Haystacks and Needles

Measuring the number of proton collisions in ATLAS and probing them for the production of new exotic particles

KTH Engineering Sciences

GIULIA RIPELLINO

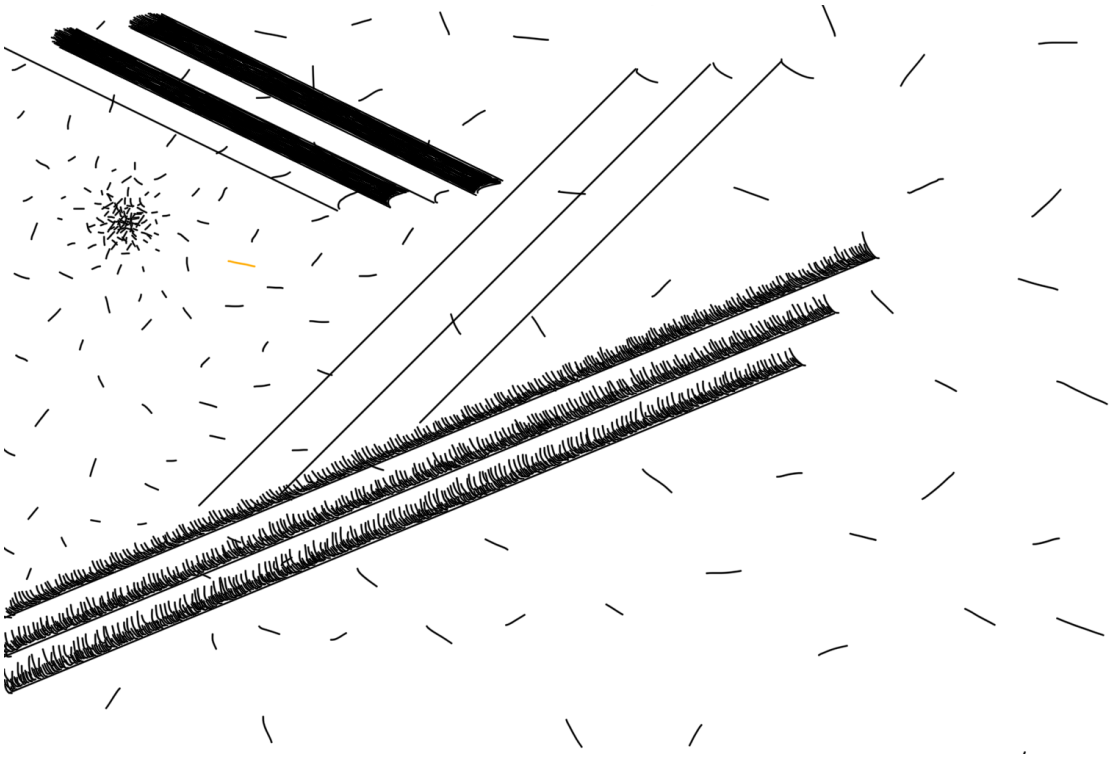

Doctoral Thesis in Physics
Stockholm, Sweden 2022

KTH Engineering Sciences

Doctoral Thesis in Physics

Haystacks and Needles

Measuring the number of proton collisions in ATLAS and probing them for the production of new exotic particles

Giulia Ripellino

Particle and Astroparticle Physics, Department of Physics,
Royal Institute of Technology, SE-106 91 Stockholm, Sweden

Stockholm, Sweden 2022

Cover image: *The One.* Katarina Bendtz

Academic Dissertation which, with due permission of the KTH Royal Institute of Technology, is submitted for public defence for the Degree of Doctor of Philosophy on Thursday 2nd of June 2022 at 1:00 p.m. in room FB41, AlbaNova Universitetscentrum, Roslagstullsbacken 21, Stockholm.

ISBN 978-91-8040-268-2
TRITA-SCI-FOU 2022:30

Printed by: Universitetsservice US-AB, Sweden 2022

Till Elvira och Miranda

Abstract

This thesis describes my work on the ATLAS luminosity measurement and on two searches for physics beyond the Standard Model, performed using $\sqrt{s} = 13\,\text{TeV}$ proton-proton (pp) collision data recorded by the ATLAS detector at CERN during Run 2 of the LHC in the years 2015-2018

ATLAS employs several detectors and methods to measure the relative luminosity and the absolute scale is determined using van der Meer scans during dedicated low-luminosity running periods in each year. The thesis presents the details of the track-counting method which is used to extrapolate the calibration to the high-luminosity physics data-taking regime. After data-quality selections, the full Run 2 pp data sample corresponds to an integrated luminosity of $139\,\text{fb}^{-1}$, with an uncertainty of 1.7 %.

The first analysis presented is a search for Supersymmetry in final states with two same-flavour opposite-charge leptons, jets, and large missing transverse momentum. The search uses $36.1\,\text{fb}^{-1}$ of pp collision data collected in 2015 and 2016. Signal scenarios involving pair-production of squarks or gluinos decaying into final states with two leptons and the lightest supersymmetric particle are targeted by the analysis. The main Standard Model backgrounds are estimated with data-driven methods and the observed data is found to be consistent with the expectation. Results are interpreted in simplified Supersymmetry models and exclude gluinos and squarks with masses up to 1.85 TeV and 1.3 TeV at 95 % confidence level, respectively.

The second analysis is a search for long-lived particles decaying into hadrons. The analysis selects events that contain multiple energetic jets and a displaced vertex and uses the full Run 2 pp collision dataset. Background estimates for Standard Model processes and instrumental effects are extracted from data. The observed yields are one event in the first signal region where 0.49 ± 0.28 are expected and zero events in the second signal region where 0.7 ± 0.4 are expected. The results are used to set limits at 95 % confidence level on Supersymmetry scenarios with long-lived electroweakinos that decay via a small R-parity-violating coupling into triplets of quarks. Electroweakly produced neutralinos with masses up to 1.3 TeV are excluded for lifetimes between 0.01 ns and 10 ns.

Sammanfattning

Elementarpartikelfysik är vetenskapen om naturens minsta beståndsdelar och drivs framåt av vår nyfikenhet och vilja att förstå det universum som vi lever i. Uppfattningen om vad som utgör dessa beståndsdelar har ändrats genom historiens gång och teorier som har antagits beskriva vår värld har gång på gång omkullkastats av nya uptäckter och teoretiska genombrott. Vår nuvarande förståelse sammanfattas av partikelfysikens standardmodell - ett teoretiskt ramverk som beskriver en uppsättning elementarpartiklar och hur dessa växelverkar med varandra och bygger upp all materia. Sammantaget är standardmodellen en fantastiskt framgångsrik teori som har testats i mycket noggranna experiment och har kunnat förutsäga en mängd fysikaliska processer. Trots detta lämnar den många fundamentala frågor obesvarade och vi vet att den kommer att bryta samman vid högre energier. Det viktigaste exemplet är kanske observationen av mörk materia som i universum finns i fem gånger så stor mängd som den materien som beskrivs av standardmodellen. Många nya teorier för fysik bortom standardmodellen har lagts fram i försök att skapa en mer komplett modell. Experimentell partikelfysik handlar idag om att med allt högre precision mäta standardmodellens observabler samt att verifiera eller falsifiera förekomsten av partiklar som de nya teorierna förutsäger. En av de mest populära teorierna är supersymmetri som dubblar partikelinnehållet genom att förutspå en ny supersymmetrisk partikel för varje existerande partikel i standardmodellen.

Studien av universums minsta beståndsdelar kräver ironiskt nog enormt stor experimentell utrustning, och teoretiska framsteg går hand i hand med teknologisk utveckling som möjligör för experiment vid allt högre energier. The Large Hadron Collider (LHC) vid CERN är världens största partikelaccelerator. Den accelererar protoner upp till nästan ljusets hastighet och kolliderar dem med varandra för att skapa nya tyngre partiklar som kan studeras. Acceleratorn ligger på gränsen mellan Schweiz och Frankrike i en 30 km lång cirkulär tunnel, cirka hundra meter under marken. Runt acceleratorn finns stora detektorer som är byggda för att observera partiklar från protonkollisionerna och mäta deras egenskaper. I denna avhandling presenteras forskning som utförts med data från ATLAS-detektorn.

Vid LHC inträffar ungefär en miljard protonkollisioner varje sekund och i varje kollision bildas ett stort antal partiklar som färdas genom ATLAS och detekteras. De allra flesta är partiklar med låg massa som är kända sedan länge och vars egenskaper är väl uppmätta i tidigare experiment. Den största utmaningen för en fysikanalys som utförs med data från ATLAS är därför att sålla i den enorma datamängden och välja ut de protonkollisioner som innehåller viktig information. Av stor vikt för alla analyser är även att känna till acceleratorns luminositet som ger ett mått på det totala antalet protonkollisoner i en given datamängd. Luminositeten används för att omvandla uppmätta frekvenser av intressanta processer till tvärsnitt, eller sannolikheter, som kan användas i jämförelser med andra experimentella resultat och teoretiska förutsägelser. I ATLAS mäts luminositeten med flera olika metoder som kombineras och jämförs för att uppnå hög precision på det slutgiltiga resultatet. I denna avhandling presenteras luminositetsmätningen för den datatagning som ägde rum mellan åren 2015-2018. Luminositeten uppmäts till $139\,\mathrm{fb}^{-1}$ med en osäkerhet på 1,7 %.

I avhandlingen presenteras även två studier som analyserar protonkollisioner i ATLAS för att söka efter tecken på fysik som inte beskrivs av standardmodellen. Båda analyserna tar avstamp i olika supersymmetriska modeller och väljer ut protonkollisioner att analysera baserat på de experimentella signaturer som förutspås av modellerna. En stor utmaning är att urskilja kollisioner med potentiell signal från kollisioner där standardmodellprocesser ger upphov till liknande signaturer och kontaminerar den utvalda datan. Dessa bakgrunder uppskattas noga med hjälp av simuleringar eller med hjälp av data. Den första presenterade analysen söker efter supersymmetriska skvarkar eller gluiner som produceras i protonkollisionerna och sedan omedelbart sönderfaller till två elektroner eller myoner som detekteras av ATLAS. Ingen signal över den förväntade bakgrunden observeras och resultaten tolkas istället med hjälp av simuleringar av supersymmetriska modeller. På så sätt kan exkluderingsgränser med 95 % konfidensintervall beräknas för gluino- och skvarkmassorna. Massor under 1,85 TeV utesluts för gluiner och massor under 1,3 TeV utesluts för skvarkar. Den andra analysen söker efter långlivade supersymmetriska partiklar som färdas en uppmätbar sträcka genom detektorn innan de sönderfaller till standardmodellpartiklar. Dessa processer är i princip fria från bakgrund, även om en liten kontaminering kan uppkomma av instrumentella effekter och felrekonstruktioner av data. Efter filtrering av datan förväntas en backgrund på $0,49 \pm 0,28$ signallika händelser och en händelse observeras. Resultatet används för att utesluta supersymmetriska neutraliner med massor upp till 1,3 GeV för partikellivslängder mellan 0,01 ns och 10 ns.

Acknowledgements

None of the work presented in this thesis would have been even remotely possible without the help and support of colleagues, friends, and family. First and foremost, I want to thank my supervisors Jonas Strandberg, Christian Ohm, and Bengt Lund-Jensen, who have all dedicated an incredible amount of time and effort to the details of my work and to strengthen my skills as a particle physicist. Jonas, we met already when I was an undergraduate student. Your way of explaining complex things in simple words opened my eyes to the beauty of particle physics and inspired me to pursue this path. Thanks for your patience through all my questions in the last few years and for your careful proofreading of all the text in this thesis. Christian, thank you for always having my back and for your persistent fights against all forms of injustice in the field. You have been extremely generous with your time and saved me countless hours by being the computer wizard I most often am not. Bengt, thank you for making it possible for me to pursue my doctoral studies at KTH, for your positivity, and for all the stories about particle physics and CERN in the past. To all three, you are true role models in how you supported me while I started a family during my time as a doctoral student. That can unfortunately not be taken for granted and I am forever grateful for having been able to do both.

I would then like to thank all present and previous colleagues in the ATLAS groups at KTH and SU, and in particular Rabia, Alex, Olle, Filip, Prim, Xuanhong, Sara, Karl, and Dave. It has been a pleasure to work with you and I am proud of our strong Stockholm teams both in luminosity and DV+jets. A special thanks to Edvin and Katarina who received me with open arms when I first started and have helped me through all the challenges that come with a career in academia. To all other friends and colleagues at KTH, SU, and the Oskar Klein Centre, thank you for the support and good times over the years.

My grateful thanks also go to my colleagues in the ATLAS collaboration and in particular in the Luminosity and SUSY groups. Working with talented people from all around the world has been highly motivating and has allowed me to greatly expand my knowledge of physics. A particular thanks to Emma, Jonathan, and Zach who watched over me in the first years as a student and helped me through the work in the Strong-2L analysis, to Valerie, Claudia, and Witold for guiding

me through the track-counting measurement, and finally to Simone, Larry, and Christian for all the help in the DV+jets analysis.

To my family and all my friends, thanks for all encouragement and help through challenging and stressful times. Babbo, thanks for bringing home science kits to the seven-year-old me, and for using up all the apples and oranges to build models of the Solar System. That is the reason why I am here today. Mamma, thanks for always standing by my side when things got a bit too overwhelming. Your support both emotionally and practically is truly immeasurable. Finally, Hector, this is where words are not enough any more, but I will try. Thank you for putting up with me through frantic thesis writing (and all the other times when I am generally frantic) and for taking care of everything at home. Mostly though, for the love. You might perhaps not understand everything in this thesis, but I hope you do understand that none of it had been possible without you.

Contents

List of abbreviations

$E_{\mathrm{T}}^{\mathrm{miss}}$	Missing transverse momentum.
pp	Proton-Proton.
BCID	Bunch-Crossing Identifier.
BCM	(ATLAS) Beam Conditions Monitor.
BSM	Beyond the Standard Model.
CKM	Cabibbo-Kobayashi-Maskawa (matrix).
CL	Confidence Level.
CR	Control Region.
CSC	Cathode Strip Chamber.
DV	Displaced Vertex.
EMB	(ATLAS) Electromagnetic calorimeter Barrel.
EMEC	(ATLAS) Electromagnetic calorimeter End-Cap.
FCal	(ATLAS) Forward Calorimeter.
FS	Flavour Symmetry.
FSR	Final State Radiation.
GMSB	Gauge Mediated Supersymmetry Breaking.
HEC	Hadronic Endcap Calorimeter.
HL-LHC	High Luminosity LHC.
HLT	High-Level Trigger.
IBL	Insertable (Pixel) B-Layer.
ID	(ATLAS) Inner Detector.
ISR	Initial State Radiation.
JER	Jet Energy Resolution.

JES	Jet Energy Scale.
JVT	Jet Vertex Tagger.
L1	Level-1 (trigger).
LAr	Liquid Argon.
LB	Luminosity Block.
LEP	Large Electron Positron collider.
LHC	Large Hadron Collider.
LLP	(BSM) Long-Lived Particle.
LRT	Large Radius Tracking.
LSP	Lightest SUSY Particle.
MBTS	(ATLAS) Minimum Bias Triggers.
MC	Monte Carlo (simulation).
MDT	Monitored Drift Tube (ionisation chamber).
MSSM	Minimal Supersymmetric Standard Model.
NLO	Next-to-Leading Order.
NLSP	Next-to-Lightest SUSY Particle.
NNLO	Next-to-Next-to-Leading Order.
PDF	Parton Distribution Function.
pdf	Probability density function.
PMT	Photomultiplier Tube.
PS	Proton Synchrotron.
PSB	Proton Synchrotron Booster.
PV	Primary Vertex.
QCD	Quantum Chromodynamics.
QED	Quantum Electrodynamics.
QFT	Quantum Field Theory.
RF	Radio Frequency.
ROI	Region Of Interest.
RPC	Resistive Plate Chamber.
RPV	R-Parity Violating.
SCT	(ATLAS) Semiconductor Tracker.
SM	Standard Model.
SPS	Super Proton Synchrotron.
SR	Signal Region.
SUSY	Supersymmetry.
TGC	Thin Gap Chambers.

TileCal	(ATLAS) Tile Calorimeter.
ToT	Time over Threshold.
TRT	(ATLAS) Transition Radiation Tracker.
vdM	van der Meer (scan).
VR	Validation Region.
WIMP	Weakly Interacting Massive Particle.
WP	Working Point.

Preface

Overview

Elementary particle physics is the study of the smallest building blocks in Nature and is driven by our curiosity about the fundamental properties of the Universe. Throughout the history of science, increasingly small particles have been discovered and the perception of which particles are elementary has evolved, turning theories that previously described the Universe into approximations of more comprehensive pictures. The best understanding to date is encapsulated in the theoretical framework of the Standard Model (SM) which explains how all matter is formed from quarks and leptons that interact via the strong, weak, and electromagnetic forces. As it stands, the SM is a very successful theory that has withstood five decades of high-precision tests and has been able to predict the outcome of a wide range of experiments. Despite all this, the SM still fails to answer several crucial questions in particle physics and cosmology, with the most obvious being the lack of an explanation for Dark Matter. Over the years, many theories for physics beyond the SM (BSM) have therefore been proposed, but so far none has been experimentally verified, making the hunt for BSM physics the main quest for particle physicists today.

The theoretical progress in particle physics goes hand-in-hand with the development of technologies that allow for experiments at increasing energies and with increasing precision. Of particular importance are accelerators, which bring particles to high energies and collide them or smash them into a target material and thereby create a high-energy environment where new heavier particles can be created and observed. This study of the smallest scales in Nature ironically requires gigantic experimental equipment, and as physicists have explored higher and higher energies, accelerators have become bigger and bigger. The Large Hadron Collider (LHC) [1] at CERN is to date the largest and highest-energy particle accelerator in the world. It accelerates and collides protons at energies on the TeV scale, and allows for precision tests of the SM as well as searches for BSM physics in an unexplored mass regime. The ATLAS detector is one of several experiments located at the LHC and is dedicated to the study of the proton-proton (pp) collision events. All studies presented in

this thesis are based on data collected by the ATLAS experiment during the LHC Run 2 which corresponds to the operational period between 2015 and 2018.

When the LHC is operating, more than 1 billion pp interactions take place in ATLAS every second. Most of these result in the production of low-mass particles with properties that are already well understood and measured to great precision by previous experiments. The main challenge in a search for BSM physics, or a precision measurement of a SM observable, is to isolate the interesting pp collision events and distinguish the signal from background processes that yield similar experimental signatures. All analyses are dependent on a measurement of the luminosity which quantifies the total number of pp interactions in a given dataset. Its value must be known to high precision in order for the analyses to be able to evaluate measured cross sections and to predict background rates. This thesis describes the overall strategy of the ATLAS luminosity measurement and goes into the details of the *track-counting* method which is one of several complementary approaches used to determine the final value.

One of the most popular BSM theories is Supersymmetry (SUSY) which matches every particle in the SM with a supersymmetric partner particle. The validity of SUSY is tested at the LHC by searching for supersymmetric particles produced in the pp collisions. This thesis presents two searches for SUSY conducted at the ATLAS experiment. The first analysis is seeking to find pp collision events in which supersymmetric *squarks* or *gluinos* are produced and then immediately decay into several SM particles, including two leptons that are detected and measured by the detector. This search is referred to as the *Strong-2L* analysis. The second search targets SUSY models in which supersymmetric particles are created in the pp collisions and travel some distance through the detector before decaying into SM particles. Such processes result in an experimental signature with a displaced vertex (DV) and several jets of SM particles. The search for DVs in multijet final states is referred to as the *DV+jets* analysis.

About this thesis

The first parts of this thesis serve to give the theoretical background and to introduce the experimental concepts necessary for the description of the presented work. Chapter 1 provides a review of the SM and a discussion about its successes and shortcomings, while BSM physics is introduced in Chapter 2. The focus is then shifted towards the experimental setup and Chapters 3 and 4 describe the LHC and the ATLAS experiment, while Chapter 5 explains how the data recorded in ATLAS is translated into meaningful information about the particles created in the pp collisions, and Chapter 6 describes the simulation of such processes. After these introductory chapters, the ATLAS luminosity measurement is presented in Chapter 7 and the track-counting method in Chapter 8. Finally, Chapters 9 and 10 present the Strong-2L and the DV+jets analyses respectively. Like any research

field, particle physics comes with a lot of abbreviations and acronyms and in order to aid the reader all abbreviations introduced in the thesis are summarised in the front matter.

Author's contributions

Research in experimental particle physics today is performed using large experiments and relies heavily on collaboration. As such, most of the work that I present in this thesis is realised as a team effort with fellow colleagues. ATLAS is a collaboration of physicists, engineers, technicians, and support staff from around the world, with almost 3000 scientific authors. The collaboration has an organisational structure where scientists work in smaller groups, dedicated to specific research areas within the scientific programme of ATLAS. Any output from the collaboration is shared by all members and is subject to rigorous internal review before being made public. Results can be released as papers submitted to journals or as so-called conference notes or public plots published directly by ATLAS. The publications are signed by all scientific authors in the collaboration. This thesis focuses on the work and the publications to which I have contributed directly.

The ATLAS organisation also contains working groups dedicated to the construction and maintenance of the detector, the operation of the detector, the data preparation and reconstruction, and the simulation. In addition, the so-called *Combined Performance* groups are dedicated to providing recommendations related to the reconstruction, identification, and calibration of specific objects which then are used by all physics analyses. The detector and performance work is essential for all final research findings of the collaboration and all members are expected to contribute to these activities.

The luminosity measurement is part of the data preparation work in ATLAS. Since 2021 I am the one of two conveners of the Inner Detector Luminosity group which is composed of around ten members. The convener position is an official ATLAS role designated to coordinate the work of the group, follow the ongoing activities closely, and set up the overall strategy for the project. As a convener I am thus responsible for the work of all other group members and the progress of the studies, and also provide supervision to other students in the group. In addition to the convener role, I work actively on several parts of the track-counting luminosity measurement. My main contributions include the performance monitoring in Section 8.5.1 as well as all simulation studies in Sections 8.2–8.4 and Sections 8.6–8.7. In addition, I regularly work with analysis frameworks produced by other team members and have generated all figures in Chapter 8. The toy simulation studies, described in Section 8.2, have resulted in ATLAS public plots published in April 2021 [2]. The track-counting measurement is included in a summary paper of the luminosity measurement in the LHC Run 2. The paper is in ATLAS review at the time of writing and has not yet been made public. An ATLAS conference note with the

Run 2 luminosity measurement was published in 2019 [3]. The results quoted in Chapter 7 correspond to the conference note.

During my doctoral studies I have been part of the SUSY working group. Since 2019 I have worked on the DV+jets analysis as part of a team consisting of around ten members. I am one of the main analysers and am responsible for the development of the tools for several parts of the analysis. My main contribution is the truth-level study described in Section 10.6.1 which has served as an important tool in the development of all background estimates. I have also developed the tools for the jet systematic studies presented in Section 10.7.2. In addition I contribute to the common analysis framework as well as the production of the data files used to generate all the analysis results. All figures in the chapter are produced by myself unless otherwise stated. The analysis has not been made public at the time of writing but is formally approved by the ATLAS SUSY group and the ATLAS internal review is ongoing.

The first two years of my doctoral studies I worked on the Strong-2L analysis in another ten-person team. My main contribution is the estimate of the flavour-symmetric background in Section 9.6.1. I also made substantial contributions to the signal region optimisation and to the estimation of the diboson background presented in Section 9.6.3. In addition, I was responsible for the production and validation of the data files used to generate all the analysis results. The Strong-2L analysis was published in The European Physical Journal C in May 2018 [4]. The figures in Chapter 9 are either produced by myself or come from the publication. An earlier iteration of the analysis, to which I also contributed, was published in The European Physical Journal C in March 2017 [5]. The Strong-2L analysis is summarised and presented in my Licenciate thesis [6] and the text in Chapter 9 is adopted from the previous text.

In addition to the luminosity work within data preparation, I have contributed to the work in the ATLAS Tracking Combined Performance group with studies of the impact parameter resolution. The work is part of the final ATLAS Run 2 recommendations for the uncertainty on the impact parameter, and some examples are included in Section 5.1.4.

Part I

Theoretical motivation

Chapter 1

The Standard Model of particle physics

The best theoretical framework available today to describe the smallest scales in Nature is the Standard Model (SM) of particle physics. It explains how matter consists of a small number of elementary particles and describes how these interact through fundamental forces and combine to make more complex structures. The model was developed in stages throughout the latter half of the 20th century, through the work of many scientists worldwide, eventually leading up to the gauge field theories of electroweak interactions [7–10] and Quantum Chromodynamics (QCD) [11, 12], as well as the Brout-Englert-Higgs mechanism to explain electroweak symmetry breaking [13–15]. These components together make up the SM.

Since its completion in the 1970s, the SM has held up to many high-precision tests and has been successful at predicting the existence of both elementary and composite particles before their experimental discovery. However, despite this success, the SM is not a complete theory of everything and fails to provide answers to several open questions in particle physics. Therefore, even if the model is correct at the energy scales that have thus far been probed, it represents a low-energy approximation of a more fundamental theory, just like many models that came before it.

This chapter introduces the SM and discusses its successes and shortcomings in order to provide the necessary background and motivation for the work presented in the thesis. An overview of the forces and particles of the SM is given in Section 1.1 and the concept of symmetries, from which the interactions emerge in the theory, is introduced in Section 1.2. Sections 1.3 and 1.4 then go into more details on the electroweak theory and QCD, and Sections 1.5 and 1.6 present methods to compute the observable interaction cross sections and particle decay rates from the theory. A brief summary of some of the main triumphs of the SM as well as some of the

Table 1.1: The four known fundamental forces of Nature. The relative strengths are approximate values for two elementary particles at a distance of 1 fm, roughly corresponding to the radius of a proton, or equivalently an energy scale of about 1 GeV. Table adapted from Ref. [17].

Interaction	**Relative strength**
Strong	1
Electromagnetic	10^{-3}
Weak	10^{-8}
Gravity	10^{-37}

open points are presented in Section 1.7. The text is based on Refs. [16] and [17], where more detailed explanations, outside the scope of this thesis, can be found.

1.1 Forces and particles in the SM

There are four known fundamental forces in the Universe; gravity, electromagnetism, the weak interaction, and the strong interaction. Gravity and the electromagnetic force are perhaps the most familiar as they are observable in everyday life, and they were also recognised long before the discovery of the other forces. Weak interactions manifest themselves in radioactive decays and are responsible for the nuclear reactions that fuel the sun and other stars. The strong force is what binds protons and neutrons together, and is thereby responsible for holding the atomic nucleus together. The SM encompasses all of these forces, except for gravity, which currently is only described by the theory of general relativity. Table 1.1 summarises the forces and their relative strengths, showing how the gravitational interaction is extremely weak in comparison to the others at a distance scale corresponding to the radius of a proton. As such, it is not relevant to the type of particle interactions that can be probed in current collider experiments.

In the SM, the forces are mediated by spin-1 gauge bosons which carry four-momentum between the interacting particles. The electrically neutral photon mediates the electromagnetic force, the charged $W^{\pm}$ and neutral Z bosons the weak force, and the eight neutral gluons the strong force. Each force carrier interacts only with particles that carry certain charges. The electromagnetic interaction couples to all electrically charged particles and the weak interaction to combinations of *weak isospin* and *weak hypercharge*. The strong interaction couples to *colour charge*. Each force is further characterised by the coupling constant which determines the relative strength of the interaction, as well as the range over which the force is effective. The coupling strength runs with energy, meaning that it depends on the four-momentum transfer Q^2 in the interaction. For the electromagnetic interaction, the coupling constant α grows stronger at large momentum transfer, or

equivalently at short distances. At zero momentum transfer, the value of α is $1/137$ and around $Q = 100\,\text{GeV}$ the value is $1/128$ [18]. The range of the interaction is inversely proportional to the mass of the force-mediating particle. Forces mediated by massless particles are therefore in principle infinite in range whereas interactions mediated by massive particles have a finite range. The photon is massless and the electromagnetic interaction is therefore infinite in range while the $W^{\pm}$ and Z bosons are massive and the weak force is a short range interaction. The long- and short-distance behaviour of the strong force is more involved and is discussed separately in Section 1.4.

All matter is built from spin-1/2 fermions which are divided into quarks and leptons based on their properties. Within each class there are three generations of particles that are essentially copies of each other but vary in lifetime and mass. For each particle, there is also a corresponding antiparticle which has the same mass and spin, but opposite values of the charges. The antiparticles are denoted by the same symbols as the particles, but with a bar over their respective symbols.

The lepton generations correspond to three types, or *flavours*, of charged leptons; the electron e, the muon μ, and the tau τ. These all carry electric charge -1 in units of the elementary charge e[1], and are paired with electrically neutral leptons called neutrinos. The neutrinos are treated as massless in the SM, but are experimentally observed to have a very small, albeit non-zero, mass. This inconsistency between theory and experiment is discussed further in Section 1.7. The charged leptons interact electromagnetically as well as weakly, while the neutrinos only interact weakly.

The quarks exist in six flavours and each generation consists of an *up-type* quark with electric charge $+2/3$ and a *down-type* quark of charge $-1/3$, in units of e. They are called up u and down d in the first generation, charm c and strange s in the second generation, and top t and bottom b in the third. The quarks differ from the leptons in that they carry colour charge and thereby couple to the strong force. Each quark can exist in three different colours referred to as red, green, and blue. Likewise, each antiquark can exist in the three corresponding anticolours. Individual quarks are not seen in Nature but are always confined to colour-neutral bound states called hadrons. These states are formed either by combining a colour with the corresponding anticolour or by combining all three colours. The hadrons thus exist in two categories; mesons consisting of a quark and an antiquark and baryons or antibaryons consisting of three quarks or three antiquarks. Some composite particles which will be mentioned in this thesis are the proton p and neutron n composed of the quarks uud and udd respectively, the pi mesons π^0 ($u\bar{u}$ or $d\bar{d}$) and $\pi^{\pm}$ ($u\bar{d}$ or $d\bar{u}$), the K^0_S meson ($d\bar{s}$ or $s\bar{d}$), and the J/ψ meson ($c\bar{c}$).

All the fermions in the SM have an intrinsic property which is called chirality and is of great importance in the weak interactions. In the massless limit, chirality

[1] $1\,e = 1.602 \times 10^{-19}\,\text{C}$

coincides with helicity which quantifies the alignment of the spin relative to the direction of motion of the particle. Particles can be right-handed or left-handed, corresponding to parallel or antiparallel alignments respectively. The weak interaction only couples to fermions corresponding to left-handed chiral states. While quarks and charged leptons exist in both left- and right-handed states, only left-handed neutrinos appear in the SM.

