distributions follow the expected behavior, noticing that at $\tau = 1\,\text{ns}$ the decays are concentrated close to the interaction point, inside the Inner Detector, while at $\tau = 30\,\text{ns}$ a long tail extends through the calorimeters and beyond. This clearly demonstrates how this search is complementary to other LLP searches, as at $\tau = 1\,\text{ns}$ the majority of decays occur inside the Inner

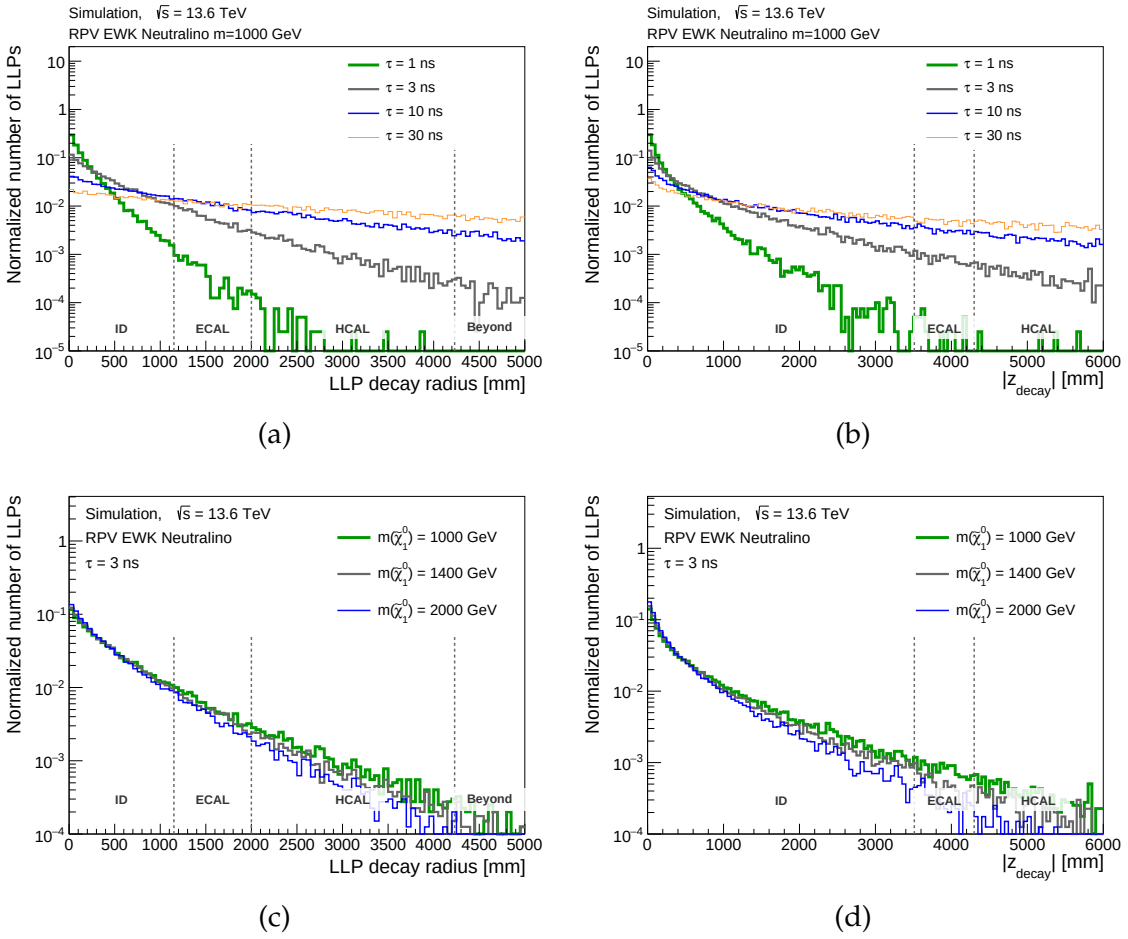

Figure 11.5: Position of the neutralino decay vertex on (a) transverse radius $R_{\text{decay}} = \sqrt{x^2 + y^2}$ and (b) longitudinal coordinates $|z_{\text{decay}}|$ at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$ for the four lifetimes. (c) Transverse radius and (d) longitudinal coordinate at the fixed lifetime $\tau = 3\,\text{ns}$ for the three mass points. The vertical dashed lines mark the nominal radial (a, c) and longitudinal (b, d) limits of the Inner Detector, the electromagnetic and the hadronic calorimeter.

Detector, the region covered by the displaced-vertices searches [123]. Instead, for the longer lifetimes, more LLPs decays falls within the calorimeter volume, with part of the $\tau = 30\,\text{ns}$ decays escaping beyond it.

The displacement of the decay is mainly dependent on the lifetime, rather than the neutralino mass. Figure 11.5(c) and (d) show the transverse decay radius and the longitudinal coordinate at the fixed lifetime $\tau = 3\,\text{ns}$ for three mass points, $m_{\tilde{\chi}_1^0} = 1000$, 1400 and 2000 GeV, with the same boundaries overlaid as in plot (a) and (b). In both coordinates the three distributions are almost indistinguishable across the Inner Detector volume, and differ only in the tails beyond the electromagnetic calorimeter, where the

heavier neutralinos decay marginally earlier likely due to the electroweak production leaving little kinetic energy for the boost of heavier LLPs. This demonstrates how the decay position is mainly dependent on τ, and the effect of the mass enters is mainly seen in the energy of the decay products.

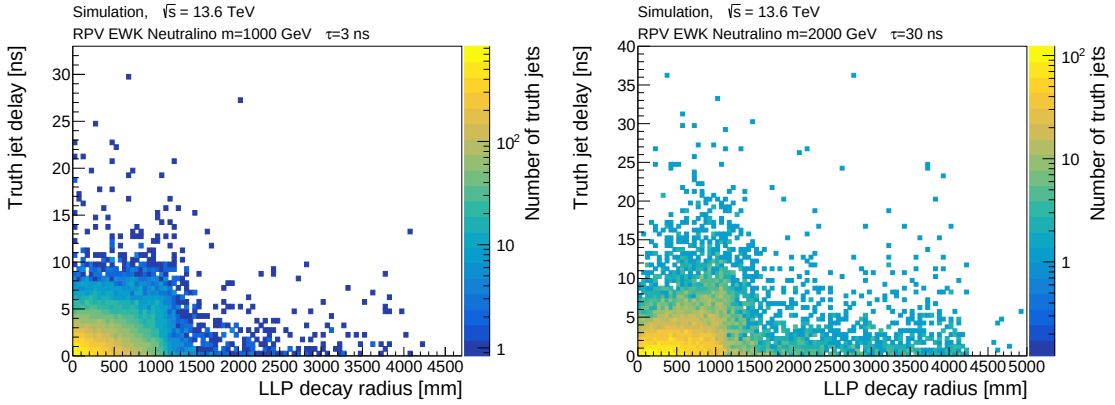

Figure 11.6: Truth jet delay as a function of the LLP decay radius, for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ grid point (left) and for a more extreme case, $m_{\tilde{\chi}_1^0} = 2000\,\text{GeV}$, $\tau = 30\,\text{ns}$ (right).

Figure 11.6 shows the truth-level jet delay as a function of the decay radius, for the reference mass point and for the heaviest, longest-lived point of the grid. The delay grows with the decay radius, and for decays within the calorimeter volume it reaches several nanoseconds, which are within the 1–15 ns jet time window used to select jets in the analysis.

11.6.2 LLP decay products and truth jets

Each neutralino decays to three quarks, which hadronize into a large number of final-state hadrons. Figure 11.7 shows the multiplicity of final-state hadrons per LLP decay. The distributions show how the heavier LLPs, as suggested earlier, produce more final-state hadrons compared to lighter particles.

Figure 11.8 shows the number of truth jets per LLP and the overall jet multiplicity per event. The truth-jet multiplicity peaks at three, as expected for the three-body decay, with lower values occurring most likely when the decay products are collimated enough to be clustered into fewer jets. The event-level jet multiplicity reflects the two LLP decays per event on top of initial-state radiation.

The p_T of the two jets that drive the trigger and the signal-region selection is shown in Figures 11.9 and 11.10. The transverse momentum of the leading and of the subleading reconstructed jet of the event are shown, for three mass points at fixed $\tau = 3\,\text{ns}$ and four lifetimes at fixed $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$.

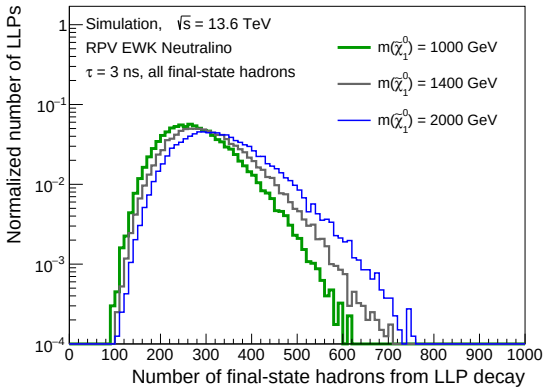

Figure 11.7: Multiplicity of final-state hadrons per LLP decay, for the three mass points at $\tau = 3\,\text{ns}$.

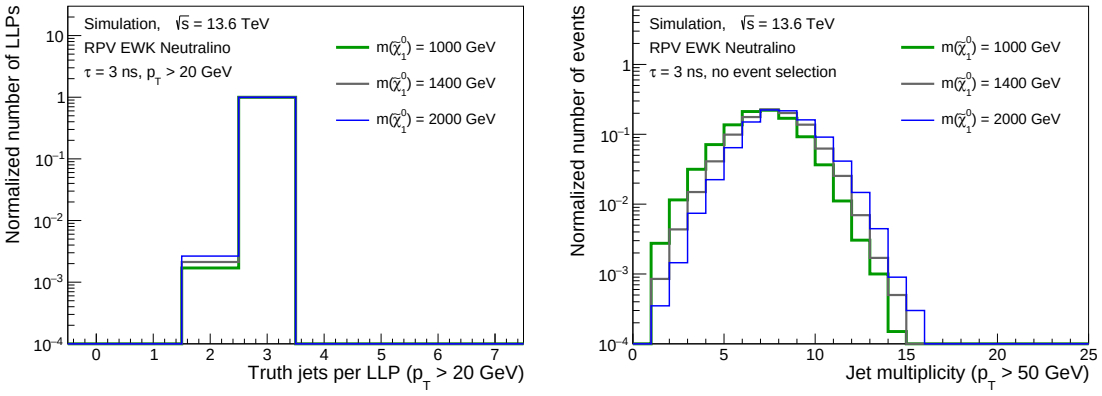

Figure 11.8: Number of truth jets with $p_\text{T} > 20\,\text{GeV}$ per LLP (left) and jet multiplicity per event for $p_\text{T} > 50\,\text{GeV}$ (right), for the three mass points at $\tau = 3\,\text{ns}$.

Across the mass scan at $\tau = 3\,\text{ns}$, the spectra harden as expected from the larger available energy, with the leading-jet distribution peaking near 400 GeV at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$ and moving to about 550 and 875 GeV at 1400 and 2000 GeV. The subleading jet follows the same pattern, peaking near 300, 425 and 650 GeV.

The lifetime scan at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, instead, exhibits an opposite behavior, as the peak of the leading-jet transverse momentum seem to decrease with lifetime. The shape of the distribution seems to changes as well. As seen at $\tau = 1\,\text{ns}$, the subleading spectrum has a peak around 400–500 GeV, while at $\tau = 30\,\text{ns}$ the distribution piles up in the 100–200 GeV region. This effect is also seen at jet-level seen for the matching-efficiency loss of Figure 11.14, showing how the many of longest-lived neutralinos decay beyond the calorimeters, so they never reconstructed into a jet.

The consequence for the selection is visible in the fraction of events above the $p_\text{T} \geq 350\,\text{GeV}$ threshold of SR1, as it will be described in Section 11.7.

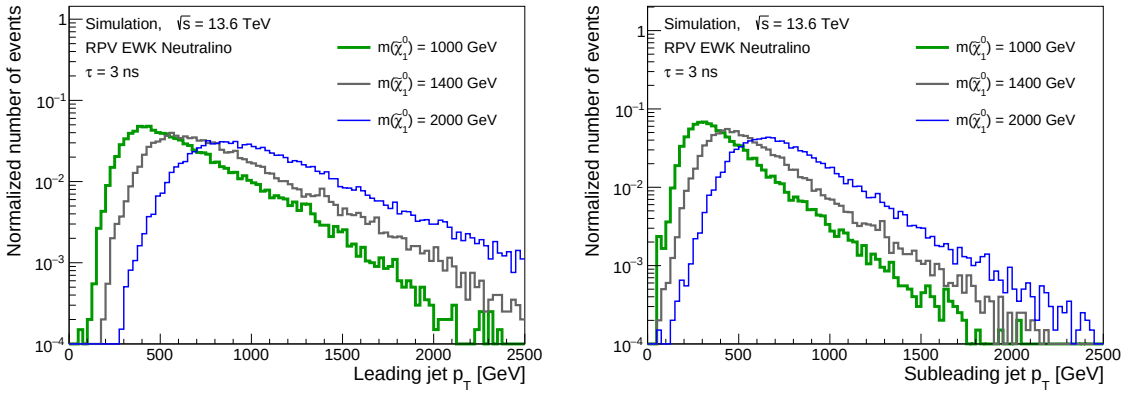

(a) (b)

Figure 11.9: Transverse momentum of the (a) leading and (b) subleading reconstructed jet of the event, for the three mass points at $\tau = 3\,\text{ns}$, from the jets with $p_\text{T} > 50\,\text{GeV}$. One entry per event, normalized to unit area.

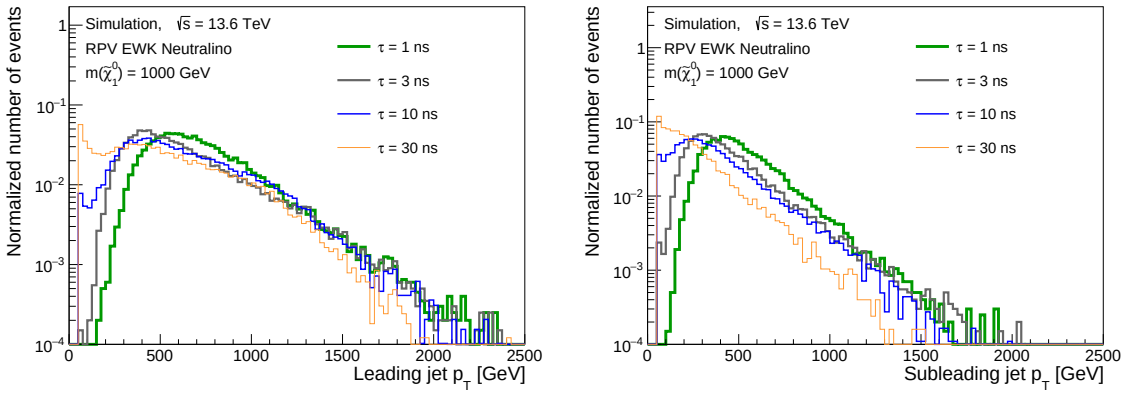

(a) (b)

Figure 11.10: Transverse momentum of the (a) leading and (b) subleading reconstructed jet of the event, for the four lifetimes at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, from the jets with $p_\text{T} > 50\,\text{GeV}$. One entry per event, normalized to unit area.

11.6.3 Truth jets in displaced decays

The standard truth-jet collection is not directly usable for this signature, because its jet direction is not the same as a reconstructed jet. This is addressed in three steps.

1. **Standard truth jets.** The standard collection (`truthjet_antikt4`) is clustered from the stable final-state particles, and the axis of each jet is the direction of the summed particle momenta.

2. **Reconstructed jets.** The axis of a reconstructed jet points from the primary vertex

(PV) to the energy deposit in the calorimeter. For a prompt jet these axes coincide. Instead, for a displaced decay, they do not: the decay products leave the decay vertex (DV), and not the PV, so the momentum direction no longer points back to the PV and the truth jet is displaced in the η–ϕ matching plane with respect to the deposit. This is shown schematically in Figure 11.11.

3. **Correction.** Custom truth jets are built per LLP using FastJet [133], by clustering only the final-state hadrons of each LLP with the anti-k_t algorithm, with a requirement of $R = 0.4$. The axis of each jet is propagated from the DV to the calorimeter reference surface, a cylinder of radius $R = 2.3\,\text{m}$ closed by planes at $|z| = 4.3\,\text{m}$, and (η, ϕ) is then recomputed as seen from the PV, the frame in which reconstructed jets are defined.

The effect of the correction is shown in Figure 11.12(a) for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ point: the mean ΔR_{min} between truth and reconstructed jets is reduced from 0.152 to 0.065, well inside the $\Delta R < 0.4$ matching cone used below. Figure 11.12(b) compares the corrected per-LLP truth jets with the standard `truthjet_antikt4` collection. The reconstructed-jet matching efficiency of the corrected per-LLP jets remains higher as the decay radius increases, especially in the region relevant for this search. The per-LLP truth jets are used as the reference for all the matching studies that follow.

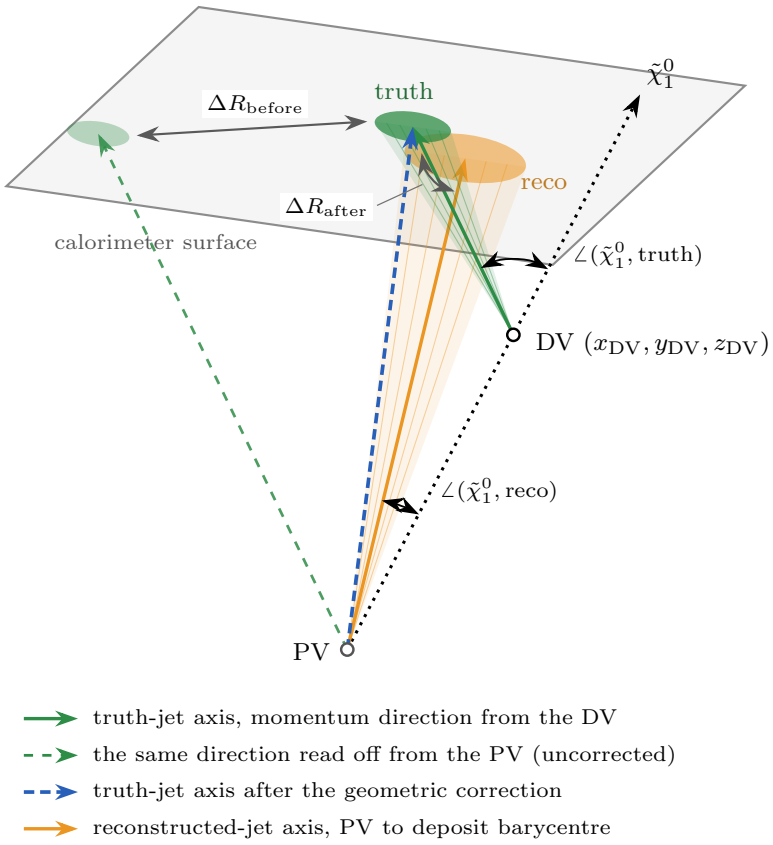

Figure 11.11: Geometric correction of the truth-jet axis for a displaced decay. The neutralino is produced at the primary vertex (PV), where the hard scatter occurs, and decays at the displaced vertex (DV). The per-LLP truth jet (green) has its axis defined by the momenta of the decay hadrons, i.e. as seen from the DV, while the reconstructed jet (orange) is defined from the PV directly to the calorimeter deposit. The uncorrected axis (dashed green), therefore places the truth jet far from the deposit in the η–ϕ matching plane. The correction thus propagates the truth axis from the DV to the calorimeter surface and re-expresses the direction from the PV (dashed blue), the frame in which reconstructed jets are defined. This leads to a reduced truth-to-reco angular separation from ΔR_{before} to ΔR_{after}.