The final SM particle is the Higgs boson, which is a massive, electrically neutral spin-0 particle. It results in the theory as a manifestation of the mechanism that gives mass to the other particles and is further described in Section 1.3. Figure 1.1 summarises the particle content of the SM. All particles are fundamental in the sense that they are described as point-like and lacking underlying substructure in the theory.

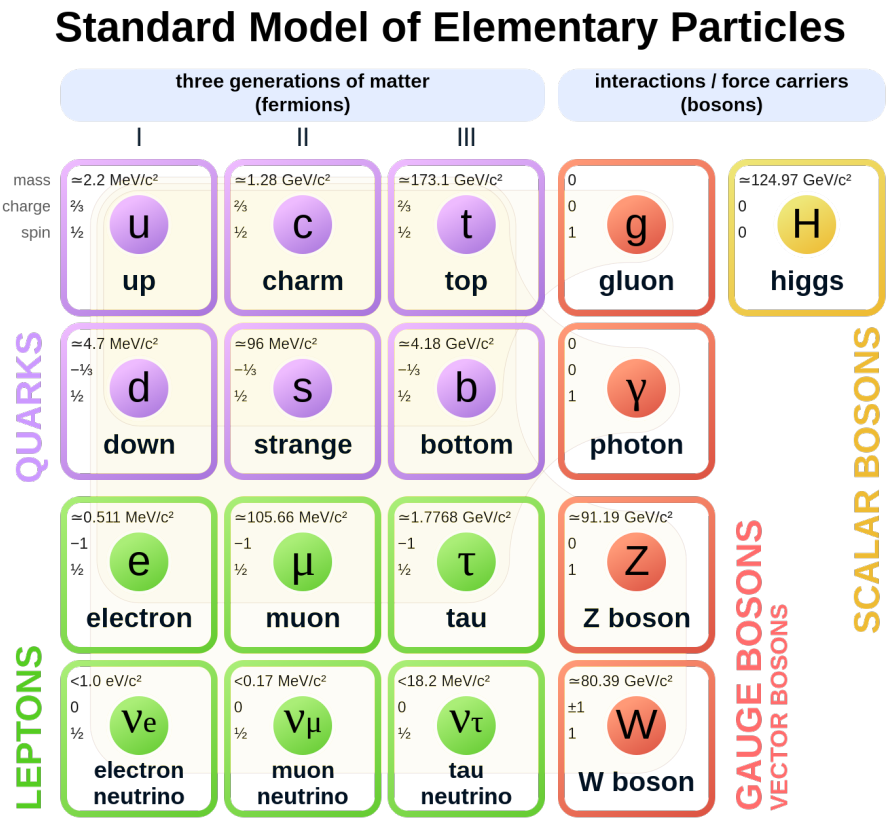

Figure 1.1: Overview of the particle spectrum of the SM [19].

1.2 Symmetries in the SM

The SM is formulated as a quantum field theory (QFT) where all particles are introduced as quanta of various fields. Fermions correspond to spin-1/2 fields and the gauge bosons correspond to spin-1, so-called vector fields. The dynamics of the system are controlled by the SM Lagrangian from which the equations of motion can be derived. This is found to be invariant under certain transformations, where each such symmetry corresponds to a quantity that is conserved in the system. The symmetries thereby lead to a set of fundamental conditions of the theory and dictate the allowed interactions between the various fields.

First of all, any QFT must be invariant under global transformations described by the Poincaré group. These include temporal and spatial translations which correspond to conservation of energy and linear momentum, as well as rotational transformations which are associated to the conservation of angular momentum. In addition, the fields of the SM are required to be invariant under a set of local phase transformations, so-called *gauge transformations*, which describe the interactions of the theory and are associated to conserved quantities via Noether's theroem [17].

The gauge principle can be illustrated for the case of quantum electrodynamics (QED), which is the QFT of the electromagnetic force and describes interactions between electrically charged fermions and photons. Before introducing the interactions, the system is described by the Lagrangian for a free spin-1/2 field ψ with mass m

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

where γ_{μ}, $\mu \in [0, 1, 2, 3]$, are the Dirac matrices. The Lagrangian is trivially invariant under a global $U(1)$ transformation where the field transforms as

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

It is however not, a priori, invariant under a local phase transformation, where the parameter θ depends on the space-time coordinate and the field transforms as

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

The invariance under the local transformation is instead assured by replacing the derivative in the Lagrangian by 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

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

The QED Lagrangian can then be written including the new vector field as

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

with

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

The expression in brackets in Eq. 1.6 corresponds to the free massive fermionic field, where the first term represents its kinetic energy and the second is the mass term. The third term in Eq. 1.6 represents an interaction between the fermionic field and the vector field A_μ, which can be identified as the photon field. The last term is the free-field Lagrangian for a massless vector field and is introduced to describe the kinematics of the photon. All terms in this expression are now invariant under the combined gauge transformations corresponding to Eqs. 1.2 and 1.5. The interaction term is proportional to the electric charge, which is the conserved quantity associated to QED. For the electron it is related to the electromagnetic coupling constant through $\alpha = e^2/4\pi$.

This principle of local gauge invariance generating the interactions between the fields as well as the associated conserved charges, can be extended to all interactions in the SM. The symmetries of the SM are described by the gauge group

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

The strong force is generated by requiring symmetry of the Lagrangian under colour charge transformations, corresponding to the symmetry group $SU(3)_C$, where C refers to the colour. The electromagnetic and weak interactions are unified into the electroweak interaction that corresponds to invariance under the $SU(2)_L \times U(1)_Y$ symmetry. Here, $SU(2)_L$ is the weak isospin group acting only on left-handed fermions and $U(1)_Y$ is the weak hypercharge group. The gauge bosons associated to the $SU(3)_C$ symmetry are the eight gluons, while three weak isospin bosons W_1, W_2, and W_3 are associated to the $SU(2)_L$ symmetry and the hypercharge boson B to $U(1)_Y$. The corresponding conserved charges are colour, third component of weak isospin I^3, and weak hypercharge Y.

The field content of the SM is presented in detail in Table 1.2. Fermion fields transform under the fundamental representation of each gauge group, while the gauge fields transform under the adjoint representation. In $SU(3)_C$ all fermions are thus represented by colour triplets, corresponding to the three colour charges, while the adjoint representation describes the eight gluons as a colour octet. The $SU(2)_L$ left-handed chiral fields appear as doublets with an up-type and down-type quark pair or a charged and a neutral lepton pair, while the right-handed chiral fields live in the singlet representation of $SU(2)_L$. The weak gauge bosons are represented by weak isospin triplets. The $U(1)_Y$ symmetry of weak hypercharge transformations is one-dimensional.

In addition to the conserved gauge charges, the SM contains accidental global $U(1)$ symmetries which arise as a consequence of imposing invariance under the symmetry transformation with a given particle content. Corresponding to the accidental symmetries are the baryon number B, lepton numbers L_l ($l = e, \mu, \tau$), and the total lepton number $L = L_e + L_\mu + L_\tau$, which are conserved in all interactions in the

Table 1.2: The field content of the SM and the transformation properties under the SM gauge groups. The three rightmost columns indicate under which representations of the gauge groups that each field transforms. For the $U(1)_Y$ symmetry the column gives the weak hypercharge Y associated with each field. For $SU(3)_C$ and $SU(2)_L$, **1** refers to the field belonging to the associated singlet representation, **2** to the doublet representation, **3** to the triplet representation, and **8** to the octet representation.

	Field label	**Content**	**Spin**	$U(1)_Y$	$SU(2)_L$	$SU(3)_C$
Quarks	Q_i	(u_L, d_L), (c_L, s_L), (t_L, b_L)	1/2	1/3	**2**	**3**
	$u_{R,i}$	u_R, c_R, t_R	1/2	4/3	**1**	**3**
	$d_{R,i}$	d_R, s_R, b_R	1/2	-2/3	**1**	**3**
Leptons	L_i	$(e_L, \nu_{e,L})$, $(\mu_L, \nu_{\mu,L})$, $(\tau_L, \nu_{\tau,L})$	1/2	-1	**2**	**1**
	$e_{R,i}$	e_R, μ_R, τ_R	1/2	-2	**1**	**1**
Gauge fields	B	B	1	0	**1**	**1**
	W	(W_1, W_2, W_3)	1	0	**3**	**1**
	G	G_a, $a \in [1, ..., 8]$	1	0	**1**	**8**
Higgs field	ϕ	(ϕ^+, ϕ^0)	0	1	**2**	**1**

SM. The lepton numbers take the value $+1$ for the leptons and -1 for the antileptons, and the baryon number takes the value $+1/3$ for the quarks and $-1/3$ for the antiquarks. For strong and electromagnetic interactions, the individual quark flavours are also conserved while flavour changing processes are allowed by the weak interaction.

1.3 Electroweak physics and spontaneous symmetry breaking

The $SU(2)_L \times U(1)_Y$ content of the SM Lagrangian predicts four gauge bosons corresponding to the W^1_μ, W^2_μ, W^3_μ, and B_μ fields. These can be compared to what is empirically known to exist; the electromagnetic interaction described by interactions with the photon in the theory of QED, and the presence of two massive charged and one massive neutral mediator of the weak force. The mass eigenstates of these physical gauge bosons are obtained by a mixing of the electroweak eigenstates according to

$$W^\pm_\mu = \frac{1}{\sqrt{2}} \left(W^1_\mu \mp i W^2_\mu \right) \tag{1.9}$$

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

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

where $W_\mu^\pm$ and Z_μ are the fields corresponding to the $W^\pm$ and Z bosons, and A_μ is the photon field as introduced in Section 1.2. The angle θ_W is known as the weak mixing angle and quantifies the amount of mixing that occurs between the neutral $SU(2)_L \times U(1)_Y$ gauge fields. Furthermore, the fermion electric charge Q can be related to the underlying $SU(2)_L$ and $U(1)_Y$ gauge symmetries through the relation

$$Q = I^3 + \frac{1}{2}Y. \tag{1.12}$$

At this point the theory treats the gauge bosons as being massless, similar to the photon field in QED. Mass terms for vector boson fields take the form $m_W^2 W_\mu W^\mu$, similar to the mass term for the fermion field introduced in Eq. 1.6. Such terms are however forbidden in the electroweak sector since they violate the $SU(2)_L \times U(1)_Y$ invariance. This is in clear disagreement with the short range property of the weak interaction, as well as the measured masses of the W and Z bosons. The solution to the problem is spontaneous symmetry breaking which is implemented by introducing the Higgs field ϕ.

The Higgs field is a complex scalar field which is invariant under $SU(2)_L \times U(1)_Y$, but has a non-zero ground state that breaks the symmetry. To generate this ground state, a scalar potential

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda \left(\phi^\dagger \phi\right)^2 \tag{1.13}$$

is added to the SM Lagrangian. The Higgs field is an $SU(2)$ doublet and it can be shown that the interactions described by Eq. 1.13 respect the $SU(2)$ gauge symmetry. Provided that $\mu^2 < 0$, the potential has non-zero minima located at $\phi = \pm v/\sqrt{2}$ with $v = \sqrt{-\mu^2/\lambda}$. The Higgs field at this equilibrium point can then be expressed as

$$\phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{1.14}$$

Now, a perturbation of the Higgs field can be performed around the minimum value, according to

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \tag{1.15}$$

where $h(x)$ corresponds to excitations of the field that represent the physically observable Higgs boson. The interactions with the gauge bosons are generated by applying the gauge principle to the Higgs field as expressed in Eq. 1.15, which also leads to mass terms proportional to v for the $W^\pm$ and Z bosons. In essence, three out of the four real, scalar degrees of freedom of the complex Higgs field are used to generate mass to the weak gauge bosons while the forth component becomes the Higgs boson. The requirement of a symmetry-breaking ground state thus leads to an $SU(2)_L \times U(1)_Y$ invariant formulation of the masses of the electroweak gauge bosons. This is what is commonly referred to as the Brout-Englert-Higgs mechanism [13–15].

The fermion masses are obtained by adding additional $SU(2)_L \times U(1)_Y$ invariant interaction terms to the SM Lagrangian between the fermions and the Higgs field, so-called Yukawa couplings. Within the quark sector these terms result in a mixing among the weak eigenstates of the quark fields to produce the observed mass eigenstates. This allows for the flavour-changing processes involving the decay of a quark of one generation into that of another generation. The amount of mixing in the quark sector is dictated by a 3×3 unitary matrix known as the Cabibbo-Kobayashi-Maskawa (CKM) matrix [20, 21].

1.4 Quantum Chromodynamics

The strong interaction is mediated by the massless gluons which exist in eight different superpositions of colour-anticolour states corresponding to the octet representation of $SU(3)$. This means that the gluons carry a net colour charge and are not neutral to the field they mediate. The gluons can therefore undergo self-interactions which lead to the concepts of *asymptotic freedom* and *confinement*. Asymptotic freedom refers to the fact that the strong coupling constant α_S becomes weaker at short distances, or equivalently, at large momentum transfer, contrary to the running of the electromagnetic coupling. Confinement is instead related to the increase of the interaction strength at large distances, or small momentum transfer. As two coloured objects move apart, the energy in the gluon field between the objects will increase and a colour neutral combination will eventually be formed. The strong interaction therefore effectively has a short range, in spite of having a massless mediator. The cut-off distance for such processes corresponds to approximately 1 fm. Figure 1.2 shows the running of the strong coupling constant.

If quarks or gluons gain enough kinetic energy to break the confinement potential, the energy in the gluon field is instead transferred into the creation of pairs of quarks and antiquarks which form new hadronic bound states, a process called hadronisation. When two coloured states move away from each other at energies that are large compared to the energy required to break the confinement of a hadron, a chain of hadronisation processes form a number of collimated hadrons called a hadronic jet.

1.5 Interaction cross sections and particle decays

Experimental particle physics is in large based on measurements of particle decay rates and cross sections. These experimentally observable phenomena represent transitions where an initial particle state evolves into a final state through some sort of interaction and can be computed by means of the SM.

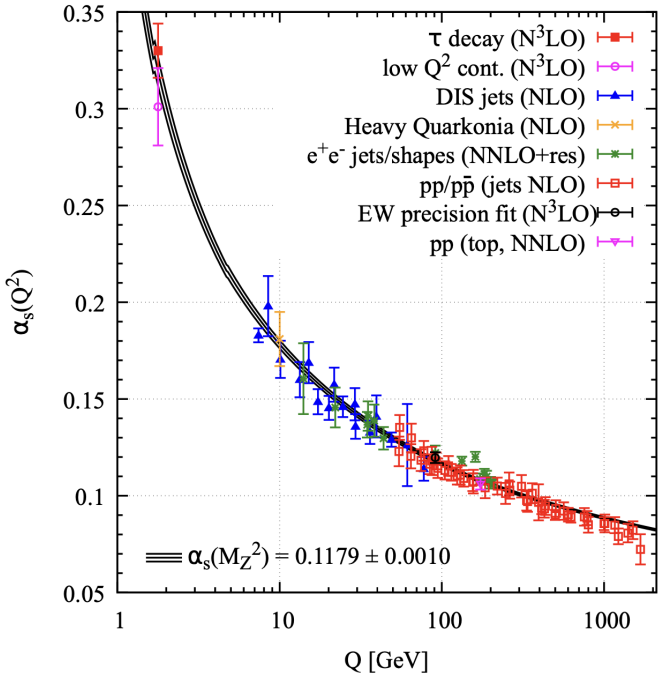

Figure 1.2: Measured values of the strong coupling constant α_S as a function of the momentum transfer Q [18].

Cross sections

The cross section represents the probability for a given process to occur. For a $2 \to 2$ process, such as the annihilation process shown in Figure 1.3, the differential cross section can be computed as

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

In this expression, p_1, p_2, m_1, and m_2 are the four-momenta and masses of the initial state particles, and p_3 and p_4 the four-momenta of the final state particles. The phase-space factor $\mathrm{d}\Phi$ describes the number of kinematically available states for the final state particles to occupy and the matrix element $|\mathcal{M}|$ describes the dynamics of the interaction.

The matrix element can be calculated from the SM Lagrangian by using perturbation theory. The procedure can be illustrated for the process $e^+ + e^- \to \mu^+ + \mu^-$ which occurs through the exchange of a photon. Figure 1.3 shows two examples of how this process can occur, as represented by Feynman diagrams. The diagram in Figure 1.3a contains two QED interaction vertices which each contribute a factor $e \propto \sqrt{\alpha}$ to the matrix element, as shown in Eq. 1.6. In total, the diagram contributes a factor α and corresponds to the lowest-order term in the perturbation series for the annihilation process. Figure 1.3b corresponds to a second-order diagram with four vertices that in total contribute with a factor α^2 to the matrix element.

In a similar manner, diagrams can be constructed representing increasing orders of α and the matrix element is the sum of all the corresponding interaction terms. Provided that the coupling constant is smaller than unity, the terms in the series become smaller and smaller and eventually negligible. This means that it is often justified to include only a limited number of Feynman diagrams in the calculation of the matrix element. For QED, the coupling constant is sufficiently small for the series to converge rapidly and the matrix element is dominated by the lowest-order term. For QCD, the perturbative method is only applicable at short distances where the strong coupling α_S is significantly smaller than unity.

When higher-order processes are added to the matrix element it will contain divergent integrals over the internal loop momenta. This issue is solved by introducing a renormalised coupling constant which absorbs the infinities. The renormalised charge will vary with the momentum transfer Q^2 and will be a function of the renormalisation scale μ_R at which the infinities are removed and the coupling is set. Renormalisation is thereby the origin for the running of the experimentally observed coupling strengths with energy.

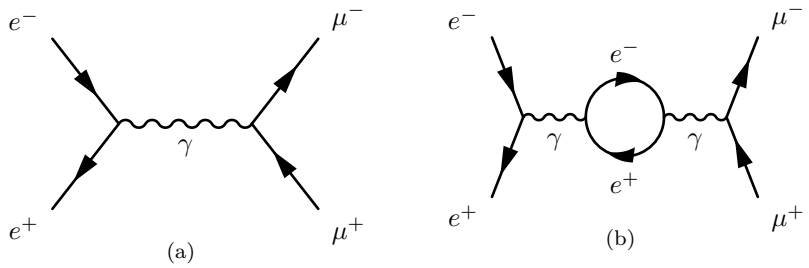

Figure 1.3: First-order (a) and second-order (b) Feynman diagrams contributing to the annihilation process $e^-e^+ \to \mu^-\mu^+$.

Decay rates

A general property of all fundamental particles is their tendency to decay into lighter particles if they are allowed to do so. Whether or not the decay is possible is dictated by the conservation laws of the SM. For instance, the electron is the lightest electrically charged particle, so conservation of energy and electric charge prevents it from decaying. Likewise, the proton is the lightest particle with non-zero baryon number and it is therefore stable.

Particles that are unstable will decay into n particles at a rate Γ given by

$$\Gamma = \frac{1}{\tau} \propto \frac{1}{2M} \int |\mathcal{M}|^2 \, d\Phi_n(P; p_1, ..., p_n) \tag{1.17}$$

where M and P are the mass and the four-momentum of the particle, $|\mathcal{M}|$ is the matrix element, p_i are the four-momenta of the decay products, and dΦ is the phase-space factor. The inverse of the decay rate is the mean proper lifetime τ of the unstable particle. Given the mean lifetime, the probability that the particle lives for a time t before decaying is given by Poisson statistics according to

$$P(t) = e^{-t/\gamma\tau}, \tag{1.18}$$

where γ is the Lorentz factor.

Figure 1.4 shows the proper lifetimes of a selection of elementary and composite particles, displaying the wide range of lifetimes in the SM. The lifetimes are dictated by the strength of the coupling and the phase-space element. All hadrons will decay through the strong interaction if they are allowed to do so and will then have a very short lifetime on the order of 10^{-23} s. If no lighter hadrons with the same quark numbers exist, the decay has to proceed through the weak force, and the particle will have a longer lifetime. Among the weak decays, interactions involving a small momentum transfer compared to the mass of the mediator will be suppressed. This

is the case for the neutron decay where the difference between the neutron mass and its decay products, the electron, proton, and antineutrino, is very small. The reason for the suppression is that the mediator is virtual, or *off-shell*, meaning that its energy and momentum do not lie on the surface defined by $E^2 - p^2 = m^2$. The notion of which particles are considered to be long-lived is typically determined by the experimental equipment used to detect the particles. A particle that travels a distance which can be resolved by the experiment before it decays is considered long-lived.

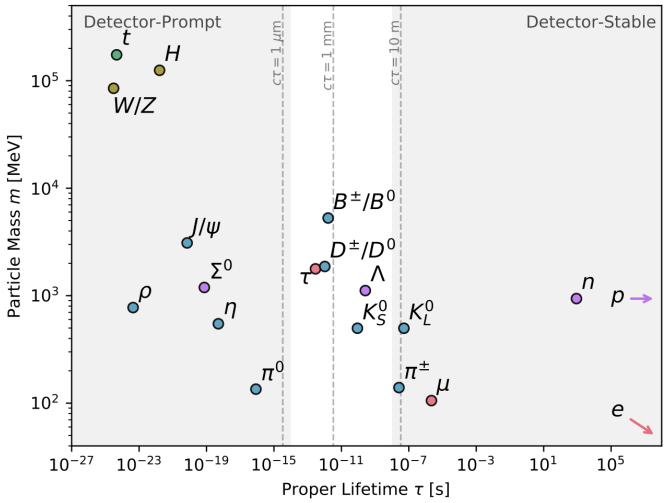

Figure 1.4: A selection of the SM particle spectrum shown as a function of mass and proper lifetime [22].

1.6 Phenomenology of proton-proton collisions

This thesis describes research within particle physics carried out using high-energy pp collisions produced at the LHC. A proton is a complex object composed of the three *valence* quarks uud, which carry most of the proton momentum, as well as *sea* quarks and gluons. These constituents are collectively called partons. At sufficiently high Q^2, the pp collisions are inelastic, and the fundamental interaction takes place between individual partons. Such processes are called hard interactions and are typically the main target of the physics analyses. The full picture is however more complex and the collision involves several additional QCD and QED processes. Figure 1.5 shows a sketch of a pp interaction. The green ellipses with pointing arrows represent the colliding protons and the filled red circle represents the hard interaction. Below is the so-called *initial state radiation* (ISR) in blue

and above is the *final state radiation* (FSR) in red. These are processes where the coloured particles present in the initial and final states radiate gluons which further split into quark-antiquark pairs. Electrically charged particles might also radiate photons, and this is represented by the yellow lines. The green ellipses represent hadronisation, where the coloured particles in the final state combine into hadrons which then further decay if they are unstable. These two processes are collectively called fragmentation. In addition to all these processes, which are related to the primary hard scatter, the *pp* collisions involve interactions between the remaining partons, illustrated by the purple ellipse in the figure.

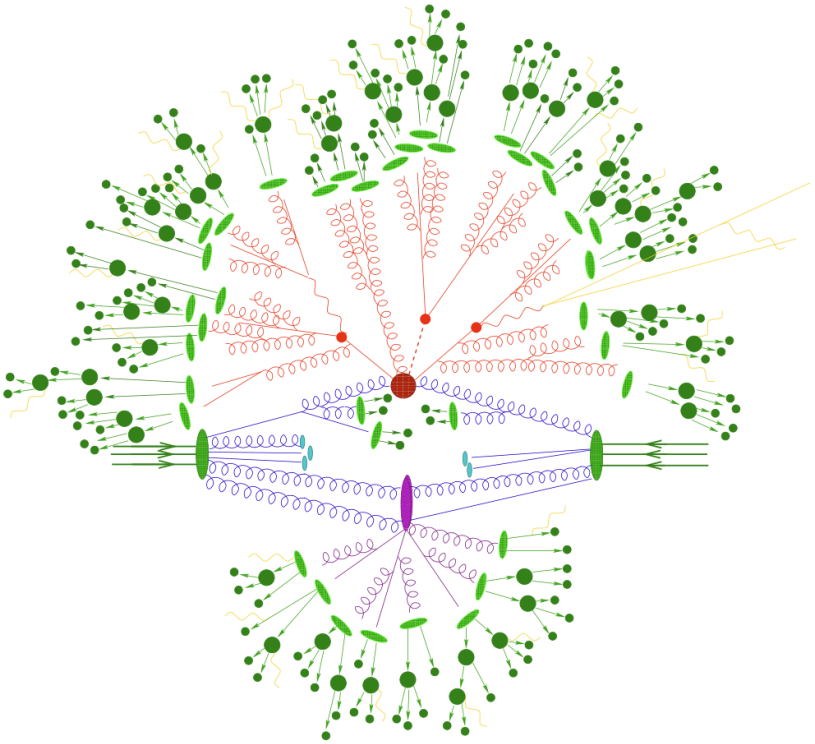

Figure 1.5: Sketch of a proton-proton collision [23]. See the text for details.

The momentum scales of the different physical processes involved in the *pp* collision are very different, ranging from the $O(1)$ GeV scale of hadronisation to the $O(1)$ TeV scale which is typical for the hard interaction at the LHC. Due to the running of the strong coupling constant, the nature of the QCD description for these various aspects varies. The hard interaction represents a scale at which QCD is perturbative and the matrix element can be calculated using the methods described

in Section 1.5, while the hadronisation scale is in the non-perturbative regime, for which processes can only be modelled phenomenologically. According the the factorisation theorem of QCD, these aspects are independent in the collinear limit, where the partons are approximately free particles moving collinearly to the hadron. The cross section for a process where two protons collide and evolve into a final state O can then be written in the factorised form

$$\sigma_{pp\to O} = \sum_{i,j} \int_0^1 \mathrm{d}x_i \int_0^1 \mathrm{d}x_j f_i(x_i, \mu_F^2) f_j(x_j, \mu_F^2) \times \hat{\sigma}_{ij\to O}\left(x_i p_i, x_j p_j, \mu_F^2, \mu_R^2\right). \tag{1.19}$$

Here, $f_i(x, Q)$ and $f_j(x, Q)$ are the so-called *parton distribution functions* (PDFs), $\hat{\sigma}_{ij\to O}$ is the partonic cross section, and the sum runs over the partons of types i and j that exist in the two protons and contribute to the interaction.

The PDFs absorb the non-perturbative aspects of the computation and parametrise the probability to find a given parton in the proton carrying a fraction x of the proton momentum, evaluated at the so-called factorisation scale, μ_F. This is an arbitrary scale introduced in the calculation of the cross section to cure infrared divergencies in the matrix element, resulting from contributions of particles with very small energies approaching zero. As a result, the PDFs run with energy and the factorisation scale defines the cut-off below which the divergencies are absorbed into the PDFs. The PDFs cannot be calculated from first principles and must be obtained from experimental measurements [24, 25]. The functions are measured at specific scales and are then extrapolated to the energy regime relevant to the physics process being calculated. This extrapolation is governed by the DGLAP evolution equations which evolve the PDFs from one scale to another [26, 27]. Examples of proton PDFs are provided in Figure 1.6.

The parton-level cross section describes the hard scatter between the initiating partons and can be calculated using the machinery of perturbation theory to a fixed order in α_S. As with all calculations involving the perturbation theory of QFT, the parton-level cross section also depends on the renormalisation scale, μ_R.

1.7 Successes and shortcomings of the SM

Over the years, the SM has proven extremely successful at predicting and explaining the outcome of particle physics experiments. The theory was formulated before the experimental observation of several particles and their subsequent discoveries have provided important proof of the validity of the model. Some examples include the discovery of the W and Z bosons in 1983 [29, 30] and the top quark in 1995 [31,32], as well as the discovery of the Higgs boson at the LHC in 2012 [33,34].

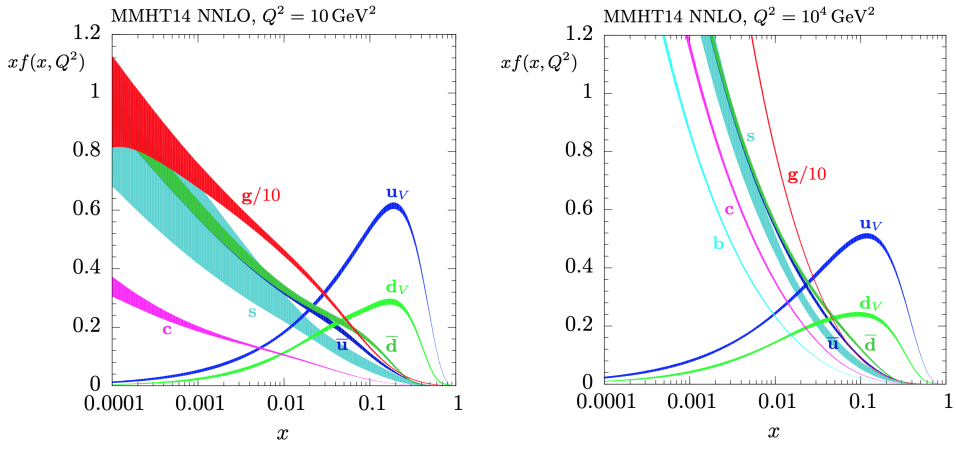

Figure 1.6: Proton PDFs evaluated at energy scales $Q^2 = 10\,\text{GeV}^2$ (a) and $Q^2 = 10^4\,\text{GeV}^2$ (b) [28]. The bands indicate the 68% confidence-level uncertainty bands. The valence-quark components are indicated with a subscript 'v' and the sea-quark components without.

With these, there is now experimental evidence for all of the particles in the SM. In addition, in principle all measurements of SM observables have so far been in excellent agreement with the values predicted by the model. Figure 1.7 shows a summary of LHC measurements of cross sections for various production processes at different energies. The agreement of these measurements with the theoretical predictions provided by the SM, spanning over ten orders of magnitude, is remarkable. The internal consistency of the model is further illustrated in Figure 1.8, which shows W-boson and top-quark masses as predicted by a fit of the SM to electroweak precision data. When all measurements except those on the Higgs mass are included, the constraints of the SM only allow for a small region of the parameter space spanned by the top and W masses. Adding the Higgs-mass measurement only shrinks this allowed area. The fact that these indirect measurements agree so well with the direct measurements illustrates the self-consistency of the SM.

It is clear that the SM accurately describes all current experimental observations of the particles and the interactions that it includes. There are however several open questions in particle physics that lead to the assumption that further extensions to the SM are needed. This section lists a few of these items.

Dark matter Most of the experimental evidence supporting the need for physics beyond the SM comes from astrophysical observations. There is compelling evidence that the majority of the matter in the Universe is composed of a non-

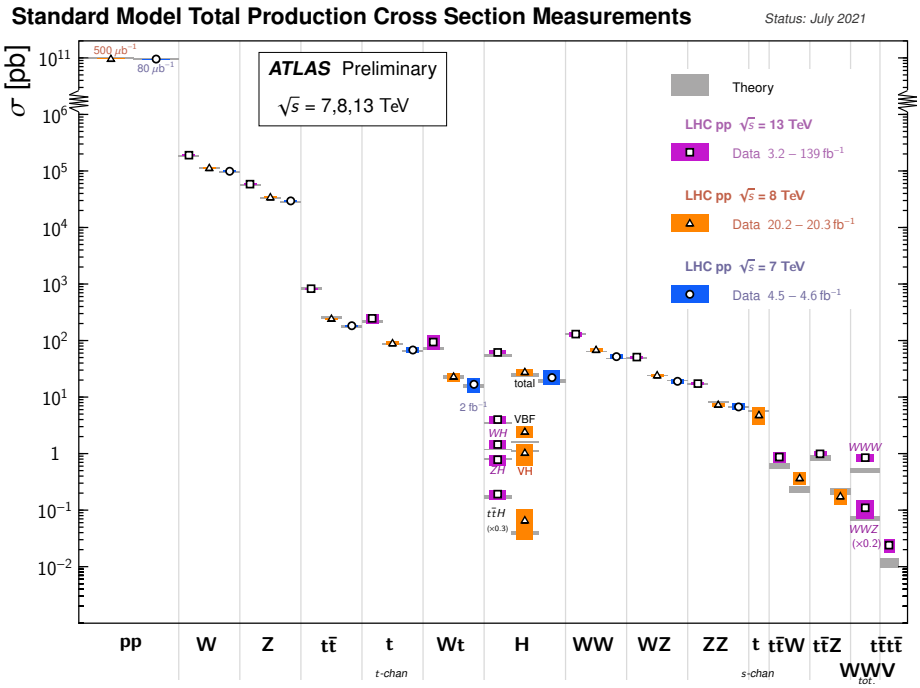

Figure 1.7: Summary of several SM total production cross-section measurements compared to the corresponding theoretical expectations [35].

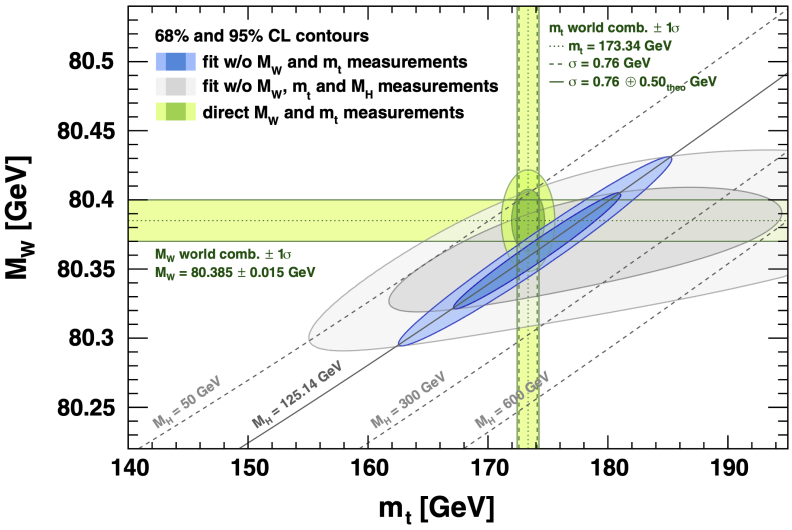

Figure 1.8: Contours at 68% and 95% CL obtained from scans of M_W versus m_t, for the global electroweak fit in comparison to the direct measurements, including Higgs-boson mass measurements in the fit (blue) or excluding them (grey) [36].

luminous, possibly weakly interacting type of matter referred to as Dark Matter, which cannot be explained by the SM. In fact, studies of the Cosmic Microwave Background show that only approximately 15 % of the matter in the Universe consists of regular SM particles, while the rest consists of Dark Matter [37]. There has so far not been any experimental proof of the nature of dark matter particles, but some guidance is provided by the astrophysical observations. The particle has to be massive, can at most interact weakly with the SM particles, and needs to be stable on time scales around the lifetime of the Universe. One of the most studied dark matter candidates is the so-called weakly interacting massive particle (WIMP), which meets all of these criteria. In most WIMP scenarios, the Dark Matter is expected to have been produced thermally in the early Universe, similar to the particles of the SM. Obtaining the correct abundance of Dark Matter today via thermal production of WIMPs requires the mass to be within the range of 100-1000 GeV. A WIMP is thus an interesting dark matter candidate that, if it exists, should appear at the electroweak scale and be accessible for searches at the LHC through its weak coupling to SM particles.

Massive neutrinos The observation of neutrino oscillations [38] provides evidence in support of neutrinos having non-zero masses, in conflict with the original formulation of the SM. It is possible to include neutrino masses in the SM [39] and revealing the mechanism that provides the masses is the motivation behind the extensive neutrino physics programme in the world.

Matter-Antimatter asymmetry The matter-antimatter asymmetry problem refers to the observation that the observable Universe only consists of matter, and that all the antimatter that must have been created at the time of the Big Bang has disappeared. A condition for this asymmery to arise is the existence of a CP-violating mechanism that differentiates between matter and antimatter [40]. While there are such mechanisms in the SM, for example in the CKM matrix, these are not enough to explain the observed asymmetry between matter and antimatter.

Hierarchy problem The hierarchy problem refers to the fact that the electroweak sector, through the scalar Higgs boson, is sensitive to the cut-off scale Λ where new physics is expected to change the high-energy behaviour of the model. In the calculation of the physical mass of the SM Higgs boson, loop-level diagrams like Figure 1.9 need to be included, which give radiative correction terms to the lowest order diagram according to

$$m_h^2 = m_{h,0}^2 + \Delta m_h^2, \tag{1.20}$$

where $m_{h,0}$ is the so-called *bare* Higgs mass of the SM Lagrangian. The corrections are proportional to the cut-off scale and can be expressed as

$$\Delta m_h^2 = -\frac{\lambda_f^2}{16\pi^2}\left[2\Lambda^2 + O\left(m_f^2 \ln\left(\frac{\Lambda}{m_f}\right)\right)\right], \tag{1.21}$$

where λ_f is the coupling parameter to the respective fermion field, and m_f is the fermion mass. Setting the cut-off scale to an energy corresponding to the Planck mass, $M_P \approx 10^{19}\,\text{GeV}$, where gravity becomes similar in strength to the other fundamental forces, these corrections grow enormous. The bare Higgs mass therefore needs to be fine-tuned in order to precisely cancel out the radiative corrections, leading to the observed Higgs mass of 125 GeV. This introduces the philosophical concerns of naturalness, which refers to the idea that models of physics should work without requiring fine-tuning, and that all such cancellations among parameter values are to have some systematic cause.

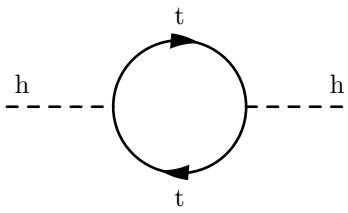

Figure 1.9: One-loop radiative correction to the mass of the Higgs boson h from the top quark t.

Grand unification The SM includes 26 free parameters that determine the masses of all the particles and the couplings of the interactions between them.

A common philosophy in physics is that any fundamental theory should consist of a minimal number of laws and as few free parameters as possible. This inspires the pursuit of a Grand Unified Theory where the electroweak and strong interactions are merged into one single force, thereby simplifying the theory. With the current formulation of the SM, the running of the three gauge couplings is found to nearly, but not quite, meet at the same point at high energies.

Chapter 2

Physics beyond the Standard Model

One of the most promising theories for physics beyond the SM (BSM), that provides a solution to the hierarchy problem and also potentially a dark matter candidate, is Supersymmetry (SUSY). It builds on the existing framework of the SM and predicts an extended particle spectrum with several new particle states. The search for these new particles is the topic of the analyses presented in Chapters 9 and 10, which are part of a wide programme of SUSY searches at the LHC. In the present chapter, an introduction to SUSY is provided in order to give the relevant background and motivation for the analyses. Section 2.1 presents the general components of SUSY and Section 2.2 then explains how the new particles can be produced in the pp collisions at the LHC. Several SUSY scenarios predict the existence of new particles which decay into SM particles through a suppressed decay which makes them long-lived. Such models will be described in Section 2.3 and are the target of the analysis presented in Chapter 10. The text is in large based on Ref. [41], where more thorough derivations and explanations can be found.

2.1 Supersymmetry

SUSY is a theoretical framework that builds on the SM and extends it by proposing a new symmetry between bosons and fermions. That is, it allows for transformations between the fields typically associated with matter and those typically associated

with the forces. It is mathematically formulated by introducing the SUSY generator $\hat{Q}$, which turns fermionic states into bosonic states and vice versa:

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

The new symmetry is a space-time symmetry and its generator is fermionic, meaning that it has spin-1/2, as opposed to the bosonic generators of the symmetries in the SM.

As a consequence of SUSY, there is a fermionic superpartner for each SM boson and a bosonic superpartner for each SM fermion. These SUSY particles, or *sparticles*, have the same internal quantum numbers as their SM counterparts except for the spin which differs by 1/2. The new superpartners of the SM leptons and quarks are spin-0 sleptons and squarks whereas the superpartners of the SM gauge bosons are spin-1/2 gauginos. The symbols for the sparticles are the same as for the corresponding particles, but with a tilde used to denote the superpartner.

One reason for SUSY being a favoured theory to provide an extension to the SM is that the introduction of the new particles potentially can solve the hierarchy problem. When computing the Higgs mass, loop corrections involving the new scalar superpartners of the fermions, as illustrated in Figure 2.1, need to be included. For each scalar S, the mass is corrected by a term given by

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

where the parameter λ_S describes the coupling of the scalar to the Higgs field and Λ is the cut-off scale of the theory. Comparing Eq. 2.2 to Eq. 1.21, it can be seen that the divergent contributions to the mass in the SM are cancelled out by the new terms, provided that SUSY is an exact symmetry where the particle and sparticle masses and couplings are identical. SUSY thereby creates a mechanism for systematic cancellation of the divergent terms and removes the need for fine-tuning amongst the parameters. An exact SUSY is however not realised in Nature, since the superpartners of the SM particles would have been found using present-day experimental equipment if they had the same mass. Any form of SUSY must therefore be spontaneously broken, similar to the electroweak symmetry breaking in the SM. The cancellation of radiative correction terms in the Higgs mass can however still approximately hold, provided that the SUSY breaking is soft [42, 43], and the corresponding SUSY-breaking mass parameters are no larger than a few TeV. There are a number of models that attempt to explain the mechanism of SUSY breaking, and the details concerning each of these impact the features and phenomenology of the model. The SUSY extension of the SM that requires the minimum amount of new particle states and new interactions is the Minimal Supersymmetric Standard Model (MSSM).

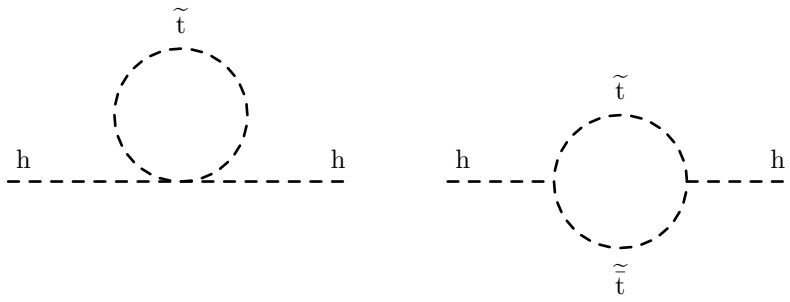

Figure 2.1: One-loop radiative corrections to the mass of the Higgs boson h from the scalar superpartner of the top quark, the stop $\tilde{t}$.

2.1.1 The Minimal Supersymmetric Standard Model

In a supersymmetric extension of the SM, every field and its superpartner can be written together as a supermultiplet on which the SUSY algebra operates. There are two kinds of supermultiplets; chiral supermultiplets consisting of a spin-1/2 field paired with a complex spin-0 field, and gauge supermultiplets composed of a spin-1 field and a spin-1/2 field.

The field content of the MSSM is summarised in Tables 2.1 and 2.2 together with the gauge transformation properties. The SM vector boson fields are part of gauge supermultiplets, while the fermions are part of chiral supermultiplets. Left-handed and right-handed fermions have different gauge transformation properties in the SM, which means that they must have their own complex scalar partner. Thus, for example, the first generation of (s)quarks corresponds to the $SU(2)_L$-doublet chiral supermultiplet containing $\tilde{u}_L, u_L$ and $\tilde{d}_L, d_L$, and the $SU(2)_L$-singlet supermultiplet containing $\tilde{u}_R^*$ and $u_R^\dagger$. The squarks and sleptons are scalar, and are therefore not inherently chiral. Their subscripts L and R instead refer to the associated SM particle. By standard convention, all chiral supermultiplets are defined in terms of left-handed fields in the mathematical formulation, which is why Table 2.1 shows the conjugates of the right-handed quarks and leptons and their superpartners.

The SM Higgs boson resides in a chiral supermultiplet together with its spin-1/2 superpartner, the higgsino. It turns out that the addition of only one Higgs chiral supermultiplet leads to gauge anomalies in the theory. These are associated with processes involving loops of particles and lead to divergent integrals in the matrix element. The anomalies in the SM cancel out exactly between the contributions from the quarks and the leptons. However, in the MSSM the addition of one higgsino contributes with extra terms that are not cancelled. This can be avoided if there are two Higgs supermultiplets with mutually opposite values of the

Table 2.1: Chiral supermultiplets in the MSSM.

	Label	**Spin-0**	**Spin-1/2**	$U(1)_Y$	$SU(2)_L$	$SU(3)_C$
Squarks/Quarks (3 generations)	Q	$(\tilde{u}_L, \tilde{d}_L)$	(u_L, d_L)	1/3	**2**	**3**
	U^c	$\tilde{u}_R^*$	$u_R^\dagger$	-4/3	**1**	$\bar{\mathbf{3}}$
	D^c	$\tilde{d}_R^*$	$d_R^\dagger$	2/3	**1**	$\bar{\mathbf{3}}$
Sleptons/Leptons (3 generations)	L	$(\tilde{\nu}_L, \tilde{e}_L)$	(ν_L, e_L)	-1	**2**	**1**
	E^c	$\tilde{e}_R^*$	$e_R^\dagger$	2	**1**	**1**
Higgs/Higgsinos	H_u	(H_u^+, H_u^0)	$(\tilde{H}_u^+, \tilde{H}_u^0)$	1	**2**	**1**
	H_d	(H_d^0, H_d^-)	$(\tilde{H}_d^0, \tilde{H}_d^-)$	-1	**2**	**1**

Table 2.2: Gauge supermultiplets in the MSSM.

	Spin-1/2	**Spin-1**	$U(1)_Y$	$SU(2)_L$	$SU(3)_C$
Gluino/Gluon	$\tilde{g}$	g	0	**1**	**8**
Winos/ *W*-bosons	$\tilde{W}_1$, $\tilde{W}_2$, $\tilde{W}_3$	W_1, W_2, W_3	0	**3**	**1**
Bino/*B*-boson	$\tilde{B}$	B	0	**1**	**1**

charges. The Higgs sector in the MSSM therefore consists of two complex scalar Higgs $SU(2)_L$-doublets and their superpartners.

Supersymmetry-breaking mechanisms

The MSSM Lagrangian is constructed by including all possible supersymmetric interaction terms that satisfy the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetries as well as the conservation of baryon and lepton number. SUSY breaking is then implemented by adding the most general set of soft SUSY-breaking terms consistent with the gauge groups. The theory however provides no explanation for the fundamental origin of the spontaneous symmetry-breaking mechanism. In total, the Lagrangian of the MSSM contains 124 physical parameters, 105 of which are genuinely new compared to the SM.

The large number of parameters can be brought down by applying physically motivated simplifications to the MSSM. Several such models assume that SUSY breaking occurs in a hidden sector of particles that have none, or very small, direct couplings to the MSSM particles. The effects of the spontaneous SUSY breaking in the hidden sector are then mediated to the MSSM by some mechanism, often involving a particle exchange. One such model is Gauge Mediated Supersymmetry Breaking (GMSB) [44] where SUSY breaking is mediated by gauge interactions,

communicated from the hidden sector via a so-called messenger sector. The particle spectrum of the GMSB contains a spin-3/2 gravitino $\tilde{G}$ which results from the symmetry breaking.

Particle spectrum of the MSSM

When the electroweak symmetry breaking and supersymmetry breaking effects are included, there can be mixing between the electroweak gauginos and the higgsinos, and within the various sets of squarks and sleptons and Higgs scalars that have the same electric charge. In the MSSM, the superpartners listed in Tables 2.1 and 2.2 are therefore not necessarily the mass eigenstates of the theory. Table 2.3 shows the full particle spectrum of the MSSM, including the extended Higgs sector.

The neutral higgsino states mix with the neutral gaugino states and the charged states mix with the charged gaugino states. The result is a series of neutral states called neutralinos numbered from the lightest to the heaviest and a series of charged states called charginos. The expected mass spectrum of these electroweakinos depends on the mixing among the particles, which can lead to more wino/bino-like or more higgsino-like states. For the squarks and sfermions, the left- and right-handed states can mix to form the physical fields. The amount of mixing depends on the SUSY breaking parameters of the theory, and in the MSSM, the mixing is only appreciable for the third generation.

The Higgs scalar fields in the MSSM correspond to eight real scalar degrees of freedom. After electroweak symmetry breaking, five of the original eight Higgs components are left, the lightest of which can be interpreted as the SM Higgs boson h^0. Two neutral states and two charged states then remain, corresponding to two more neutral and two charged Higgs bosons.

R-parity

The most generic versions of a supersymmetric theory include terms which violate lepton and baryon number. These take the form

$$W_{\Delta \mathrm{L}=1} = \frac{1}{2}\lambda_{ijk} L_i L_j E^c_k + \lambda'_{ijk} L_i Q_j D^c_k + \mu_i L_i H_u \tag{2.3}$$

$$W_{\Delta \mathrm{B}=1} = \frac{1}{2}\lambda''_{ijk} U^c_i D^c_j D^c_k, \tag{2.4}$$

where the field labels have been introduced in Table 2.1, i, j, k are generation indices, and λ_{ijk}, λ'_{ijk}, λ''_{ijk}, μ_i are the couplings for each interaction term. The chiral supermultiplets carry baryon number $B = +1/3$ for Q_i and $B = -1/3$ for U^c_i and D^c_i, and lepton number $L = +1$ for L_i and $L = -1$ for E^c_i. All other fields have $B = L = 0$. The terms in Eqs. 2.3 and 2.4 therefore violate the lepton and baryon number by one unit respectively. The values of the couplings in these

Table 2.3: Additional particle content of the MSSM, including the extended SM Higgs sector.

	Spin	**Gauge eigenstate**	**Mass eigenstate**
Squarks	0	$\tilde{u}_L$, $\tilde{u}_R$, $\tilde{d}_L$, $\tilde{d}_R$	same
		$\tilde{s}_L$, $\tilde{s}_R$, $\tilde{c}_L$, $\tilde{c}_R$	same
		$\tilde{t}_L$, $\tilde{t}_R$, $\tilde{b}_L$, $\tilde{b}_R$	$\tilde{t}_1$, $\tilde{t}_2$, $\tilde{b}_1$, $\tilde{b}_2$
Sleptons	0	$\tilde{e}_L$, $\tilde{e}_R$, $\tilde{\nu}_e$	same
		$\tilde{\mu}_L$, $\tilde{\mu}_R$, $\tilde{\nu}_\mu$	same
		$\tilde{\tau}_L$, $\tilde{\tau}_R$, $\tilde{\nu}_\tau$	$\tilde{\tau}_1$, $\tilde{\tau}_2$, $\tilde{\nu}_\tau$
Neutralinos	1/2	$\tilde{B}$, $\tilde{W}^3$, $\tilde{H}_u^0$, $\tilde{H}_d^0$	$\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$, $\tilde{\chi}_3^0$, $\tilde{\chi}_4^0$
Charginos	1/2	$\tilde{W}^1$, $\tilde{W}^2$, $\tilde{H}_u^+$, $\tilde{H}_d^-$	$\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^\pm$
Gluino	1/2	$\tilde{g}$	same
Higgs bosons	0	H_u^0, H_d^0, H_u^+, H_d^-	h^0, H^0, A^0, $H^\pm$

terms are tightly constrained by experiment, and in particular by the observed stability of the proton. If both λ' and λ'' are non-zero, this can result in proton decay, where a squark mediates the decay into a lepton and pion as shown in Figure 2.2. The decay time of the proton is measured experimentally to be in excess of 10^{29} years [18]. Therefore, at least one of the two couplings must be suppressed for each combination of generations.

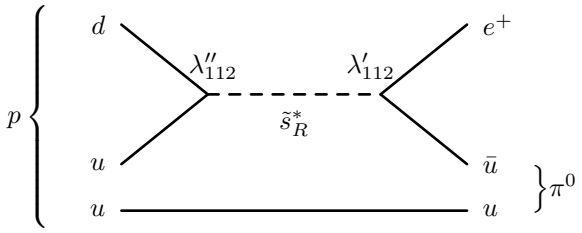

Figure 2.2: Diagram showing the proton decay $p \to e^+\pi^0$ mediated by a strange squark.

The terms in Eqs. 2.3 and 2.4 are not included in the MSSM. This is enforced by adding a new symmetry, R-parity, which has the effect of eliminating the possibility of baryon and lepton number violating terms in the Lagrangian. R-parity is a

multiplicative quantum number defined as

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

where B is the baryon number, L is the lepton number, and s is the spin of the particle. All SM particles have positive R-parity while all sparticles have negative R-parity.

In SUSY models with R-parity conservation, the lightest SUSY particle (LSP) is required to be stable since it cannot decay into a SM particle without violating the conservation of R-parity. This particle, provided that it does not interact electromagnetically or strongly, provides a good WIMP-like candidate for a dark matter particle. In the MSSM, the lightest neutralino is often viewed as the dark matter candidate, and in GMSB models it is the gravitino.

2.2 Searches for Supersymmetry at the LHC

At the LHC, the search for SUSY proceeds both directly by searching for the decay products of new particles produced in the pp collisions, and indirectly through precision measurements of flavour or electroweak observables, which can be analysed for radiative effects of new particles. These approaches are complementary and equally important.

SUSY particles can be produced in the pp collisions via the electroweak force through interactions similar to those shown in Figure 2.3, or via the strong force through interactions like those in Figure 2.4. The predicted cross sections for various production modes are shown in Figure 2.5 for pp centre-of-mass energies of 8 TeV and 13-14 TeV, corresponding to the LHC operating energies. Strong production has the largest cross section and gluino and squark pairs are therefore expected to be the most abundantly produced SUSY particles at the LHC if allowed by their masses.

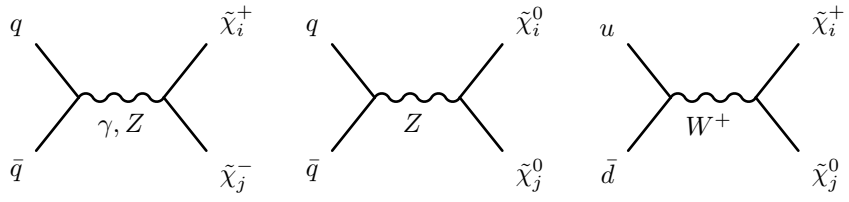

Figure 2.3: Diagrams showing the electroweak production of sparticles in pp collisions.

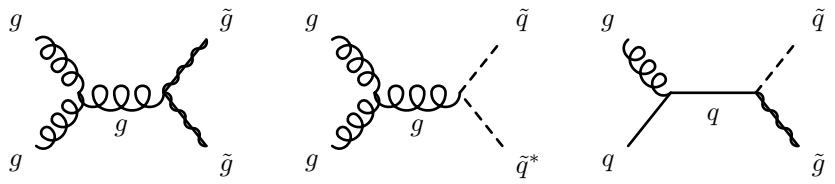

Figure 2.4: Diagrams showing the production of gluinos and squarks in pp collisions via the strong interaction.

Typically, the searches for SUSY exploit the topology of the SUSY signals in order to define the event selections of the analysis. The vast parameter space of the MSSM however presents a challenge for making predictions of the phenomenology of the SUSY interactions. This has led to a general search approach in simplified scenarios involving a reduced number of parameters, each predicting a given production and decay topology, which can be realised in several different SUSY models. The simplified models consist of one diagram describing the production of sparticles and their decay. All sparticles not directly involved in the process are effectively decoupled, for example by setting their masses to an energy which is well beyond the reach of the experiment. The masses of the remaining sparticles in the simplified process are then free to be tuned directly.

The analyses often make use of grids with simulations generated by varying the masses of the sparticles in the targeted simplified model. Such grids allow for a fast scan of the possible signals and for the prediction of the likelihood to detect them in the experiment. If the experimental results are in agreement with the SM prediction, the grids are used to set exclusion limits, typically at 95 % confidence level (CL), on the masses of the sparticles involved in the simplified models. These limits provide help for the design of future analyses and for the comparison of results between different searches and experiments. However, they are based on all assumptions of the considered simplified models and should therefore not be interpreted to mean that no other SUSY model can exist with the excluded masses.

2.3 Long-lived particles

The decay rate of a SM particle was described in Section 1.5, together with several mechanisms which can suppress the decay and give the particle a longer lifetime. Similar mechanisms can also cause BSM particles to be long-lived. This section will list a few models which predict SUSY particles with long lifetimes. Similar mechanisms can also be realised in other BSM models, including dark matter models,

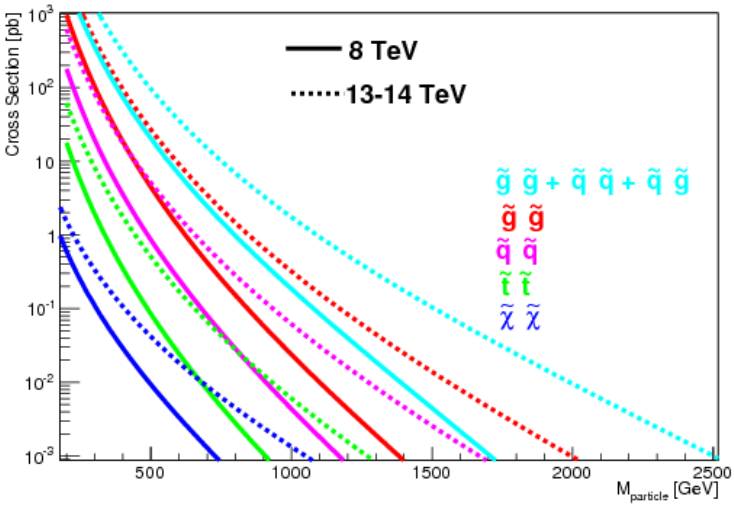

Figure 2.5: Cross sections for SUSY particle production in pp collisions at centre-of-mass energies at 8 TeV and 13-14 TeV, as a function of the sparticle mass [45]. The coloured particle cross sections are evaluated at 8 TeV and 13 TeV while the electroweak pure higgsino cross sections are evaluated at 8 TeV and 14 TeV.

magnetic monopoles, and hidden valley models [22], resulting in the same kind of experimental signatures.

R-parity violating SUSY

In the MSSM, it is assumed that R-parity is an exact symmetry, thereby setting the R-parity-violating (RPV) couplings λ, λ', and λ'' to zero. It is however possible to relax this requirement within the experimental limits that exist on B- and L-violating processes and build models of SUSY where R-parity is violated. The couplings are still expected to be very suppressed, meaning that a SUSY particle which decays through an RPV coupling should have a relatively long lifetime.

Split SUSY

In Split SUSY models, the requirement of naturalness is relaxed, and SUSY is no longer the solution to the hierarchy problem. The SUSY breaking can then occur at a much larger energy scale compared to the electroweak-scale symmetry breaking in the MSSM. All the scalars in are then significantly heavier than the SUSY fermions, with the exception of one Higgs boson, which is fine-tuned to be light. In this kind

of models, the gluino often has a long lifetime which arises due to the fact that it only can decay through a squark which, by construction, is much heavier. The decay will therefore be suppressed and the gluino lifetime will depend on the mass splitting [46].

GMSB SUSY

In GMSB models, the mechanism that gives rise to long-lived particles depends on the mass scale of the SUSY-breaking mediator. In this model, the next-to-lightest SUSY particle (NLSP) is typically the lightest neutralino which decays into the gravitino and a SM particle at a rate suppressed by the messenger mass scale. Depending on the fundamental scale of supersymmetry breaking, the NLSP can have a measurable lifetime relevant for long-lived particle searches [47, 48].

Part II

Experimental setup

Chapter 3

The Large Hadron Collider

The LHC [1] is a circular particle accelerator, with a 27 km circumference, operating at CERN. It is contained in an underground tunnel located at an average distance of 100 m below ground, on the border between Switzerland and France. The machine is nominally used for proton-proton (pp) collisions but can also be run in heavy-ion configurations wherein proton-lead or lead-lead collisions take place. For the pp configuration, the LHC is designed for collisions at a centre-of-mass energy up to $\sqrt{s} = 14\,\text{TeV}$. The following sections introduce the design of the LHC and provide a more detailed discussion about the pp operation. Additionally, the concepts of centre-of-mass energy and luminosity, which dictate the physics potential of the collider, are introduced.

3.1 LHC design and operating principle

The LHC tunnel hosts two adjacent, parallel beam pipes kept at ultra-high vacuum which, during operation, contain counter-rotating beams of protons guided around the ring by a strong dipole magnetic field. The two beams are prepared by a system of smaller accelerators, making up the LHC injector chain, and are sorted into packs of protons called bunches. Acceleration is achieved as the beam repeatedly traverses a number of radio frequency (RF) cavities that provide an oscillating longitudinal electric field. The dipole magnetic field is increased synchronously with the energy of the particles and once the maximum field is achieved, the coasting beams are brought into collision.

3.1.1 Machine layout

The LHC is composed of eight octants, as shown in Figure 3.1, consisting of straight elements connected by arcs. In four of the straight sections, labelled Point 1, 2, 5, and 8, the beam pipes are intersected, causing the beams to collide. At these points, large underground caverns are built and host the experiments ATLAS [49], ALICE [50], CMS [51], and LHCb [52] which are dedicated to the study of the physics processes that occur in the collisions. ATLAS and CMS are general-purpose detectors with broad research programs, whereas the ALICE and LHCb detectors are built specifically for the study of heavy-ion collisions and b-hadron physics, respectively. The remaining points contain various service elements necessary for the operation of the accelerator.

The system of RF cavities that accelerates the beams to their collision energies is placed in Point 4. A total of 16 RF cavities are housed in four cylindrical cryomodules which enable them to work in a superconducting state. Each cavity can reach a maximum voltage of 2 MV and is tuned to oscillate at a frequency of 400 MHz.