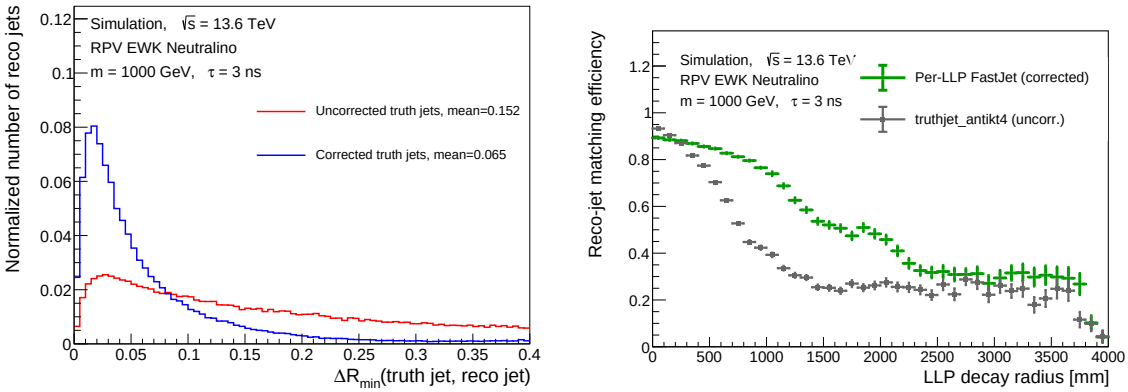

(a) (b)

Figure 11.12: Validation of the per-LLP truth-jet definition, for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ benchmark. (a) Minimum ΔR between truth and reconstructed jets before and after the geometric correction of the truth-jet axis. (b) Reconstructed-jet matching efficiency as a function of the LLP decay radius, for the corrected per-LLP truth jets and for the standard truth-jet collection.

11.6.4 Truth-to-reconstruction jet matching

Reconstructed jets carry no information about which neutralino, if any, produced them. Establishing that correspondence, or *matching*, is what allows the reconstruction of the signature to be characterized. For example, it helps defining the region of the detector in which a displaced decay is still reconstructible, and it separates efficiency lost to the decay geometry from efficiency lost to the selection thresholds.

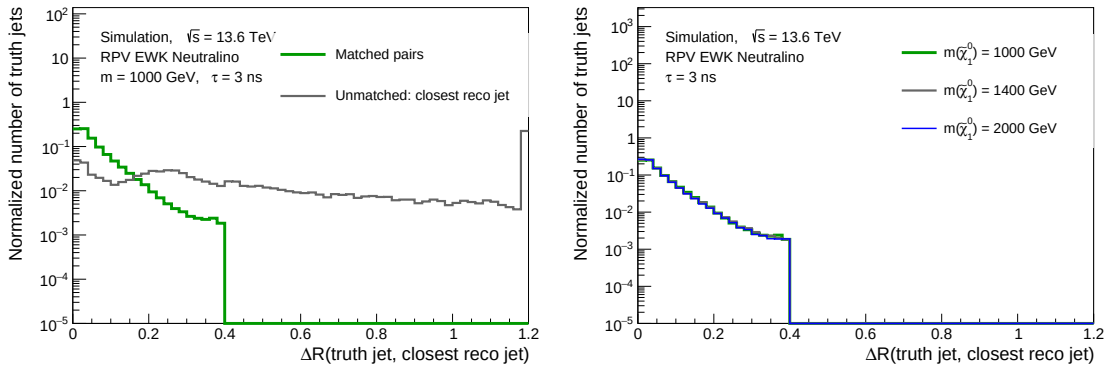

Figure 11.13: ΔR between each truth jet and the closest reconstructed jet, split into matched and unmatched truth jets for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ benchmark (left), and for matched pairs at the three mass points (right).

Truth matching is performed in the frame in which reconstructed jets are defined, that is, as seen from the PV: the truth-jet direction entering the matching is the geometrically corrected axis of Section 11.6.3. The reconstructed jets, on the other hand, enter with their nominal (η, ϕ). The angular distance is therefore evaluated between the apparent truth axis and the reconstructed axis in the same η–ϕ plane. Candidate pairs are formed from truth jets with $p_\text{T} > 20\,\text{GeV}$ and reconstructed jets with $p_\text{T} > 50\,\text{GeV}$ and $|\eta| < 2.8$, and are required to satisfy $\Delta R \leq 0.4$. All pairs inside the cone are ordered by increasing ΔR and accepted only if neither jet has already been used. Figure 11.13 shows the resulting ΔR for matched pairs, together with the ΔR to the closest reconstructed jet for unmatched truth jets, irrespective of the cone. Matched pairs peak at small ΔR, well inside the cone, while for unmatched truth jets the closest reconstructed jet sits far outside it; the shape of the matched peak is independent of the neutralino mass, confirming the choice of $\Delta R \leq 0.4$.

Figure 11.14 shows the efficiency as a function of the LLP decay radius, comparing the three mass points (1000, 1400 and 2000 GeV) at fixed lifetime and the three lifetimes (, 3, 10 and 30 ns) at fixed mass.

The efficiency is high for decays occurring before and within the first calorimeter layers and degrades as the decay moves outward. At a given decay radius the three lifetimes agree more, suggesting that the degradation is dependent mainly on the decay

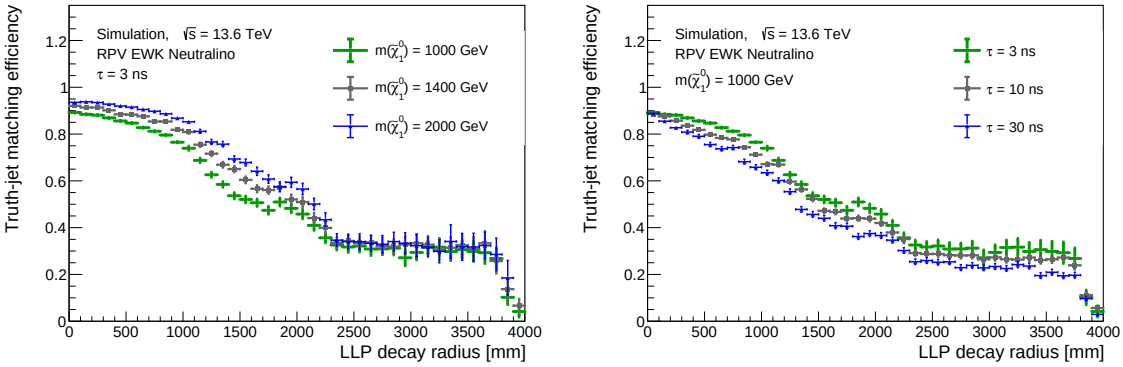

(a) (b)

Figure 11.14: Truth-jet matching efficiency as a function of the LLP decay position: as a function of the decay radius for (a) the three mass points at $\tau = 3\,\text{ns}$ and (b) three lifetimes at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$.

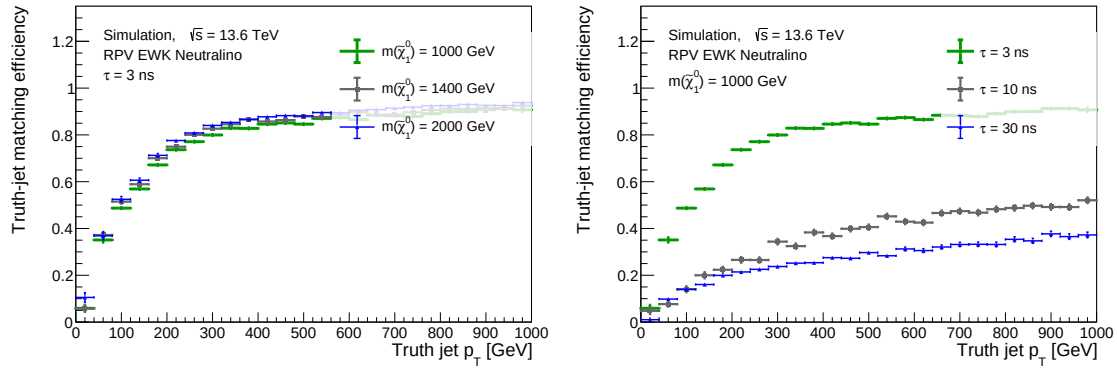

Figure 11.15: Truth-jet matching efficiency as a function of the truth-jet p_{T}, for the three mass points at $\tau = 3\,\text{ns}$ (left) and for three lifetimes at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$ (right).

position itself. The limits of the different subdetectors can be clearly seen in the stepwise changes in matching efficiency. Figure 11.15 shows the efficiency as a function of the truth-jet p_{T}, which increases with the jet momentum as harder jets are more likely to survive reconstruction thresholds. Notice in Figure 11.15(b) how the matching efficiency falls with the lifetime. This suggests it is dependent specifically on where the longest-lived neutralinos decay, given their decay-radius distribution extends well beyond the calorimeters (Figure 11.5), so a large fraction of them decay too far in the detector be reconstructed at all.

Figure 11.16 show the number of matched reconstructed jets per LLP compared with the number of truth jets. The matched multiplicity appears to shift toward lower values with respect to the truth one, as expected from deposits merging more collimated jets

into one, and from decays occurring too far into the detector to be reconstructed and thus matched. On the other hand, Figure 11.16(c) shows the rate of reconstructed jets without a truth partner as a function of the reconstructed p_{T}. The concentration at lower p_{T} originates from pileup and initial-state radiation jets, falling well below the $p_{\mathrm{T}} \geq 350\,\mathrm{GeV}$ threshold of Signal Region SR1, that will be defined in Section 11.7.

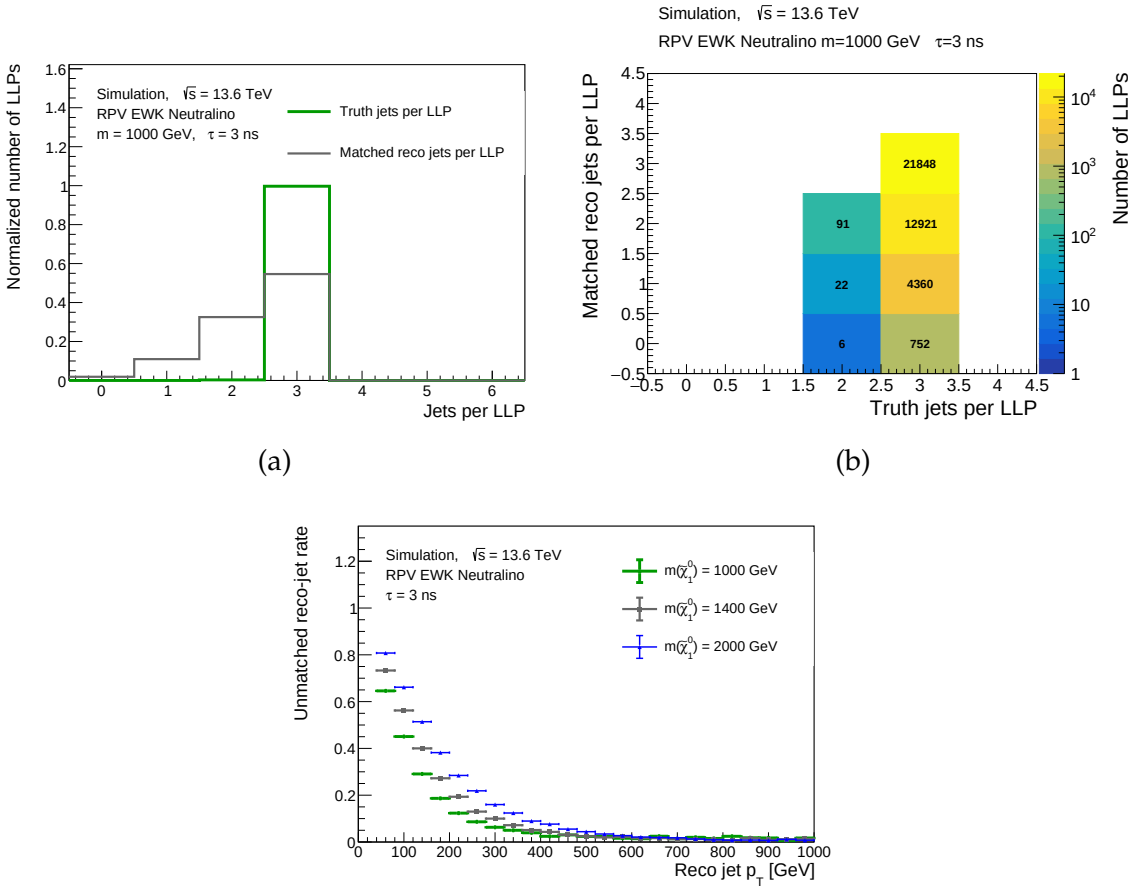

Figure 11.16: (a) Number of truth jets and of matched reconstructed jets per LLP and (b) their correlation, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark. (c) Rate of reconstructed jets not matched to any truth jet, as a function of the reconstructed jet p_{T}, for the three mass points at $\tau = 3\,\mathrm{ns}$.

The matching behavior is illustrated at the level of single events in Figures 11.17 and 11.18, which show one representative well-matched and one representative poorly matched event of the reference sample, visualizing the r–z view of the displaced-decay geometry. In the well-matched event of Figure 11.17 both neutralinos decay at small radii, and all six truth jets find a reconstructed partner within the cone, with a small ΔR between the matched pairs. In the poorly matched event of Figure 11.18 the outermost

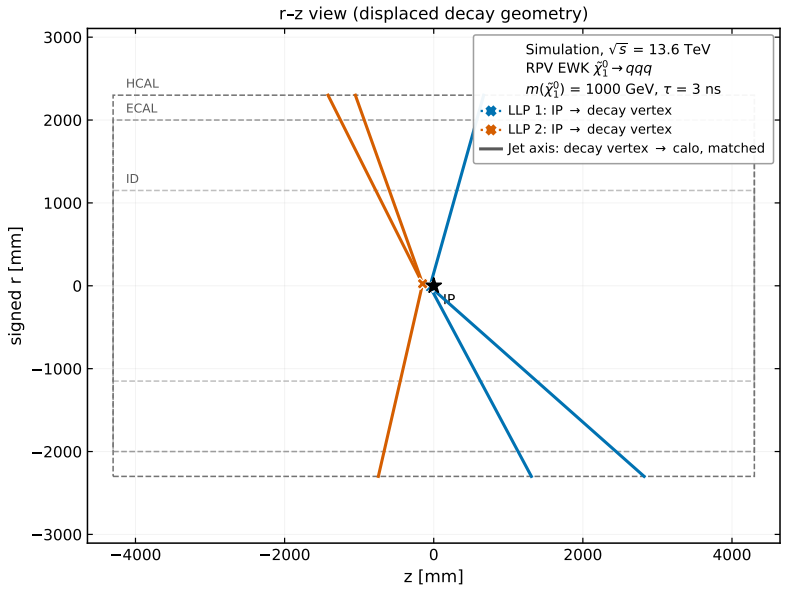

Figure 11.17: Event display of a representative well-matched event for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ benchmark, in the r–z view of the displaced-decay geometry. All six truth jets are matched to a reconstructed jet.

decay occurs at a radius of about 1.7 m, already within the calorimeter volume. Here, the decay products of each neutralino seem to merge into a single calorimeter deposit, so only three reconstructed jets are found, of which two are matched. Here, four of the six truth jets are left unmatched. The event thus displays the merging and efficiency-loss mechanisms at large decay radii quantified in Figure 11.14.

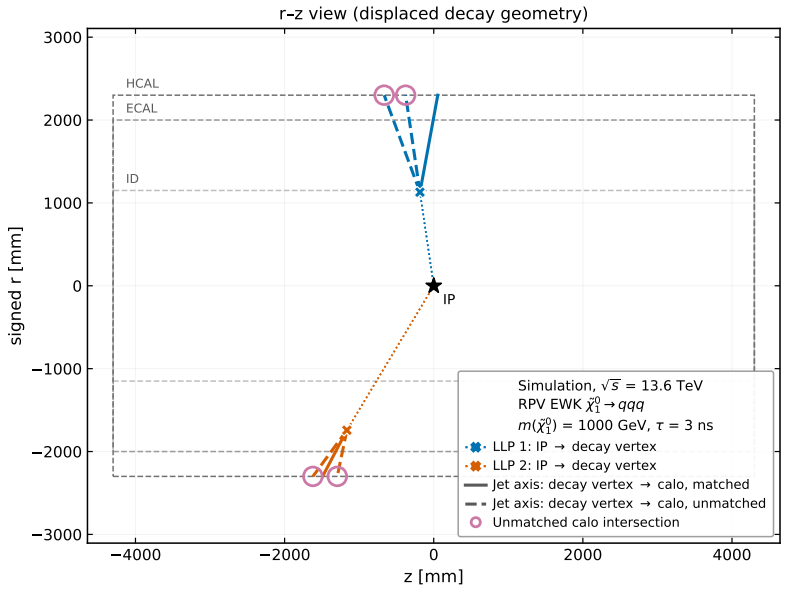

Figure 11.18: Event display of a representative poorly matched event for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ benchmark, in the r–z view of the displaced-decay geometry. The outermost decay occurs at a radius of about 1.7 m, the decay products merge into single deposits, three reconstructed jets are found of which two are matched, and four of the six truth jets remain unmatched.

11.7 Signal region definition

All signal regions are built on a common baseline selection. First of all, an event must pass one of the dedicated delayed-jet triggers of Section 11.4. Jets entering the selection are then required to pass the following cuts:

- $p_\text{T} > 100\,\text{GeV}$,
- $|\eta| < 2.5$,
- $t_\text{jet} > 1\,\text{ns}$,
- $t_\text{jet} < 12.5\,\text{ns}$,

with the delay subsequently required to be *significant* as defined in Eq. (11.4). The lower p_T bound of 100 GeV removes the soft jets from pileup and initial-state radiation identified in Figure 11.16(c). The $|\eta|$ requirement restricts the selection to the region of the calorimeter in which the jet timing can well measured, while the minimum delay of 1 ns removes the bulk of the prompt jets. The upper limit set to 12.5 ns, half the 25 ns bunch spacing, restricts the baseline signal that is so late as to be attributed to the following bunch crossing, and simultaneously bounds the region in which out-of-time pileup from the neighboring crossing could occur. This is an tighter offline cut compared to the already existing 15 ns upper limit set by the trigger chain.

Three signal regions are defined, following the design principle of one inclusive

single-jet selection and two multiplicity-based ones:

SR1 (inclusive): At least one jet with $p_{\mathrm{T}} \geq 350\,\mathrm{GeV}$ passing the t/σ_t requirement. Being a single-jet requirement, SR1 makes no assumption about the decay multiplicity or topology of the LLP decay. This SR can be target a broad range of models producing at least one hard, delayed jet. Hence, it is designed to be the model-independent result of the search.

SR2 (two jets, same hemisphere): At least two jets with $p_{\mathrm{T}} > 150\,\mathrm{GeV}$, both passing the t/σ_t requirement, in the same detector hemisphere. This targets models in which the LLP decays to two or more jets, as in the RPV neutralino model, where the LLPs are produced back-to-back. The idea is that the probability for several jets from Standard Model processes to be mismeasured in time simultaneously is small, so requiring a second delayed jet suppresses the mistiming background far more strongly than the signal.

SR3 (three jets): At least three jets with $p_{\mathrm{T}} > 100\,\mathrm{GeV}$, all passing the t/σ_t requirement, anywhere in the detector. No hemisphere requirement is applied. This is the region targeting most directly the RPV neutralino model, in which each LLP decays to three quarks, and it extends the multiplicity argument of SR2 by one further jet, with no requirement for time coincidence. The background is further suppressed, at the cost of an increased loss in signal acceptance.

A scan is performed over the five working points in delay significance $t/\sigma_t > 3$, 5, 7, 9 and 12, applied to all three region, and also in terms of jet transverse momentum. The final working point will be chosen once the background estimates of Section 11.5 are complete.

11.7.1 Signal acceptance scans

The two quantities that define each signal region, namely the jet p_{T} threshold and the delay significance t/σ_t, are optimized together because they are coupled by the resolution parameterization of Eq. (11.3). This is done with a two-dimensional scan, in which each region is analyzed over every point of a grid in $(p_{\mathrm{T}}, t/\sigma_t)$ and for every mass–lifetime point of the signal grid.

Figure 11.19 shows such a scan, for SR1 at the benchmark point, with the acceptance normalized to the *total* number of generated events, hence including the trigger efficiency for in this region and mass point. The plot shows highest acceptance at the lower end of both spectra, as expected, where the decays of an LLP with this mass and lifetime most are located, and are surviving the cuts.