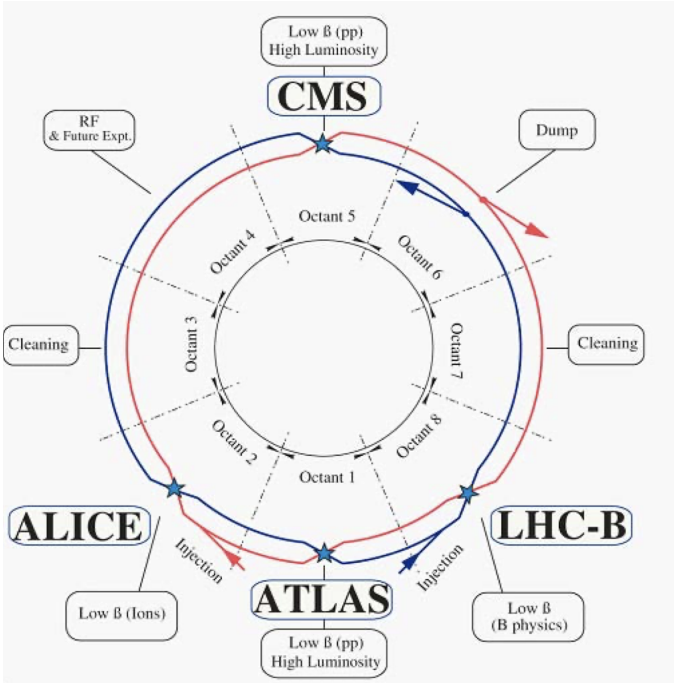

Figure 3.1: Layout of the LHC and its two counter-rotating beams [1]. Beam 1 is illustrated in blue and rotates counter-clockwise while beam 2 is in red and rotates clockwise. At the center of each octant is a straight section which houses the experimental caverns or LHC beam facilities.

The dipole magnets are placed in the arcs connecting the straight sections. A total of 1232 superconducting dipole magnets, which are 15 m long and produce a field with a design strength of 8.3 T, are used to steer the beam. The magnets are cooled by liquid helium and are operated at a temperature of 1.9 K, with both beam pipes housed in the same cryostat. In addition, a total of 392 superconducting quadrupole magnets are used to focus the beams, with stronger magnets close to the intersection points. Under nominal conditions the beams do not collide head on, but with a small angle to avoid unwanted collisions near the interaction point. Besides the dipole and quadrupole magnets, Point 8 hosts kicker magnets which can ramp up quickly and are used to divert the beams out of the LHC ring in a safe manner.

3.1.2 Injection chain

The LHC relies on a series of pre-acceleration steps that bring initial low-energy protons to an energy of 450 GeV before they are injected into the LHC. These steps are referred to as the LHC injector chain and form the heart of the CERN accelerator complex illustrated in Figure 3.2. The protons are initially sourced from hydrogen atoms which are released at the start of the linear accelerator Linac2 where the protons are brought up to an energy of 50 MeV. This is followed by three circular accelerators of increasing size, the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS), where the energy is brought to 1.4 GeV, 25 GeV and 450 GeV, respectively. The beams are then injected into the LHC at two points, one for the counter-clockwise beam, and one for the clockwise beam. A standard LHC fill takes on the order of 4 minutes per ring and the maximum LHC energy is reached in about 20 minutes.

3.1.3 Bunch structure

The protons arrive at the LHC in bunches which are initially prepared in the injection chain and are then kept in their final bunch structure by the oscillations of the RF cavities. When the beam has reached the required energy, an ideally timed proton with exactly the right energy will not be accelerated. By contrast, protons with slightly different energies arriving to the cavities earlier or later will be accelerated or decelerated so that they stay close to the desired energy. The relationship between the RF oscillations and the bunch structure is illustrated in Figure 3.3.

The oscillation frequency of of the RF cavities defines the boundaries in which proton bunches can lie. These boundaries are called RF buckets and, along with the circumference of the LHC, dictate the number of proton bunches that can potentially fit in the LHC. In principle, the 400 MHz frequency allows for 35640 RF buckets, but in practice only a tenth of them are used. This is because the

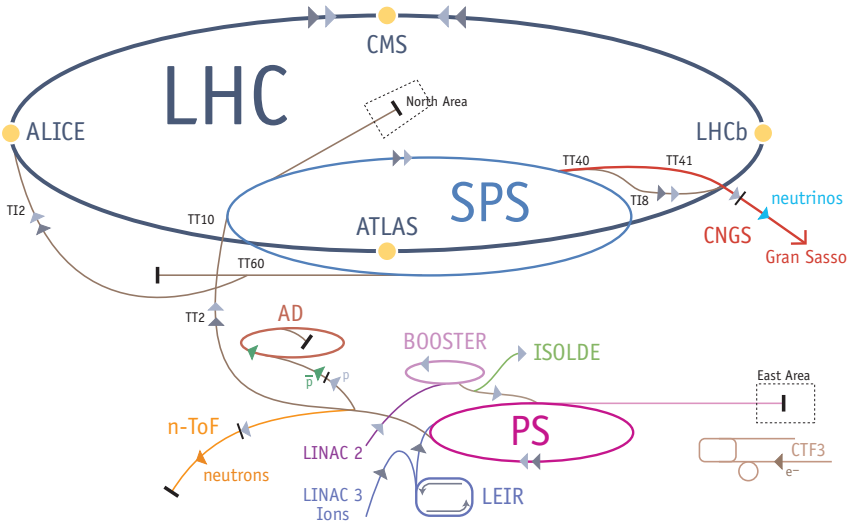

Figure 3.2: The LHC accelerator complex [53]. Protons are extracted from hydrogen and are accelerated in steps in Linac2, the PSB (Booster), the PS, and the SPS before being injected into the LHC.

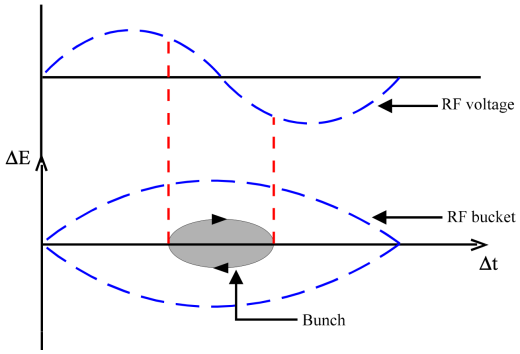

Figure 3.3: Illustration of the proton bunch structure in the LHC [54]. The amplitude of the RF field is illustrated at the top. The boundary of the RF bucket is defined by a full period of the RF oscillation and the particle bunch formation, depicted in grey, occurs at the central node of the oscillation.

experiments are designed for a maximum bunch-crossing rate of 40 MHz and higher frequencies might result in too high radiation levels and difficulties in separating the data between the bunch crossings. Typically, the proton bunches are therefore filled with a spacing in time of 25 ns in the LHC. The position of a proton bunch along the ring with a 25 ns spacing is called the bunch-crossing identifier (BCID) and runs from 0 to 3563.

Proton bunches are filled in two main patterns into the LHC; either as individual bunches where there are empty BCIDs before and after the filled bunch, or as bunch trains consisting of patterns of several bunches in consecutive 25 ns slots with longer gaps between these patterns. To fill an accelerator with proton bunches, several fills of the preceding accelerator in the injection chain are needed. The exact scheme in which the different accelerators are sequentially filled with proton bunches effectively determines the final bunch pattern in the LHC. When the PS is filled it holds 72 proton bunches and 12 empty bunches with a 25 ns spacing. An SPS fill then consists of three or four PS fills and the LHC consists of 12 of these variable-length SPS fills. This is illustrated in Figure 3.4. The total number of bunches is always lower than the 3564 BCIDs potentially allow for. This is in part due to the non-trivial bunch-filling schemes used, but also because there needs to be a 3 µs abort gap without any filled BCIDs to allow the kicker magnets to ramp up when aborting the beam.

3.2 Performance goals and limitations

The LHC project was formally approved by CERN in 1994 as the next step in a long history of previous accelerators, all with increasing collision energies. In the years leading up to the LHC, several important discoveries in particle physics had been made at hadron collider experiments. The W and Z bosons were discovered at the SPS, and the top quark was discovered in at the Tevatron [55], which both collided protons and antiprotons at energies up to 900 GeV and 1.96 TeV respectively. In addition, the Large Electron Positron (LEP) collider [56, 57], which was in operation at CERN between the years of 1989 to 2000, had provided extremely precise measurements of several electroweak observables. The LHC was planned as the successor for LEP at CERN and its design was steered by the goal to have an accelerator able to probe any new physics discovered at LEP or the Tevatron, and to provide the necessary power to search for still-elusive hints of SM or BSM physics. In particular, in case of a non-discovery of the predicted Higgs boson at the previous experiments, the LHC should be powerful enough to produce Higgs bosons at sufficient rate, a requirement expected, and now also confirmed, to be met by an $O(10)$ TeV pp collider.

The physics potential of a collider depends in large on two main figures of merit; the centre-of-mass energy available to probe smaller scales and produce heavier parti-

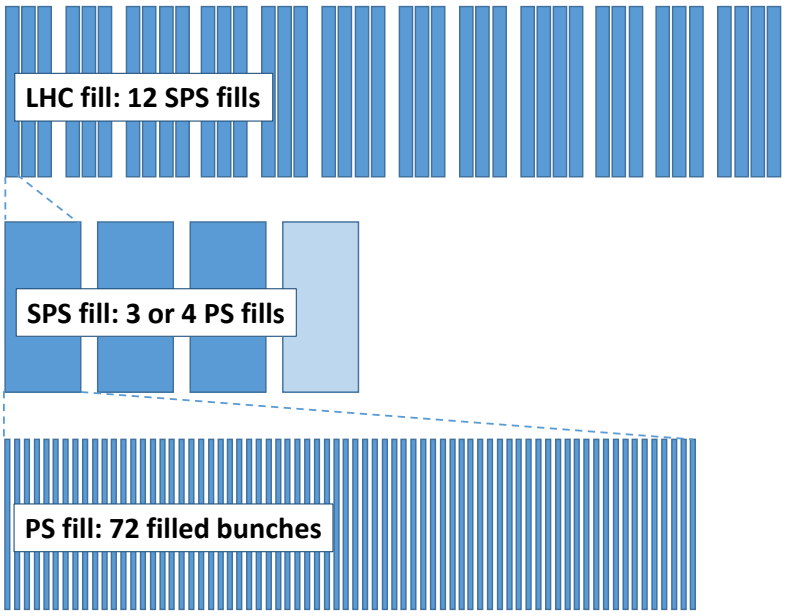

Figure 3.4: A schematic view of the 25 ns filling scheme of the LHC.

cles, and the luminosity which in essence is a measure of the amount of collisions the collider produces.

3.2.1 Centre-of-mass energy

For a circular hadron collider, the centre-of-mass collision energy is limited by the size of the accelerator and the strength of the magnetic fields used to keep the circulating charged particles in orbit. This can be seen by first considering the expression for the relativistic cyclotron frequency, ω, of a particle moving in a circular orbit in a magnetic field,

$$\omega = \frac{qB}{m\gamma}, \tag{3.1}$$

where m and q are the rest mass and electric charge of the particle, B is the magnitude of the magnetic field, and γ is the relativistic Lorentz factor. This can then be related to the kinetic energy according to

$$E_{\text{kin}} \propto mv^2 = m(\omega R)^2 = \frac{q^2 B^2 R^2}{m\gamma^2}, \tag{3.2}$$

where R is the radius of the orbit. In order to reduce the costs in civil engineering and real-estate works for its construction, the LHC is housed in the same tunnel as was used for the LEP accelerator. The beam orbit was thus already prescribed, and the main challenge to meet the $O(10)$ TeV design goal was the development of a magnet system strong enough to bend the particles along the existing tunnel.

3.2.2 Luminosity

The luminosity is determined by the rate of particle collisions produced by the collider. Assuming the production of a general final state O at the LHC, the luminosity $\mathcal{L}$ is the process-independent proportionality factor between the rate $R_{pp\to O}$ and the production cross section $\sigma_{pp\to O}$:

$$R_{pp\to O} = \mathcal{L} \cdot \sigma_{pp\to O} \tag{3.3}$$

The luminosity is a function of the LHC beam parameters and in the simplified case of uniform bunch population, round beams, and equal beam parameters, it can be expressed according to [58]:

$$\mathcal{L} = \frac{N_{\text{b}}^2 n_{\text{b}} f_{\text{r}}}{4\pi\sigma^2} F. \tag{3.4}$$

Here, N_{b} is the number of protons per bunch, n_{b} the number of colliding bunches, f_{r} the revolution frequency of the bunches, and σ is the transverse beam width. The

factor F accounts for the luminosity reduction due to the beam crossing angle which leads to a varying transverse bunch size at the collision point. The beam parameters determining the transverse beam size at the interaction point are the emittance ϵ and the β-function. The emittance is a measure of the average spread of the beam in position and momentum space. A low-emittance beam is one where the particles are confined to a small distance and have nearly the same momentum. This is an intrinsic beam parameter and cannot be changed by the focusing properties of the accelerator. The amplitude function, β, determines the maximum amplitude a single particle trajectory can reach at a given position in the ring and β^* is its value at the interaction point. It is determined by the focusing properties of the quadrupole magnets. For Gaussian beams, the transverse beam size at the interaction point is given by $\sigma^2 = \sqrt{1 - 1/\gamma^2} \cdot \epsilon\beta^*$ where γ is the Lorentz factor of a beam particle.

The running conditions of the LHC vary with time and the luminosity therefore has a time dependence. The definition in Eqs. 3.3 and 3.4 is referred to as the *instantaneous luminosity* and is usually expressed in units of $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The integrated luminosity, denoted by $L = \int \mathcal{L}(t)\mathrm{d}t$, refers to the instantaneous luminosity accumulated over a certain time interval, e.g a data-taking year. It is measured in units of cm^{-2}, or equivalently in inverse barns b^{-1}.

The luminosity can be increased by increasing the number of protons per bunch, the number of colliding bunches, or by reducing the transverse size of the beam at the collision point. The latter can be done by using a lower emittance beam, or by squeezing the beam more with the focusing magnets i.e reducing β^*. Limitations of the beam intensity instead come from undesired beam-beam interaction effects and electron clouds which build up around the beam and can be harmful for the machine and the experiments [1]. Typical values of the beam parameters are presented in Section 3.3. The design instantaneous luminosity of the LHC is $1 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [1]. During an LHC fill, the instantaneous luminosity typically decreases as a function of time, mainly due to the decay in beam intensity as the protons collide, emittance growth, and increase of the bunch length. This natural decay is approximately exponential with a time constant on the order of 6 hours [58].

The typical instantaneous luminosities at which the LHC operates are so high that it is very likely that multiple proton pairs will collide in the same bunch crossing. The multiple interactions are referred to as *in-time* pileup and must be separated from the main interaction of interest for the physics analysis. The average number of inelastic collisions per bunch crossing is denoted by the symbol $\langle\mu\rangle$. In addition, the experiments can also be sensitive to *out-of-time* pileup from pp collisions occurring in bunch crossings just before or after the collision of interest. When detectors are sensitive to several bunch crossings or their electronics integrate over more than 25 ns, these collisions can affect the signal in the collision of interest.

3.3 Operational history

The LHC began its planned research program in the spring of 2009. It operates in so-called runs corresponding to periods of continuous data taking, interleaved by long shut-downs for maintenance and upgrades of the machine and the experiments. At the time of writing, there have been two runs of the LHC; Run 1 which took place during the years 2009–2012 and Run 2 which took place during 2015–2018. In Run 1, the accelerator was operated at centre-of-mass energies of 7 TeV and 8 TeV and with a bunch spacing of 50 ns. A total of 5.61 fb^{-1} of 7 TeV data and 23.3 fb^{-1} of 8 TeV data was delivered in 2011 and 2012, respectively [59]. In Run 2, the LHC was operating at a centre-of-mass energy of 13 TeV and with a bunch spacing of 25 ns. Thanks to increased instantaneous luminosity, 156 fb^{-1} of data was delivered to the experiments during these years. The integrated luminosities for each of the data-taking years between Run 1 and Run 2 is shown in Figure 3.5.

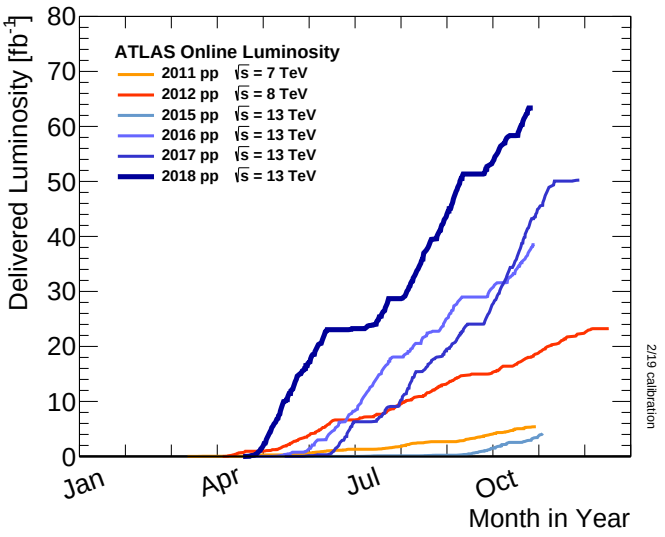

Figure 3.5: Integrated luminosity versus day delivered to ATLAS during stable beams for pp collisions [60].

The LHC will resume operation in spring 2022, marking the start of its Run 3 which is expected to last until the end of 2025. A small increase in centre-of-mass energy to 13.6 TeV is also expected, while the instantaneous luminosity will be kept similar to Run 2. After Run 3 the statistical gain in running the accelerator without a significant luminosity increase will become marginal. Therefore, to maintain scientific progress, the LHC will be upgraded to run at up to 10 times higher instantaneous luminosity, with a predicted accumulated dataset of around 3000 fb^{-1}. This project

Table 3.1: Selected LHC parameters for pp collisions at $\sqrt{s} = 13\,\text{TeV}$ in 2015-2018. The values shown are representative of the best accelerator performance during normal physics operation. In 2017, the LHC was run in two modes: standard 25 ns bunch train operation with long trains, and '8b4e', denoting a pattern of eight bunches separated by 25 ns followed by a four bunch-slot gap. Values are given for both configurations.

Parameter	**2015**	**2016**	**2017**	**2018**
Max number of colliding bunch pairs, n_b	2232	2208	2544/1909	2544
Bunch spacing (ns)	25	25	25/8b4e	25
Typical bunch population (10^{11} protons)	1.1	1.1	1.1/1.2	1.1
β^* (m)	0.8	0.4	0.3	0.3–0.25
Peak luminosity ($10^{33}\,\text{cm}^{-2}\text{s}^{-1}$)	5	13	16	19
Peak $\langle\mu\rangle$	~ 16	~ 41	$\sim 45/60$	~ 55

is referred to as the High Luminosity LHC (HL-LHC) and is predicted to start operation in 2029. To handle the increased luminosity, pileup, and the radiation of the HL-LHC environment, the detectors at the LHC will also need to be upgraded, and this will take place in the long shut-down after Run 3.

The work described in this thesis is based on the Run 2 dataset. Table 3.1 shows an overview of typical best beam parameters for each year in the run, together with the total delivered integrated luminosity. Figure 3.6 shows the distribution of the average number of interactions per bunch crossing. All the Run 2 running took place with long trains of bunches with 25 ns bunch spacing within the trains, except for the second part of 2017, where a special filling pattern with eight filled bunches separated by 25 ns followed by a four bunch-slot gap was used. This beam pattern suppresses the formation of electron clouds compared to the standard beam, and was necessary to cope with a temporary LHC vacuum issue. The luminosity was levelled by a partial beam separation at the beginning of such 8b4e LHC fills to give a maximum $\langle\mu\rangle$ of 60.

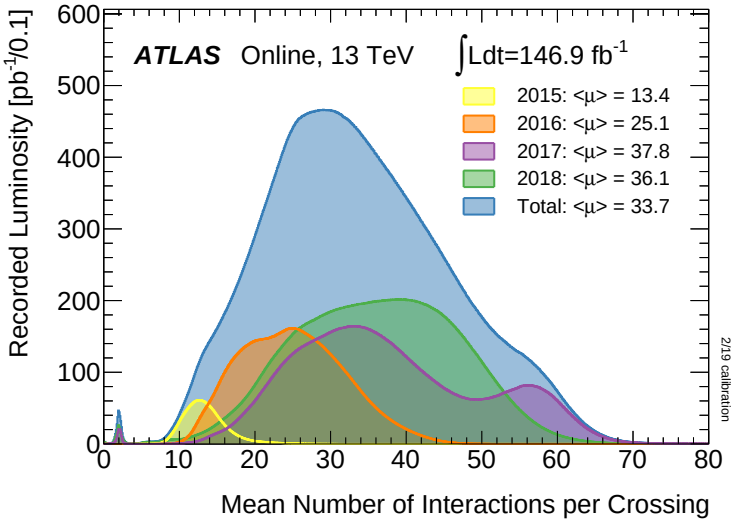

Figure 3.6: Average number of interactions per bunch crossing weighted by the integrated luminosity shown for 2015, 2016, 2017, and 2018 separately, as well as the sum of the four years [60].

Chapter 4

The ATLAS experiment

The ATLAS detector is hermetic and covers nearly the full 4π solid angle around the interaction point in order to exploit the full physics potential of the *pp* collisions at the LHC. The detector consists of various subsystems that are dedicated to the identification of particles and to the measurement of their kinematic properties. These are layered around the interaction point and are cylindrically symmetric since the *pp* interactions in the LHC have no preferred direction in the plane transverse to the beam line. ATLAS is thus roughly cylindrical in shape, with the LHC beam pipe as its longitudinal center axis, and is in total 44 m long and 25 m in diameter.

An overview of the ATLAS detector is shown in Figure 4.1. It consists of concentric detector layers in the so-called *barrel* and is sealed by two *endcap* structures where the detector layers form disks. Closest to the beam pipe is the Inner Detector (ID) which is designed to reconstruct the trajectories of charged particles, called tracks, and to locate the *pp* interaction point. The ID is submerged in a magnetic field from a solenoid magnet that surrounds it and allows for a momentum measurement from the bending of charged particle tracks. Outside the solenoid are the electromagnetic and hadronic calorimeters that have the purpose of stopping particles in order to measure their energy. The outermost subdetector is the Muon Spectrometer which identifies muons and measures their momenta. This system is installed inside and around a set of toroidal magnets which allow for a precise muon momentum measurement. In addition to these systems, the ATLAS detector contains subdetectors that are designed specifically for the measurement of the luminosity.

The ATLAS detector subsystems are designed to tolerate the high-intensity radiation doses from the LHC and use fast, radiation-hard or shielded electronics capable of providing a distinct readout for a 25 ns bunch-crossing rate. However, neither the data acquisition system nor the resources for doing offline analysis are capable of handling the full amount of data. Therefore, a trigger system is required to select only the most interesting events to be written to disk and analysed further offline.

In the following sections, the coordinate system used to describe the ATLAS geometry is introduced and the magnet systems and all detector subsystems are discussed in more detail. Particular focus is dedicated to the ID, since both the luminosity work and the search for long-lived particles described in this thesis are heavily dependent on the reconstruction of charged particle tracks. The trigger and data acquisition systems are also discussed. A complete description of the ATLAS experiment can be found in Ref. [49].

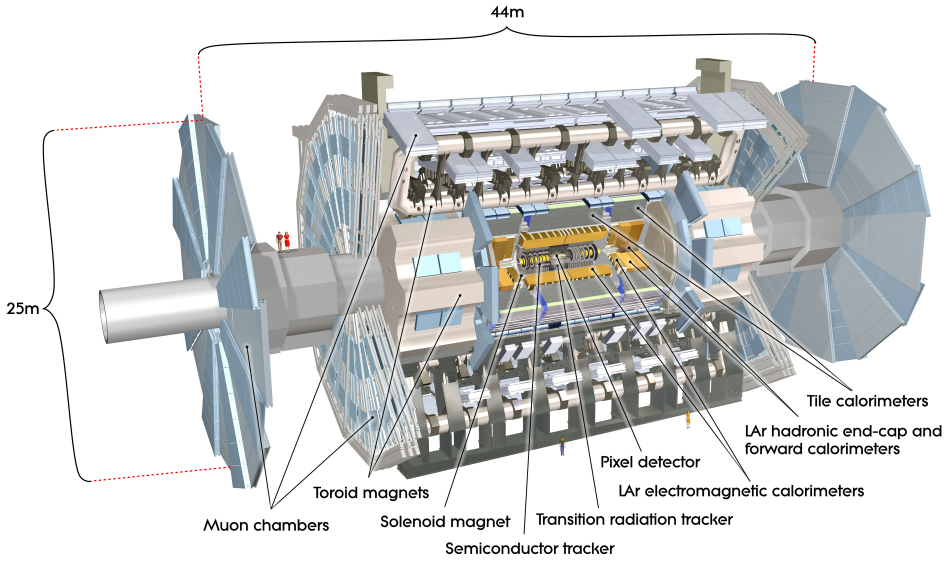

Figure 4.1: Overview of the ATLAS detector, with labelled subsystems and magnets [49].

4.1 Coordinate system

The coordinate system used in ATLAS is a right-handed system with the origin at the nominal pp collision point, as shown in Figure 4.2a. Its x-axis points towards the center of the LHC ring, its y-axis points upwards, and its z-axis points along the counter-clockwise beam direction. The sides associated with positive and negative values of z are referred to as the A and C sides of the detector, respectively. A cylindrical coordinate system is also used, with the same z-axis, the radius R measured from the origin in the x-y plane, and the azimuthal angle ϕ measured from the positive x-axis in the x-y plane. In addition, the polar angle θ is frequently employed and is measured from the positive z-axis.

When describing systems of particles or locations inside the ATLAS detector, the pseudorapidity $\eta = -\ln(\tan\theta/2)$ is often preferred over the polar angle, since the particle flow is approximately uniform per unit in η. Positive or negative values of η then correspond to the positive and negative directions along the z-axis, while $\eta = 0$ corresponds to the transverse plane, as shown in Figure 4.2b. Objects at large absolute value of η are referred to as *forward*, while objects at small values are referred to as *central*. A metric commonly used to describe the distance between two systems of particles in the detector is $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

The four-momentum of a particle in ATLAS is typically described in the coordinates $(p_{\mathrm{T}}, \eta, \phi, m)$ where $p_{\mathrm{T}} = p\sin\theta$ is the transverse momentum and m the mass of the particle. The transverse momentum, the azimuthal angle, and the mass are invariant under boosts along the z-axis and the pseudorapidity changes approximately only by an additive constant. The difference in pseudorapidity between two particles is therefore invariant under boosts along the z-axis. This choice of coordinates is motivated by the nature of the pp collisions at the LHC. When two protons collide, the fundamental interaction takes place between individual partons, which do not typically have perfectly balanced momenta resulting in a boost of the interaction products along the beam line.

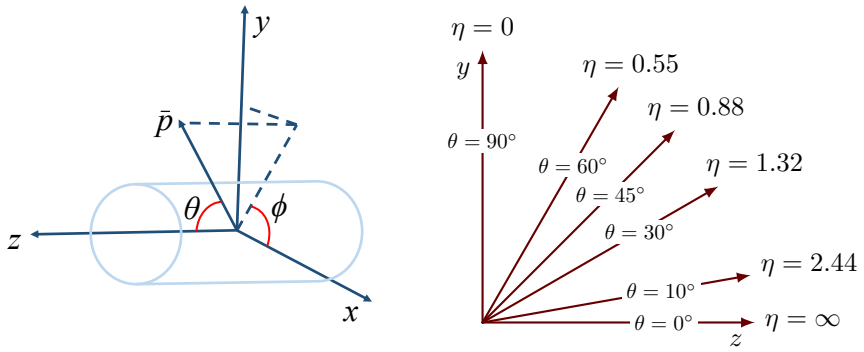

(a) (b)

Figure 4.2: (a): Coordinate system used in ATLAS. (b): Relation between the pseudorapidity η and the polar angle θ.

4.2 Magnet system

A charged particle in a magnetic field will follow a curved path where the bending radius depends on the momentum of the particle and the strength of the magnetic

field. For a uniform field B, the momentum of the particle transverse to the field is given by

$$p_{\mathrm{T}} = qBr, \tag{4.1}$$

where q is the charge and r is the bending radius. ATLAS exploits this relationship and uses powerful magnetic fields to bend the trajectories of the high-energy charged particles created in the pp collisions in order to provide accurate momentum measurements. The ATLAS magnet system [61] is shown in Figure 4.3 and consists of four superconducting magnets; an inner solenoid, a barrel toroid, and two endcap toroids. These are cooled with liquid helium and are kept at an operating temperature of about 4.5 K.

The solenoid encloses the ID and provides a 2 T axial magnetic field that causes the trajectories of charged particles to bend in the transverse plane. It is important that the solenoid interferes minimally with particles in order to allow for a high-resolution measurement of their energies in the calorimeters. The magnet is made of superconducting wire and shares a vacuum chamber with the electromagnetic calorimeter in order to avoid extra material. It stretches 5.8 m along the z-axis and between 2.5-2.6 m in diameter.

The toroidal magnets consist of eight equally spaced superconducting coils in the barrel, symmetrically spaced around the ϕ direction, and two end-cap structures. These provide an azimuthal magnetic field which causes the trajectories of the muons to bend in the r-z plane. The magnetic field has a complex structure and is mapped by magnetic field sensors. In the barrel, the field varies between 0.2-2.5 T and in the endcap toroids it varies between 0.2-3.5 T. The inner and outer diameters of the barrel toroid magnet system are 9.4 m and 20.1 m, respectively, and the magnet system spans 25.3 m along the beam direction.

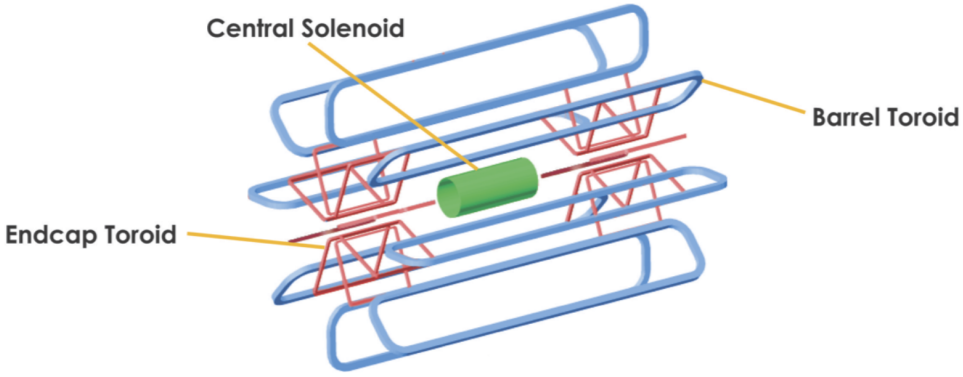

Figure 4.3: Overview of the ATLAS magnet system [61]. The solenoid magnet is shown in green, the barrel toroid system in blue, and endcap toroid magnets in red.