An extended selection efficiency maps for all grid points of Figure 11.20, shown for the three regions at the nominal p_{T} threshold of each, 350 GeV for SR1, 150 GeV for SR2

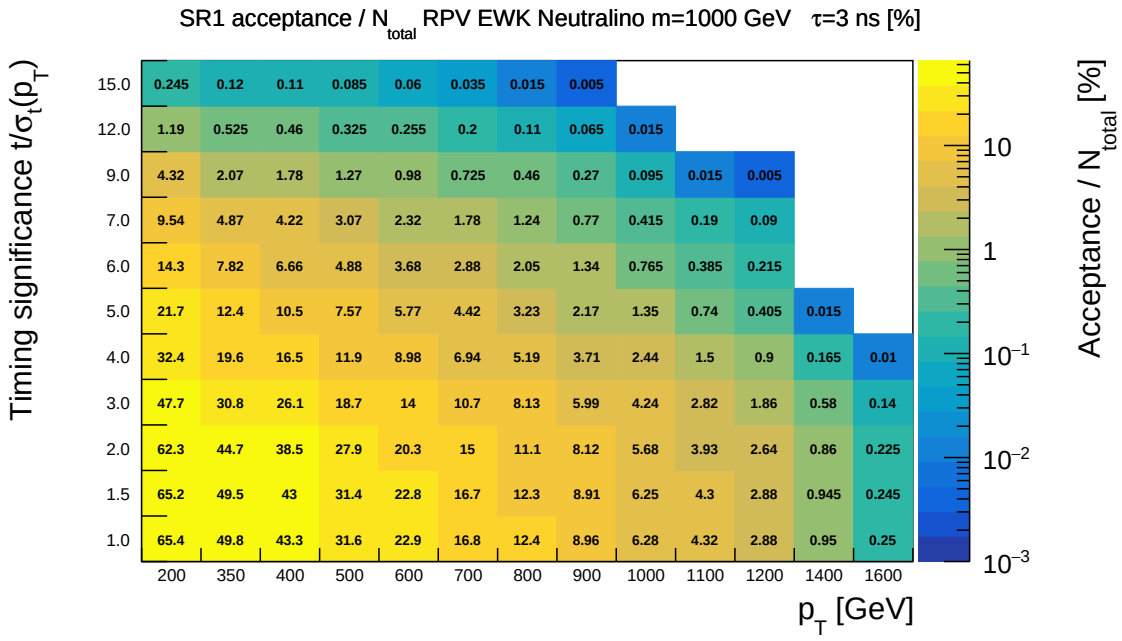

Figure 11.19: SR1 acceptance in the plane of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark. The acceptance is expressed as a percentage of the generated events and therefore includes the efficiency of the delayed-jet triggers.

and 100 GeV for SR3, and $t/\sigma_t > 3$. In all three regions the efficiency peaks at $\tau = 10\,\mathrm{ns}$, the lifetime at which the largest fraction of decays occurs inside the calorimeters, and falls away at both shorter and longer lifetimes.

After looking at the total acceptance of the regions, which includes the trigger efficiency, it is then necessary to study the acceptance relative to only the *triggered* events, as shown in the following plots. The way this efficiency is built up is shown by the cutflows of Figure 11.21, which follow the acceptance of triggered events through the selection steps of each region at $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, for the four available lifetimes.

The effect of kinematic requirements depends strongly on the region. In SR1 the single $p_{\mathrm{T}} > 350\,\mathrm{GeV}$ jet requirement keeps between 70% and 80% of the triggered events at every lifetime. Here, the choice of a 350 GeV threshold is justified by looking at the leading jet p_{T} distribution in Figure 11.9a: the plots shows the peak of the distribution for the lowest mass point is located around exactly around 350 GeV. On the other and, the multiplicity and topology requirements of SR2 and SR3 appear to cut more on the longest-lived samples, where the decay products are likely often lost or merged. After that, cuts in delay significance are performed stepwise, and the separation between the lifetimes grows: at the loosest working point, $t/\sigma_t > 3$, the four lifetimes lie within roughly a factor of two of one another, but by $t/\sigma_t > 12$ the $\tau = 1\,\mathrm{ns}$ sample has

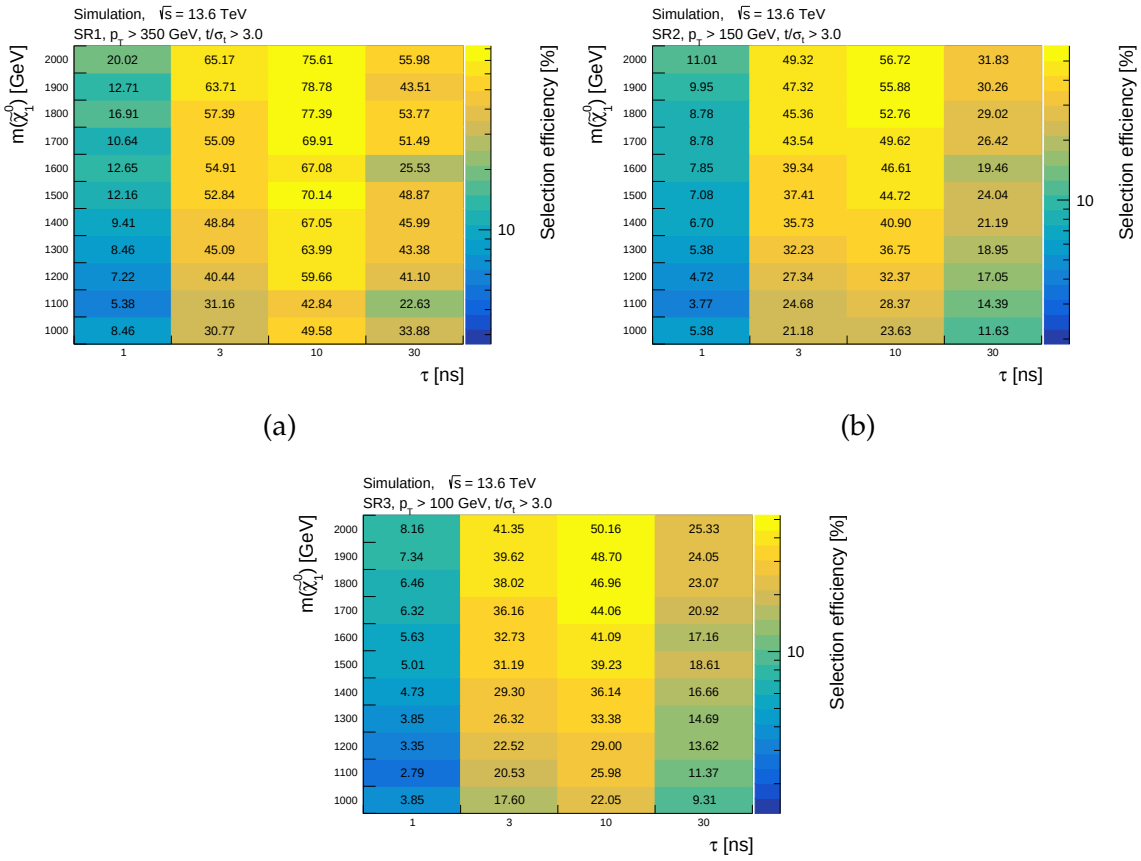

Figure 11.20: Selection efficiency of (a) SR1, (b) SR2 and (c) SR3 across the signal grid, in the plane of neutralino proper lifetime and mass, at the nominal p_{T} threshold of each region, 350 GeV for SR1, 150 GeV for SR2 and 100 GeV for SR3, and $t/\sigma_t > 3$. The efficiency is normalized to the total number of generated events and therefore includes the efficiency of the delayed-jet triggers.

been removed entirely while the $\tau = 10$ and 30 ns samples retain of order 10% of the triggered events in SR1. For SR2 and SR3, the multiplicity requriements are expected to reduce most background. For this reason, lower p_{T} thresholds of 150 and 100 GeV, respectively, were chosen, as they sit comfortably below the peaks of the subleading jet p_{T} distribution shown in Figure 11.9b.

A more detailed scan on the triggered events is shown for SR1 at the benchmark point in Figure 11.22. The corresponding scans for SR2 and SR3 are available in Appendix B. The three regions share the cuts described above, and the maps differ only in the extent of the populated area, which shrinks as the required jet multiplicity grows.

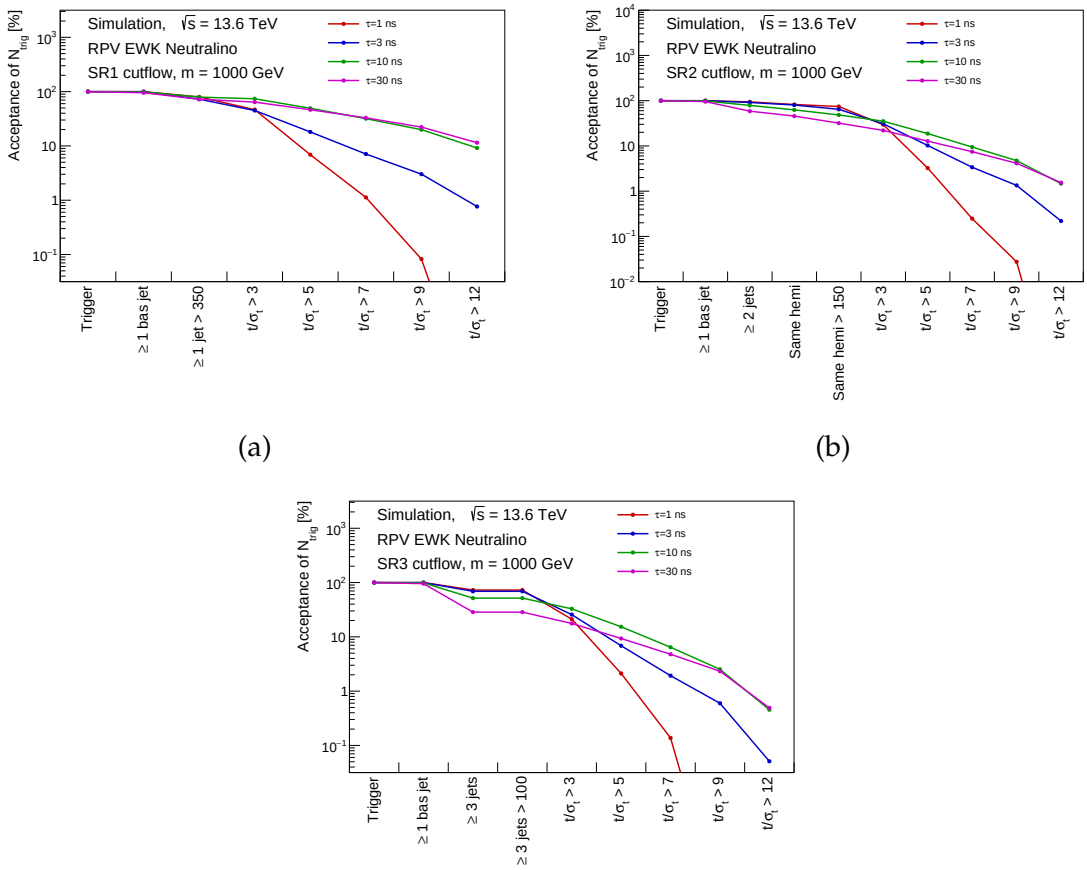

Figure 11.21: Cutflows of the (a) SR1, (b) SR2 and (c) SR3 selections at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, for the four lifetimes. The acceptance is normalized to the number of events passing the delayed-jet triggers and shown on a logarithmic scale. Each chain proceeds from the trigger through the baseline jet requirement and the jet multiplicity and topology requirements of the region, and then through the five delay-significance working points: (a) one jet with $p_\text{T} > 350\,\text{GeV}$, (b) two jets with $p_\text{T} > 150\,\text{GeV}$ in the same detector hemisphere, and (c) three jets with $p_\text{T} > 100\,\text{GeV}$ with no hemisphere requirement.

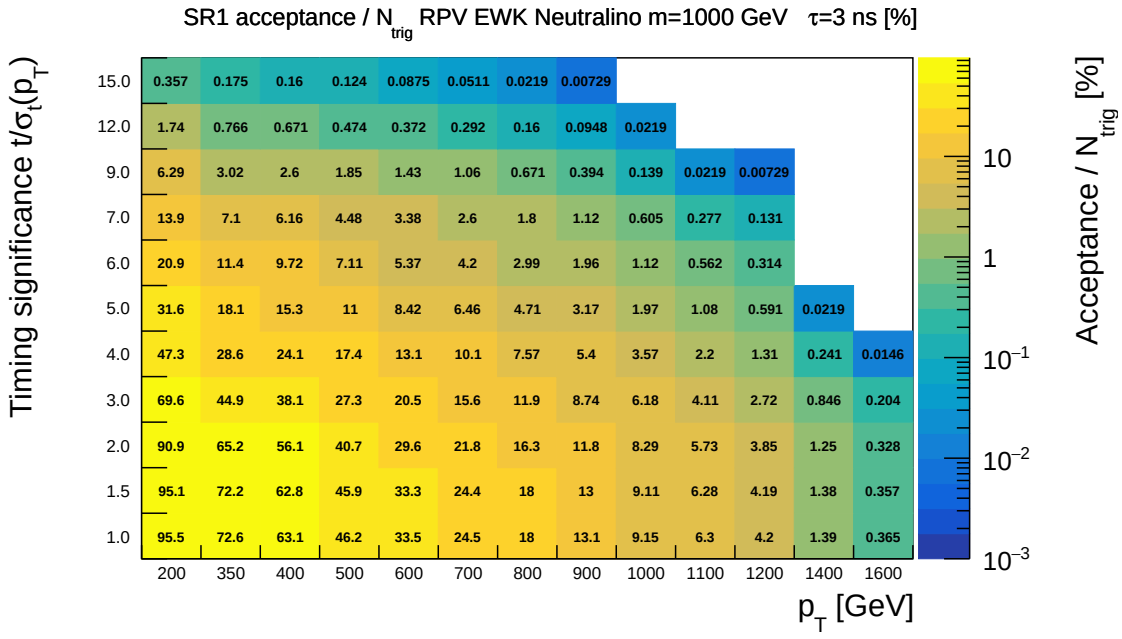

Figure 11.22: SR1 acceptance in the plane of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark. The acceptance is expressed as a percentage of the events passing the delayed-jet triggers.

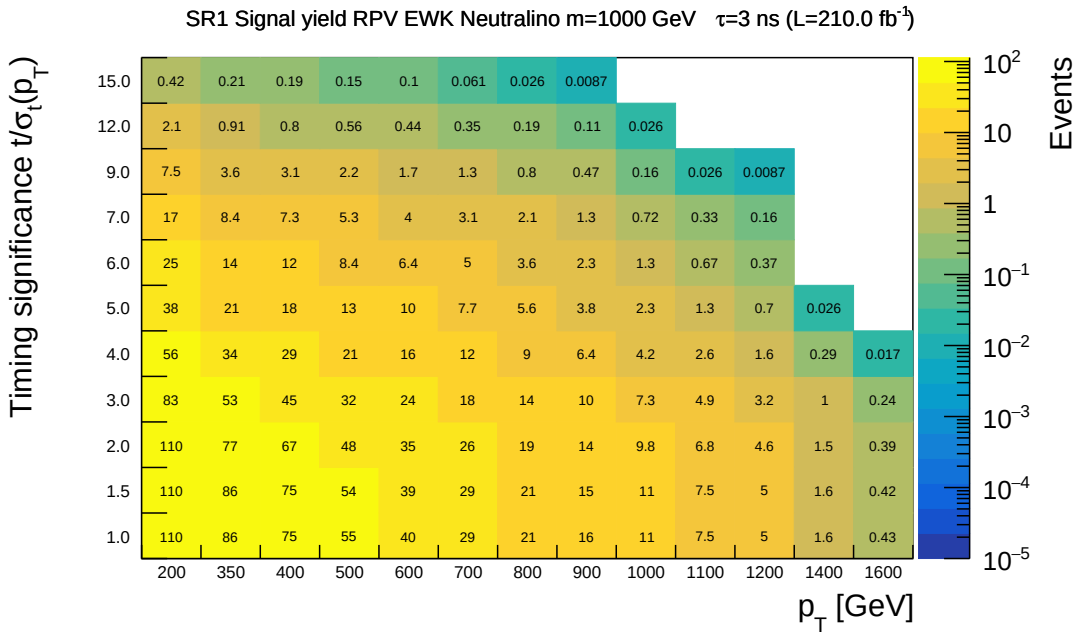

Figure 11.23: Expected signal yields in SR1 as a function of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark and $\mathcal{L}_{\mathrm{int}} = 210\,\mathrm{fb}^{-1}$. The map is obtained from the acceptance normalized to N_{gen} with Eq. (11.5), and is therefore the same scan as Figure 11.22 rescaled by the production cross section and the integrated luminosity.

11.7.2 Expected yields

Combining the acceptances computed so far with the production cross sections σ of Section 11.2.2 and the target integrated luminosity $\mathcal{L}_{\mathrm{int}}$ of $210\,\mathrm{fb}^{-1}$ gives the expected signal yields of Table 11.3, quoted at the loosest working point, $t/\sigma_t > 3$. The expected yield of a region at a given grid point is

$$N_{\mathrm{exp}} = \sigma\,\mathcal{L}_{\mathrm{int}}\,\frac{N_{\mathrm{sel}}}{N_{\mathrm{gen}}}, \tag{11.5}$$

where N_{gen} is the total number of generated events of the sample and N_{sel} the number of events passing the delayed-jet trigger and the full offline selection of the region. The trigger efficiency is therefore included in the yields. The corresponding yields across the $(p_{\mathrm{T}}, t/\sigma_t)$ plane are shown in Figure 11.23 for SR1 (the corresponding scans for SR2 and SR3 are given in Appendix B). The yields confirm the expectation introduced in the premise of Section 11.1. Even before background subtraction, they reach about 86 events at the most favorable grid point and fall below one event over most of the grid.

The yield seems to be correlated to the lifetime and peaks at $\tau = 10\,\mathrm{ns}$, reflecting what was found in the efficiency maps. This demonstrated the sensitivity of the search

in the lifetime range that vertex-based searches do not reach. Another effect seen is the mass dependence: at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$ the yields climb as high as 86 events, at 1400 GeV they range from about one to 13 events, and at 2000 GeV they remain below one event in every region and at every lifetime. Sensitivity at this luminosity is therefore confined to the mid-lower end of the mass grid. Finally, the three regions give relative comparable yields, so the choice between them can be decided by their background levels rather than by these yields, once the backgrounds are estimated.

Table 11.3: Expected signal yields in the three signal regions for $\mathcal{L}_{\text{int}} = 210\,\text{fb}^{-1}$, computed with Eq. (11.5) at $t/\sigma_t > 3$ and at the nominal p_{T} threshold of each region, namely 350 GeV for SR1, 150 GeV for SR2 and 100 GeV for SR3. Yields are obtained from the simulated samples weighted by the total production cross section of the four degenerate higgsino channels ($\tilde{\chi}_1^0\tilde{\chi}_1^+$, $\tilde{\chi}_1^+\tilde{\chi}_1^-$, $\tilde{\chi}_1^0\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0\tilde{\chi}_1^-$) and include the efficiency of the delayed-jet triggers.

$m_{\tilde{\chi}_1^0}$ [GeV]	τ [ns]	SR1	SR2	SR3
1000	1	14.65	9.31	6.67
1000	3	53.30	36.69	30.48
1000	10	85.88	40.93	38.18
1000	30	58.68	20.14	16.13
1400	1	1.85	1.32	0.93
1400	3	9.60	7.02	5.76
1400	10	13.18	8.04	7.10
1400	30	9.04	4.17	3.27
2000	1	0.23	0.12	0.09
2000	3	0.73	0.55	0.46
2000	10	0.85	0.64	0.56
2000	30	0.63	0.36	0.28

Tightening the delay significance reduces the yields heavily, and by $t/\sigma_t > 12$ even the best performing mass point, $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$ at $\tau = 10\,\text{ns}$, yields only about 10 events in SR1 and 0.55 in SR3. The final working point is consequently a trade-off that cannot be fixed from the signal alone, as it requires the background estimates of Section 11.5, which are still in progress at the time of writing. For this reason, no significance figure is shown in this thesis and the selection is presented as an optimized but not yet finalized scan.

The same yields are shown across the full grid in Figure 11.24, which extends Table 11.3 to all available mass points. It can be seen how the yield contours run almost horizontally in mass and peak in the same lifetime column in all three regions. The maps apply the nominal p_{T} threshold of each region, 350 GeV for SR1, 150 GeV for SR2 and 100 GeV for SR3, which are the same working points as Table 11.3.

Overall, the computed expected yield suggests that a certain sensitivity can be retained at mass points over the exclusion limit: assuming a background of 1 ± 1, the yields shown in Figure 11.24 act as a preliminary significance scan. Here, in SR1 at least three events are expected up to $m_{\tilde{\chi}_1^0} = 1700\,\text{GeV}$ at $\tau = 10\,\text{ns}$, and up to $m_{\tilde{\chi}_1^0} = 1600\,\text{GeV}$ at $\tau = 3\,\text{ns}$ (Figure 11.24(a)).