4.3 Inner Detector

The ID [62, 63] is composed of three subsystems; the Pixel detector closest to the beam pipe, the Semiconductor Tracker (SCT), and the outermost Transition Radiation Tracker (TRT). Figure 4.4 shows all three subsystems in the barrel and endcaps, and Figure 4.5 shows further information about the physical locations in r, z, and η of each of the components that make up the various ID subsystems. The ID measures 6.2 m in length and 2.1 m in diameter and the coverage extends up to $|\eta| = 2.5$ for the Pixel detector and the SCT and up to $|\eta| = 2.0$ for the TRT.

The modules in the various layers in the ID produce an electrical signal, called a hit, when a charged particle travels through their active material. In the ATLAS tracking software, which is described in Section 5.1.1, all these hits are considered in order to reconstruct the trajectory of the particle. Several crucial aspects of particle identification depend on information from the tracking and it is essential to reconstruct the particle trajectories with very high momentum, angle, and position resolution. At the same time, the ID must be light-weight in order for passing particles not to lose their energy. It is therefore constructed to take accurate position measurements so that tracks can be reconstructed using only a few hits. The design resolution on the transverse momentum of a particle measured in the ID is given by

$$\frac{\sigma_{p_\mathrm{T}}}{p_\mathrm{T}} = 0.05\,\% \cdot p_\mathrm{T}[\mathrm{GeV}] \oplus 1\,\%. \tag{4.2}$$

Here, the momentum-dependent term results from the intrinsic resolution on the measurement points and the constant term results from multiple scattering, where the charged particles undergo many small coulomb interactions with nuclei in the detector material. Table 4.1 summarises the element sizes and intrinsic resolutions, while the following sections go into more details of each subsystem.

Subdetector	**Element size** (µm)	**Intrinsic resolution** (µm)
Pixel	50×400	10×115
	50×250 (IBL)	8×40 (IBL)
SCT	80	17×580
TRT	4000	130

Table 4.1: Summary of the main characteristics of the ID subdetectors. The intrinsic resolution of the Pixel and SCT is reported along r-ϕ and z, and for TRT along r-ϕ. The innermost Pixel layer, called IBL, has a finer resolution compared to the other layers, which is quoted separately. For SCT and TRT the element sizes refer to the spacing of the readout strips and the diameter of the straw tubes, respectively.

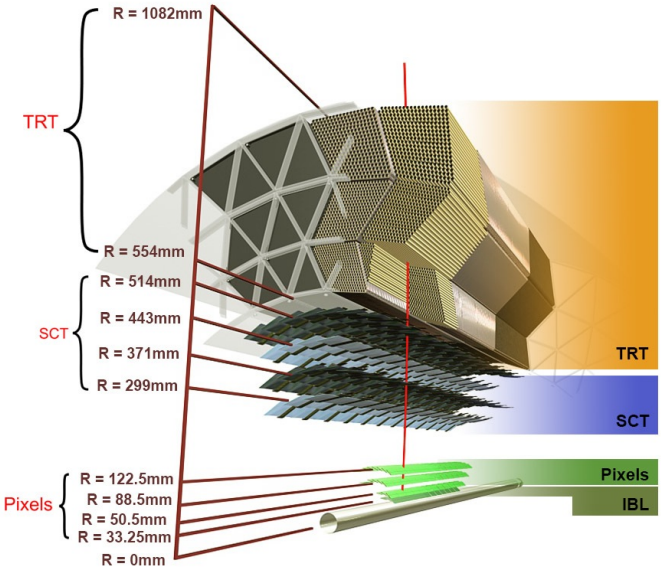

(a)

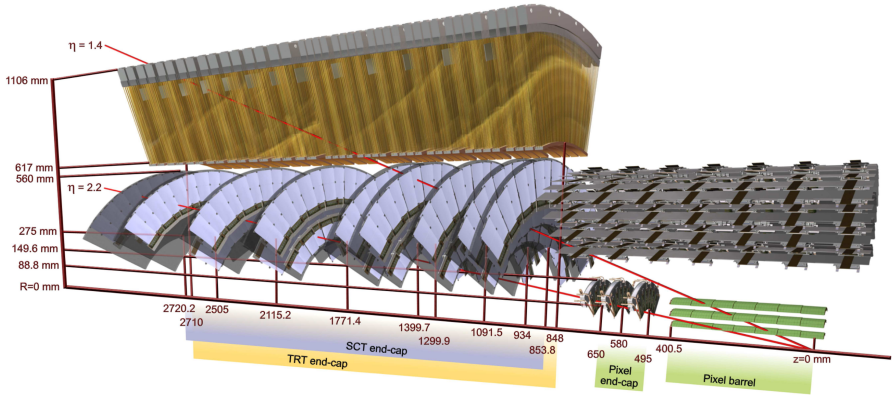

(b)

Figure 4.4: Cut-away view of the barrel [64] (a) and endcap [49] (b) portions of the ATLAS ID, with each of the three subdetectors indicated along with their envelopes in R and z.

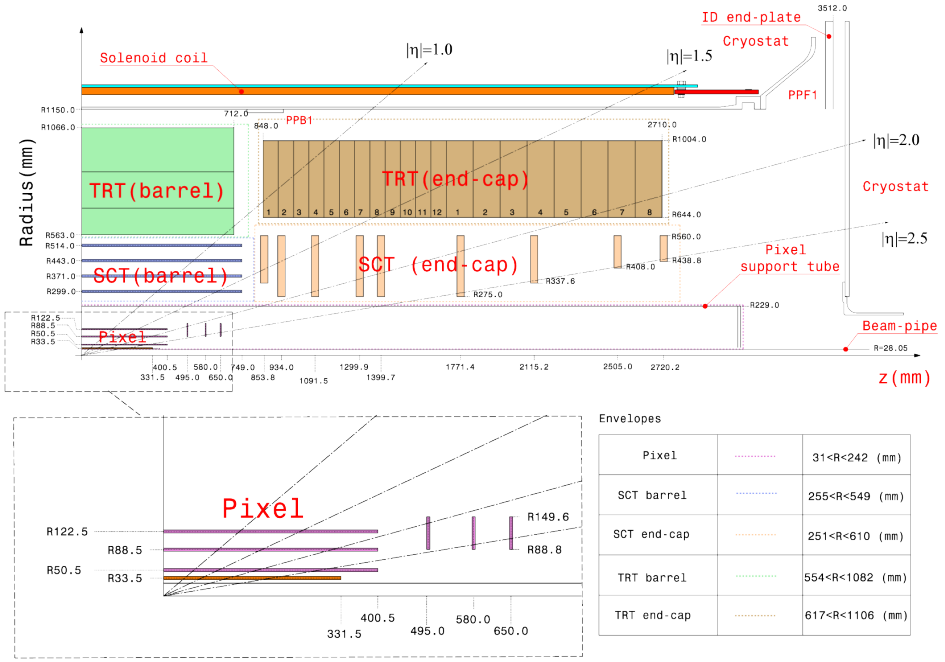

Figure 4.5: Detailed drawing of one quadrant of the ID in the r-z plane [65]. Positions in r, z, and η are shown for all individual layers and are given in mm.

4.3.1 The Pixel detector

Closest to the interaction point, the flux of particles is very high and a high granularity is needed in order to reconstruct individual tracks. This is ensured by the silicon pixel sensors that form the inner layers of the ID. The pixel detector is composed of 2024 silicon pixel modules with a total of approximately 92 million readout channels. In the barrel, the modules are arranged in four coaxial layers placed at increasing radius. These are complemented in each endcap by three disks placed along the longitudinal axis. The design ensures the existence of four measurement points for each track over the full covered η-range. The innermost layer, called the Insertable B-Layer (IBL) [66], was inserted for the LHC Run 2 in order to improve tracking by reducing the distance between the interaction point and the first layer. Each module contains several pixels with a size of $50\,\mu\text{m} \times 250\,\mu\text{m}$ in the IBL and $50\,\mu\text{m} \times 400\,\mu\text{m}$ in the other layers, with the larger side along the z-coordinate.

The pixel sensors are n-type silicon wafers with a depletion voltage applied. When a charged particle passes through a sensor it produces electron-hole pairs inside the material. The electrons then drift in the electric field towards the mounted read-out system. A hit occurs when the resulting current becomes large enough to pass a predefined threshold. The charge deposition is typically shared between adjacent pixels and the relative pulse height allows for a position measurement with a resolution better than the pixel size. This way, a resolution of about $8\,\mu\text{m} \times 40\,\mu\text{m}$ is achieved in the IBL and $10\,\mu\text{m} \times 115\,\mu\text{m}$ in the other three layers.

4.3.2 The Semiconductor Tracker

Further away from the interaction point, the density of tracks is lower and silicon microstrip modules are used instead of pixel modules. The SCT consists of 4088 modules with a total of approximately 6.3 million readout channels, arranged in four barrel layers and nine wheels in each endcap. The barrel modules have a rectangular shape with an area of $64.0\,\text{mm} \times 63.6\,\text{mm}$ and a strip pitch of $80\,\mu\text{m}$, while the endcap modules have a trapezoidal shape with a strip pitch varying between $56.9\,\mu\text{m}$ and $94.2\,\mu\text{m}$. In the barrel, the strips run parallel to the z-axis and in the endcap they run radially. Each module is composed of two layers of sensors glued back to back with a relative stereo angle of 40 mrad. This design allows for a position determination with a precision better than the strip length in the z-coordinate in the barrel and the r-coordinate in the endcap. The modules provide a measurement point with a resolution of $17\,\mu\text{m}$ in the r-ϕ coordinate and $580\,\mu\text{m}$ in the z-coordinate in the barrel and in the radial coordinate in the endcap. Depending on where in η and ϕ a given track goes though the SCT, between four and nine hits are provided.

4.3.3 The Transition Radiation Tracker

The TRT is the outermost of the ID subdetectors and uses a classical technique for detecting charged particles. It is made of approximately 351,000 straws filled with a Xe-based gas mixture and a gold plated tungsten wire as anode in the middle of the tube. When the straws are traversed by a charged particle the gas is ionised and a voltage causes the free electrons to move towards the anode. The acceleration of the electrons causes further ionisation close to the anode and an avalanche of electrons develops which generates a signal on the wire. A position measurement is achieved by converting the electron drift time to the wire to a distance.

The straw tubes are 4 mm in diameter and are organised in modules interleaved with polypropylene fibres in the barrel and foils in the endcaps. The barrel consists of 73 modules with tubes arranged along the z-axis and the endcaps consist of 160 layers with radially arranged tubes. On average, 36 one-dimensional measurement points with a resolution of 130 µm in the r-ϕ coordinate are provided per track.

In addition to the tracking, the TRT can also provide particle identification. Particles passing through the TRT produce transition radiation as they traverse the dielectric material embedding the tubes. The amount of transition radiation is proportional to the Lorentz boost factor $\gamma = E/m$ of the particle and leads to an increased signal. All signals are therefore discriminated against two adjustable low and high thresholds and the amount of high-threshold hits along the track is counted. This provides discrimination power for separating electrons from heavier charged particles.

4.4 Calorimeter system

The basic principle of the calorimeter system is to stop particles produced in the pp interactions and measure the energy that they deposit in the active material. The calorimeter system consists of the electromagnetic calorimeter designed to absorb electrons and photons, the hadronic calorimeter designed to absorb hadrons, and an additional forward calorimeter that extends the coverage at large pseudorapidity. When particles pass through the calorimeter material they create a shower of secondary particles. The showers created by electrons and photons are fully contained within the electromagnetic calorimeter while the hadrons only deposit some energy in the electromagnetic calorimeter and are not fully stopped until they reach the hadronic calorimeter. Both systems are sampling detectors in which absorbers, where the shower of secondary particles is initiated, are interleaved with layers of an active detector material measuring the energy. The calorimeters are separated in barrel, endcap, and forward structures covering a pseudorapidity region up to $|\eta| = 4.9$. Figure 4.6 shows the full calorimeter system. The electromagnetic calorimeter barrel has a length of 6.4 m and stretches from 1.4 m to 2 m in radius,

while the hadronic calorimeter barrel is 5.8 m long with inner and outer radii of 2.3 m and 4.3 m respectively.

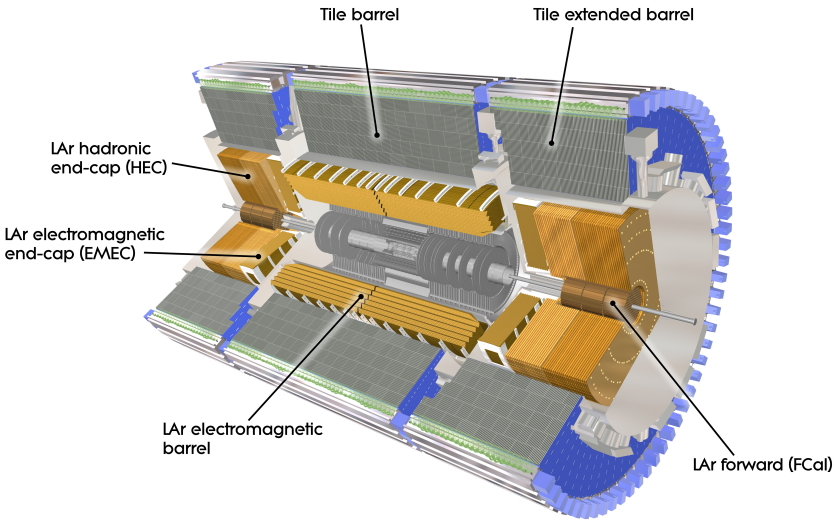

Figure 4.6: Overview of the calorimeter system of the ATLAS detector [49].

4.4.1 The electromagnetic calorimeter

The electromagnetic calorimeter [67] consists of one barrel part (EMB) and two endcaps (EMEC) where Liquid Argon (LAr) is used as the active material and lead as the absorber. A basic unit consists of a LAr filled gap between two parallel lead plates. The main energy loss mechanisms relevant for photons and electrons are bremsstrahlung, pair production, and ionisation. When a photon or an electron traverses the calorimeter it will interact with the material leading to photon conversions into electron-positron pairs and bremsstrahlung where electrons radiate photons. This creates a shower of electrons and photons. The size of the shower is determined by the radiation length X_0 of the material. Both bremsstrahlung and pair production have a cross section proportional to the atomic number Z of the material and the shower therefore develops predominantly in the lead. The electrons in the shower ionise the argon and an electric signal is created by the drift of ionisation electrons towards an anode readout electrode.

The EMB covers up to $|\eta| = 1.475$ and consists of three layers with various thickness. Its first layer is around $4\,X_0$ deep at $|\eta| = 0$ and has a very fine segmentation in η. This provides precise energy and positional measurements that are used for the identification of electrons and photons. The middle layer is $16\,X_0$ deep and

absorbs most of the electromagnetic shower, while the third layer is $2\,X_0$ deep and is designed to collect the tail of the shower. The lead is structured symmetrically around the z-axis in an accordion shape, as shown in Figure 4.7. This design allows for full azimuthal coverage and assures that particles pass through the same amount of material in all ϕ-directions. Figure 4.8 shows the thickness of the material in terms of the radiation lengths before and inside the electromagnetic calorimeter barrel.

Each EMEC module consists of two co-axial wheels. The outermost wheel covers up to $|\eta| = 2.5$ and consists of three layers with similar thickness and design as the barrel layers. This allows for precise energy measurements over the full range covered by the ID. The innermost wheel covers up to $|\eta| = 3.2$ and consists of two layers with coarser granularity. In the region $|\eta| < 1.8$, a so-called presampler

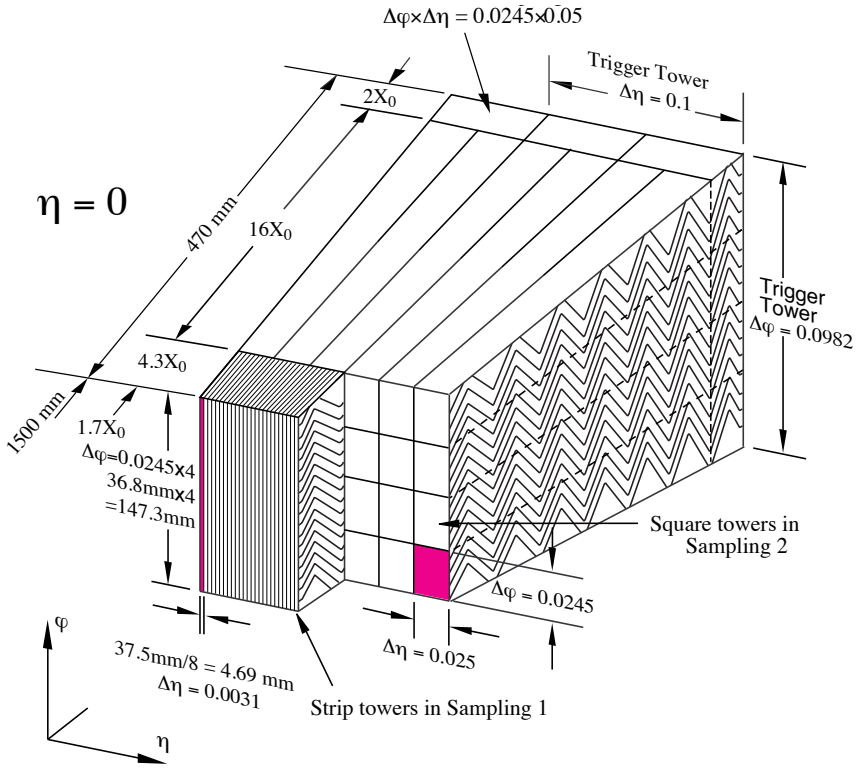

Figure 4.7: Sketch of the cell structure in the different sampling layers of the LAr calorimeter at central η [67].

detector is used to correct for the energy lost by electrons and photons due to material interactions occurring before the calorimeter. It is a single LAr layer, with widths of 1.1 cm and 0.5 cm in the barrel and endcap respectively.

The design energy resolution of the electromagnetic calorimeter depends on the energy and can be described by

$$\frac{\sigma_E}{E} = \frac{10\,\%}{\sqrt{E[\mathrm{GeV}]}} \oplus 0.7\,\%. \tag{4.3}$$

Here, the energy-dependent term results from statistical fluctuations in the electromagnetic showering and the constant term represents the intrinsic detector resolution.

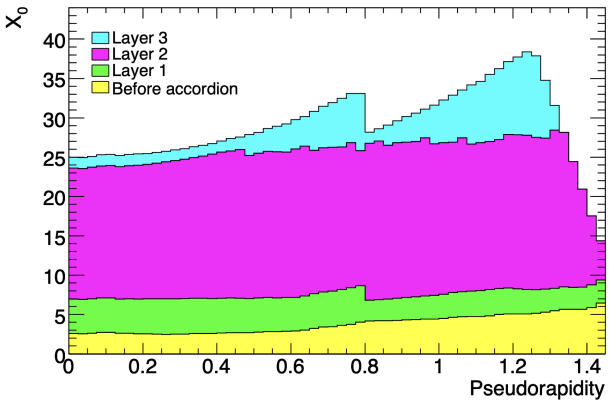

Figure 4.8: Cumulative amount of material, in units of the radiation length X_0, as a function of $|\eta|$, in front of and in the electromagnetic calorimeter barrel [49].

4.4.2 The hadronic calorimeter

The hadronic calorimeter consists of the Tile sampling calorimeter (TileCal) [68] in the barrel and the LAr Hadronic Endcap Calorimeter (HEC) [67]. These use different technologies motivated by the physics performance requirements as a function of the pseudorapidity. The development of the hadronic shower is considerably different from the development of the electromagnetic showers, with strong interactions contributing to the production of secondary particles. The scale of the shower is determined by the nuclear absorption length λ which is larger than the radiation length, explaining why more material is needed to contain the shower. Figure 4.9 shows the thickness of the material in terms of nuclear absorption lengths for the entire detector.

The TileCal consists of a central barrel structure and two extended barrels that together cover up to $|\eta| = 1.7$. It is composed of a steel absorber interleaved with scintillating plastic tiles as active material. Photomultiplier tubes (PMTs) measure

the light emitted from the scintillators and produce the readout signal. The total depth corresponds to 7.4 λ in both the central and extended barrels. The design energy resolution of the TileCal can be described by

$$\frac{\sigma_E}{E} = \frac{50\,\%}{\sqrt{E[\mathrm{GeV}]}} \oplus 3\,\%. \tag{4.4}$$

The HEC is made of two independent wheels covering the region $1.5 < |\eta| < 3.2$. These consist of absorbing copper plates interleaved with LAr as active material.

4.4.3 The Forward Calorimeter

The Forward Calorimeter (FCal) [67] serves as an extension of the electromagnetic and hadronic calorimeters into the region $3.2 < |\eta| < 4.9$. It is composed of two endcaps made of three modules each. The first uses copper as absorber and is optimised for precise electromagnetic energy deposit measurements, while the other two use tungsten and are optimised for hadronic interactions. All three modules use LAr as active material and the total depth of each endcap is about 10 λ. The design energy resolution of the FCal is given by

$$\frac{\sigma_E}{E} = \frac{100\,\%}{\sqrt{E[\mathrm{GeV}]}} \oplus 10\,\%. \tag{4.5}$$

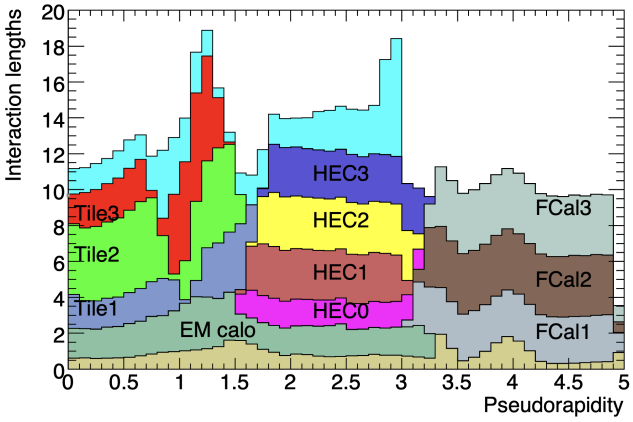

Figure 4.9: Cumulative amount of material, in units of the nuclear absoption length λ, in front of and in the hadronic calorimeter [49]. From the bottom is the ID (unlabelled), followed by the electromagnetic and hadronic calorimeters, and the material in front of the first active layer of the muon spectrometer (unlabelled).

4.5 The Muon Spectrometer

Muons penetrate very deep into matter, and in ATLAS they typically cross both the ID and the calorimeter system with only a minimal energy loss. This is due to their high mass, which makes them less susceptible to emit bremsstrahlung than electrons. At the same time they do not interact strongly, meaning that no hadronic interaction with the detector material takes place. The Muon Spectrometer [69] is therefore the outermost subdetector in ATLAS, reaching from 5 m to 11 m in radius, and is dedicated to the measurement of charged particle tracks. The system is shown in Figure 4.10 and uses chambers with different types of detectors organised in three layers in the barrel and three layers in the endcaps, providing three measurement points across the full covered η-range. For a muon reconstructed only with information from the muon system the design p_{T} resolution is 10 % at 1 TeV.

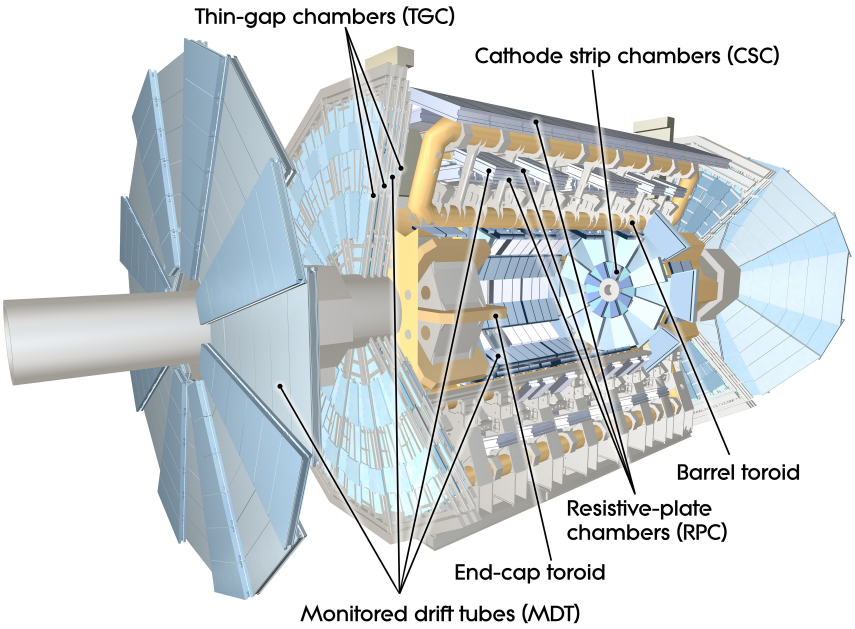

Figure 4.10: Overview of the muon system of the ATLAS detector [49].

The majority of the muon system consists of Monitored Drift Tube (MDT) ionisation chambers that cover the range up to $|\eta| = 2.7$ and are responsible for the precision tracking. These consist of between three to eight layers of aluminium tubes filled with an Ar-based gas mixture. The chambers provide a two-dimensional measurement in the bending plane of the muons with an average resolution of 35 µm. At large pseudorapidity, between $2.0 < |\eta| < 2.7$ where the flux of particles is higher, the MDTs in the innermost layer are replaced by Cathode Strip Chambers (CSCs).

In these chambers, closely spaced anode wires are stretched between two cathodes and the ionisation electrons drift towards the closest wire. The CSC chambers provide a three-dimensional measurement point with a resolution of 40 µm in the bending plane and 5 mm in the azimuthal plane.

In order to get a fast signal for triggering, Resistive Plate Chambers (RPCs) are installed in the barrel up to $|\eta| = 1.05$ and Thin Gap Chambers (TGCs) are installed in the endcaps between $1.05 < |\eta| < 2.4$. The RPCs are parallel-plate gaseous detectors that combine an adequate position resolution with a very fast response time. The TGCs are based on a multi-wire technology and have a small cathode-anode distance in order to give a very fast timing information. Both systems also provide a three-dimensional measurement point used to complement the information from the MDTs in the azimuthal plane.

4.6 Luminosity detectors

ATLAS contains two subdetectors in the forward region which are used for the measurement of the luminosity; LUCID [70] and the Beam Conditions Monitor (BCM) [71]. These are designed to have a sufficient time resolution to measure the number of interactions per bunch crossing and to have a high acceptance for inelastic *pp* collisions.

The LUCID detector is the only ATLAS subdetector that is dedicated entirely to the luminosity measurement. It is a Cherenkov detector consisting of two modules positioned symmetrically around the beam pipe at about ±17 m from the ATLAS interaction point, at a pseudorapidity $|\eta| = 5.6$. Each module consists of 16 PMTs surrounding the beam pipe, with thin quartz windows which are used as active medium. As charged particles produced from the proton collisions pass through the quartz, Cherenkov light is emitted and is read out by the PMTs. A hit is registered if the signal is above a pre-defined threshold. The LUCID detector is read out with dedicated electronics which provide hit counts for every colliding bunch. To monitor and, if needed, adjust the gain of the photomultipliers, radioactive bismuth-207 (Bi-207) sources are mounted directly in front of the PMT quartz-windows for a subset of the PMTs. The amplitude of the signal resulting from monoenergetic internal conversion electrons serves as reference for potential gain adjustments.

The BCM detector is comprised of two groups of four diamond sensors, arranged in a cross pattern on either side of the interaction point, located at $z = \pm 184$ cm and $R = 5.5$ cm. It is primarily used to monitor particle fluxes near the beam pipe in order to protect the innermost subsystems of ATLAS and trigger a beam abort in case of excessive radiation levels. The detector has fast readout electronics which also allows it to be used as a luminosity detector.

4.7 Trigger and data acquisition

With a collision frequency of 40 MHz and around 30 interactions per bunch crossing, huge amounts of data are produced every second at the LHC. A bunch crossing with all its associated processes that is captured and recorded by ATLAS is called an event and represents around 1.5 MB of raw data. It is impossible to fully process and store all this information at the rate dictated by the collision frequency and for every event a fast decision must be taken, determining whether or not it should be kept for analysis. One of the most important aspects of ATLAS is therefore the trigger and data acquisition systems, which make the online event selection and handle the readout of the data for the selected events.

In ATLAS, the trigger system is composed of two steps, the Level-1 Trigger (L1) and the High-Level Trigger (HLT). The data from each crossing is held in buffers while the L1 uses low-granularity information from dedicated subsystems in the calorimeters and the muon system to search for high-p_{T} physics objects and events with large momentum imbalance in the transverse plane. This first step is completed in about 2.5 µs, and reduces the event rate down to 100 kHz. The L1 calculations are done on custom-built hardware placed underground next to the detector cavern.

If an event is selected by the L1, so-called Regions of Interest (ROIs) in η and ϕ that contain the interesting activity are passed on to the HLT. The reduced event rate allows for a more detailed analysis and more complex calculations similar to the event reconstruction described in Chapter 5. A software-based system performs a partial event reconstruction on the available data in the ROI with full granularity and with information from the ID included. This step reduces the event rate further down to 3.5 kHz with a latency of 40 ms. The final step in the HLT employs an event filter with simplified versions of the software used for offline event reconstruction and runs on ordinary computers. In this step the event rate is further reduced to approximately 1 kHz. Events that are selected are sent for full reconstruction at the CERN on-site computing center called Tier-0. The reconstructed data is then distributed to computing centres around the world and is available for analysis.

The decisions taken by the trigger chain are steered by a predefined trigger menu that specifies all types of events to be kept for analysis. Event types that occur very frequently, such that it would require too much of the total trigger bandwidth to record all events passing a given threshold, can be prescaled so that they are recorded at a fraction of the actual rate. When these events are used for analysis they are weighted by the prescaling rate. Since this is associated with an effective loss of luminosity it is however desirable for analyses to use unprescaled triggers when possible.