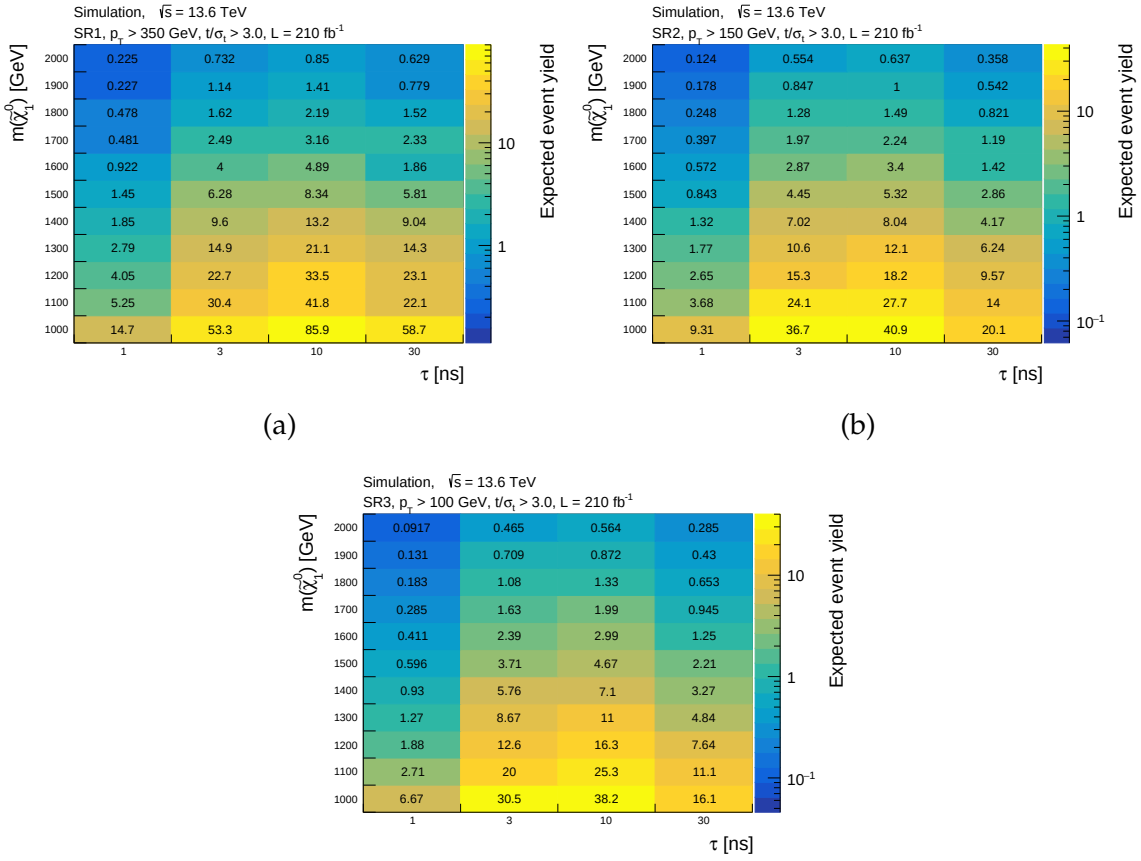

(a) (b)

(c)

Figure 11.24: Expected signal yields of (a) SR1, (b) SR2 and (c) SR3 across the signal grid, in the plane of neutralino proper lifetime and mass, for $\mathcal{L}_{\text{int}} = 210\,\text{fb}^{-1}$ at the nominal p_{T} threshold of each region, 350 GeV for SR1, 150 GeV for SR2 and 100 GeV for SR3, and $t/\sigma_t > 3$. The yields are obtained with Eq. (11.5) from the acceptances weighted by the summed production cross section of the four degenerate higgsino channels.

11.8 Next steps

With a signal this small, the sensitivity of the search is set almost entirely by how far the background can be suppressed, so the four background sources introduced in Section 11.5, all of which are being estimated from data, remain to be performed. Their

completion will set the delay-significance working point, which unfortunately cannot be chosen from the signal alone, and will allow a significance figure of merit to be estimated for each region and the most sensitive one to be identified. Nevertheless, assuming an expected background reduced to 1 ± 1, the yields suggest that the search is sensitive in regions beyond the exclusion limits set by previous analyses. Beyond that, the final GRL and integrated luminosity with its uncertainty are still to be settled, and no systematic uncertainties have yet been evaluated.

The second stage of the analysis, in which the events surviving the preselection are re-reconstructed with cell-level information retained, will then allow the final selection to use timing quantities that are not available in the standard derivation format. On the signal side, the two additional models mentioned in Section 11.1, strong production of the same long-lived neutralinos and heavy scalars decaying to a pair of long-lived scalars, are still to be generated and will extend the interpretation beyond the electroweak RPV benchmark. Once the background estimates are complete and the analysis is unblinded, the search will further push ATLAS limits on long-lived particles using the arrival time of calorimeter deposits, in a lifetime range where the existing vertex-based searches lose sensitivity.

Concluding remarks

The main result of this thesis is that it firmly established the LAr Digital Trigger, a readout chain built for the Phase-I trigger upgrade, as a luminosity detector, through three different methods that exploits observables that the same system already computes. First, the dispersion of the digitized samples (Chapter 8) is a linear and stable proxy for the instantaneous luminosity, reaching on the `physics_Main` stream a resolution relative to the ATLAS online luminosity of 0.04% at $\langle\mu\rangle = 58.6$ (Figure 8.31) and 0.21% at $\langle\mu\rangle = 120.6$ (Figure 8.33), both approved by the ATLAS collaboration [116] and published [117]; its mean ratio to LUCID, 1.000 ± 0.016, matches the 0.990 ± 0.014 of track counting and the 0.992 ± 0.014 of the LAr gap currents (Figure 8.35). The second method, which uses the deconvolution of the mean ADC sample pedestal per BCID with the bipolar LAr pulse (Eq. (9.1)) as kernel, was shown to reconstruct the LUCID per-bunch luminosity with a residual RMS of 3.4% with the constrained inverse matrix, and 3.1% with the constrained FFT (Table 9.2). The LATOME Pulse Rate Monitor, free of any trigger dependence, works as a hit-counting luminometer with 0.83% resolution in nominal conditions and 0.79% at $\langle\mu\rangle = 120.6$ (Figure 10.10), and remains linear at van der Meer luminosities (Figure 10.14), where the RMS method does not.

For the search of Chapter 11, a per-LLP truth-jet definition with a geometric axis correction reduced the mean ΔR_{min} to the closest reconstructed jet from 0.152 to 0.065 (Figure 11.12), and three signal regions built on the resolution-scaled delay significance of Eq. (11.4) were defined and scanned in the $(p_{\text{T}}, t/\sigma_t)$ plane. Their efficiency peaks at $\tau = 10\,\text{ns}$ (Figure 11.20) and the expected yields in $210\,\text{fb}^{-1}$ reach about 86 events at $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, falling below one event at 2000 GeV (Table 11.3).

While many of the results obtained in this thesis are remarkable, several obstacles made the work challenging, and at times precluded the results. The main obstacle in the luminosity study was the access to data: the MON stream, the only source of raw samples, is prescaled by 60,000, leaving the feasibility study with a run-to-run spread of $\pm$0.010–0.013 against the $\pm$0.003–0.007 of the established methods (Section 8.3.12), while `physics_Main`, the most precise stream (Table 8.2), occupies terabytes per run, so the long-term stability study could not be repeated with it. The RMS method is also bounded by the electronic noise floor of Eq. (8.1), which makes it insensitive to emittance scans (Section 8.3.9) and loses its correlation almost entirely at vdM conditions (Section 8.3.10). The per-bunch measurement, instead, is limited by the kernel, measured in 2025 delay runs for 2026 data and truncated by the 32 recordable samples, so the 3%-level residual, obtained with a kernel extracted from LUCID itself, shows what an adequate kernel achieves rather than an independent measurement (Section 9.5). For

the PRM, rates span order 10 Hz to 10^3 Hz across layers, making full-detector averaging detrimental (Figure 10.7). Furthermore, a refined optimization of the energy thresholds needs to be implemented to improve the sensitivity. Finally, the delayed-jet physics search is blinded, and its four background estimates are still in progress (Section 11.5), while no systematic uncertainty has been evaluated, which leaves the analysis unfinished at the time of the writing of this thesis.

It is worth noting that the Digital Trigger was studied as a *relative* luminometer (Section 7.2.1): the absolute scale is set by the van der Meer calibration of Section 7.4 and was not derived here, as every estimator was calibrated empirically against the ATLAS online luminosity. The quoted figures are therefore resolutions with respect to a reference, not uncertainties on an integrated luminosity of the kind broken down in Table 7.2, and all three methods were demonstrated offline rather than online or in firmware. Futhermore, for the search, the contributions are the ntuple framework, the truth studies of Section 11.6 and the signal-region definition of Section 11.7; the background estimation and the unblinding lie outside this thesis, and no exclusion limit is quoted.

Relevance for the HL-LHC

At the HL-LHC, ATLAS will operate at up to $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ with $\langle\mu\rangle$ of about 200 [93, 94]. This is a regime in which these methods are shown to perform very well: the pileup term of the sample dispersion grows with μ while the noise floor stays fixed (Section 8.2.1). The highest pileup studied here is $\langle\mu\rangle = 120.6$, in run 509891, which has the conditions closest to HL-LHC available in Run 3 data. There the RMS method reaches 0.21% in relative resolution and the PRM reaches 0.79% with the EMB presampler alone. This demonstrates that both methods work at high μ, though the run differs from nominal conditions in instantaneous luminosity and number of colliding bunches as well as in μ. None of the methods has been tested at $\langle\mu\rangle = 200$, so the linearity in that regime remains an open question. This is especially relevant for the PRM: it counts pulses above an certain energy threshold, so overlapping pulses could change the E_T spectrum relative to that threshold, which in turn could introduce a non-linearity. A further significant point is that nothing new has to be installed: the LAr Phase-I components are retained through the HL-LHC era and the Digital Trigger becomes part of the Phase-II Level-0 calorimeter trigger [97], so the samples, their BCID labels and the PRM counters all survive into Run 4. Few ATLAS systems resolve individual bunch crossings (Section 7.3), which the samples of the DT do, while the PRM counters are integrated over the orbit.

Thee three methods outlined in this thesis are relative: the absolute scale comes from the vdM calibration, and every estimator was calibrated against the online luminosity. The results therefore do not translate into a reduction of the Run 2 uncertainty budget

of Table 7.2. The two terms they could address are the long-term stability, which contributes 0.12–0.23% per year and 0.08% to the combined Run 2 result, and the calibration transfer, which contributes 0.50% [109]. For the calibration transfer, the PRM, which retains sensitivity down to vdM luminosities, could be cross-calibrated to LUCID in the same fills, as is done for track counting, but no such study was performed.

Future work

In the context of the luminosity measurement, the most direct step is to move the sample-based estimators online or into the LATOME firmware during LS3, which removes both the statistics limit of the MON stream and the size limit of reprocessing `physics_Main` data offline. For the first vdM program of Run 4, an unprescaled stream fed by a random L1 item gated on colliding BCIDs should be commissioned, so that the vdM sensitivity of the RMS method can finally be tested (Section 8.3.10). Furthermore, a ML-based estimation of the luminosity using the RMS is also possible, in order to probe the maximal capabilities of this method. The per-bunch measurement, on the other hand, needs per-scieta delay runs taken contemporaneously with the data to be deconvolved and a readout window extended beyond 32 samples, alongside the learned inverse outlined at the end of Section 9.5. The PRM energy thresholds of Eq. (10.1) should be optimized to improve the overall sensitivity, and, at the same time, fixed per layer and per scieta measurements should be taken from the best-performing region in each regime rather than using a full-detector average. Furthermore, the linearity of the PRM up to $\langle\mu\rangle = 200$ should be studied in simulation.

For the physics search, the four background estimates of Section 11.5 have to be finalized, which fixes the delay-significance working point. The cell-level event-picking stage then has to be performed, and once the Good Run List, luminosity and systematics are settled, the analysis can be unblinded around $\tau = 10\,\text{ns}$, where the signal yield is highest and vertex-based searches [123] lose acceptance.

Part V

Appendices

Appendix A

Data quality and performance studies on the Digital Trigger

Maintaining optimal performance of the calorimeter system requires ongoing calibration, monitoring, and optimization [136]. The LAr Data Quality (DQ) team is responsible for these tasks throughout the data-taking chain, from real-time monitoring in the control room to detailed offline assessments and the final certification of datasets for physics analysis. The remainder of this chapter first describes the LAr DQ operations that define the framework within which this work was carried out (Section A.1), and then four specific studies conducted during the 2024 and 2025 data-taking campaigns: the optimization of the DT optimal filtering coefficients (Section A.2), the processing of *MON data* (Section A.3), the investigation of *FCal correlations* (Section A.4), and the study of *pedestal jumps* (Section A.5).

A.1 LAr data quality operations

The LAr DQ procedures are designed to promptly identify, diagnose, and mitigate any issues that could compromise detector performance, thereby maximizing the fraction of usable data for physics. ATLAS data are monitored at several stages of the acquisition and processing chain to detect as early as possible any problem that could compromise their quality. Most of the monitoring infrastructure is common to both the online and offline environments, but the level of detail in the monitoring procedure evolves with the refinement of the analysis [137].

- **Real-time and online monitoring:** During data taking, shifters in the ATLAS control room monitor detector-level parameters and reconstructed event characteristics as data are being collected. They use dedicated software tools and real-time histograms to identify hardware or software issues as they arise, enabling immediate action to minimize the impact on the collected data. If anomalies or pathologies are observed, such as abnormal timing, noise bursts, or channel failures, the DQ team coordinates with calibration experts and detector specialists to investigate further and, where possible, apply corrections or masking in real time.

- **Trigger streams for DQ assessment:** The ATLAS trigger system organizes data into streams for both physics and calibration purposes. For LAr DQ, four main calibration streams are used: the `Express` stream (a fraction of the recorded data reconstructed

promptly, often within minutes to hours after data taking), `CosmicCalo` (events from empty bunch crossings), `LArCells` (partially built collision events containing only LAr cells with high-energy deposits), and `LArCellsEmpty` (partial events from empty bunches). These streams facilitate targeted and efficient monitoring of detector performance.

- **Offline data quality assessment:** After a run ends, the DQ team performs a detailed offline evaluation of the data, starting with the `Express` data stream. Custom algorithms identify issues such as problematic cells or noise bursts. The results are summarized into status flags stored in central databases, that guide the following processing and analysis. If problems are found, calibration constants may be updated, problematic regions or cells masked, or, in rare cases, affected data excluded (*vetoed*) from physics analyses.

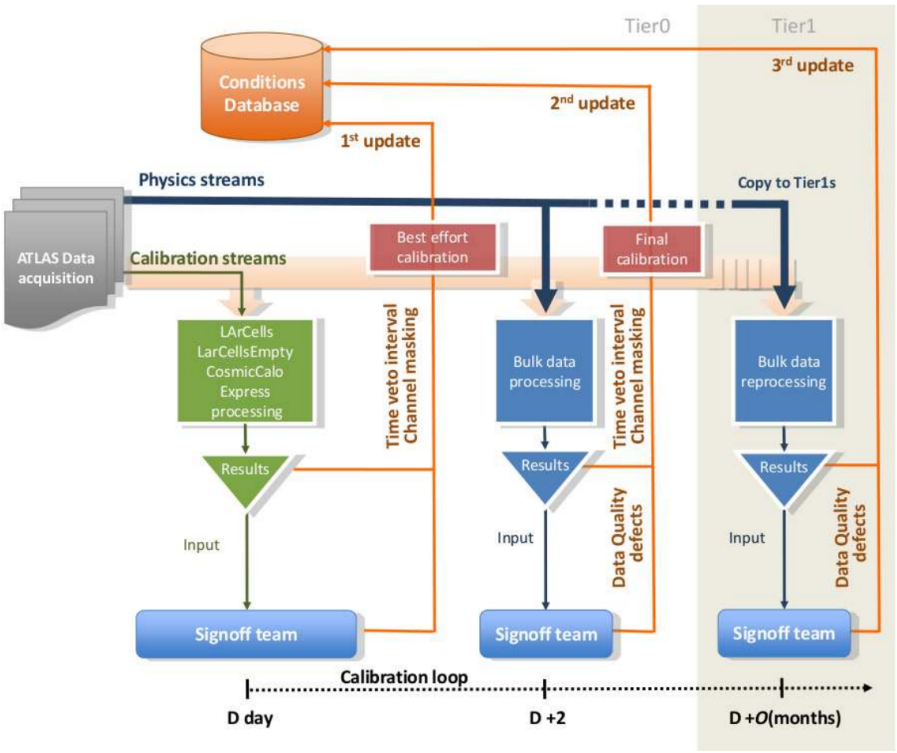

Figure A.1: The ATLAS DQ calibration and sign-off loop.

- **Final sign-off of runs:** The LAr DQ team is also responsible for the final certification of the data. This involves a two-stage review process: initial assessment by the LAr DQ experts, followed by sign-off from the broader ATLAS DQ coordination group. Only data that pass these checks are certified for eventual use in physics analyses.

The effectiveness of these procedures is demonstrated by the high fraction of usable data achieved consistently throughout LHC operations. The latest results from Run 3,

from the 2026 data-taking campaign [136], indicate that the LAr calorimeter continues to deliver data of excellent quality, with efficiencies above 99%, as evidenced in Figure A.2.

ATLAS nominal pp Run-3: 2026

Trigger	Inner Tracker			Calorimeters		Muon Spectrometer					Magnets		Global	
L1+HLT	Pixel	SCT	TRT	LAr	Tile	MDT	RPC	TGC	MM	sTGC	Solenoid	Toroid	Lumi. calib.	Other
99.6	99.7	98.6	99.7	99.7	100	99.9	99.6	100	100	100	100	100	99.7	99.5

Good for physics: 96.5% (28.1 fb^{-1})

Luminosity weighted good data quality efficiencies (in %) in 2026 during stable beam operations of *pp* physics runs at $\sqrt{s}$ = 13.6 TeV and an average number of collisions per bunch crossing of 60, corresponding to an integrated luminosity of 28.1 fb^{-1}, for 29.2 fb^{-1} *pp* data recorded. Technical runs such as luminosity calibration scans are not accounted for in the efficiencies. Runs taken at very low pile-up are also ignored.
When the stable beam flag is raised, the tracking detectors initiate a "warm start", which involves ramping up the high-voltage and activating the pre-amplifiers for the Pixel and SCT systems. The inefficiency due to this, as well as the DAQ inefficiency, are not included in the table above, but accounted for in the ATLAS recording efficiency.
A large fraction of the SCT inefficiency is due to scans for radiation damage monitoring during nominal operations.
The good-for-physics luminosity is 28.1 fb^{-1} for all analyses.

Figure A.2: The overall efficiency of ATLAS, and by subdetector, during the 2026 proton–proton data-taking campaign, with LAr reaching above 99% data quality efficiency.

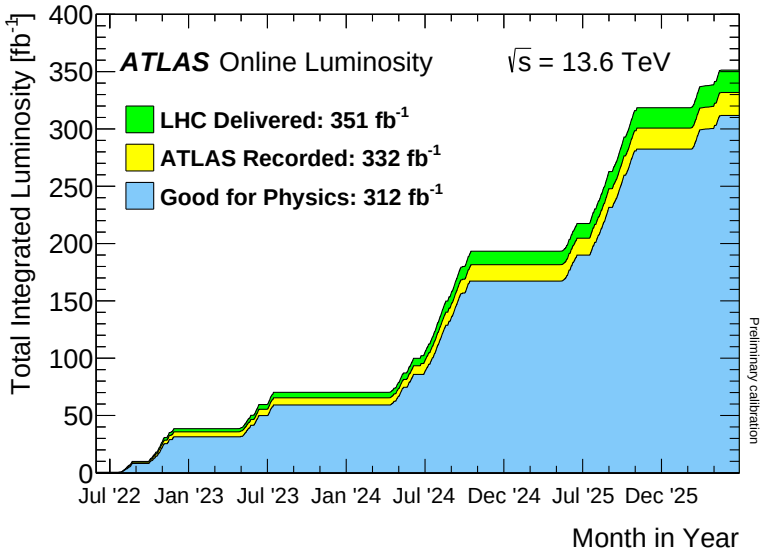

Figure A.3: The integrated luminosity of proton–proton collision data for Run 3.

The continuous refinement of monitoring algorithms and investigation procedures ensures that the LAr calorimeter consistently delivers data of the highest quality, enabling ATLAS to achieve its physics goals. This is essential to the success of the ATLAS physics program, and the high amount of "Good for Physics" data recorded by the experiment, as seen in Figure A.3.