The ATLAS data taking is divided in runs, which typically correspond to one LHC fill. Each run is split into Luminosity Blocks (LBs), which are small sections of time where the instantaneous luminosity is approximated to be constant. Data events within a LB are recorded with the same triggers and under the same detector

conditions. A typical LB length in ATLAS is 60 s. The integrated luminosity then corresponds to the sum

$$L = \sum_{\text{LBs}} \mathcal{L}_{\text{LB}} \cdot \Delta t_{\text{LB}} \tag{4.6}$$

where $\mathcal{L}_{\text{LB}}$ is the constant luminosity inside the LB and Δt_{LB} is the duration of the LB, and the sum runs over all the LBs in the considered time window. The data used for analysis is subject to offline quality criteria which require all reconstructed physics objects to be of good data quality. LBs where essential detector components were malfunctioning during data taking are therefore not used. After the quality requirements, 139 fb^{-1} are available for data analysis in Run 2.

Chapter 5

Physics building blocks

Particles produced in the pp collisions in the LHC travel outwards through the ATLAS detector and leave different signatures in the various subsystems. When an event is triggered, the electronic channels of the subdetectors are read out, and the corresponding digitised signals make up the raw data from the bunch crossing. Before any physics analysis can be performed on the event, the data needs to be translated into well-defined representations of the underlying physical objects. A variety of different algorithms is deployed for this so-called *event reconstruction*, which all attempt to collect signals that belong to the same particle and use those to reconstruct its kinematic properties.

The event reconstruction starts with the formation of low-level objects which are then used as input to identify the final high-level physics objects. Low-level objects include clusters constructed from energy deposits in the calorimeters, and tracks and vertices built from measurements in the ID. The high-level physics objects used for the work described in this thesis include electrons, photons, muons, and jets. ATLAS is designed to produce a unique signature for each of these particles by combining tracking and calorimeter information. In addition, the missing transverse momentum $E_{\mathrm{T}}^{\mathrm{miss}}$, which quantifies the momentum imbalance in the transverse plane, is used to infer the presence of particles that pass through the detector without interacting.

The following sections describe the ATLAS tracking and vertexing procedures, the identification and reconstruction of electrons, photons, muons, and jets, as well as the definition and computation of the missing transverse momentum. The track-counting luminosity measurement and the DV+jets analysis, described in this thesis, are heavily dependent on the reconstruction of charged particle tracks in the ID. More emphasis is therefore dedicated to the ATLAS tracking and vertexing procedures.

5.1 Inner Detector tracking and vertexing

The reconstruction of charged particle tracks is one of the most complex parts of the ATLAS event reconstruction. In a pp collision at $\sqrt{s} = 13\,\text{TeV}$, on average 15 charged particles are expected to be created [72]. Combined with the Run 2 average of 34 simultaneous pp interactions, this means that around 500 charged particles will be produced in each bunch crossing. All the corresponding individual tracks must be reconstructed with a high efficiency and with a low rate of fake tracks resulting from spurious combinations of hits in the ID. In addition, the corresponding pp interaction points must be determined to high precision in order to be able to separate the hard-scatter event of interest for physics analysis from the multiple pileup interactions. This is done by clustering tracks together to find their common point of origin in the vertex reconstruction.

The basic idea of the tracking is to fit the hits in the ID to a parametrisation of the trajectory followed by a charged particle in the field from the solenoid magnet. For a homogeneous magnetic field, and in the absence of any material, the path is described by a helix with its axis parallel to the direction of the field. In general, five parameters are needed to describe the helical track. The representation used by ATLAS is given by

$$\tau = (d_0, z_0, \phi_0, \theta_0, q/p)\,, \tag{5.1}$$

where all parameters are expressed at the so-called *perigee*, corresponding to the point of closest approach to a chosen reference point. The impact parameters, d_0 and z_0, are the projections of the distance of closest approach to the reference point in the transverse plane and along the z-axis, respectively. The angles ϕ_0 and θ_0 are the azimuthal and polar angles of the track at the perigee, and the curvature is described by the ratio of the particle charge and momentum q/p. Figure 5.1 shows a graphical view of the track parameters. In practice, the helical path is slightly altered by inhomogeneities in the magnetic field as well as interactions between the charged particles and the material in ATLAS, which result in energy losses as well as multiple scattering. All these effects must be taken into account by the tracking algorithms.

In the following, Sections 5.1.1–5.1.4 describe the standard tracking and vertexing procedures used in ATLAS, and provide a discussion of the parameters representing the tracks. The standard tracking is optimised for particles originating from the pp interaction point and is not optimal for the reconstruction of massive LLPs which decay with a significant displacement and result in tracks that do not point back to the pp interaction point. Such particles are predicted by a wide range of BSM models, as described in Section 2.3. In analyses searching for LLPs, a special event reconstruction is therefore applied with an extra Large Radius Tracking (LRT) and a secondary vertexing algorithm. This special reconstruction is presented in Sections 5.1.5–5.1.6.

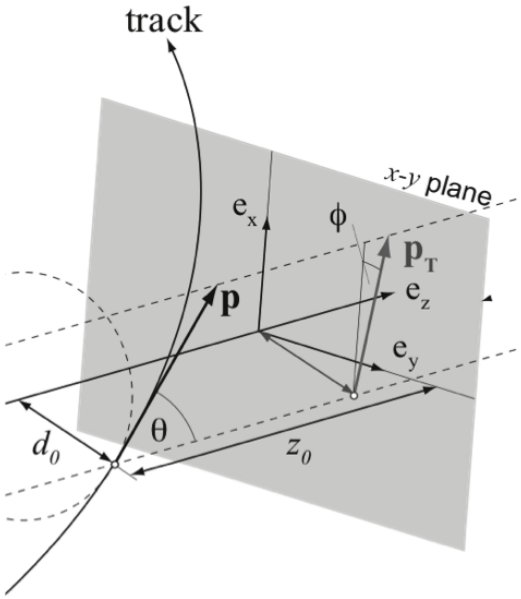

Figure 5.1: Illustration of the track parameters at the perigee [73]. In this scenario, the reference point corresponds to the detector origin.

5.1.1 Standard track reconstruction

In ATLAS, the tracking is performed first *inside-out* and then *outside-in* with respect to the components of the ID [74]. The inside-out tracking starts from track segments in the silicon detectors and extrapolate into the TRT. This is optimised for the reconstruction of so-called *primary tracks* originating from the *pp* interactions. The outside-in tracking starts in the TRT and extrapolates into the SCT and the Pixel detector. It is employed to increase the acceptance to so-called *secondary particles* produced at a greater distance from the interaction point, such as electrons originating from photon conversions in the detector material and the decays of *b*-hadrons. Both procedures however assume that the tracks point back to the *pp* interaction point.

Inside-out tracking

The inside-out tracking procedure can be divided into various steps which will be described separately below.

Spacepoint creation: In the first step, all hits in the silicon detectors are translated into three-dimensional spacepoints by clustering together neighbouring pixels and strips that have recorded a signal over threshold. The spacepoints represent the intersection of charged particles crossing the detector layers, and also account for the uncertainty of the measured hit positions. For the IBL and the Pixel detector, the spacepoints correspond directly to the hits in the detector. For the SCT, the single-sided hits on the back-to-back readout strips on each layer must be combined using the stereo-angle information to give the three-dimensional spacepoint information. Detector elements which are known to be malfunctioning and not sensitive at the time of data taking are counted as providing a spacepoint, in order to not lose tracking efficiency.

Track seeding: The track reconstruction algorithm starts with the identification of so-called seeds corresponding to triplets of hit in separate silicon layers, which are compatible with originating from a charged particle track. An initial set of track parameters for each seed are estimated assuming a perfect helical trajectory and using the detector origin as reference point. These are then used to define search windows in which to look for additional hits to be added to the seeds.

Track candidate building: Starting from the seeds, a set of track candidates are constructed by adding hits within the search windows using a combinatorial Kalman filter [75, 76]. This is a recursive procedure where the track parameters are propagated through successive detector layers taking into account the magnetic field map and the material in the detector. The energy loss of the particle is estimated by assuming the mass of the pion. Hits compatible with the predicted trajectory in each layer are added to the track candidate and the track parameters are updated with the new information. If several hits on the same layer are compatible, multiple track

candidates are formed and proceed independently through the remaining layers. The final track parameters are determined by a global least-squares fit of the track prediction to all hits associated to the track candidate. This final fit also returns the track covariance matrix which describes the uncertainties on the track parameters and their correlations.

Ambiguity solving: The seeded track finding results in a very high number of track candidates, many of which share hits, are incomplete, or describe fake tracks. This is solved by a scoring algorithm which ranks the tracks in their likelihood to describe the real trajectories of charged particles. The scoring scheme favours high-p_{T} tracks and applies positive scores for unshared hits and good fit quality in terms of the chi-squared per degree of freedom χ^2/n_{dof}, and negative scores for holes or shared hits. Holes are defined as intersections of the reconstructed track trajectory with a sensitive detector element that does not contain a matching hit. Shared measurements are defined as hits that are assigned to multiple track candidates. At the end, only tracks that pass a threshold for the scoring scheme will remain.

Track extension: Tracks that are kept are extended into the TRT with a new Kalman fit. Similar to the creation of track candidates, a road search is performed into the TRT volume starting from the position estimation of the track candidate. Compatible TRT hits are added to the tracks which are eventually refit with a global least-squares fit if the extension is successful.

Outside-in tracking

The final step in the tracking uses any left-over hits from the inside-out pass and builds on standalone TRT track segments in regions seeded by energy deposits in the electromagnetic calorimeter. These seeds are extended into the SCT and Pixel detector, building track candidates with the same procedure as for the inside-out tracking, i.e using a combinatorial Kalman filter, an ambiguity solver, and a final global least-squares fit.

5.1.2 Primary vertex reconstruction

The tracks built with the above procedure are used to reconstruct the points in space where the pp interactions took place. These points are called primary vertices and are effectively found by looking for points where at least two tracks intersect close to the beam axis using an iterative procedure [77]. For tracks to be considered in the construction of a vertex they must pass a set of additional quality requirements on the hit multiplicities in the various ID layers [78]. The procedure of primary vertex reconstruction then consists of the following steps:

- A seed position for the first vertex is selected. Initially this point corresponds to the detector origin in x and y and to the global maximum of the z_0-distribution of the tracks in z.

- The tracks and the seed are used to fit the best vertex position with a least-squares based adaptive vertex fitter [79]. The fit is an iterative procedure, and in each iteration less compatible tracks are down-weighted and the vertex position is recomputed.
- After the vertex position is determined, tracks that are incompatible with the vertex are removed and allowed to be used in the determination of other vertices.
- The procedure is repeated until no tracks are left in the event or no additional vertex can be found in the remaining set of tracks.

The output of the vertex reconstruction algorithm is a set of three dimensional vertex positions and their covariance matrices. The primary vertex for which the sum of the p_{T} of the associated tracks is the highest is often referred to as the Primary Vertex (PV) of the event. Other vertices are associated to pileup interactions. The reconstruction of primary vertices is therefore essential for the separation of the hard interaction of interest for analysis from additional pileup interactions.

Once the primary vertices are reconstructed they are used to detemine the so-called *beamspot* which corresponds to the region in the detector where the *pp* interactions between the two beams take place. The beamspot reconstruction is based on an unbinned maximum-likelihood fit to the spatial distribution of primary vertices collected from many events [80]. This is carried out during data taking, typically every two minutes [81]. The centre of the beamspot is then used as an additional measurement during the primary vertex fit, to help constrain the transverse position. In addition, all tracks are recomputed and their parameters are expressed with the beamspot as the reference point. Figure 5.2 shows the size of the beamspot during the LHC Run 2. In the transverse plane, it is on the order of 10 µm and in the longitudinal plane it is on the order of 40 mm. The transverse size is determined by the focusing of the LHC beams near the interaction region and by the geometric emittance of the beams, while the longitudinal size is determined by the bunch length and by the angle at which the two beams are brought into collision.

5.1.3 Offline track selection

In the track reconstruction, the design of the quality criteria applied to the tracks is a trade-off between achieving a high tracking efficiency and reducing the rate of fake tracks. Requirements designed for a strong rejection of fake tracks are referred to as *tight*, while those designed to achive a higher tracking efficiency at the cost of a larger fake-track rate are referred to as *loose*. The criteria applied in the ambiguity solver are purposely left loose and for most analyses the selection of tracks needs to be tightened offline. The standard ATLAS track selection for analyses is called *TightPrimary* [82]. Table 5.1 shows a list of TightPrimary track selection criteria together with the requirements in the track reconstruction, referred to as *Loose*. In order to put the hit-multiplicity requirements into perspective it can be noted that

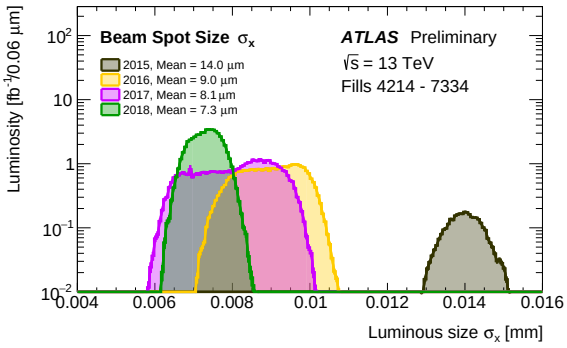

(a)

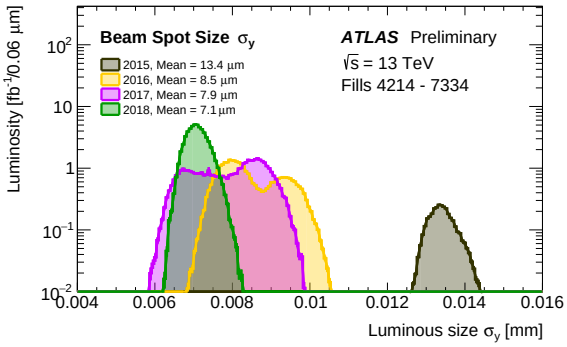

(b)

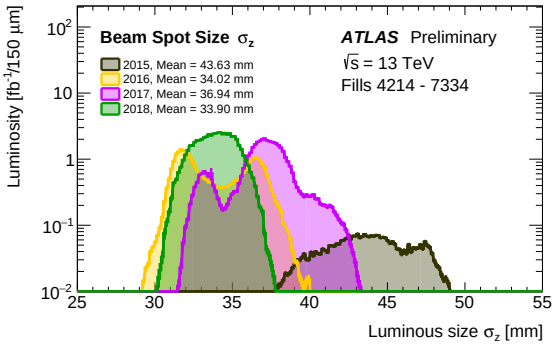

(c)

Figure 5.2: The luminosity weighted distribution of the size in x (a), y (b), and z (c) of the beamspot in ATLAS in each data-taking year in Run 2 [81]. The data points are the result of a maximum likelihood fit to the spatial distribution of primary vertices collected over a two minute period.

Table 5.1: List of criteria applied in the online track reconstruction (Loose) and for the TightPrimary selection. $N_{\text{hits}}^{\text{Pix}}$ and $N_{\text{hits}}^{\text{SCT}}$ denote the total number of hits in the Pixel and single-sided hits in the SCT detectors, respectively, and include the number of inactive sensors. $N_{\text{hits}}^{\text{IBL}}$ and $N_{\text{hits}}^{\text{B-layer}}$ denote the number of hits in the IBL and first layer of the Pixel detector, respectively. $N_{\text{holes}}^{\text{Pix}}$ and $N_{\text{holes}}^{\text{SCT}}$ denote the number of holes in the Pixel and SCT detectors, respectively. $N_{\text{shared}}^{\text{Pix}}$ and $N_{\text{shared}}^{\text{SCT}}$ denote the number of shared hits in the Pixel and shared single-sided hits in the SCT detectors, respectively.

Criterion	**Loose**	**TightPrimary**
p_{T}	500 MeV	500 MeV
$\lvert\eta\rvert$	2.5	2.5
$N_{\text{hits}}^{\text{Si}} = N_{\text{hits}}^{\text{Pix}} + N_{\text{hits}}^{\text{SCT}}$	≥ 7	≥ 9 if $\lvert\eta\rvert \leq 1.65$, ≥ 11 if $\lvert\eta\rvert > 1.65$
$N_{\text{hits}}^{\text{IBL}} + N_{\text{hits}}^{\text{B-layer}}$	-	> 0 if hits are expected both layers
$N_{\text{holes}}^{\text{SCT}}$	≤ 2	≤ 2
$N_{\text{holes}}^{\text{Pix}}$	≤ 1	0
$N_{\text{shared}}^{\text{Pix}} + \frac{1}{2}N_{\text{shared}}^{\text{SCT}}$	≤ 1	≤ 1

a track on average consists of one IBL hit, three Pixel hits, eight single-sided hits in the SCT and about 34 hits in the TRT.

5.1.4 Track parameters

The result of the standard tracking and vertexing is the final collection of reconstructed standard tracks with parameters expressed with the beamspot as reference point. The impact parameters define the origin of the track, the angles ϕ and θ its direction at the origin, and the ratio q/p together with the angle θ gives the p_{T} measurement. Figure 5.3 shows distributions of the track parameters for the TightPrimary track selection derived from simulated minimum bias events[1].

The track impact parameters are important quantities for constraining and reducing background processes and to distinguish between primary and secondary tracks. When applying requirements on these parameters it is often necessary to take the associated uncertainty into account. Figure 5.4 shows the d_0 and z_0 uncertainties as a function of the pseudorapidity and p_{T} of the track, as derived from minimum bias simulation. The uncertainty increases at large η where a track passes through more

[1]Minimum bias refers to events collected or simulated without any particular restriction on the particles in the final state and include all possible hard-scatter processes at the given collision energy. Simulation of minimum bias events is discussed in Section 6.4

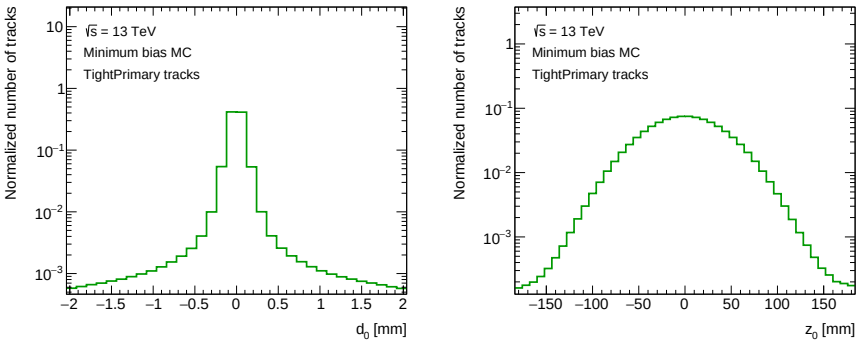

(a)

(b)

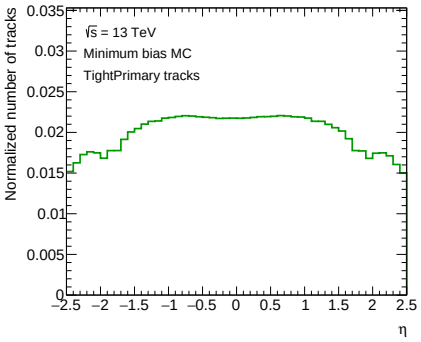

(c)

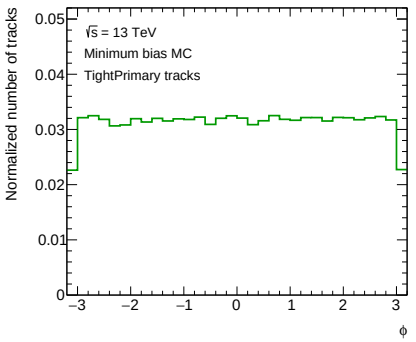

(d)

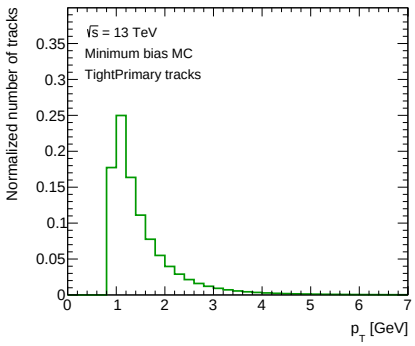

(e)

Figure 5.3: Normalised distributions of d_0 (a), z_0 (b), η (c), ϕ (d), and p_{T} (e) for tracks selected with the TightPrimary requirements, with the addition of a stricter p_{T} requirement at 900 MeV, derived from minimum bias simulation. The minimum bias events include all possible hard-scatter processes at the given collision energy. See Section 6.4 for further description of minimum bias samples.

material and thereby undergoes more multiple scattering. Likewise, it increases at small p_{T} where the effect of the multiple scattering is larger. Figure 5.5 shows the uncertainty as a function of p_{T} only, where the distribution measured in simulation is compared to data. Such comparisons are very important for analyses which heavily rely on the impact parameters and are used to derive smearing factors that can be applied to the simulation to take it closer to the data. In general, it is not possible to measure the impact parameter resolution directly using the convariance matrix of the track, because it is convolved with the resolution on the reference point. The results shown in Figures 5.4 and 5.5 are therefore derived with an iterative unfolding method to separate the resolution on the reference point from the measurement of the impact parameter resolution itself [83].

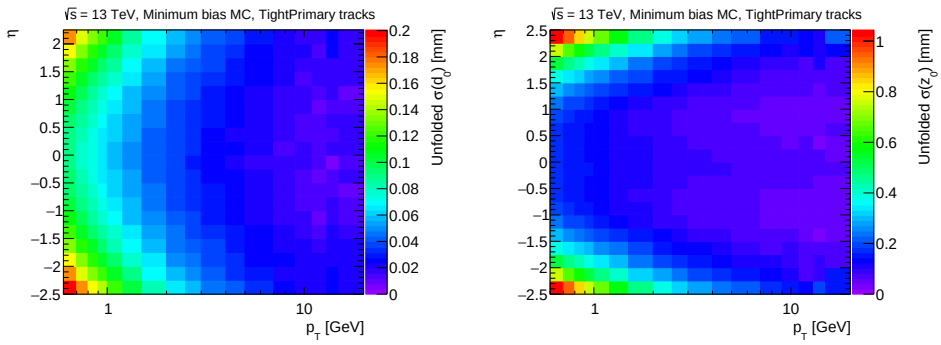

(a) (b)

Figure 5.4: Unfolded d_0 (a) and z_0 (b) resolution, as a function of the track pseudorapidity and p_{T}, derived from minimum bias simulation. See Section 6.4 for a description of minimum bias samples.

Since the uncertainties have a large dependence on the kinematics of the tracks, requirements are often placed on the impact parameter significances when requiring a track to come from the pp interaction point. These are defined as

$$d_0\text{-significance} = d_0/\sigma_{d_0} \tag{5.2}$$

$$z_0\text{-significance} = z_0/\sigma_{z_0}, \tag{5.3}$$

and ensure that tracks measured with higher uncertainty are not punished harder by the requirements. In addition, for the track-counting luminosity measurement, which depends on measurements of all primary tracks, the impact parameter resolution is corrected for the beamspot width in x- and y-direction through an increase of the uncertainty on d_0 according to

$$\sigma_{d_0,\mathrm{corr}} = \sqrt{\sigma_{d_0}^2 + \sigma_{bs}^2}, \tag{5.4}$$

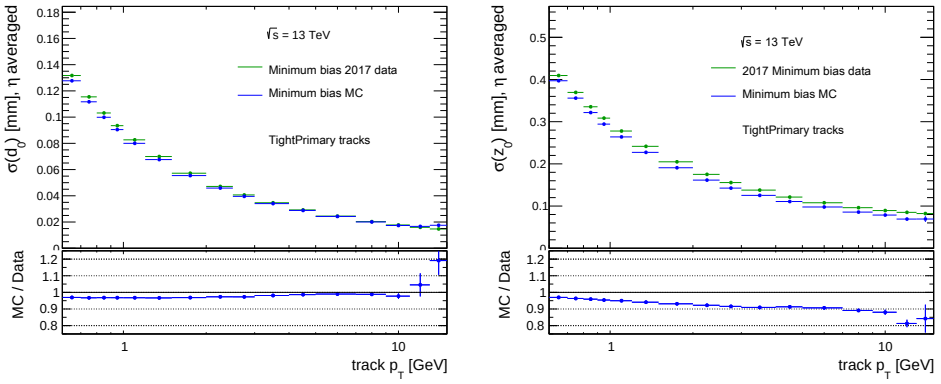

(a) (b)

Figure 5.5: Unfolded d_0 (a) and z_0 (b) resolution, as a function of the track p_{T}, derived from minimum bias simulation and data collected in 2017, using the TightPrimary track selection. See Section 6.4 for a description of minimum bias samples.

with the beamspot uncertainty given by

$$\begin{aligned}\sigma_{bs} = &\sin(\phi_{\mathrm{trk}}) \left(\sin(\phi_{\mathrm{trk}}) \cdot \sigma^2_{bs,x} - \cos(\phi_{\mathrm{trk}}) \cdot \sigma_{bs,xy}\right) + \\ &\cos(\phi_{\mathrm{trk}}) \left(\cos(\phi_{\mathrm{trk}}) \cdot \sigma^2_{bs,y} - \sin(\phi_{\mathrm{trk}}) \cdot \sigma_{bs,xy}\right) \ .\end{aligned} \tag{5.5}$$

Here, $\sigma_{bs,x}$ and $\sigma_{bs,y}$ are the widths of the beamspot in the x and y directions, $\sigma_{bs,xy}$ is their correlation, and ϕ_{trk} is the azimuthal angle of the track at the perigee with respect to the beamspot position. The d_0-significance computed with $\sigma_{d_0,\mathrm{corr}}$ allows for a compatibility check of the origin of a given track and the area corresponding the beamspot in the transverse plane. The corrected d_0-significance distribution is shown in Figure 5.6.

5.1.5 Large Radius Tracking

For analyses targeting LLPs, a special event reconstruction is employed which includes LRT [84]. This runs after the standard tracking and makes use of hits that are leftover to seed and reconstruct large radius tracks using an inside-out approach. Compared to the standard tracking algorithm, LRT uses a slightly extended region to search for hits for seeding and less strict requirements on the track impact parameters and the hit multiplicities. The strategy used to extend seeds into track candidates is also different from the standard reconstruction and prevents the creation of multiple track candidates from the different possible combinations of points by using a sequential Kalman filter approach instead of the combinatorial Kalman

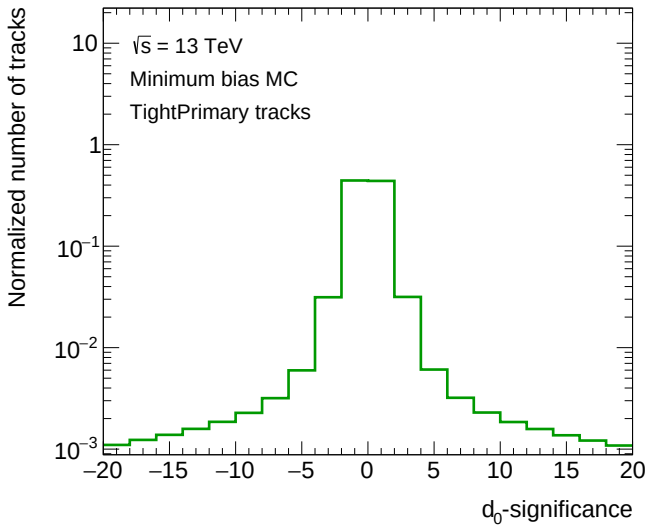

Figure 5.6: Normalised d_0-significance distribution for tracks selected with the TightPrimary requirements, with the addition of a stricter p_{T} requirement at 900 MeV, derived from minimum bias simulation, using the TightPrimary track selection. The d_0-significance is computed including the correction for the size of the beamspot. See Section 6.4 for a description of minimum bias samples.

Table 5.2: Comparison of some of the key track reconstruction parameters for the standard tracking and LRT. Table adapted from Ref. [84].

Selection	**Standard tracking**	**LRT**
Minimum p_{T} (MeV)	500	900
Maximum $\lvert d_0 \rvert$ (mm)	10	300
Maximum $\lvert z_0 \rvert$ (mm)	250	1500
Maximum $\lvert \eta \rvert$	2.7	5
Maximum shared silicon modules	1	2
Minimum unshared silicon hits	6	5
Minimum silicon hits	7	7
Seed extension	Combinatorial	Sequential

filter. Once the track candidates are created, ambiguity resolution and TRT extension proceed as with standard tracking but with relaxed constraints on the track candidates. A summary of the main differences between standard tracking and LRT is given in Table 5.2. Figure 5.7 shows the track reconstruction efficiency for standard tracking, LRT, and a combination thereof, as a function of the radius of production of displaced particles. At small radii, the standard tracking is very efficient and few hits are left over to form tracks in the LRT step. Tracks which are produced far from the interaction point are unlikely to be reconstructed by the standard tracking and the addition of LRT increases the reconstruction efficiency substantially beyond radii of 50 mm.

At the Run 2 pileup conditions, running LRT increases the time taken to reconstruct a data event by a factor of around 2.5. It is therefore applied only to a few percent of the data used for physics analysis, filtered out based on quantities computed during the standard reconstruction. The filter allows for a separate storage of the selected raw data events for further processing with less stringent requirements on the computational resources.

5.1.6 Secondary vertex reconstruction

In addition to the primary vertexing, various secondary vertexing algorithms exist in ATLAS, which all target different types of decay topologies. One of those is designed to search for LLPs [85] and is employed by the DV+jets analysis. The algorithm uses tracks from the standard and LRT configurations, and imposes no specific restrictions on the track orientation, which means that a wide phase space of outgoing track combinations can be reconstructed. The steps for secondary vertex formation are quite different from primary vertexing and include several steps which are described below.

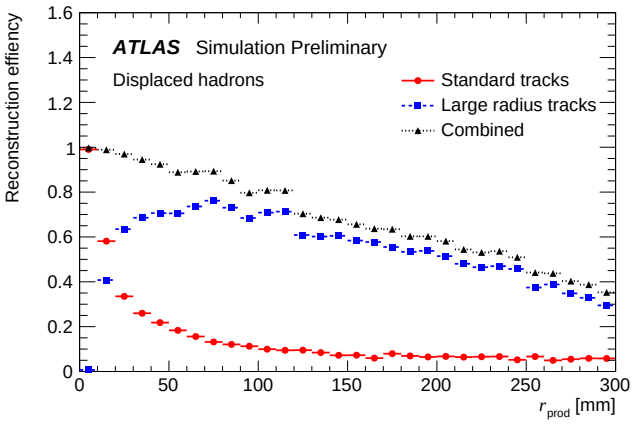

Figure 5.7: Track reconstruction efficiency as a function of the radius of production of displaced particles, derived for decays of long-lived gluinos with a mass of 1.4 TeV [84].

Seed track selection: Due to the high multiplicity of tracks in the typical Run 2 LHC collisions, it is not feasible to consider all standard and LRT tracks for secondary vertexing. Instead, tracks are preselected for vertex seeding with a series of quality criteria on the hit multiplicities, and are required to fulfil $p_{\mathrm{T}} > 1\,\mathrm{GeV}$, and to not be associated to any primary vertex in the event. These tracks are referred to as *selected tracks*.