A.2 DT OFCs optimization

As described in Section 6.2.3, the OFCs are central to the energy and timing reconstruction, and are carefully optimized to extract the true signal from digitized calorimeter waveforms. For the 2024 data-taking campaign, a dedicated three-step procedure was implemented to optimize the DT OFCs, ensuring the best possible reconstruction under continuously evolving LHC conditions.

- **Step 1: OFC scan and identification of problematic SCs**

  Before refining the OFCs, a comprehensive scan was undertaken to identify "bad" SCs that could degrade performance with the current OFC set. This process utilized calibration data and simulations to systematically evaluate the response of each SC. The criteria for flagging problematic SCs included examining the mean and RMS of the reconstructed transverse energy (E_T) and timing (t) values from simulated pulses; each SC was simulated ten times with 150 GeV input using both calibration and physics waveforms. Additional checks focused on the pulse slope, specifically searching for sudden changes in ADC values or visibly deformed pulse shapes, such as the one shown in Figure A.4a, as well as identifying cases where the reconstructed amplitude or the OFCs themselves were anomalously large or small, as in Figure A.4b.

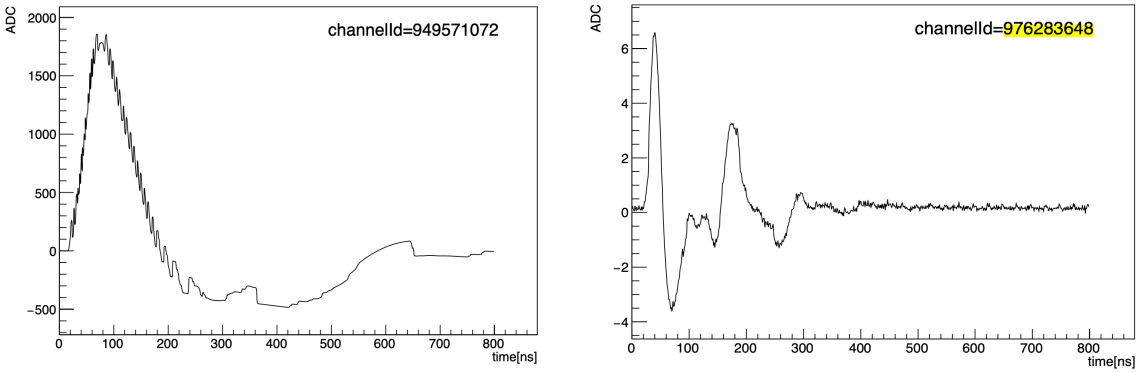

(a) (b)

Figure A.4: Examples of a pulse classified as distorted (a) and "off amplitude" (b) that resulted in the offending channels to be labeled as "bad".

SCs failing one or more of these criteria were flagged for further investigation and compiled into a list of problematic channels. As a result of this scan, 249 SCs were flagged, of which 54 were masked, as listed in Table A.1; the 54 masked SCs are 24 fewer than in the previous year.

- **Step 2: Optimization of the phases**

  With the problematic SCs identified and masked, the next step was to optimize the *phase* for each remaining SC. The optimal phase for each SC was computed, leveraging both

Table A.1: The final list of "bad" SCs for the 2024 campaign. Several labels were used: *masked*, *usable*, *recoverable*, *patch* and *no issue*. At the end of the scan many of the flagged SCs were recovered after investigation.

Criteria	Title in list (cause)	Flag	N of SCs	Total
OFC	`BadOFC`	masked	13	30
		usable	17	
Undershoot	`offamplitude`	masked	1	2
	`Deformed_Pulse`	no issue	1	
Pulse slope	`Deformed_Pulse`	masked	12	18
		recoverable	5	
		usable	1	
Mean & RMS	`offamplitude`	masked	1	5
	`tinySpike`	usable	1	
	`Badreconstruction`	no issue	3	
Amplitude	`LowAmp`, `HighAmp`, `Amp`	masked	27	194
		patch	24	
		usable	115	
		no issue	28	
Total				249

E_T and t reconstruction from the same run as the OFC scan (run 473915, 1970 colliding bunches, 19 April 2024). For each SC, the phase that maximized the reconstructed transverse energy per entry was selected, as exemplified in Figure A.5. The optimization was performed with a 1 ns resolution, scanning 49 possible phases for each OFC. The resulting best OFC phases for each SC were documented and stored for implementation in the online databases used for data taking.

- **Step 3: Validation of new OFCs and phases**

  The final step was to validate the new OFC set and the optimized phases. E_T and timing reconstructions were performed with the new coefficients and compared to the reference online reconstruction. Validation was carried out using the express runs 473915 (April 19, 2024, with 1970 colliding bunches distributed in 15 bunch trains). Two versions of the new OFC constants were tested: a set with the timing coefficients b_i of Eq. (6.1) (OFCb) tuned per SC and a set with a flat weight of 1.5 for the HEC OFCb. The validation confirmed an improved t reconstruction with the new optimized OFCs, especially in the EMB and HEC, as presented in Figure A.6 and good energy reconstruction across all LAr subdetectors, without an appreciable enhancement compared to the reference OFC set, as evidenced in Figure A.7.

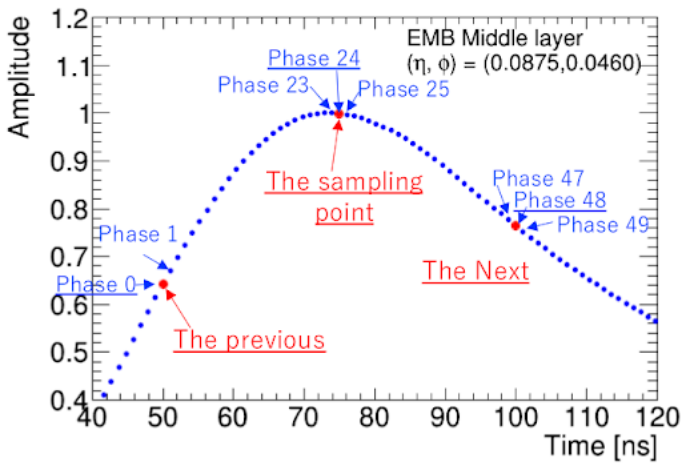

Figure A.5: The best phase for the OFCs was chosen as the one that yields the highest amplitude, in the case of the example, phase 24.

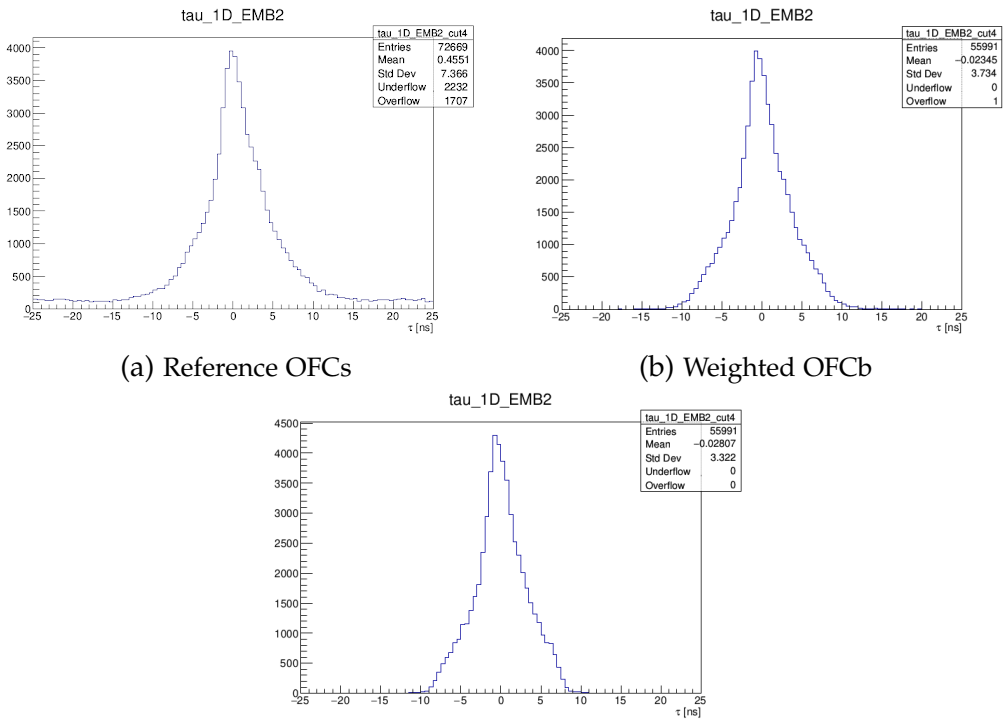

(a) Reference OFCs

(b) Weighted OFCb

(c) Flat 1.5 HEC OFCb

Figure A.6: The result of the time (t) reconstruction in EMB middle layer with the different sets of OFCs with run 474462. A narrower t distribution was achieved with the new sets of optimized OFCs in (b) and (c)

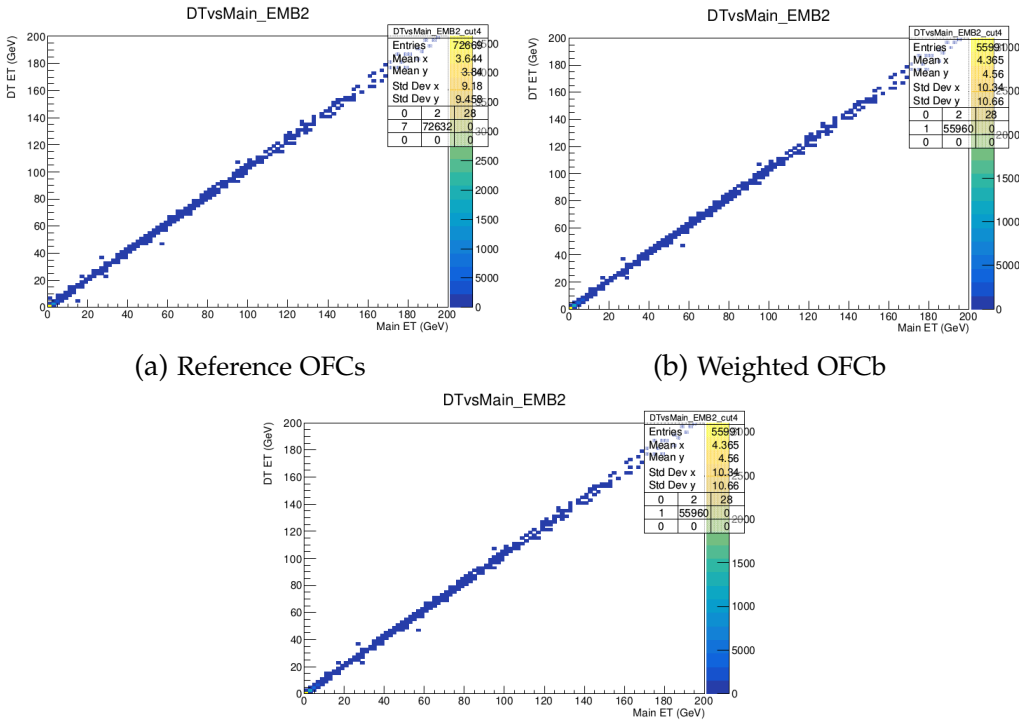

(a) Reference OFCs

(b) Weighted OFCb

(c) Flat 1.5 HEC OFCb

Figure A.7: The results of the DT E_{T} reconstruction in EMB middle layer compared to the main readout E_{T} with run 474462 did not yield a significant improvement in linearity with the main readout.

A.3 Processing of MON data

The MON data stream is a special, unofficial data path within the LAr system, designed to provide access to raw digitized calorimeter signals for commissioning and monitoring. Unlike the standard monitoring data flow, which routes through the SWROD and is optimized for trigger performance and event selection, the MON stream operates independently and outputs data directly from the LATOME modules via a 10G UDP server.

A key feature of the MON stream is its ability to deliver extended information: for each event, it provides between 6 to 32 *raw* and *baseline-corrected* samples (also named *samples_ADC_BAS*) per channel, far exceeding the granularity and detail available in the main trigger path. This makes the MON stream the only source where raw samples can be inspected, which is invaluable for data quality studies.

The stream is configured without a trigger type requirement and typically with a prescale factor of 60,000 (i.e., only one out of every 60,000 events is stored). This greatly reduces the available statistics, making the stream unsuitable for studying rare issues. The MON data are then processed by a custom pipeline that automatically handles each

run, creating ntuples with event, bunch crossing and luminosity block information.

Access to the raw samples in the MON stream enables several critical analyses, including baseline correction checks, or studies of noise such as the "FCal correlations" (see Section A.4). These analyses are essential for understanding and correcting issues in the baseline, which can impact energy reconstruction, especially given the long LAr pulse duration (400–500 ns) relative to the LHC bunch spacing and the effects of out-of-time pileup. Notably, baseline corrections, applied every six seconds to mitigate baseline shifts at the start of bunch trains, as illustrated in Figure 6.5, can only be directly observed and validated using the MON data.

A.4 FCal correlations

During the commissioning and monitoring phases of the DT system, unexpected correlations associated with SC noise were observed in the FCal. Initially, this noise was believed to affect only the FCal, particularly in specific regions. However, later studies revealed a broader pattern, extending to additional sections of the calorimeter system.

Examination of the MON data uncovered that these correlations were not exclusive to the FCal. Instead, similar structures appeared in other regions of the detector. The correlations manifested in a distinctive "petal" pattern when visualized in physical detector coordinates, as seen in Figure A.8. This structure corresponded to alternating values of the SC azimuthal index ($i\phi$), specifically distinguishing between even and odd channels, but it did not uniformly appear across all values of the pseudorapidity index ($i\eta$, the scieta of Chapter 8). Importantly, this petal structure was only observed in the Digital Trigger and was entirely absent from the main data readout.

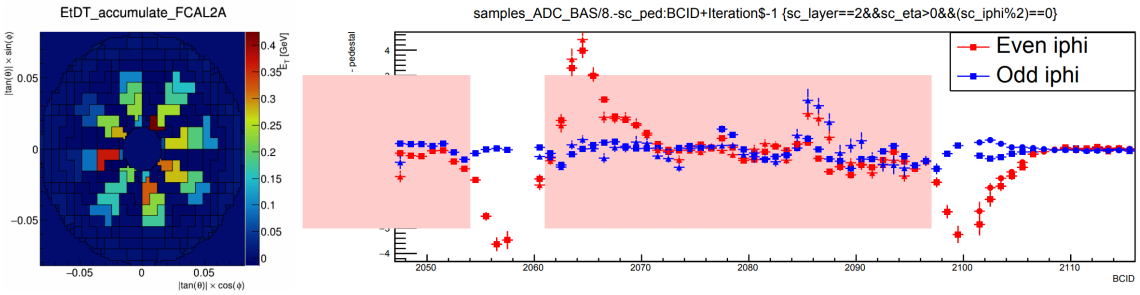

Figure A.8: An example of petal-shaped digital noise from the correlations on the left-hand side. On the right, it is shown how the baseline correction seems to be be applied on a iphi-dependency.

The emergence of these correlations prompted a detailed investigation into the underlying cause. The evidence pointed toward an incorrect application of the baseline correction as the primary culprit. Detailed studies on the MON data were initiated

to understand how the baseline correction was behaving across different regions and conditions.

The issue was investigated in run 474462 (April 26, 2024, with 1946 bunches). During this run, noticeable "FCal correlations" emerged, particularly on the FCalC side of the detector, while FCalA remained unaffected, as evidenced in Figure A.9. The correlations were especially prominent in even $i\phi$ slices, consistent with the petal structure observed earlier.

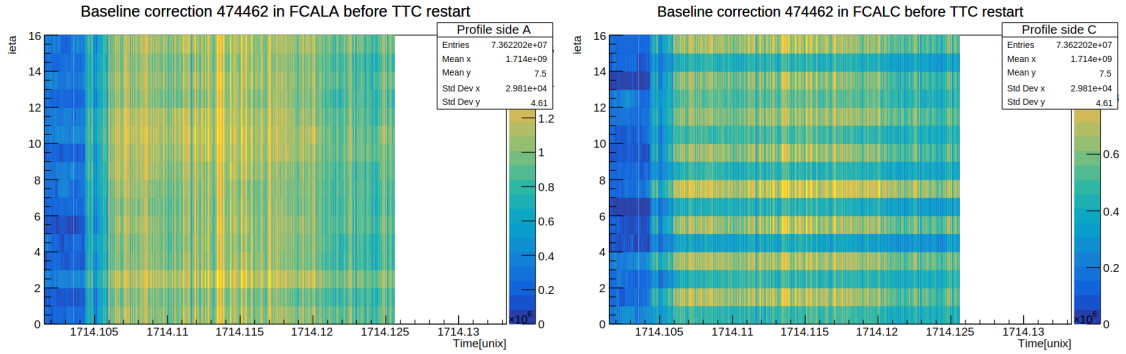

(a) (b)

Figure A.9: The baseline correction as a function of $i\eta$ and time for FCalA (a) and FCalC (b) showing how the baseline correction was applied properly only to FCalA. On FCalC it was only applied to odd-numbered $i\phi$ SCs, as evidenced by the strip pattern in the plot on the right.

During the run, the issue persisted until approximately LB 640, at which point a TTC restart [1] was performed, as seen in the dip in recorded luminosity of the run in Figure A.10. This intervention resolved the correlations, as demonstrated in Figure A.11, confirming that the baseline correction problem was isolated to FCalC and strongly tied to even $i\phi$ Super Cells.

In contrast, the baseline correction in the EMBA and EMECC regions appeared to function properly throughout the same period. Additional checks during runs without collisions also confirmed that the baseline correction was not always reliably set to zero, revealing further inconsistencies.

The investigation ultimately identified a firmware issue as the root cause of the incorrect baseline correction application. The discovery and correction of this firmware fault led to the disappearance of the observed correlations. Since then, such petal-structured correlations have no longer been observed in following runs.

[1] The TTC system in ATLAS distributes clock, trigger, and synchronization signals to detector electronics. A TTC restart re-synchronizes all components by issuing commands like the BCR, ensuring proper timing alignment across the experiment. It is a mechanism used to resolve, among others, configuration and other digital issues in the ATLAS partition during data taking [138].

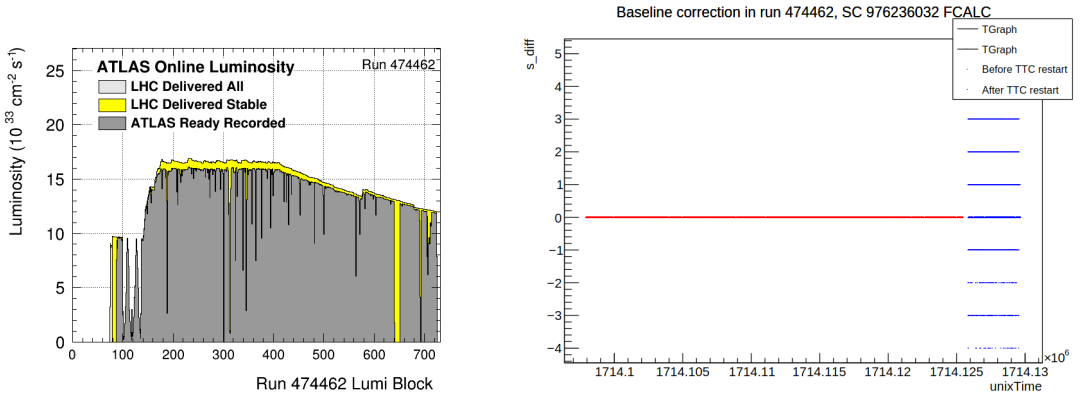

(a) (b)

Figure A.10: (a) The luminosity profile of run 474462. Around LB 640 a TTC restart was performed, as seen by the drop of *ATLAS Ready Recorded* luminosity. (b) A plot showing how the baseline correction was not being applied before the TTC restart ($s_{diff} = \frac{samples_ADC_BAS}{8} - samples$).

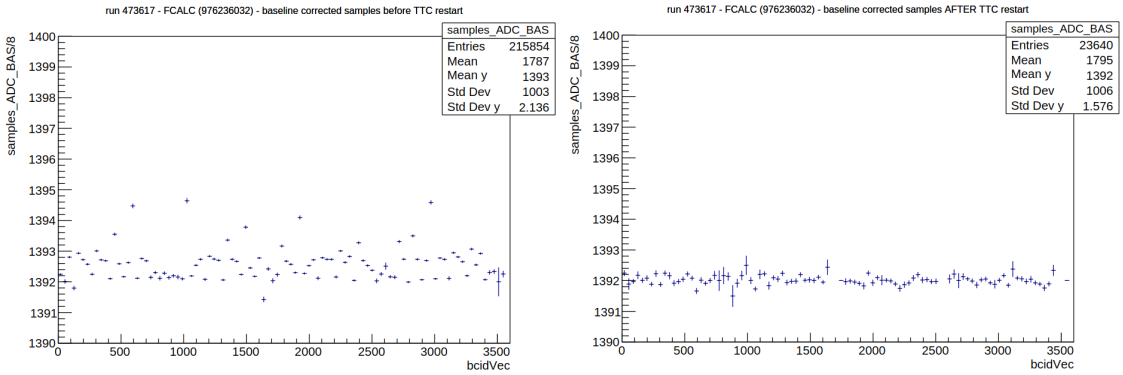

(a) (b)

Figure A.11: The profile of the baseline corrected samples [in ADC counts] before (a) and after the TTC restart (b) as a function of bunch crossing. The baseline shift is evident before the TTC restart due to the baseline correction not being applied properly.

A.5 Pedestal jumps study

Pedestal jumps, also referred to as ADC jumps, constitute a persistent and operationally significant anomaly in the liquid argon system that affect the operation of the whole ATLAS detector. These events are characterized by sudden shifts in the pedestal, or *baseline*, of the digitized samples from the Digital Trigger, such as the one shown in the example in Figure A.12. They occur randomly across various parts of the detector, often appearing in groups of SCs, and can cause disruptive conditions, most notably leading to *Fixed Frequency Trigger Veto* (FFTV) that hinder data acquisition and could

even damage the detector.

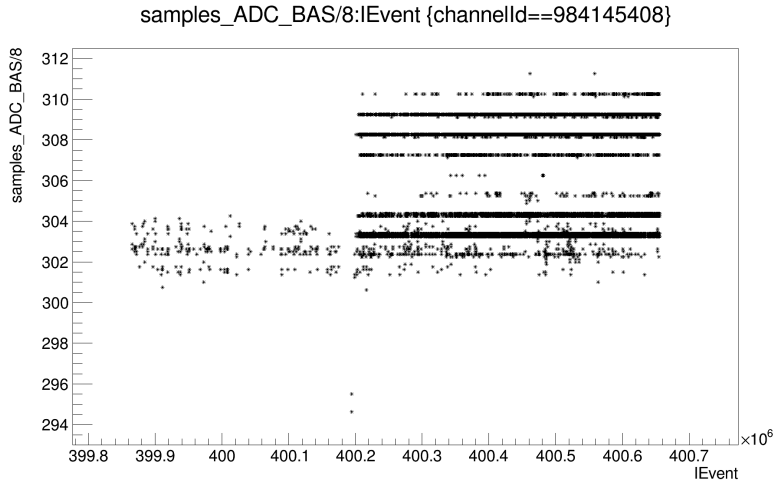

Figure A.12: An example of an instance of a pedestal jump during a stable beams run that caused a FFTV. The samples experience a sudden jump and remain at high counts.

The FFTV is a protective mechanism originally developed for the ATLAS SCT to prevent damage to the silicon-module bond wires from mechanical resonances that can be excited by triggers occurring at specific frequencies in the magnetic field of the Inner Detector [139]. In the ATLAS trigger system, regularly spaced triggers can arise from various detector anomalies, such as noise bursts or the pedestal jumps described in this section, and pose risks both to the SCT hardware and to the synchronization of the readout electronics. The FFTV is implemented in the TTC Interface Module (TIM) firmware of the SCT partition [140, 141], and works by identifying and vetoing triggers that fall within these dangerous frequency patterns, thereby preserving the integrity of data collection and ensuring stable detector operation, albeit at the cost of suspending data taking while the veto is active.

Currently, there is no known method to predict or systematically prevent pedestal jumps. Masking the affected SCs remains the sole immediate operational mitigation, underlining the importance of fully characterizing these jumps to maintain detector efficiency during data-taking periods. Efforts to study pedestal jumps have focused on developing robust algorithms for their detection and characterization:

- Duration and length of jumps;
- ADC count variations;
- Topological patterns (e.g., grouping of affected SCs or location in the detector);
- Periodicity and rates;

Distinguishing pedestal jumps from genuine physics activity (such as real pulses or pileup) poses a particular challenge. For this reason, the study focuses on *empty* data, obtained from empty bunch crossings, using background bunch group keys such as `BCRVeto`, `AllWithoutCalreq`, or `Empty` [142].

- *BCRVeto:* This group designates a short interval of bunch crossings around the LHC BCR signal, during which triggers are generally vetoed. The purpose is to allow proper distribution of the BCR signal to the on-detector electronics and to avoid collecting potentially problematic data in this narrow window.
- *Empty:* The Empty bunch group consists of bunch crossings where there are no proton bunches present in either direction. These crossings are typically used for collecting noise, cosmic rays, or calibration triggers and serve as a clean reference sample, free from pp collisions.
- *AllWithoutCalreq:* This group refers to all bunch crossings except those specifically reserved for calibration requests. It is used to select data that excludes calibration triggers.

Preprocessing steps to extract the data include:

1. Filtering for empty bunch data to filter out physics-induced signals
2. Removing invalid or "bad" SCs from the sample
3. Calculating pedestals and pedestal noise for each channel and run

Initial approaches: pull distribution

The initial analytic technique focused on the pull distribution, defined as:

$$\text{Pull} = \frac{\text{samples} - \text{pedestal}}{\text{noise}} \tag{A.1}$$

A "clean" example is shown in Figure A.13, where the plotted pull values of an SC follow a Gaussian distribution. The idea is to set a cutoff for extreme values and study the resulting outliers, that given the extreme pull value, should present issues in the pedestal that *could* come from an ADC jump.

Although plotting this distribution highlighted channels with anomalous behaviors, it did not effectively reveal pedestal jumps. Rather, most results showed two channels in the EMB with highly-deviating values on the third sample, suggesting a firmware issue limited to this run. This behavior is shown in Figure A.14.

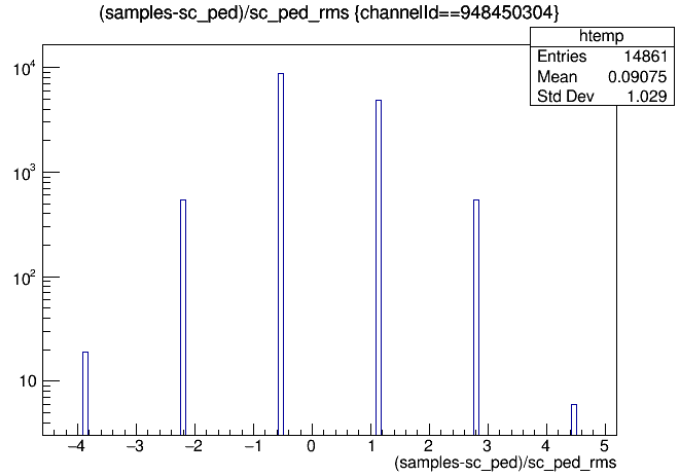

Figure A.13: An example of an SC pull distribution for the clean data from run 480620. The approach using a pull distribution should yield a Gaussian distribution from which to pick the outliers.

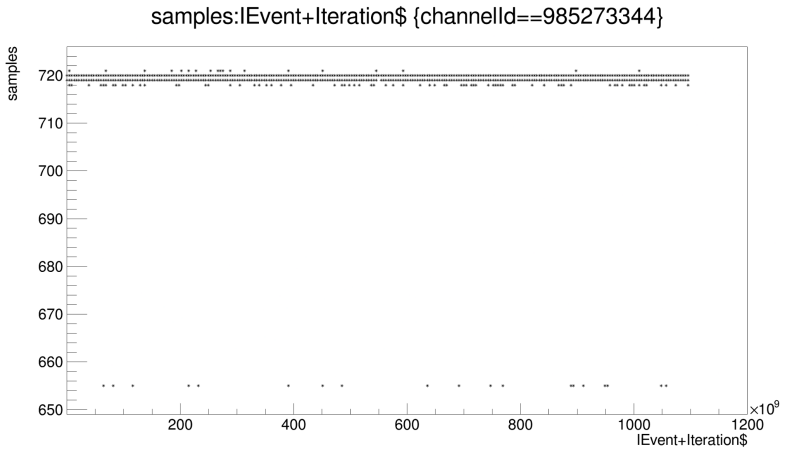

Figure A.14: One of the offending channels (985273344) from run 480620 showing a consistently deviating value on the 3rd sample, suggesting a firmware issue.

Ratio of samples RMS to pedestal noise

A next strategy involved evaluating the ratio of the RMS of the samples to the pedestal noise from the database:

$$\frac{\text{RMS(samples)}}{\text{pedestal noise (database)}} \tag{A.2}$$

Empty data and exclusion of bad channels remained crucial. Several runs were investigated, and the method identified several outlier channels, and highlighted some issues in the pedestal. One example of this is the investigation of run 484074. The distribution of the calculated ratios is illustrated in Figure A.15.

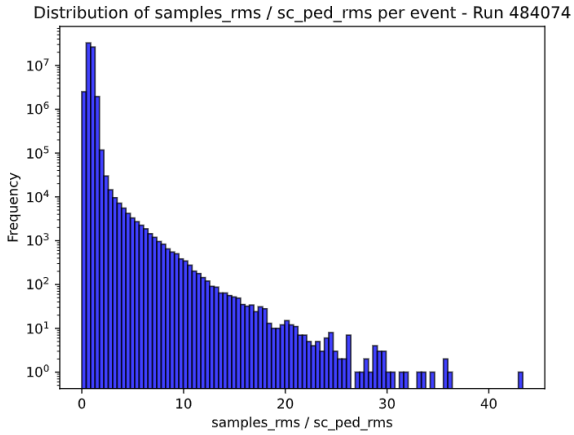

Figure A.15: The distribution of the ratio between the samples RMS and the corresponding pedestal noise from the database. A cutoff around 30 was set to isolate potential outliers.

Figure A.16 shows how the results from the outliers of this distribution, filtered out from a cutoff value of 30, revealed issues in the pedestal not related to ADC jumps. Clear identification of pedestal jumps, however, was still elusive, likely due to the high prescale factor in the MON data, where it is unlikely to observe a rare jump. This prompted a change in the strategy, and suggested that unprescaled data was needed.

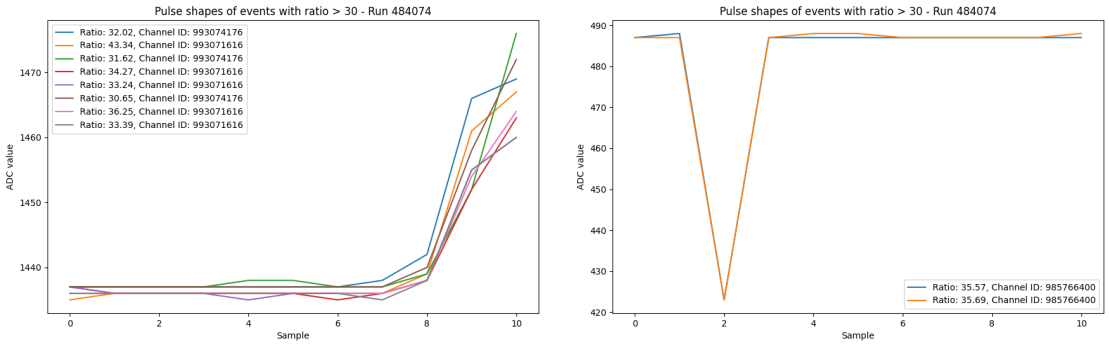

(a) (b)

Figure A.16: The two types of deviating channels revealed by the method. The plot on the left shows a constant increase in the ADC value, suggesting it might be the "anticipation" of an incoming physics pulse. The second shows a dip in the pedestal value at every 3rd sample, suggesting a firmware issue.

Usage of SWROD data

A breakthrough emerged with the use of SWROD raw data in specific runs where pedestal jumps had been previously confirmed, such as in run 489909, where a pedestal

jump caused an FFTV on LB 410 on 4 SCs, as evidenced in Figure A.17.

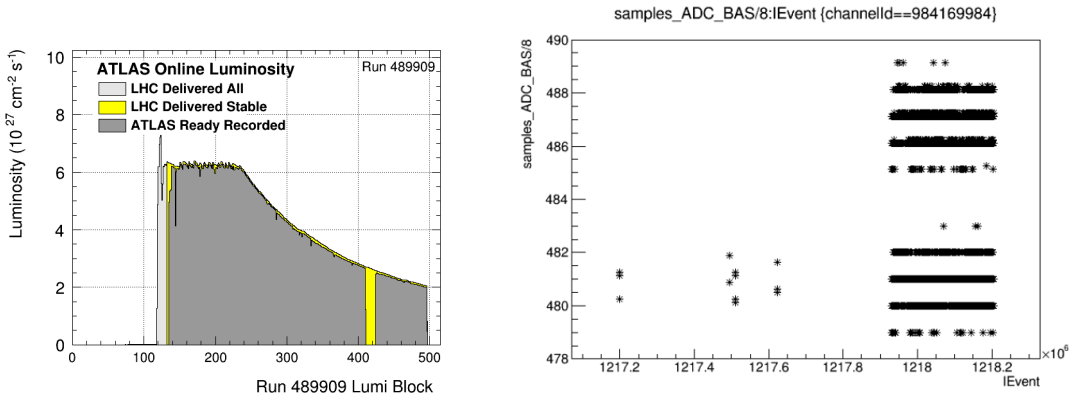

(a) (b)

Figure A.17: The FFTV caused by the ADC jump can be clearly seen in the online luminosity profile of run 489909 in LB 410, where ATLAS could not record any data due to the protection mechanism.

The refined algorithm included:

- Selecting events with a high RMS-to-noise ratio (e.g., > 3)
- Further filtering for a sustained high mean after the second sample, exceeding the channel's pedestal by a multiple of its noise (this characteristic is shown in the pedestal jump in Figure A.17b).

By plotting the distribution of the ratio between the RMS of the samples and the electronic noise of the corresponding channel, a large number of outliers was revealed. However, after investigation, it was found that all of the most extreme outliers were caused by invalid samples with default value -0.125 ADC. This is illustrated clearly in Figure A.18. This prompted the decision to look at a smaller interval of outliers.

Focusing on outliers with ratio between 3–7, and selecting events with a sustained high value of ADC after the second samples successfully flagged events corresponding to known pedestal jumps in specific SCs. The results are shown in Figure A.19.

Investigation of baseline correction test run 493045

Run 493045, a dedicated baseline correction study using SWROD data, was used for further testing, given that no prescale was applied, and no cuts to get empty bunch crossing data were needed (as a test run with no collisions in the LHC, all the events recorded are "empty"). A new data selection criteria to skim the large amount of data

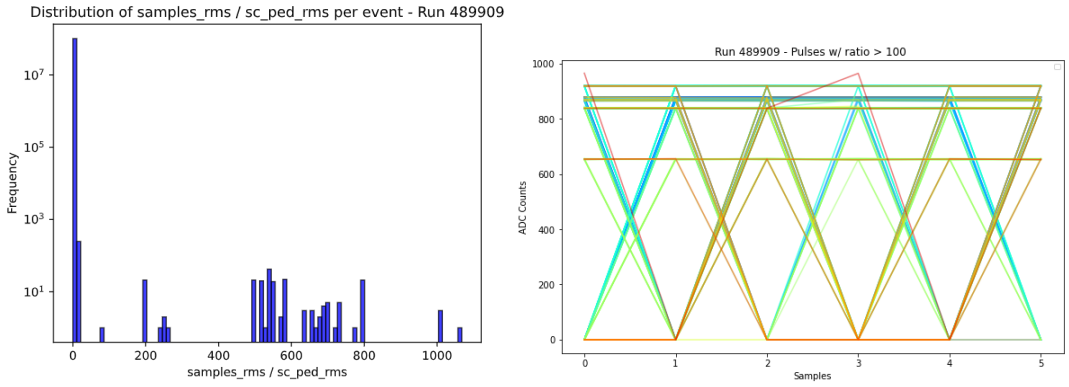

(a) (b)

Figure A.18: The distribution of the ratio (a) showing a large number of outliers, most of them coming from invalid samples (-0.125 ADC). As evidenced in (b), the pulses have samples with recurring values of -0.125.

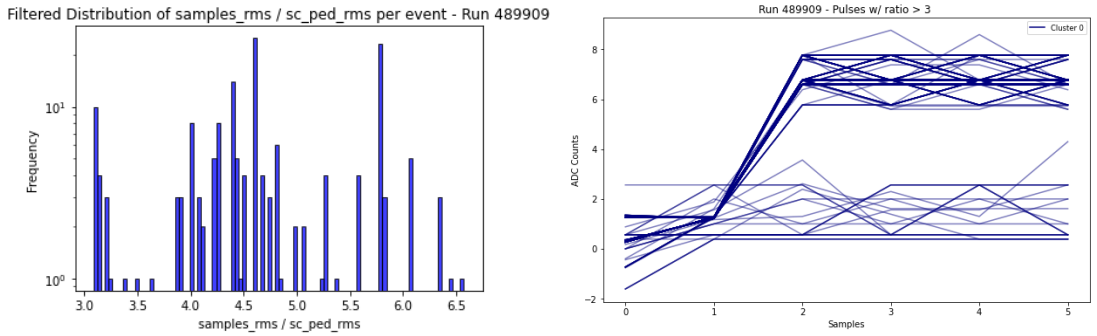

(a) (b)

Figure A.19: (a) The distribution of the events with ratio between 3 and 7, and (b) the pulses with subtracted pedestal showing ADC jumps coming from these events.

was decided:

$$\text{samples_maxmin} > 6\,\text{ADC} \tag{A.3}$$

$$\text{Sum}\left(|\text{samples_jumps}| > 6\right) = 1 \quad (\text{Single jump} > 6) \tag{A.4}$$

Once the data was skimmed, the RMS-to-noise ratio method was applied. Once again, the ratio between 3 and 7 was chosen, and it revealed a large number of ADC jumps, as demonstrated in Figure A.20.

This approach effectively highlighted localized pedestal jumps, particularly in certain ADCs. After selecting events from the run that passed the selections described above, it was important to check where these events came from. In a distribution *per LTDB*, it was revealed that most of these events with ADC jumps indeed come from specific ADCs. This was demonstrated, for instance, for LTDB 105 and 120, where most

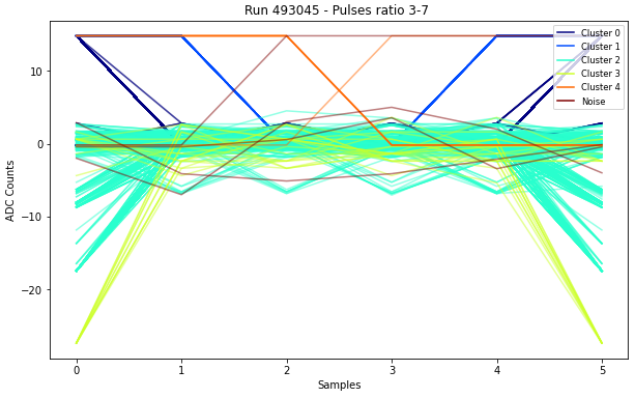

Figure A.20: The results of the new method on the SWROD data revealed a high number of ADC jumps, as seen in the baseline-subtracted pulses illustrated in the plot.

of the jumps come from ADC 64 and 46, respectively. The corresponding jumps were also plotted, with an example of the behavior over time of one of the offending SCs. This is illustrated in Figures A.21 and A.22.