Two-track seed finding: The vertex finding starts by considering all possible pairs of selected tracks and assessing their compatibility with a least-squares vertex fit [79]. Pairs corresponding to a fit with $\chi^2/n_{\mathrm{dof}} < 5$ are retained as seed vertices and are subject to additional requirements in order to ensure that the vertex position and the tracks are consistent with an LLP decay. Firstly, at least one track in each two-track vertex must have $d_0 > 2\,\mathrm{mm}$. Secondly, the two-track vertices are filtered based on the hit pattern of the associated tracks. Since tracks typically travel outwards from the vertex, the tracks are required to have no hits in tracker layers with smaller radii than the vertex position and to have a hit in the first tracking layer beyond the vertex. All track pairs passing these requirements are labelled as *compatible*.

Multi-track vertex formation: The vertexing algorithm continues by merging two-track vertex seeds to create multi-track vertices. This is done with a graph theory technique [85] which decides which tracks should be clustered together based on the compatibility of all selected track pairs. The extracted group of tracks are simultaneously fitted to a single multi-track vertex. A single track can still be associated to different vertices, and in that case the track is included in the best-fitting vertex.

Vertex merging: After the first multi-track vertex formation, an additional merging step is necessary to recombine vertices that are artificially split by the graph theory technique. Vertices are considered for merging if the two positions are within 10σ, where σ is the uncertainty on the difference between the two vertex positions. A set of tests, involving the extrapolation of tracks from one vertex to the other, are then used to determine if the vertices should be merged or not. Finally, any remaining vertices within 1 mm are forced to merge, and the merged vertices are refit with all tracks.

Track attachment: The last step of the vertexing algorithm is to attach tracks satisfying looser quality requirements, in order to recover efficiency lost by the tight requirements on the selected tracks. In this step, the impact parameters with respect to the secondary vertices are calculated for all unused tracks satisfying $p_{\mathrm{T}} > 1\,\mathrm{GeV}$, including tracks which are already associated to primary vertices. Tracks with transverse and longitudinal impact parameter significances with respect to the target secondary vertex less than 5 are associated to the vertex. Compared to the initial vertex formation, a relaxed hit-pattern consistency check is employed which only requires the presence of hits in the outer layers in order to accept potential back-scattered tracks. The tracks added to the vertices in this step are referred to as *attached tracks*.

Figure 5.8 shows the efficiency of the secondary vertexing as a function of the vertex radius for three signal models predicting LLP decays inside the ID. The efficiency steeply falls off below radii of 2 mm due to the requirement of $|d_0| > 2\,\mathrm{mm}$ applied to the selected tracks. Above this cut-off, the efficiency degrades with increasing radius. This degradation is partly inherited from the LRT efficiency which falls with radius, and partly caused by the nature of the vertices. For a given mean proper lifetime, the average LLP boost increases with decay radius, resulting in more collimated decay products and a vertex topology which is more challenging to reconstruct.

5.2 Object identification and reconstruction

5.2.1 Electrons and photons

Photons and electrons are reconstructed by combining information from the ID and the electromagnetic calorimeter [86]. In a simplified picture, an electron would leave a track in the ID and all its energy in the calorimeter while a photon, being electrically neutral, would pass through the ID without interacting and only leave a shower in the electromagnetic calorimeter. The reconstruction however needs to take into account energy losses through Bremsstrahlung for the electrons and the possibility of photon conversions into electron-positron pairs in the ID material.

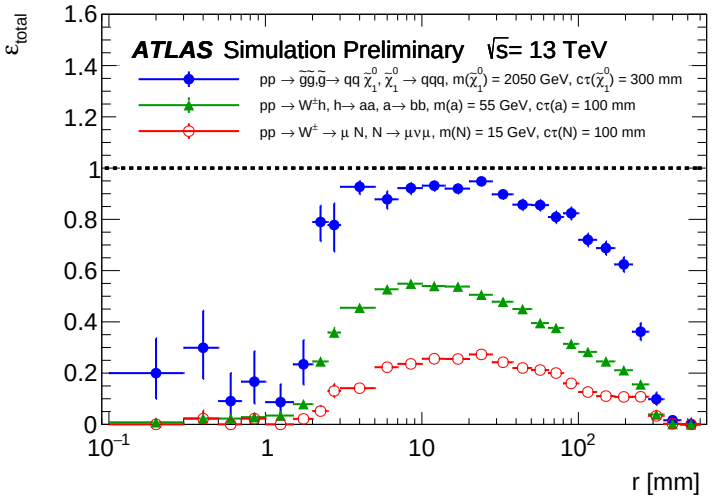

Figure 5.8: Vertex reconstruction efficiency as a function of the parent particle decay radius for three LLP models [85]. The blue points correspond to the gluino signal model used in the DV+jets analysis described in Section 10.

The reconstruction of electrons and photons begins with the identification of energy deposits in the electromagnetic calorimeter. Collections of topologically connected cells with an energy deposit larger than a predefined threshold form so-called seed clusters [87]. These are used to seed a new tracking pass which takes into account the increased energy losses of a photon compared to the default pion hypothesis. For electrons and converted photons, the seed clusters are matched to the new ID tracks by extrapolating each track to the middle layer of the calorimeter and identifying nearby clusters. A cluster is required to match the track in both position and momentum. Seed clusters matched to two oppositely charged tracks that can be fit to a conversion vertex are associated to converted photons. Finally, any seed cluster without an associated track is identified as an unconverted photon.

The p_{T} of central electrons is determined though a combination of the calorimeter energy measurement and the track measurement, while the η and ϕ are taken entirely from the track. In the forward region where the ID lacks coverage, the electron energy is determined with reduced precision and the center of the cluster is used to determine the angular coordinates of the electron.

The quality of each reconstructed electron candidate is assessed using an algorithm that assigns a likelihood that the candidate is a true electron. This is based on input from different variables including track quality and cluster shape. The electrons are classified as *Loose*, *Medium*, or *Tight* [88] which correspond to different likelihood criteria. A tighter category corresponds to a higher purity, or

lower mis-identification rate, but also to a lower identification efficiency. A similar classification is applied to photons.

In addition to the identification criteria described above, most analyses require electrons to fulfil isolation requirements, to further discriminate between signal and background. The isolation variables quantify the energy of the particles produced around the electron candidate and allow to disentangle promptly produced electrons from other, non-isolated electron candidates such as electrons originating from converted photons produced in hadron decays or electrons from heavy-flavour hadron decays, which typically are embedded in jets.

5.2.2 Muons

Muons passing though the ATLAS detector will typically leave tracks in the ID, small energy deposits in the calorimeters, and tracks in the muon system. They are thus reconstructed by combining information from the Muon Spectrometer with information from the ID, and in some cases also with information from the calorimeters [89]. The nominal reconstruction method forms so-called *combined muons* by matching tracks in the muon detector with tracks in the ID. Tracks are first reconstructed independently in the ID and in the Muon Spectrometer and matching tracks are then combined with an overall track fit.

While the combined muons have the highest purity, three other types of muons are also used in ATLAS; *stand-alone*, *segment-tagged*, and *calorimeter-tagged* muons. Stand-alone muons are formed by muon system tracks extrapolated to the beam line with energy losses and multiple scattering in the calorimeters taken into account. This algorithm is mainly used in the $2.5 < |\eta| < 2.7$ region where the ID lacks coverage. Segment-tagged muons are reconstructed starting from tracks in the ID and looking for matching track segments in the first muon system layer. These muons recover inefficiencies for low-p_{T} muons that might lose all their energy before passing through all muon system layers. Calorimeter-tagged muons are built from ID tracks matched to a calorimeter energy deposit compatible with a muon. This algorithm is used at very central $|\eta|$ where the muon system lacks coverage.

The reconstructed muons are subject to different quality and isolation criteria and are categorised as *Loose*, *Medium*, *Tight*, and *High-p_T* [89]. These selections are defined to cope with various analysis requirements. One of the key variables for their discrimination is the q/p significance, which quantifies the consistency between the ID and muon system measurements of the momentum. Another important discriminator is the χ^2/n_{dof} of the combined fit. The Medium selection is the default and uses combined muons in the central part complemented by stand-alone muons above $|\eta| = 2.5$.

5.2.3 Jets

A jet describes a collimated spray of particles resulting from the hadronisation of an initial quark or gluon, and can be identified in ATLAS by large energy deposits in the calorimeters. Several algorithms can be used to collect calorimeter energy deposits to form jets and the characteristics of the jets are heavily algorithm-dependent. The standard jet reconstruction algorithm used in ATLAS is called anti-k_t [90] and is infrared and collinear safe. Infrared safety refers to the independence of the jet definition under the inclusion of soft radiation, while collinear safe algorithms have the property of reliably including a split parton within a given jet.

The algorithm begins with a collection of clusters in the electromagnetic and hadron calorimeters defined by topologically connected cells with energy deposits above theshold [87]. The clusters are merged into jets by considering their transverse momenta and their separation in terms of the distance measure

$$d_{ij} = \min\left(p_{\mathrm{T},i}^{-2}, p_{\mathrm{T},j}^{-2}\right)\frac{\Delta R_{ij}^2}{D^2}. \tag{5.6}$$

Here, $p_{\mathrm{T},i}$ and $p_{\mathrm{T},j}$ are the transverse momenta associated to two clusters, ΔR_{ij} is their angular separation, and D is an arbitrary distance parameter. The algorithm starts by taking the cluster with the highest p_{T} as seed cluster i and computes the distance d_{ij} to the closest cluster j. This is compared to a second distance measure given by

$$d_{iB} = p_{\mathrm{T},i}^{-2}, \tag{5.7}$$

which quantifies the distance from the seed to the beam axis. If the distance d_{ij} is smaller than the distance d_{iB}, the two clusters are merged into a new cluster which takes the role as seed. Otherwise the seed i is called a jet and its associated clusters are removed from the collection. The clustering process continues until all clusters are associated with a jet. In the anti-k_t algorithm, the inverse dependence on the p_T produces jets which are conical in shape and have energetic cores and soft edges. The distance parameter D defines the size of the cone and the minimum distance between two jets and is commonly set to 0.4.

The energies of the jets are measured from the associated calorimeter energy deposits and need to be calibrated to account for several experimental effects. These include energy lost in inactive areas of the detector, leakage of energy from the jet out of the calorimeter, and energy deposition in jets from pileup interactions. A multi-step calibration procedure is applied with correction factors derived both from simulations and data using a well-measured reference object [91, 92]. In addition, the so-called *Jet Vertex Tagger* (JVT) algorithm is used to provide an indication of how likely a jet is to be due to a pileup interaction or not [93]. This is based on the tracks that are associated with the reconstructed jets and quantifies the fraction of tracks in the jet that originate from the PV.

The jet reconstruction algorithm cannot identify which type of parton initiated the jet. In the case of jets containing b-hadrons these can be identified using tracking

information in a process referred to as b-tagging. Due to the long lifetime of the b-quark, the b-hadrons typically have a decay length of a few millimetres. This is long enough to reconstruct a secondary vertex that can be used to identify them.

5.2.4 Missing transverse momentum

The ATLAS experiment is designed to infer the presence of particles that do not give a signature in the detector through the missing transverse momentum. This quantity is defined as the negative vectorial sum of the transverse momenta of all reconstructed particles and detector energy deposits in a given event,

$$E_{\mathrm{T}}^{\mathrm{miss}} = -\left[\sum \mathbf{p}_{\mathrm{T}}^{\mathrm{jet}} + \sum \mathbf{p}_{\mathrm{T}}^{\mathrm{electron}} + \sum \mathbf{p}_{\mathrm{T}}^{\mathrm{muon}} + \sum \mathbf{p}_{\mathrm{T}}^{\mathrm{photon}} + \sum \mathbf{p}_{\mathrm{T}}^{\mathrm{soft}}\right]. \quad (5.8)$$

Here, the so-called *soft term* $\sum \mathbf{p}_{\mathrm{T}}^{\mathrm{soft}}$ includes all energy deposits that are not associated to the reconstructed jets, electrons, muons, and photons. In a pp collision, the interacting partons have negligible transverse momentum, and the sum of all end product momenta in the transverse plane must therefore be zero. Any missing transverse momentum can thus be attributed to invisible particles. The soft term can be calculated based either on calorimeter or track measurements. A track-based soft term is more robust against pileup since the tracking information can be used to separate the additional vertices from the PV, but is insensitive to the momenta of neutral particles.

5.2.5 Removal of overlapping objects

The object reconstruction methods are run independently and it is common for calorimeter energy deposits and tracks to be shared between jets and particles of different types. In order to prevent double counting in such situations, a process called overlap removal is used which specifies the requirements for two objects to be considered overlapping and which ones to remove. The process is performed on *baseline* objects which have looser selections than the final *signal* objects used for analysis. This allows for a tuning of the final signal objects to best match the experimental signature targeted by the analysis, while leaving the overlap removal process unchanged. In the Strong-2L analysis the overlap procedure consists of all the steps listed below, while the DV+jets analysis only applies the first two.

- **Electron-Muon:** Electrons are removed if they share an ID track with a muon.
- **Electron-Jet:** A jet is discarded if it resides within $\Delta R = 0.2$ of an electron. If the jet is b-tagged the electron is removed instead. This ensures that electrons originating from heavy-flavour decays are discarded. All remaining electrons with ΔR satisfying $\Delta R < \min(0.04 + (10\,\mathrm{GeV})/p_{\mathrm{T}}, 0.4)$ of a jet are then removed from the event. Here, the p_{T}-dependent cone size is designed to reject low-p_{T}

heavy-flavour muons while preserving muons resulting from the decay of high-p_{T} particles, which are closely aligned with the other products of the decay.

- **Muon-Jet:** Identical to Electron-Jet.
- **Photons:** A photon is removed if it resides within $\Delta R = 0.4$ of a baseline electron or muon, and any jet within $\Delta R = 0.4$ of any remaining photon is discarded.

Chapter 6

Simulation of proton-proton collisions

The advancements in particle physics rely on comparisons between observed data and predictions from the theory. For the research conducted at the LHC, this entails producing simulations of the pp collisions, down to the level of individual particles, which are used as a detailed theoretical reference for the measurements performed at the experiments. The task is tackled by using Monte Carlo (MC) event generators which simulate the collisions in several steps.

This chapter first describes the various calculations necessary for the event simulation and then presents the MC event generators relevant to the work presented in this thesis. Finally, the simulation of the ATLAS detector and its response to the pp collision products, as well as the modelling of pileup, is presented.

6.1 Event simulation

The phenomenology of a pp collision was described in Section 1.6 and involves several physics processes at different energy scales. The hard interaction is typically the main target of the physics analyses at the LHC, but the event structure is significantly more complex and the description of the full final state necessitates involved multi-particle calculations.

Hard-scatter interaction: Any collider simulation is built around the hard primary interaction which is specified by the user, and is generated using the matrix element and according to the available phase space. The computation of higher orders requires a choice of factorisation and renormalisation scales, typically set around $\mu_F = \mu_R \approx 100\,\text{GeV}$. Both the factorisation scale and the renormalisation

scale are non-physical parameters and any observable should not depend on them. This would be true if the matrix element was computed from the entire perturbation series, which in practice is impossible. However, the more terms of the series that are included, the less the observable depends on the scales. Most MC event generators available today compute the matrix element to next-to-leading order (NLO). The computation takes experimentally measured PDFs as input to model the initial momenta of the interacting partons. There are a number of PDF sets available, all based on precision measurements at various energy scales.

Parton showering: The initial state radiation (ISR) and final state radiation (FSR) can be treated perturbatively. A complete matrix element calculation using all Feynman diagrams is however very difficult for higher orders [94]. Instead, the parton shower approach is usually used, which simulates a successive emission of quarks and gluons from the partons in the final or initial states. In this process the partons successively lose energy [95]. For FSR, the partons are evolved from the hard scatter through the parton showering until they are below the energy scale at which the hadronisation process takes place. For the simulation of ISR, a backward evolution is instead used from the initial high-energy scale of the hard scatter to the final scale of the ancestor partons provided by the PDFs.

Fragmentation: When the partons have reached the hadronisation scale, the confining nature of QCD takes over and colour neutral particles are formed. This happens in the non-perturbative regime and the hadronisation, as well as the hadron decays, are described by phenomenological models. The main models used in the MC simulation are the Lund string model [96] and the cluster hadronisation model [97].

Underlying event: The event also contains the beam remnants corresponding to the remaining partons. These may interact and lead to the so-called *underlying event*. The physics processes of the underlying event are driven by low-energy phenomena and must therefore be simulated using phenomenological modes. The employed models are characterised by tuneable parameters which are optimised by dedicated measurements and comparisons to observed data in the experiments [98].

6.2 MC event generators

MC generators are computational algorithms that rely on repeated random sampling to determine the probabilistic results of the calculations discussed above. There are many different generators which all have different strengths. Some, so-called *general purpose* generators can perform all of the steps, while other so-called *matrix element* generators only perform the calculations related to the hard interaction. These must be interfaced with separate parton shower generators in order to provide the full functionality needed to simulate the events. There are also several so-called *afterburners* which are run at the end of the event simulation, supplementing the hadronisation step in order to improve its modelling.

For each physics process of interest the generator which best describes the experimental observation is typically selected for use in the analyses. Different physics processes may therefore be simulated using different generators. In addition, the generators can be tuned by varying parameters in the simulations to create outputs that most closely match the experimental data.

The generators relevant to the various simulated samples used in this thesis are MadGraph5_aMC@NLO [99, 100] and Powheg Box [101–103], which are matrix element generators capable of calculating the matrix element to NLO, and Pythia [104, 105] and Sherpa [106, 107] which are general purpose generators. Pythia is mainly used a as a parton shower generator interfaced to one of the matrix element generators, or to model so-called minimum bias interactions described in Section 6.4. Sherpa on the other hand is used for the full event simulation, with the matrix element calculated up to NLO. In addition to these generators, the afterburner EvtGen [108] is used in some samples for the simulation of b-hadron decays in the final state.

6.3 Detector simulation

The end result of the steps simulated by the MC event generators is a collection of four-vectors corresponding to all stable particles after hadronisation. This collection can be used to study physics processes in the pp collisions at so-called *truth-level*, without the effects of the detector response and its geometric acceptance. The simulations are however mainly used to make predictions for the data collected in ATLAS, and must therefore be analysed after the full reconstruction process, described in Chapter 5. To accomplish this, the interactions of the truth-level particles with the material in the various detector subsystems are simulated using a model of the ATLAS detector in the Geant4 [109] software. The ATLAS model contains all the active detector elements and the magnetic field, as well as any inactive material such as support structures or cabling. In the software, the interaction of the particles with the detector material is simulated and the energy deposition in every part of the detector is calculated. New particles created in the material interactions might also be appended to the simulated event. The simulated energy depositions are eventually turned into hits using custom algorithms developed for each of the ATLAS subdetectors. This step is referred to as digitisation and results in simulated raw detector data that can be reconstructed using the same techniques as used for the data.

6.4 Pileup simulation

The events in ATLAS contain both in-time and out-of-time pileup, as described in Section 5.1.2, and these components must also be simulated. This is done separately

by simulating so-called minimum bias events and overlaying them as additional energy deposits in the detector before the digitisation step. The minimum bias events include all possible hard-scatter processes at a given collision energy. These are generated with parameters tuned to match the overall production of the LHC collisions using datasets collected with the ATLAS Minimum Bias Triggers (MBTS) [110]. For the in-time pileup, a number of minimum bias events are overlaid based on a sampling of the pileup distribution for a given data-taking period, shown in Figure 3.6. For out-of-time pileup, the effect of the additional interactions depends strongly on the detector technology, and different time windows are used to include the appropriate level of interactions in each detector.

6.5 The truth record

Both the truth-level information and the reconstructed-level information are useful for the analyses in ATLAS. While the reconstructed objects only represent the final state, the truth record from the generator contains all the particles that were ever generated in the event. Every interaction in the truth record is encoded as a list of incoming particles, an interaction vertex, and a list of outgoing particles. The particles have kinematic properties as well as a so-called PDG ID [111] which encodes the particle type, and the vertices have temporal and spatial properties corresponding to when and where the interaction happened. In addition, both types of objects have links to each other so that a particle can be traced to its production or decay vertex, and further to the next linked particle. The full ancestry of any particle in the truth record can therefore in principle be tracked all the way back to the initial pp interaction.

All particles and vertices in the truth record have unique barcodes defining them. These barcodes also contain information about where the particle was created, with barcodes below 200,000 corresponding to particles from the generator and barcodes above corresponding to particles from GEANT4. This feature will be exploited heavily by the analysis described in Chapter 10, where particles from material interactions contribute substantially to the backgrounds.

In order to save disk space, the truth record stored in the final output files is typically reduced by applying requirements to the kinematic properties of the particles. For primary particles from the hard-scatter interaction, these requirements are very loose so that the truth objects corresponding to the reconstructed objects are always present. Stricter requirements are placed on secondary particles, such as the decay products from hadronic interactions. For the minimum bias simulation used to overlay pileup events on the hard scatter, the truth record is usually not stored at all. This allows considerable disk space savings, as well as a reduction of CPU time when reading the events from disk. However, it also entails that no truth-level information about pileup particles is available. For studies which heavily depend

on the physics of the pileup interactions, dedicated samples may be produced where the full truth record is stored, including also the overlaid minimum bias events.

Several of the reconstructed objects have links to the truth particles from which they are built. Of large importance for the work in this thesis are the links from the ID tracks. Reconstructed tracks are matched to truth particles using a hit-based track-to-truth particle association. Each of the hits in the ID is associated to the truth particle which contributed the largest energy deposition to the hit in the simulation. The hits are then weighted according to their importance in the track reconstruction. If the hits are from the Pixel detector, the weight is set to 10, if they are from the SCT, the weight is set to 5 and if they are from the TRT, the weight is set to 1. A weighted matching probability is then defined using the ratio of the number of hits which are common to a given track and the corresponding truth particle, and the number of hits which form the track, according to

$$p_{\text{match}} = \frac{10 \cdot N_{\text{common}}^{\text{Pixel}} + 5 \cdot N_{\text{common}}^{\text{SCT}} + 1 \cdot N_{\text{common}}^{\text{TRT}}}{10 \cdot N_{\text{track}}^{\text{Pixel}} + 5 \cdot N_{\text{track}}^{\text{SCT}} + 1 \cdot N_{\text{track}}^{\text{TRT}}}, \tag{6.1}$$

where $N_{\text{common}}^{\text{Pixel/SCT/TRT}}$ is the number of tracks common to the track and the respective truth particle in the Pixel, SCT, and TRT detectors and $N_{\text{track}}^{\text{Pixel/SCT/TRT}}$ is the number of hits which form the track.

Every track receives a link to the truth particle which contributes the most to the truth match probability and is assigned the corresponding value of p_{match} which quantifies the quality of the match. If the truth particle is dropped from the truth record, the link is marked as invalid. A track can thus have an invalid truth link even if it has a non-zero match probability. For such tracks, it is not possible to study the truth origin. The truth match probability is often used to identify fake tracks by defining a threshold below which tracks are considered as fakes. Similarly, invalid truth link are often used to identify tracks from pileup interactions in standard samples which do not contain the full truth record.

Part III

Luminosity measurement

Chapter 7

ATLAS luminosity measurement

While the instantaneous luminosity depends on the configuration and steering of the LHC, it is up to each experiment to determine the delivered luminosity at a particular interaction point. In principle all analyses at the ATLAS experiment are dependent on the luminosity measurement, but the required precision varies considerably from case to case. This chapter describes the strategy of the ATLAS luminosity measurement, starting with a motivation for the needed precision. The focus is on the measurement in the LHC Run 2, corresponding to the data used for the studies in this thesis.

7.1 Importance of the luminosity measurement

A precise measurement of the luminosity is an essential part of most SM precision measurements and in particular for direct cross-section measurements where the luminosity in some cases even is the dominant uncertainty. For a precision test of the SM, the experimental accuracy should match the uncertainty affecting the theoretical predictions for the physical process under study. The precision of the measured cross section is determined by the quantities in Eq. 3.3 which can be rewritten as

$$\sigma_{pp\to O} = \frac{R_{pp\to O}}{\mathcal{L}} = \frac{\mathrm{d}N_{pp\to O}}{\mathrm{d}t} \cdot \frac{1}{\mathcal{L}} = \frac{N_{pp\to O}}{L}. \tag{7.1}$$

Here, the precision on the measured number of events $N_{pp\to O}$ depends on the knowledge of detector acceptances and biases, reconstruction efficiencies, and background

subtraction. For analyses with very precise understanding of these experimental effects, the final precision of the cross-section measurement might be dominated by the luminosity uncertainty. The precision on the theoretical cross section depends on the partonic cross section and the PDFs, as explained in Section 1.6. The calculation of the partonic cross section is purely a theoretical issue where the inclusion of higher orders of perturbation theory makes the prediction more accurate, while the knowledge of the PDFs relies on both theory and experiment.

An example of a very precise measurement is the W and Z inclusive production cross-section measurement in ATLAS with the 2011 7 TeV dataset [112]. The precision is limited by systematic uncertainties, and the total experimental uncertainty is around 0.5 %, excluding the luminosity uncertainty of 1.8 %. At the same time, the theoretical predictions are very precise with a dominant uncertainty associated with the knowledge of the PDFs of a few percent. An absolute luminosity determination with significantly better precision can therefore provide an important test of the self-consistency of the SM as well as valuable constraints on the PDFs. This kind of measurements drives the aim to obtain a luminosity estimate with an uncertainty below a couple of percent.

In searches for new physics such as the Strong-2L and DV+jets analyses, the delivered luminosity has to be known in order to evaluate background levels and determine the search sensitivity. The results from such analyses are however typically dominated by the statistical precision and the they only require a moderately precise luminosity determination. Besides the use in physics analyses, the monitoring of the real-time instantaneous luminosity is of great importance for the trigger to set the prescales and for the accelerator diagnostics.

7.2 Overview of the luminosity measurement

The determination of the luminosity amounts to measuring the event rate $R_{pp\to O}$ in Eq. 7.1. Selecting inelastic pp collisions as a the reference process, the luminosity for each colliding bunch can be expressed as

$$\mathcal{L}_{\mathrm{b}} = \frac{\mu \cdot f_{\mathrm{r}}}{\sigma_{\mathrm{inel}}}, \tag{7.2}$$

where f_{r} is the bunch revolution frequency and σ_{inel} is the inelastic pp cross section. The *pileup parameter* μ is the average number of inelastic interactions, and should be understood as the mean of the Poisson probability distribution of the actual number of interactions N_{int} per bunch crossing

$$P(N_{\mathrm{int}}) = \frac{e^{-\mu} \cdot \mu^{N_{\mathrm{int}}}}{N_{\mathrm{int}}!}. \tag{7.3}$$

Summing over the colliding crossings of proton bunches in the LHC ring, the resulting luminosity can be expressed in terms of the average expected number of

inelastic interactions per bunch crossing $\langle\mu\rangle$ according to

$$\mathcal{L} = \sum_{\text{BCIDs}} \frac{\mu_{\text{BCID}} \cdot f_{\text{r}}}{\sigma_{\text{inel}}} = \frac{\langle\mu\rangle \cdot f_{\text{r}} \cdot n_{\text{b}}}{\sigma_{\text{inel}}}, \tag{7.4}$$

where n_{b} is the number of filled BCIDs.

ATLAS has several detectors sensitive to the luminosity. These so-called *luminometers* typically measure a fraction of all inelastic interactions per bunch crossing and Eq. 7.2 can thus be rewritten as

$$\mathcal{L}_{\text{b}} = \frac{\epsilon\mu f_{\text{r}}}{\epsilon\sigma_{\text{inel}}} = \frac{\mu_{\text{vis}} f_{\text{r}}}{\sigma_{\text{vis}}}. \tag{7.5}$$

Here, ϵ is the efficiency for one inelastic pp collision to be recorded by the luminometer and the *visible interaction rate* μ_{vis} is the average number of pp collisions per bunch crossing measured by the luminometer. The *visible cross section* σ_{vis} is the calibration constant that relates the measurable quantity μ_{vis} to the absolute bunch luminosity $\mathcal{L}_{\text{b}}$ and is different for each luminometer.

The measurement of the visible interaction rate is referred to as the relative luminosity measurement and the determination of the visible cross section is referred to as the calibration of the absolute luminosity scale. In ATLAS, the calibration is done in so-called *van der Meer* (vdM) scans. The procedure determines the absolute luminosity from the measurable parameters of the colliding bunches and is performed in low-μ LHC fills with special optics settings. A *calibration transfer* method is then used to port the calibration from low μ to high μ.

The luminometers provide either *online* or *offline* luminosity measurements, or both. The online luminosity is the first measurement of the instantaneous luminosity delivered to ATLAS. It is published in real time during data taking, independently for every BCID and LB, and is reported back to the LHC. Each luminometer typically provides several *algorithms* which correspond to different ways of measuring the luminosity with the detector, for example using different PMTs in the case of LUCID. The offline measurement corresponds to the final integrated luminosity estimate used by the analyses in ATLAS. Its determination involves a more careful treatment of the data from several different luminosity detectors and algorithms.

7.3 Luminometers in ATLAS

ATLAS employs a battery of different detectors and algorithms to measure the luminosity. The redundancy is important for validation cross checks, monitoring of the stability of the measurement over time, and for the estimate of the systematic uncertainty on the integrated luminosity. The two primary luminometers are LUCID and BCM, described in Section 4.6. These both provide per-bunch, online and

offline estimates based on the number of hits in the detectors. The measurement provided by LUCID is used as the nominal ATLAS luminosity estimate in the LHC Run 2.

ATLAS also uses the calorimeter system as well as the ID to measure the luminosity. The calorimeter measurement is based on the currents drawn in the different parts of the detectors. In particular, the PMT current drawn in the TileCal modules as well as the current drawn across the LAr gaps in the EMEC and FCal modules, is used for this purpose. These measurements provide the average particle rate over longer time scales, and are therefore not on a per-bunch level. Nevertheless, the techniques provide highly valuable online and offline luminosity monitoring and is totally independent from the LUCID and BCM measurements. The ID measurements are based on a counting of objects such as tracks, vertices, and hits in the Pixel detector. This method is not suitable as an online luminometer because it relies on events recorded at a finite trigger rate. Track counting is however of large importance for the offline measurement and is the main topic of Chapter 8.

7.4 Relative luminosity measurement

The relative luminosity measurements amounts to determining the visible interaction rate μ_{vis}, which is proportional to the luminosity up to an overall scale factor according to Eq. 7.5. The algorithms used in ATLAS to estimate μ_{vis} can be grouped into *event-based*, *hit-based*, and *rate-based* algorithms.