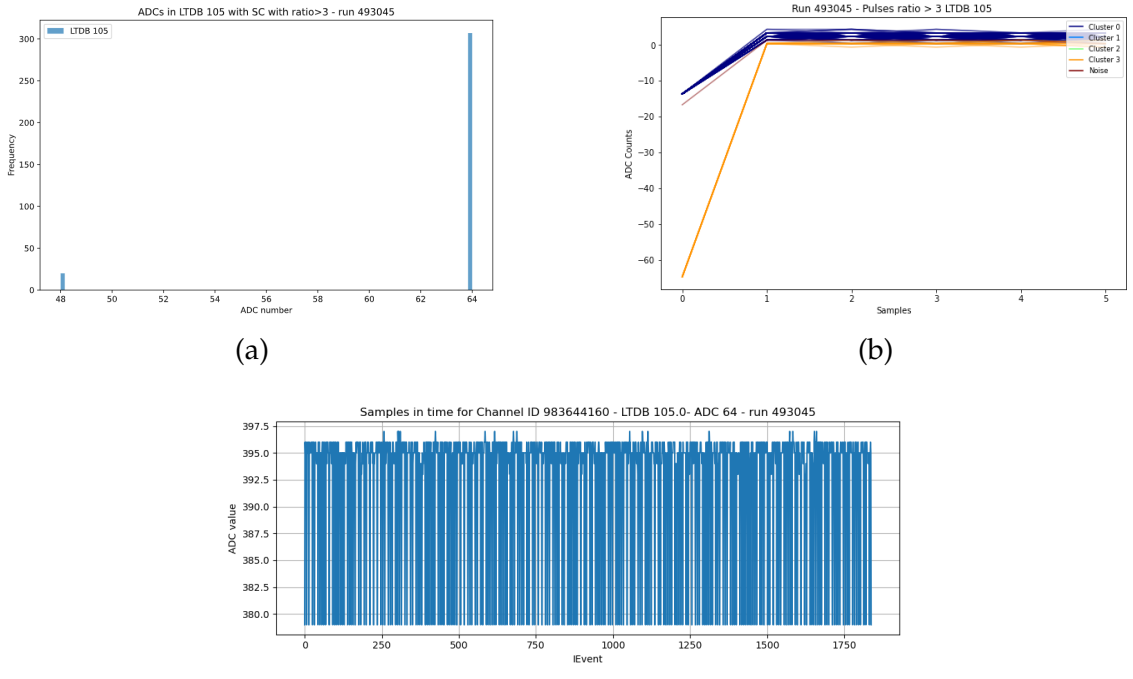

(a) (b)

(c)

Figure A.21: The jumps coming from LTDB 105 are localized mostly in ADC 64 (a). The jumps are shown in (b), and an example of an offending SC jumping over time is shown in (c).

The observation that pedestal jumps are localized within individual ADCs indicates

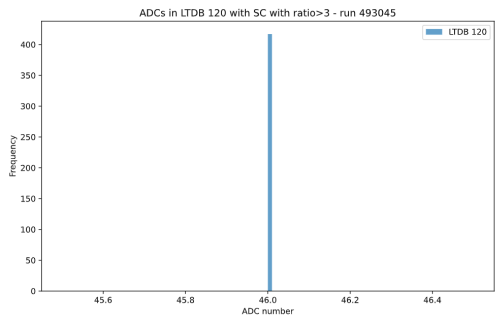

(a)

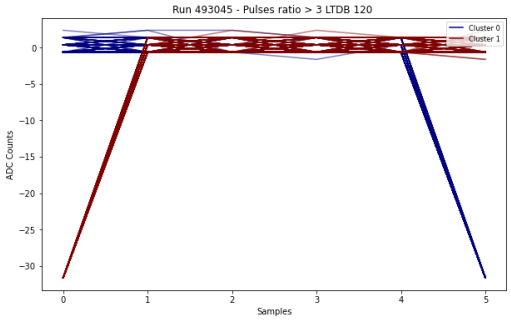

(b)

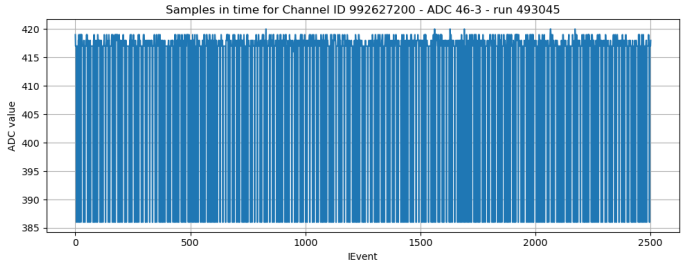

(c)

Figure A.22: All the jumps coming from LTDB 120 are localized mostly in ADC 46. The jumps are shown in (b), and an example of an offending SC jumping over time is shown in (c)

that the origin of the issue lies in the ADC electronics themselves. Because the ADCs are installed on the front-end LTDBs, these components cannot be accessed or replaced during routine operation and require a long shutdown for intervention, so any recommendations or improvements for the ADCs will be considered as part of the Phase-II upgrades for the High-Luminosity LHC.

Appendix B

Two-dimensional signal-region optimization scans

This appendix collects the two-dimensional scans for SR2 and SR3 at the benchmark point ($m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$) that complement the SR1 scans shown in Section 11.7: the acceptance relative to the triggered events (Figures B.1 and B.2) and the corresponding expected signal yields (Figures B.3 and B.4).

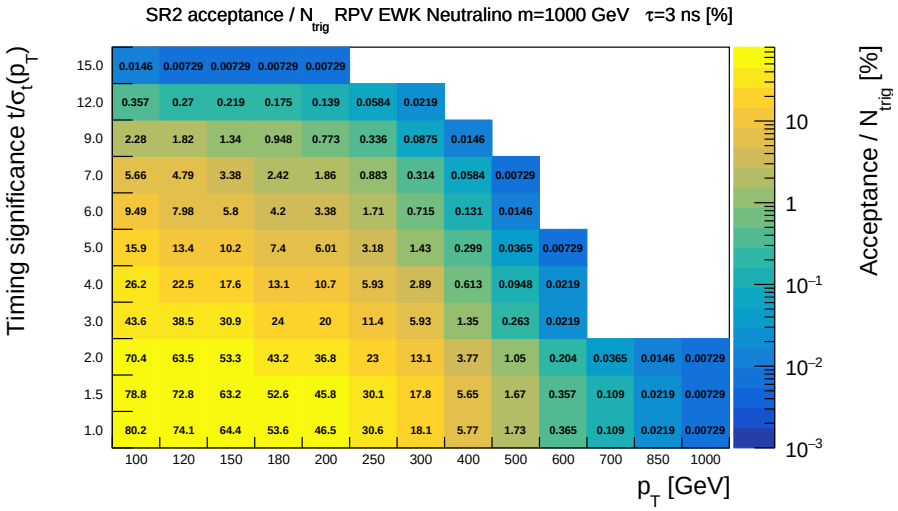

Figure B.1: SR2 acceptance in the plane of the jet p_T threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\text{GeV}$, $\tau = 3\,\text{ns}$ benchmark, normalized as in Figure 11.22. Both jets are required to exceed the threshold and to lie in the same detector hemisphere.

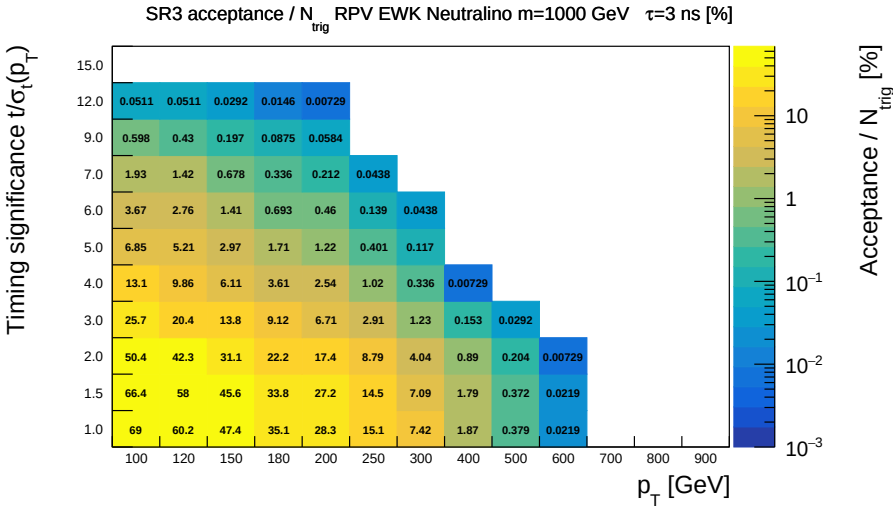

Figure B.2: SR3 acceptance in the plane of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark, normalized as in Figure 11.22. All three jets are required to exceed the threshold, with no hemisphere requirement.

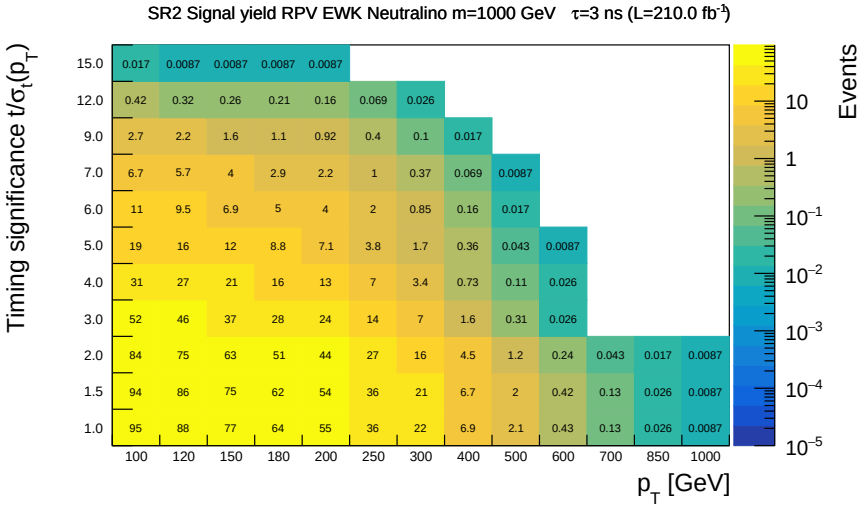

Figure B.3: Expected signal yields in SR2 as a function of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark and $\mathcal{L}_{\mathrm{int}} = 210\,\mathrm{fb}^{-1}$. The map is obtained from the acceptance normalized to N_{gen} with Eq. (11.5), and is therefore the same scan as Figure B.1 rescaled by the production cross section and the integrated luminosity.

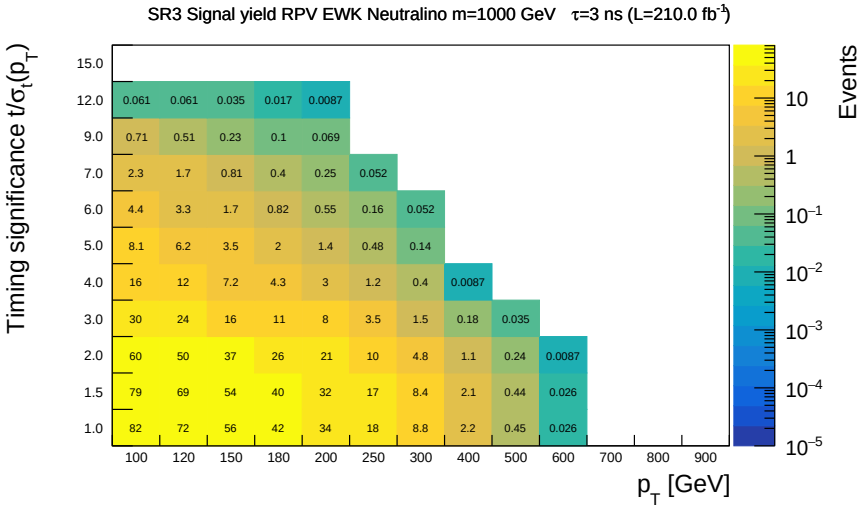

Figure B.4: Expected signal yields in SR3 as a function of the jet p_{T} threshold and the delay significance t/σ_t, for the $m_{\tilde{\chi}_1^0} = 1000\,\mathrm{GeV}$, $\tau = 3\,\mathrm{ns}$ benchmark and $\mathcal{L}_{\mathrm{int}} = 210\,\mathrm{fb}^{-1}$. The map is obtained from the acceptance normalized to N_{gen} with Eq. (11.5), and is therefore the same scan as Figure B.2 rescaled by the production cross section and the integrated luminosity.

Bibliography

[1] Glashow, Sheldon L. "Partial Symmetries of Weak Interactions". In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.

[2] Weinberg, Steven. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

[3] Salam, Abdus. "Weak and Electromagnetic Interactions". In: *Elementary Particle Theory: Relativistic Groups and Analyticity, Proceedings of the Eighth Nobel Symposium*. Ed. by Nils Svartholm. Stockholm: Almqvist & Wiksell, 1968, pp. 367–377.

[4] Englert, F. and Brout, R. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321.

[5] Higgs, Peter W. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508.

[6] Guralnik, G. S., Hagen, C. R., and Kibble, T. W. B. "Global Conservation Laws and Massless Particles". In: *Phys. Rev. Lett.* 13 (1964), pp. 585–587. DOI: 10.1103/PhysRevLett.13.585.

[7] Cush. *Standard Model of Elementary Particles Anti*. 2018. URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles_Anti.svg.

[8] LHCb Collaboration. "Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \to J/\psi K^- p$ decays". In: *Phys. Rev. Lett.* 115 (2015). 48 pages, 18 figures including the supplementary material, v2 after referee's comments, now 19 figures, p. 072001. DOI: 10.1103/PhysRevLett.115.072001. arXiv: 1507.03414. URL: https://cds.cern.ch/record/2033887.

[9] Particle Data Group. "Review of Particle Physics". In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[10] Noether, Emmy. "Invariante Variationsprobleme". In: *Nachr. Ges. Wiss. Göttingen, Math.-Phys. Kl.* 1918 (1918), pp. 235–257.

[11] Peskin, Michael E. and Schroeder, Daniel V. *An Introduction to Quantum Field Theory*. Reading, MA: Addison-Wesley, 1995. ISBN: 978-0-201-50397-5.

[12] Cabibbo, Nicola. "Unitary Symmetry and Leptonic Decays". In: *Phys. Rev. Lett.* 10 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.

[13] Kobayashi, Makoto and Maskawa, Toshihide. "CP Violation in the Renormalizable Theory of Weak Interaction". In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652.

[14] Gross, D. J. and Wilczek, F. "Ultraviolet Behavior of Non-Abelian Gauge Theories". In: *Phys. Rev. Lett.* 30 (1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343.

[15] Politzer, H. David. "Reliable Perturbative Results for Strong Interactions?" In: *Phys. Rev. Lett.* 30 (1973), pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346.

[16] Lee, Lawrence, Ohm, Christian, Soffer, Abner, and Yu, Tien-Tien. "Collider searches for long-lived particles beyond the Standard Model". In: *Prog. Part. Nucl. Phys.* 106 (May 2019), pp. 210–255. ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2019.02.006. URL: http://dx.doi.org/10.1016/j.ppnp.2019.02.006.

[17] Buckley, Andy et al. "General-purpose event generators for LHC physics". In: *Phys. Rept.* 504 (2011), pp. 145–233. DOI: 10.1016/j.physrep.2011.03.005. arXiv: 1101.2599 [hep-ph].

[18] Gieseke, S. *Event Generators and Monte Carlo Methods*. https://indico.desy.de/event/5064/sessions/7828/. Lectures at the DESY Monte Carlo School 2012. Mar. 2012.

[19] Collins, John C., Soper, Davison E., and Sterman, George. "Factorization for Short Distance Hadron-Hadron Scattering". In: *Nucl. Phys. B* 261 (1985), pp. 104–142. DOI: 10.1016/0550-3213(85)90565-6.

[20] UA1 Collaboration. "Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV". In: *Phys. Lett. B* 122 (1983), pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.

[21] UA2 Collaboration. "Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider". In: *Phys. Lett. B* 122 (1983), pp. 476–485. DOI: 10.1016/0370-2693(83)91605-2.

[22] UA1 Collaboration. "Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/c^2 at the CERN SPS Collider". In: *Phys. Lett. B* 126 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.

[23] UA2 Collaboration. "Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ Collider". In: *Phys. Lett. B* 129 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.

[24] CDF Collaboration. "Observation of Top Quark Production in $\bar{p}p$ Collisions with the Collider Detector at Fermilab". In: *Phys. Rev. Lett.* 74 (1995), pp. 2626–2631. DOI: 10.1103/PhysRevLett.74.2626. arXiv: hep-ex/9503002.

[25] D0 Collaboration. "Observation of the Top Quark". In: *Phys. Rev. Lett.* 74 (1995), pp. 2632–2637. DOI: 10.1103/PhysRevLett.74.2632. arXiv: hep-ex/9503003.

[26] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[27] CMS Collaboration. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].

[28] "Comprehensive analysis of the $B^0 \to K^{*0}\mu^+\mu^-$ decay". In: *Phys. Rev. Lett.* 137 (2026), p. 021802. DOI: 10.1103/PhysRevLett.137.021802. arXiv: 2512.18053.

[29] Planck Collaboration. "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[30] Super-Kamiokande Collaboration. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

[31] SNO Collaboration. "Direct Evidence for Neutrino Flavor Transformation from Neutral Current Interactions in the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301. arXiv: nucl-ex/0204008.

[32] Georgi, H. and Glashow, S. L. "Unity of All Elementary Particle Forces". In: *Phys. Rev. Lett.* 32 (1974), pp. 438–441. DOI: 10.1103/PhysRevLett.32.438.

[33] Super-Kamiokande Collaboration. "Search for proton decay via $p \to e^+\pi^0$ and $p \to \mu^+\pi^0$ with an enlarged fiducial volume in Super-Kamiokande I-IV". In: *Phys. Rev. D* 102.11 (2020), p. 112011. DOI: 10.1103/PhysRevD.102.112011. arXiv: 2010.16098 [hep-ex].

[34] Martin, Stephen P. "A Supersymmetry Primer". In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998), pp. 1–98. DOI: 10.1142/9789812839657_0001. arXiv: hep-ph/9709356.

[35] Resseguie, Elodie. "Searches for Electroweak SUSY: Motivation and Models". In: *Electroweak Physics at the Large Hadron Collider with the ATLAS Detector: Standard Model Measurement, Supersymmetry Searches, Excesses, and Upgrade Electronics*. Cham: Springer International Publishing, 2020, pp. 121–129. ISBN: 978-3-030-57016-3. DOI: 10.1007/978-3-030-57016-3_6. URL: https://doi.org/10.1007/978-3-030-57016-3_6.

[36] Arkani-Hamed, Nima and Dimopoulos, Savas. "Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC". In: *JHEP* 06 (2005), p. 073. DOI: 10.1088/1126-6708/2005/06/073. arXiv: hep-th/0405159 [hep-th].

[37] Giudice, G. F. and Romanino, A. "Split supersymmetry". In: *Nucl. Phys. B* 699 (2004). [Erratum: Nucl. Phys. B 706 (2005) 65], pp. 65–89. DOI: 10.1016/j.nuclphysb.2004.08.001. arXiv: hep-ph/0406088 [hep-ph].

[38] Giudice, G. F. and Rattazzi, R. "Theories with gauge-mediated supersymmetry breaking". In: *Phys. Rept.* 322 (1999), pp. 419–499. DOI: 10.1016/S0370-1573(99)00042-3. arXiv: hep-ph/9801271 [hep-ph].

[39] Barbier, R. et al. "R-parity violating supersymmetry". In: *Phys. Rept.* 420 (2005), pp. 1–202. DOI: 10.1016/j.physrep.2005.08.006. arXiv: hep-ph/0406039 [hep-ph].

[40] Neutelings, Izaak. *Signatures of long-lived particlesSignatures of long-lived particles*. 2026. URL: https://tikz.net/bsm_longlived/.

[41] Evans, Lyndon and Bryant, Philip. "LHC Machine". In: *JINST* 3.08 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001. URL: https://cds.cern.ch/record/1129806.

[42] Arduini, G. et al. "LHC Upgrades in preparation of Run 3". In: *JINST* 19.05 (2024), P05061. DOI: 10.1088/1748-0221/19/05/P05061.

[43] Brüning, O, Burkhardt, H, and Myers, S. "The Large Hadron Collider". In: *Prog. Part. Nucl. Phys.* 67 (2012). Published in "Progress in Particle and Nuclear Physics" (2012), 10.1016/j.ppnp.2012.03.001, pp. 705–734. DOI: 10.1016/j.ppnp.2012.03.001. URL: https://cds.cern.ch/record/1443022.

[44] Lefèvre, Christiane. "The CERN accelerator complex". 2008. URL: https://cds.cern.ch/record/1260465.

[45] Bailey, R and Collier, Paul. *Standard Filling Schemes for Various LHC Operation Modes*. Tech. rep. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/691782.

[46] ATLAS Collaboration. *Non-collision backgrounds as measured by the ATLAS detector during the 2010 proton-proton run*. Tech. rep. Geneva: CERN, 2011. URL: https://cds.cern.ch/record/1383840.

[47] Herr, Werner and Muratori, B. "Concept of luminosity". In: (2006). DOI: 10.5170/CERN-2006-002.361. URL: https://cds.cern.ch/record/941318.

[48] ATLAS Collaboration. "The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3". In: *JINST* 19.05 (2024), P05063. DOI: 10.1088/1748-0221/19/05/P05063. arXiv: 2305.16623. URL: https://cds.cern.ch/record/2859916.

[49] ATLAS Collaboration. *ATLAS Luminosity Public Results for Run 3*. 2026. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3.

[50] Dominguez, Daniel, Moles, Kevin Patrice, and Mehlhase, Sascha. *ATLAS detector schematics*. 2021. URL: https://cds.cern.ch/record/2777214.

[51] ATLAS Collaboration. "The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions". In: *Eur. Phys. J. C* 85 (2025), p. 606. DOI: 10.1140/epjc/s10052-025-14062-8. arXiv: 2402.05858. URL: https://cds.cern.ch/record/2888751.

[52] ATLAS Collaboration. *ATLAS inner detector: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331063.

[53] ATLAS Collaboration. "Alignment of the ATLAS Inner Detector in Run-2". In: *Eur. Phys. J. C* 80.12 (2020), p. 1194. DOI: 10.1140/epjc/s10052-020-08700-6. arXiv: 2007.07624. URL: https://cds.cern.ch/record/2724037.

[54] ATLAS Collaboration. *ATLAS Insertable B-Layer: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 2010. URL: https://cds.cern.ch/record/1291633.

[55] ATLAS IBL Collaboration. "Production and Integration of the ATLAS Insertable B-Layer". In: *JINST* 13.05 (2018), T05008. DOI: 10.1088/1748-0221/13/05/T05008. arXiv: 1803.00844.

[56] ATLAS Collaboration. *ATLAS pixel detector: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1998. URL: https://cds.cern.ch/record/381263.

[57] ATLAS Collaboration. "Operation and performance of the ATLAS semiconductor tracker in LHC Run 2". In: *JINST* 17 (2022), P01013. DOI: 10.1088/1748-0221/17/01/P01013. arXiv: 2109.02591. URL: https://cds.cern.ch/record/2780336.

[58] ATLAS Collaboration. "The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance". In: *JINST* 3 (2008), P02013. DOI: 10.1088/1748-0221/3/02/P02013. URL: https://cds.cern.ch/record/1094549.

[59] ATLAS Collaboration. "The ATLAS Trigger System for LHC Run 3 and Trigger performance in 2022". In: *JINST* 19 (2024), P06029. DOI: 10.1088/1748-0221/19/06/P06029. arXiv: 2401.06630. URL: https://cds.cern.ch/record/2886407.

[60] ATLAS Collaboration. *ATLAS calorimeter performance: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. URL: https://cds.cern.ch/record/331059.

[61] Pequenao, Joao. "Computer Generated image of the ATLAS calorimeter". 2008. URL: https://cds.cern.ch/record/1095927.

[62] ATLAS Collaboration. *ATLAS muon spectrometer: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[63] Bauer, F. et al. "Construction and Test of MDT Chambers for the ATLAS Muon Spectrometer". In: *Nucl. Instrum. Meth. A* 461 (2001), pp. 17–20. DOI: 10.1016/S0168-9002(00)01156-6. URL: https://cds.cern.ch/record/2145043.

[64] Wotschack, Jorg. *ATLAS muon chamber construction parameters for CSC, MDT, and RPC chambers*. Tech. rep. CERN, Apr. 2008.

[65] Stelzer, Bernd. *The New Small Wheel Upgrade Project of the ATLAS Experiment*. Tech. rep. Geneva: CERN, 2016. DOI: 10.1016/j.nuclphysbps.2015.09.182. URL: https://cds.cern.ch/record/1958265.

[66] Pequenao, Joao. "Computer generated image of the ATLAS Muons subsystem". 2008. URL: https://cds.cern.ch/record/1095929.

[67] ATLAS Collaboration. *Performance of the ATLAS RPC detector and Level-1 muon barrel trigger at TeV*. 2019. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/MDET-2019-01/.

[68] ATLAS Collaboration. *ATLAS magnet system: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997. DOI: 10.17181/CERN.9O5C.VDTM. URL: https://cds.cern.ch/record/338080.

[69] Rodriguez Vera, Ana Maria and Antunes Pequenao, Joao. "ATLAS Detector Magnet System". General Photo. 2021. URL: https://cds.cern.ch/record/2770604.

[70] Buttinger, W. *The ATLAS Level-1 Trigger System*. Tech. rep. Geneva: CERN, 2012. DOI: 10.1088/1742-6596/396/1/012010. URL: https://cds.cern.ch/record/1456546.

[71] ATLAS TDAQ Collaboration. "New Software Based Readout Driver for the ATLAS Experiment". In: *IEEE Trans. Nucl. Sci.* 68.8 (2021), pp. 2233–2240. DOI: 10.1109/TNS.2021.3085970. arXiv: 2010.14884 [physics.ins-det].

[72] ATLAS Collaboration. *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/616089.

[73] ATLAS Collaboration. *ATLAS Computing: Technical Design Report*. Tech. rep. CERN-LHCC-2005-022, ATLAS-TDR-17. CERN, 2005. URL: https://cds.cern.ch/record/837738.

[74] Fabjan, Christian Wolfgang and Gianotti, F. "Calorimetry for Particle Physics". In: *Rev. Mod. Phys.* 75 (2003), pp. 1243–1286. DOI: 10.1103/RevModPhys.75.1243. URL: https://cds.cern.ch/record/692252.

[75] ATLAS Collaboration. *ATLAS liquid-argon calorimeter*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: 10.17181/CERN.FWRW.FOOQ. URL: https://cds.cern.ch/record/331061.

[76] Wigmans, Richard. *Calorimetry: Energy Measurement in Particle Physics*. Oxford University Press, Sept. 2017. ISBN: 9780198786351. DOI: 10.1093/oso/9780198786351.001.0001. URL: https://doi.org/10.1093/oso/9780198786351.001.0001.

[77] Groom, D. and Klein, S. "Passage of particles through matter". In: *Eur. Phys. J. C* 15 (Mar. 2000), pp. 163–173. DOI: 10.1007/BF02683419.

[78] ATLAS Collaboration. *ATLAS tile calorimeter*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: 10.17181/CERN.JRBJ.7O28. URL: https://cds.cern.ch/record/331062.

[79] ATLAS Collaboration. "Operation and performance of the ATLAS tile calorimeter in LHC Run 2". In: *Eur. Phys. J. C* 84.12 (2024), p. 1313. DOI: 10.1140/epjc/s10052-024-13151-4. arXiv: 2401.16034. URL: https://cds.cern.ch/record/2887689.

[80] ATLAS Collaboration. "ATLAS Liquid Argon Calorimeter 2m prototype". 1990. URL: https://cds.cern.ch/record/2254084.

[81] ATLAS Collaboration. *ATLAS Public Results - Public Liquid-Argon Calorimeter Plots on Detector Status*. 2024. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LArCaloPublicResultsDetStatus.

[82] Krieger, P. "The ATLAS liquid argon calorimeter". In: (2006). URL: https://cds.cern.ch/record/1015504.

[83] Bélanger, G et al. *The ATLAS Liquid Argon Forward Calorimeter*. Tech. rep. Geneva: CERN, 2005. URL: https://cds.cern.ch/record/919047.

[84] Aubert, B et al. "Development and Construction of Large Size Signal Electrodes for the ATLAS Electromagnetic Calorimeter". In: *Nucl. Instrum. Meth. A* 539 (2005), pp. 558–594. DOI: 10.1016/j.nima.2004.11.005. URL: https://cds.cern.ch/record/737956.

[85] Colas, J. et al. "Response Uniformity of the ATLAS Liquid Argon Electromagnetic Calorimeter". In: *Nucl. Instrum. Meth. A* 582 (2007), pp. 429–455. DOI: 10.1016/j.nima.2007.08.157. arXiv: 0709.1094 [physics.ins-det].

[86] ATLAS Collaboration. "Drift Time Measurement in the ATLAS Liquid Argon Electromagnetic Calorimeter using Cosmic Muons". In: *Eur. Phys. J. C* 70 (2010). Comments: 30 pages, 25 figures, accepted by EPJC, pp. 755–785. DOI: 10.1140/epjc/s10052-010-1403-6. arXiv: 1002.4189. URL: https://cds.cern.ch/record/1244611.

[87] Buchanan, N J et al. "Design and implementation of the Front End Board for the readout of the ATLAS liquid argon calorimeters". In: *JINST* 3 (2008), P03004. DOI: 10.1088/1748-0221/3/03/P03004. URL: https://cds.cern.ch/record/1099777.

[88] Achenbach, R et al. *The ATLAS Level-1 Calorimeter Trigger*. Tech. rep. Geneva: CERN, 2008. DOI: 10.1088/1748-0221/3/03/P03001. URL: https://cds.cern.ch/record/1080560.

[89] Majewski, S. *Electronic Readout of the ATLAS Liquid Argon Calorimeter: Calibration and Performance*. Tech. rep. Geneva: CERN, 2010. URL: https://cds.cern.ch/record/1270533.

[90] Cleland, W.E. and Stern, E.G. "Signal processing considerations for liquid ionization calorimeters in a high rate environment". In: *Nucl. Instrum. Meth. A* 338.2 (1994), pp. 467–497. ISSN: 0168-9002. DOI: 10.1016/0168-9002(94)91332-3. URL: https://www.sciencedirect.com/science/article/pii/0168900294913323.

[91] Aleksa, M (CERN) et al. *ATLAS Liquid Argon Calorimeter Phase-I Upgrade*. Tech. rep. Final version presented to December 2013 LHCC. CERN, 2013. URL: https://cds.cern.ch/record/1602230.

[92] ATLAS Collaboration. *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*. Tech. rep. CERN-LHCC-2013-018, ATLAS-TDR-023. Geneva: CERN, 2013. URL: https://cds.cern.ch/record/1602235.

[93] HiLumi LHC Project. *High-Luminosity Large Hadron Collider (HL-LHC) schedule*. 2026. URL: https://hilumilhc.web.cern.ch/about-hilumi-lhc/.

[94] ATLAS Collaboration. *ATLAS Phase-II Upgrade Scoping Document*. Tech. rep. CERN-LHCC-2015-020, LHCC-G-166. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2055248.

[95] ATLAS Collaboration. *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*. Tech. rep. CERN-LHCC-2017-021, ATLAS-TDR-030. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2285585.

[96] ATLAS Collaboration. *Technical Design Report for the ATLAS Inner Tracker Strip Detector*. Tech. rep. CERN-LHCC-2017-005, ATLAS-TDR-025. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2257755.

[97] ATLAS Collaboration. *Technical Design Report for the Phase-II Upgrade of the ATLAS LAr Calorimeter*. Tech. rep. CERN-LHCC-2017-018, ATLAS-TDR-027. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2285582.

[98] ATLAS Collaboration. *Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter*. Tech. rep. CERN-LHCC-2017-019, ATLAS-TDR-028. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2285583.

[99] ATLAS Collaboration. *Technical Design Report: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*. Tech. rep. CERN-LHCC-2020-007, ATLAS-TDR-031. Geneva: CERN, 2020. URL: https://cds.cern.ch/record/2719855.

[100] ATLAS Collaboration. *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*. Tech. rep. CERN-LHCC-2017-020, ATLAS-TDR-029. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2285584.

[101] ATLAS Collaboration. *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - Event Filter Tracking Amendment*. Tech. rep. CERN-LHCC-2022-004, ATLAS-TDR-029-ADD-1. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2802799.

[102] Aad, G. et al. "The Phase-I Trigger Readout Electronics Upgrade of the ATLAS Liquid Argon Calorimeters". In: *JINST* 17 (2022). 56 pages, 41 figures, 6 tables, P05024. DOI: 10.1088/1748-0221/17/05/P05024. arXiv: 2202.07384. URL: https://cds.cern.ch/record/2804119.

[103] *Report on the Physics at the HL-LHC, and Perspectives for the HE-LHC*. Tech. rep. CERN-2019-007. Geneva: CERN, 2019. DOI: 10.23731/CYRM-2019-007.

[104] ATLAS Collaboration. *LAr Super Cells Baseline Correction Performance*. 2023. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LArCaloPublicStableBeam2023BC.

[105] ATLAS Collaboration. *LAr Calorimeter Public Results: Digital Trigger, Stable Beams 2022-2024*. 2024. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsLAr.

[106] ATLAS Collaboration. *L1Calo Trigger Public Results*. 2023. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults.

[107] Grafström, P. and Kozanecki, W. "Luminosity determination at proton colliders". In: *Prog. Part. Nucl. Phys.* 81 (2015), pp. 97–148. DOI: 10.1016/j.ppnp.2014.11.002.

[108] ATLAS Collaboration. "Precision measurement and interpretation of inclusive W^+, W^- and Z/γ^* production cross sections with the ATLAS detector". In: *Eur. Phys. J. C* 77.6 (2017), p. 367. DOI: 10.1140/epjc/s10052-017-4911-9. arXiv: 1612.03016. URL: https://cds.cern.ch/record/2238565.

[109] ATLAS Collaboration. "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC". In: *Eur. Phys. J. C* 83 (2023), p. 982. DOI: 10.1140/epjc/s10052-023-11747-w. URL: https://cds.cern.ch/record/2844887.

[110] Avoni, G et al. "The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS". In: *JINST* 13.07 (2018), P07017. DOI: 10.1088/1748-0221/13/07/P07017. URL: https://cds.cern.ch/record/2633501.

[111] Cindro, V et al. "The ATLAS beam conditions monitor". In: *JINST* 3 (2008), P02004. DOI: 10.1088/1748-0221/3/02/P02004. URL: https://cds.cern.ch/record/1094819.

[112] ATLAS Collaboration. *The LUCID-2 detector*. 2026. URL: https://lucid3.web.cern.ch/.

[113] ATLAS Collaboration. "Improved luminosity determination in pp collisions at $\sqrt{s}$ = 7 TeV using the ATLAS detector at the LHC". In: *Eur. Phys. J. C* 73 (2013), p. 2518. DOI: 10.1140/epjc/s10052-013-2518-3. arXiv: 1302.4393. URL: https://cds.cern.ch/record/1517411.

[114] ATLAS Collaboration. *Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s}$ = 13 TeV*. Tech. rep. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2110140.

[115] ATLAS Collaboration. "Readiness of the ATLAS Tile Calorimeter for LHC collisions". In: *Eur. Phys. J. C* 70 (2010). Accepted for publication in EPJC (latest version has minor editorial modifications following requests from EPJC referees), pp. 1193–1236. DOI: 10.1140/epjc/s10052-010-1508-y. arXiv: 1007.5423. URL: https://cds.cern.ch/record/1282535.

[116] ATLAS Collaboration. *LAr Calorimeter Public Results: Luminosity with the Digital Trigger, Stable Beams 2025*. 2024. URL: https://twikiaiext.cern.ch/twiki/bin/view/AtlasPublic/LArCaloPublicStableBeam2025.

[117] Pozzi, Ruben. "Performance of the new ATLAS LAr Calorimeter Digital Trigger and its possible application to luminosity measurement". In: *Nucl. Instrum. Meth. A* 1092 (2026), p. 171851. DOI: 10.1016/j.nima.2026.171851.

[118] Wache, M. and Lehmann Miotto, G. "A Persistent Back-End for the ATLAS Online Information Service (P-BEAST)". In: *J. Phys.: Conf. Ser. (CHEP 2013)*. Vol. 513. 2014, p. 012031. DOI: 10.1088/1742-6596/513/1/012031. URL: https://cds.cern.ch/record/1432912.

[119] Richardson, William Hadley. "Bayesian-Based Iterative Method of Image Restoration". In: *J. Opt. Soc. Am.* 62.1 (1972), pp. 55–59. DOI: 10.1364/JOSA.62.000055.

[120] Lucy, L. B. "An iterative technique for the rectification of observed distributions". In: *Astron. J.* 79 (1974), p. 745. DOI: 10.1086/111605.

[121] ATLAS Liquid Argon Collaboration. *LArWebApp: Liquid Argon SuperCell Rates Monitor*. URL: https://atlasop.cern.ch/LArWebApp/.

[122] ATLAS Collaboration. "Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector". In: *Eur. Phys. J. C* 82 (2022), p. 606. DOI: 10.1140/epjc/s10052-022-10489-5. arXiv: 2201.02472 [hep-ex].

[123] ATLAS Collaboration. "Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *JHEP* 06 (2023), p. 200. DOI: 10.1007/JHEP06(2023)200. arXiv: 2301.13866 [hep-ex]. URL: https://cds.cern.ch/record/2847989.

[124] ATLAS Collaboration. "Search for displaced photons produced in exotic decays of the Higgs boson using 13 TeV pp collisions with the ATLAS detector". In: *Phys. Rev. D* 108.3 (2023). 48 pages in total, author list starting in page 31, 11 figures, 2 tables. All figures including auxiliary figures are available at http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/SUSY-2019-14, p. 032016. DOI: 10.1103/PhysRevD.108.032016. arXiv: 2209.01029. URL: https://cds.cern.ch/record/2826224.

[125] ATLAS Collaboration. *Search in diphoton and dielectron final states for displaced production of Higgs or Z bosons with the ATLAS detector in $\sqrt{s} = 13$ TeV pp collisions*. Tech. rep. 1. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/SUSY-2020-28/. 2023, p. 012012. DOI: 10.1103/PhysRevD.108.012012. arXiv: 2304.12885. URL: https://cds.cern.ch/record/2856888.

[126] Alwall, J. et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[127] Fuks, Benjamin. "Beyond the Minimal Supersymmetric Standard Model: from theory to phenomenology". In: *Int. J. Mod. Phys. A* 27 (2012), p. 1230007. DOI: 10.1142/S0217751X12300074. arXiv: 1202.4769 [hep-ph].

[128] Bierlich, Christian et al. "A comprehensive guide to the physics and usage of PYTHIA 8.3". In: *SciPost Phys. Codebases* 2022 (2022), p. 8. DOI: 10.21468/SciPostPhysCodeb.8. arXiv: 2203.11601 [hep-ph].

[129] ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. Tech. rep. ATL-PHYS-PUB-2014-021. CERN, 2014. URL: https://cds.cern.ch/record/1966419.

[130] Ball, Richard D. et al. "Parton distributions with LHC data". In: *Nucl. Phys. B* 867 (2013), p. 244. DOI: 10.1016/j.nuclphysb.2012.10.003. arXiv: 1207.1303 [hep-ph].

[131] Group, LHC SUSY Cross Section Working. *SUSY Cross Sections for 13, 14, 33 and 100 TeV pp Collisions*. 2026. URL: https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections.

[132] Cacciari, Matteo, Salam, Gavin P., and Soyez, Gregory. "The Anti-k_t jet clustering algorithm". In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[133] Cacciari, Matteo, Salam, Gavin P., and Soyez, Gregory. "FastJet User Manual". In: *Eur. Phys. J. C* 72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[134] ATLAS Collaboration. "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1". In: *Eur. Phys. J. C* 77 (2017), p. 490. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex].

[135] Patzwahl, Marcel. *Feasibility studies for triggering on delayed jets with the ATLAS experiment at the LHC*. July 2023.

[136] ATLAS Collaboration. *ATLAS Data Quality Public Results*. 2024. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/DataQualityResults.

[137] ATLAS Collaboration. "Monitoring and data quality assessment of the ATLAS liquid argon calorimeter". In: *JINST* 9 (2014), P07024. DOI: 10.1088/1748-0221/9/07/P07024. arXiv: 1405.3768 [physics.ins-det].

[138] Gällnö, P. "Timing, trigger and control distribution and dead-time control in ATLAS". In: 2000. DOI: 10.5170/CERN-2000-010.415. URL: https://cds.cern.ch/record/479703.

[139] Barber, T. J., Charlton, D. G., Murray, W., Villani, G., Warren, M., and Weidberg, A. R. "Resonant bond wire vibrations in the ATLAS semiconductor tracker". In: *Nucl. Instrum. Meth. A* (2005). URL: https://cds.cern.ch/record/908860.

[140] Barr, A. J., Postranecky, M., and Warren, M. R. M. *A Fixed-Frequency-Trigger Veto for the ATLAS SCT*. Tech. rep. Geneva: CERN, 2006. URL: https://cds.cern.ch/record/1010263.

[141] Butterworth, J. M., Lane, J., Postranecky, M., and Warren, M. "TTC Interface Module for ATLAS Read-Out Electronics: Final production version based on Xilinx FPGA devices". In: *Proceedings of the 10th Workshop on Electronics for LHC and Future Experiments (LECC 2004)*. 2004. URL: https://www.hep.ucl.ac.uk/atlas/sct/tim/Production_TIM_LECC2004_Paper.pdf.

[142] ATLAS Collaboration. "Operation of the ATLAS trigger system in Run 2". In: *JINST* 15.10 (2020), P10004. DOI: 10.1088/1748-0221/15/10/P10004. arXiv: 2007.12539. URL: https://cds.cern.ch/record/2725146.