Event-counting algorithms determine the fraction of bunch crossings in which a certain detector registers an event satisfying a given selection requirement. For instance, a bunch crossing might be counted as an event when there is at least one hit in the detector. Using Poisson statistics, the rate of bunch crossings containing an event can be used to obtain an estimate of μ_{vis}. The Poisson probability for observing no event in a given bunch crossing is $P_0(\mu_{\text{vis}}) = e^{-\mu_{\text{vis}}}$, and the probability of observing an event is thus

$$P_{\text{evt}}(\mu_{\text{vis}}) = \frac{N_{\text{evt}}}{N_{\text{BC}}} = 1 - P_0(\mu_{\text{vis}}) = 1 - e^{-\mu_{\text{vis}}}, \tag{7.6}$$

where the raw event count N_{evt} is the number of bunch crossings during a given time interval in which at least one hit is recorded, and N_{BC} is the total number of bunch crossings during the same interval. This equation can be solved for μ_{vis} to obtain

$$\mu_{\text{vis}} = -\ln\left(1 - \frac{N_{\text{evt}}}{N_{\text{BC}}}\right). \tag{7.7}$$

The event-based counting is applicable mainly in the regime where N_{evt} is small. With increasing μ_{vis}, the probability to have no signal in the detector decreases and the fraction of no-signal events approaches zero. In the limit where $N_{\text{evt}}/N_{\text{BC}} \approx 1$,

the algorithm completely loses sensitivity, which is is referred to as *saturation*. The issue can be mitigated by limiting the acceptance of the detector, for example by restricting the measurement to only one or a few of the LUCID PMTs.

To continue beyond the limit where event-based algorithms saturate, it is necessary to move to hit-based algorithms, where the number of hits in a given detector is counted. In this case the expression for μ_{vis} becomes

$$\mu_{\text{vis}} = -\ln\left(1 - \frac{N_{\text{hit}}}{N_{\text{BC}} N_{\text{ch}}}\right), \tag{7.8}$$

where N_{hit} and N_{BC} are the total numbers of hits and bunch crossings during a time interval, and N_{ch} is the number of detector channels [113]. This provides more information about the interaction rate, and increases the luminosity value at which the algorithm saturates [58]. However, hit-counting algorithms are typically more sensitive than event-counting methods to instrumental imperfections such as threshold effects and instrumental noise.

LUCID and BCM provide both event-based and hit-based luminosity measurements with several algorithms corresponding to different counting conditions. In an EventOR algorithm, a bunch crossing is counted as containing an event if there is at least one hit on either the A side or the C side of the detector. This corresponds to the expression in Eq. 7.7. For an EventAND algorithm, a bunch crossing is counted if there is at least one hit in each of the detector sides. This coincidence condition can be satisfied either from a single *pp* interaction or from individual hits on either side of the detector from different *pp* interactions in the same bunch crossing. The expression for μ_{vis} then becomes more complicated [58]. For each data-taking year, the most stable algorithm is chosen as the nominal algorithm, while the others are used for cross checks and evaluation of systematic uncertainties. Table 7.1 shows the nominal LUCID algorithms used for the luminosity measurement in Run 2. The 2015 and 2018 algorithms correspond to event counting with a single PMT on the C side, while the BiHitOr algorithms used in 2016 and 2017 correspond to hit counting using bismuth-calibrated PMTs on both the A side and C side of LUCID.

The formulae in Eqs. 7.7 and 7.8 based on Poisson statistics are inherently non-linear with the measured quantities N_{evt} or N_{hit}. Therefore, it is not possible to sum up the measurement over several BCIDs and then compute μ_{vis}, unless the value is the same for all BCIDs. At the LHC the intensity and emittance can vary between 10-20 % from bunch to bunch [58]. It is therefore essential that the LUCID and BCM detectors can provide bunch-by-bunch measurements of the number of hits.

The track-counting and calorimeter luminosity measurements are based on the proportionality between the rate of a given process and μ_{vis}. The track-counting method relies on data collected at a finite trigger rate and typically does not provide the luminosity per bunch crossing. The calorimeters do not have the possibility to

Table 7.1: LUCID algorithms used for the baseline luminosity determination in each year of Run 2 data taking, together with the visible cross sections σ_{vis} determined from the absolute vdM calibration (see Section 7.5), the peak $\langle\mu\rangle$ value, and the fraction of bunch crossings $f_{\text{no-hit}}$ which do not have a hit in any of the considered PMTs at this $\langle\mu\rangle$ value. For the HitOR algorithms, this fraction represents an average over all the contributing PMTs.

Year	**Algorithm**	σ_{vis} **(mb)**	**Peak** $\langle\mu\rangle$	$f_{\text{no-hit}}$
2015	PMT C9	6.540	16	0.270
2016	BiHitOR	6.525	41	0.035
2017	BiHitOR	6.706	60	0.006
2018	PMT C12	6.860	55	0.009

resolve individual bunch crossings and therefore also measure the bunch-averaged luminosity. Both methods are intrinsically linear in μ, and so the averaging does not introduce any bias.

7.5 Absolute calibration of the luminosity scale

In ATLAS, the calibration of the absolute luminosity scale is performed with the use of vdM beam-separation scans [58,114] in specific LHC fills. The scan is performed separately in the x and y directions, in discrete steps with increasing separation. For each scan point, the value of μ_{vis} is measured by the dedicated luminosity detectors and the resulting scan curve is used to determine σ_{vis}. Figure 7.1 shows an example scan curve measured with LUCID in a vdM session in July 2017.

For beams colliding at zero crossing angle, the bunch instantaneous luminosity for two colliding beams can be expressed as

$$\mathcal{L}_{\text{b}} = f_{\text{r}} n_1 n_2 \int \rho_1(x,y)\rho_2(x,y)\mathrm{d}x\mathrm{d}y, \tag{7.9}$$

where f_{r} is the machine revolution frequency, $n_1 n_2$ is the bunch-population product, and $\rho_i(x,y)$ is the normalised particle density in the transverse plane of each beam at the interaction point. The main assumption of the vdM method at the LHC is that the proton densities can be factorised into independent horizontal and vertical components such that

$$\rho(x,y) = \rho_x(x)\rho_y(y). \tag{7.10}$$

Under this assumption the bunch luminosity can be rewritten as

$$\mathcal{L}_{\text{b}} = \frac{f_{\text{r}} n_1 n_2}{2\pi\Sigma_x\Sigma_y}, \tag{7.11}$$

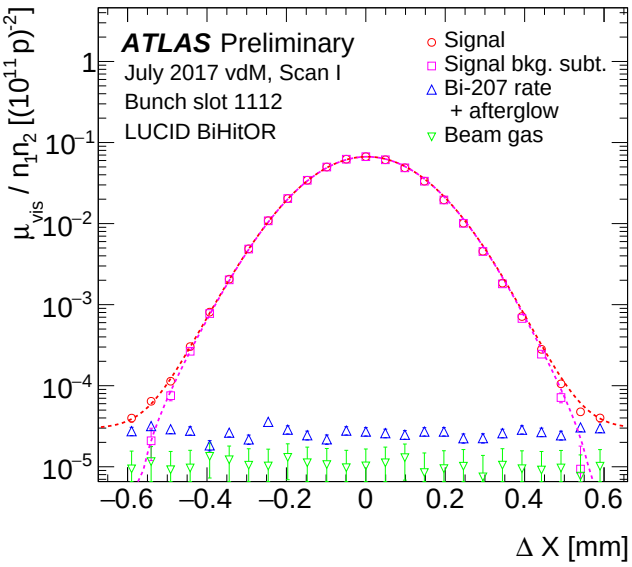

Figure 7.1: Visible interaction rate μ_{vis} per unit bunch population product $n_1 n_2$, as a function of the beam separation Δx in the horizontal plane, as measured by the LUCID BiHitOR algorithm for BCID 1112 in a vdM scan performed in July 2017 [3]. The points correspond to the scan points. At zero separation, the beams are fully overlapping, and the μ_{vis} is maximal. With increasing separation the luminosity decreases because the overlapping area gets smaller. The blue and green triangles show the constant background from e.g. the radioactive bismuth decays and beam halo.

where the effective width of the beam overlap region is given by

$$\frac{1}{\Sigma_x} = \sqrt{2\pi} \int \rho_{1x}(x)\rho_{2x}(x)\mathrm{d}x, \tag{7.12}$$

and equivalently for Σ_y. In the vdM scan these quantities are measured from the scan curve according to

$$\Sigma_x = \frac{1}{\sqrt{2\pi}} \frac{\int R(\Delta x)\mathrm{d}\Delta x}{R\left(\Delta x^{\max}\right)}, \tag{7.13}$$

where $R(\Delta x)$ is the luminosity in arbitrary units, measured during a horizontal scan when the two beams are separated by the distance Δx. The instantaneous luminosity $R\left(\Delta x^{\max}\right)$ corresponds to the peak of the scan curve which is at $\Delta x^{\max} \approx 0$. If the form of $R(\Delta x)$ is Gaussian, Σ_x is equal to the standard deviation of the distribution, but the method is valid for any functional form of $R(\Delta x)$. The normalisation of $R(\Delta x)$ cancels out in Eq. 7.13 and any quantity proportional to the luminosity can therefore be used to determine the scan curve.

The calibration of a given algorithm is given by equating Eqs. 7.5 and 7.11 resulting in an expression for the visible cross section according to

$$\sigma_{\mathrm{vis}} = \mu_{\mathrm{vis}}^{\max} \frac{2\pi\Sigma_x\Sigma_y}{n_1 n_2}, \tag{7.14}$$

where the visible interaction rate $\mu_{\mathrm{vis}}^{\max}$ corresponds to the peak of the scan curve, and the bunch population product $n_1 n_2$ is measured by the accelerator instrumentation [115]. Since the quantities entering Eq. 7.14 are different for each colliding bunch pair, it is essential to perform a bunch-by-bunch analysis to determine σ_{vis}, in practice limiting the vdM absolute calibration in ATLAS to the LUCID and BCM luminosity algorithms. All other luminometers are calibrated by normalising the measurement to the LUCID luminosity.

In order to limit the systematic uncertainties associated to the method, the vdM scan is performed with special beam conditions optimised for the purpose. Compared to normal physics running, the β^* and beam emittances are increased and special care is taken in the LHC injector chain to produce beams with Gaussian-like transverse profiles in x and y to minimise non-factorisation effects in the scans. In addition, the beams are collided with zero crossing angle and filling schemes with few and widely spaced colliding bunch pairs are used to avoid collisions between incoming and outgoing bunches away from the interaction point. Furthermore, lower bunch densities are used in order to minimise beam-beam effects. These conditions typically result in a pileup of $\langle\mu\rangle \approx 0.5$. During the LHC Run 2, up to three vdM scans were performed for each data-taking year.

7.6 Calibration transfer

The calibration of LUCID in vdM conditions needs to be transferred from low to high μ. LUCID suffers from significant non-linearity with μ, and requires a correction of $O(10\%)$ in the physics data-taking regime. The non-linearity arises when many small signals from secondary particles combine to form a signal in one tube above the PMT threshold. When the number of interactions per bunch crossing increases, the probability also increases that two or more secondary particles will give a combined signal that is above the threshold value. While saturation effects tend to underestimate the value of μ, this effect tends to overestimate it.

The non-linearity is measured by comparing the LUCID luminosity measurement to the track-counting measurement in a long physics run with a significant reduction in $\langle\mu\rangle$ over time. This run should ideally happen close in time to the vdM scan in order to limit the impact of changing data-taking conditions. Figure 7.2 shows, in red, the ratio of the track-counting and LUCID BCID-integrated measurements for the 2016 run selected for the calibration transfer, as a function of the uncorrected LUCID value. The ratio is parametrised by a first order polynomial and the fitted parameters are used to correct the LUCID measurement in all other runs according to

$$\langle\mu_{\text{corr}}\rangle = p_0 \langle\mu_{\text{uncorr}}\rangle + p_1 \left(\langle\mu_{\text{uncorr}}\rangle\right)^2 . \tag{7.15}$$

In this expression, μ_{uncorr} is the uncorrected and μ_{corr} the corrected LUCID μ value and the factors p_0 and p_1 correspond to the parameters obtained from the linear fit to the ratio.

The calibration transfer procedure implicitly assumes that track counting suffers from no significant non-linearity between the low-luminosity vdM and high-luminosity physics regimes. This assumption is probed using both internal studies comparing track-counting results with different track selections, and external studies comparing track counting to independent detectors. Figure 7.2 shows, in black, the ratio of $\langle\mu\rangle$ between track counting and TileCal. The ratio is flat which indicates that the non-linearity seen in the ratio to LUCID does not originate from track counting. A systematic uncertainty associated to the calibration transfer is evaluated from the ratio between the TileCal and track-counting luminosity measurements in the vdM run and a close-in-time high-luminosity physics fill. A difference in the ratio observed in the two regimes indicates a non-linearity between their relative responses, which is assigned as a systematic uncertainty on the correction applied to LUCID at high luminosity [3].

7.7 Long-term stability

The vdM calibration, followed by the calibration transfer to the physics regime, determines the absolute LUCID luminosity scale at any point during the data-

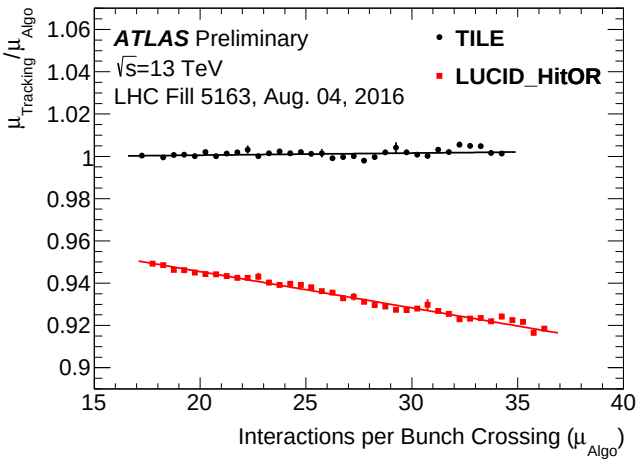

Figure 7.2: Ratios of instantaneous luminosities measured by track counting and LUCID (red), and track counting and TileCal (black), as a function of the $\langle\mu\rangle$ measured by LUCID or TileCal, in LHC fill 5163 in 2016 [3]. The lines show linear fits to the points. The TILE integrated luminosity agrees well with track counting and the ratio is close to one. The LUCID integrated luminosity shows an increasing deviation from track counting and the ratio is around 0.92 at $\mu = 35$.

taking year. To quantify any possible drifts over the year, the LUCID integrated luminosity estimates, corrected for the μ dependence, are compared for each physics run with those from other subdetectors. In this comparison, all luminosity estimates are normalised to agree with LUCID in a long reference run close to the vdM scan.

As an example, Figure 7.3 shows fractional differences between the run-integrated luminosities measured from track counting, EMEC, FCal, and TileCal with respect to the baseline LUCID algorithm chosen for 2018. The values for each run are shown as a function of the cumulative delivered luminosity fraction, ranging from zero at the start of the year to one at the end. The run-to-run agreement between the various luminosity measurements is generally at the percent level or better for the bulk of the data. Various short- and long-term trends are visible but all ratios are in agreement, meaning that the variation can be attributed to the LUCID measurement. At the start of the data-taking year LUCID systematically underestimates the luminosity compared to the other detectors. Part of this is attributed to dependencies on the number of colliding-bunch pairs which was increased gradually over a period of weeks in the beginning of the year. However, the integrated luminosity fraction affected by the problems is only a few percent of the total. A long-term stability uncertainty is evaluated conservatively from these studies by considering a stability band which encloses the bulk of the differences between LUCID and any of the other luminosity measurements [3].

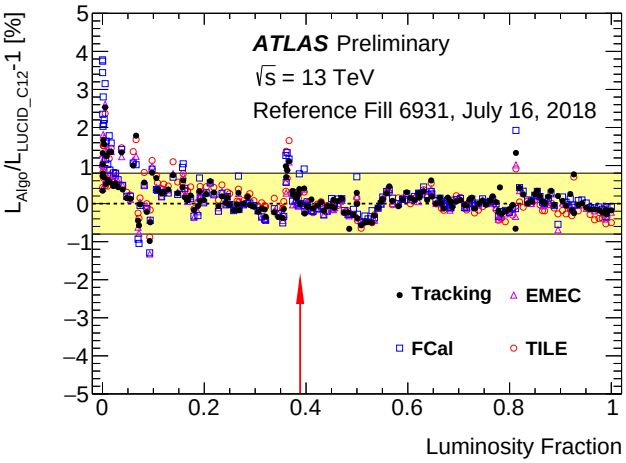

Figure 7.3: Fractional differences in run-integrated luminosity between LUCID and the track-counting, TileCal, EMEC, and FCal measurements, plotted as a function of the cumulative delivered luminosity normalised to the 2018 total integrated luminosity [3]. The luminosity measurements from the other detectors have been normalised to that of LUCID in the reference run indicated by the red arrow. The assigned long-term stability uncertainty is shown by the yellow band.

7.8 Uncertainties in the luminosity measurement

In addition to the systematic uncertainties associated to the calibration transfer and the long-term stability, the offline luminosity measurement also needs to account for uncertainties related to the LUCID measurement itself as well as the vdM calibration. The LUCID measurement needs to be corrected for background contributions related to afterglow and beam-halo interactions. The afterglow is caused by photons from nuclear de-excitation, which is induced by the hadronic cascades initiated by pp collision products. A corresponding correction can be derived by studying the luminosity signal in the BCID immediately preceding a colliding bunch pair [113]. The beam-halo background is a combination of beam-gas interactions and halo particles recorded by LUCID in time with the main beam. This background can be estimated from the luminosity signal in un-paired bunch crossings where only one beam is filled [113].

For the vdM scan a number of additional effects need to be taken into account compared to description in section 7.5, leading to corrections and additional uncertainties. Some of the dominant uncertainties are due to beam-beam effects, determination of the length scale, non-factorisation effects, and the bunch-to-bunch and scan-to-scan consistency [3]. Beam-beam effects are caused by electromagnetic interactions between the two beams which lead to a non-linear distortion of the

intended beam separation. The beam-beam deflection at each scan point, and its impact on the corresponding beam separation, can be calculated analytically from the measured bunch currents and beam widths [116][1]. Beam-beam deflections are taken into account on a bunch-by-bunch basis, by correcting the nominal separation at each scan step by the predicted displacement.

The vdM formalism assumes that the particle densities in each bunch can be factorised into independent horizontal and vertical components. Non-factorisation corrections are evaluated in each scan session from a combined fit to the beam-separation dependence of the LUCID luminosity and the position and shape of the beamspot. The beamspot is also used to measure the length scale which dictates the actual beam separation produced by a particular setting of the steering magnets. This is measured in a dedicated scan by displacing both beams transversely keeping the beams well centred on each other in the scanning plane, and measuring the actual displacement of the beamspot.

The final uncertainties relate to the consistency of the calibration between different scans and different bunches. Every year, typically two vdM sessions are held with a few scans in each plane. Results from the different scans in one year are compared to set a scan-to-scan uncertainty. The relative spread of σ_{vis} values from bunch to bunch within one scan is also used to set a bunch-by-bunch σ_{vis} consistency uncertainty.

7.9 Run 2 luminosity results

The final Run 2 luminosity measurements per year, including the uncertainties, are listed in Table 7.2. Most ATLAS physics analyses include the full Run 2 dataset and therefore only require the luminosity uncertainty on the four years combined. This is measured to be in total $139\,\text{fb}^{-1}$ with an uncertainty of 1.7 %. Since some of the uncertainty sources are correlated between all years, some are correlated between a subset of years, and some are uncorrelated, the relative error on the total is not simply a weighted sum of the relative errors on the individual years.

The measured average number of interactions per bunch crossing as a function of date in Run 2 is shown in Figure 7.4. In 2015, the $\langle\mu\rangle$ is rather stable at around 15, and in 2016 there is a slow increase from around $\langle\mu\rangle = 20$ at the start of the year to $\langle\mu\rangle = 30$ at the end of the data taking. This follows along with improvements in the beam parameters from the accelerator. The $\langle\mu\rangle$ for 2017 is generally higher than for 2015 and 2016, with $\langle\mu\rangle \approx 35$ at the start of the year, and is increased in a step to $\langle\mu\rangle \approx 50$ at the end of September. This increase is associated to the switch from standard 25 ns trains to 8b4e running, discussed in Section 3.3. In 2018, the

[1] In the preliminary Run 2 luminosity measurement [3], the corrections are computed from simulations [117].

Table 7.2: Summary of the integrated luminosities and uncertainties for each individual year of the Run 2 *pp* data sample at $\sqrt{s} = 13$ TeV (with 2015 and 2016 combined). Table adapted from [3].

Data sample	**2015+2016**	**2017**	**2018**
Integrated luminosity (fb^{-1})	36.2	44.3	58.5
Total uncertainty (fb^{-1})	0.8	1.0	1.2
Uncertainty contributions (%):			
vdM calibration	1.1	1.5	1.2
Calibration transfer	1.6	1.3	1.3
After-glow and beam-halo subtraction	0.1	0.1	0.1
Long-term stability	0.7	1.3	0.8
Tracking efficiency time-dependence	0.6	0.0	0.0
Total uncertainty (%)	2.1	2.4	2.0

pileup is rather stable at $\langle\mu\rangle \approx 45$, with slightly higher values in the first few weeks of data taking.

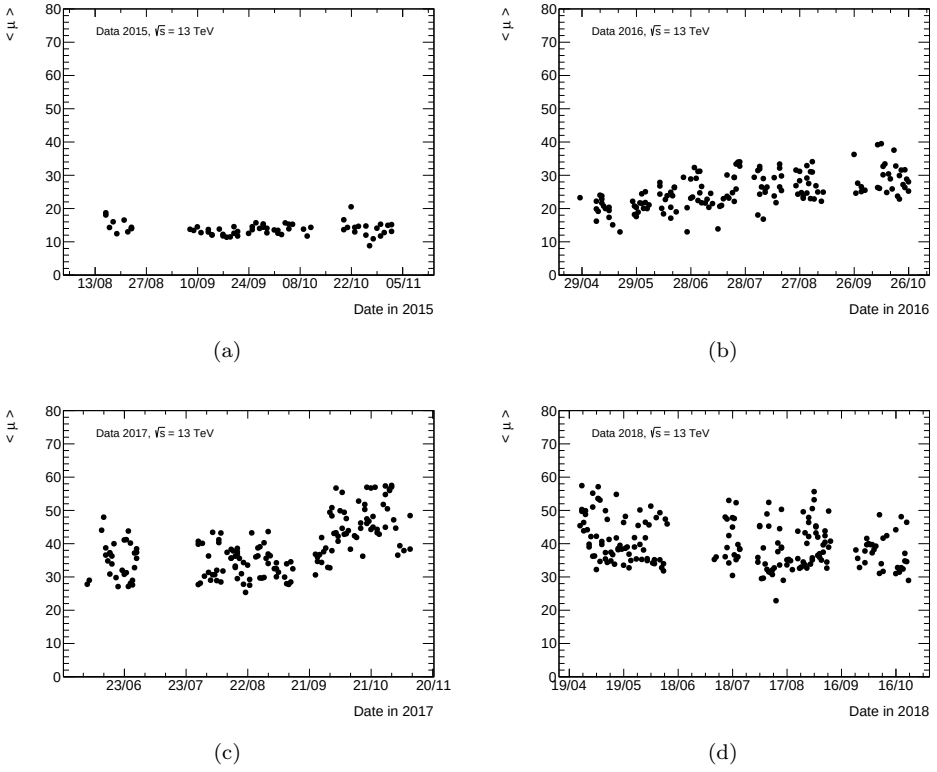

Figure 7.4: Average number of interactions per bunch crossing per run throughout 2015 (a), 2016 (b), 2017 (c), and 2018 (d). The measurements are based on the LUCID detector, using the preferred algorithm for each year.

Chapter 8

Track-counting luminosity measurement

Track counting is a central component in the ATLAS offline luminosity measurement, where it is used to determine the calibration transfer and the long-term stability. The measurement is based on the proportionality between the average number of reconstructed charged particle tracks per bunch crossing and the pileup parameter μ. This chapter discusses the methodology and presents results from the Run 2 track-counting measurement. An overview of the measurement strategy and the detailed studies presented in the chapter is given in the first section.

8.1 Overview of track counting

The relative luminosity measurement amounts to measuring the visible interaction rate μ_{vis}, as described in Section 7.2. In track counting, it is calculated from the average number of tracks in events selected using a random trigger. For such an event selection, the average number of tracks is in theory proportional to μ. The linearity of the track-counting measurement is illustrated using a toy simulation model in Section 8.2. In practice, the linearity is affected by the tracking efficiency and the presence of fake tracks, that may vary as a function of the pileup. These effects are studied as a function of pileup in fully simulated samples in Section 8.4, and as a function of the bunch-train structure in Section 8.6. The tracks used for the measurement are selected using requirements on the track quality, optimised for a high selection efficiency and to reduce the non-linear effects. Section 8.1.1 describes the event and track selections further.

The track-counting measurement is calibrated to an absolute luminosity scale by normalising to the LUCID luminosity in a process referred to as *anchoring*. This

is described in Section 8.1.3. The final deliverable consists of text files listing the measured track-counting luminosity in each LB, which are used in the ATLAS offline luminosity analysis to determine the calibration transfer and the long-term stability. In this procedure, the uncertainty related to the track-counting measurement is determined by the comparison to other luminometers. Internal monitoring tools are also in place to study and ensure the stability of the track-counting measurement and are described in Section 8.5. In addition, studies are performed in order to assess the statistical and systematic errors associated to track counting. These are presented in Sections 8.3 and 8.7 respectively.

8.1.1 Data acquisition and track reconstruction

A central part of track counting is the selection of events and tracks for the measurement. The event selection must include colliding-bunch crossings where no pp interaction takes place, so-called *empty* bunch crossings, as the measurement otherwise will be biased towards high luminosity. This excludes the use of normal physics triggers that select events based on the presence of high-p_{T} objects. The data used for the track-counting measurement is instead recorded with a trigger that randomly accepts events from colliding-bunch crossings, providing an unbiased sampling. For each triggered event, the information from the IBL, Pixel, and SCT detectors is read out and saved for further reconstruction. All other ATLAS sub-detectors are ignored. This partial read-out mode is chosen to reduce the disk footprint in order to allow for a high read-out rate. In Run 2, the rate was typically at 200 Hz during normal physics data taking and several kHz during vdM scans.

The partial read-out mode requires the tracks to be reconstructed using hits from the IBL, Pixel, and SCT detectors only. Tracks used in the track-counting luminosity measurement are therefore reconstructed with special settings without any dependence on TRT hits. In the standard track reconstruction, the TRT hits are most important for the outside-in tracking pass that serves to reconstruct secondary particles. The luminosity measurement counts primary particles and the removal of the TRT hits is therefore not an issue. Another difference compared to standard tracking is that the minimum p_{T} requirement is raised to 900 MeV.

After reconstruction, the tracks are subject to a selection for good track quality. Various working points (WPs), corresponding to different sets of selection requirements, are studied for the track-counting measurement. The first WP, *Tight*, is the ATLAS standard track selection TightPrimary described in Section 5.1.3, with the modified p_{T} requirement. Based on Tight, three other track selections are studied; *TightLumi*, *TightModLumi*, and *TightModSiPlusLumi*. These all contain a requirement on the maximum d_0-significance, which serves to reduce the rate of fake tracks. TightModLumi and TightModSiPlusLumi restrict the acceptance to $|\eta| < 1$ and allow for one hole per track in the Pixel detector. In TightModSiPlusLumi, the requirement on the number of Silicon hits is further tightened compared to the other WPs. Table 8.1 lists the criteria applied for each of the WPs. The

Table 8.1: List of track selection criteria applied for the Tight, TightLumi, TightModLumi, and TightModSiPlusLumi WPs. Criteria marked in gray are the same as in the standard TightPrimary selection.

Criterion	**Tight**	**TightLumi**	**TightModLumi**	**TightModSiPlusLumi**
$\lvert\eta\rvert$	< 2.5	< 2.5	< 1.0	< 1.0
$N_{\text{hits}}^{\text{Si}}$	≥ 9 if $\lvert\eta\rvert \leq 1.65$ ≥ 11 if $\lvert\eta\rvert > 1.65$	≥ 9 if $\lvert\eta\rvert \leq 1.65$ ≥ 11 if $\lvert\eta\rvert > 1.65$	≥ 9	≥ 10
$N_{\text{holes}}^{\text{Pix}}$	$=0$	$=0$	≤ 1	≤ 1
$\lvert d_0\rvert / \sigma_{d_0}$	–	< 7	< 7	< 7

nominal WP used for the Run 2 track-counting measurement is TightModLumi and the other selections are primarily used for studies of systematic uncertainties and long-term stability. Section 8.4 discusses the implications on the tracking efficiency and the fake rate of each track selection requirement.

The track-counting results presented in this chapter correspond to the full Run 2 dataset of $139\,\text{fb}^{-1}$. In addition, a number of special data sets, recorded at lower values of μ and with varying beam conditions, are used for dedicated studies, and will be presented in the text.

8.1.2 Relative luminosity measurement

In track counting, the visible interaction rate is calculated from the average number of tracks per event, with the average computed over all filled and colliding BCIDs in each LB. The bunch-integrated track-counting visible interaction rate is thus given by

$$\langle\mu\rangle_{\text{vis}}^{\text{TC}} = \langle N_{\text{trk}}\rangle_{\text{LB}} = \frac{N_{\text{trk}}^{\text{LB}}}{N_{\text{evt}}^{\text{LB}}} = \frac{\sum_{i=1}^{N_{\text{evt}}^{\text{LB}}} N_{\text{trk}}^{i}}{N_{\text{evt}}^{\text{LB}}}, \tag{8.1}$$

where N_{trk}^{i} is the number of reconstructed tracks in each event, $N_{\text{trk}}^{\text{LB}}$ the total number of tracks in the LB, and $N_{\text{evt}}^{\text{LB}}$ the total number of events recorded in the LB. The track-counting measurement can in principle also be performed for each BCID separately. Bunch-specific measurements are however only produced for special studies.

The statistical error on the measurement is taken as the standard error on the mean according to

$$\Delta\langle\mu\rangle_{\text{vis}}^{\text{TC}} = \frac{\sigma\left(N_{\text{trk}}^{\text{LB}}\right)}{\sqrt{N_{\text{evt}}^{\text{LB}}}}\,, \tag{8.2}$$

where the standard deviation is given by

$$\sigma\left(N_{\mathrm{trk}}^{\mathrm{LB}}\right)=\sqrt{\frac{1}{N_{\mathrm{evt}}^{\mathrm{LB}}}\sum_{i=1}^{N_{\mathrm{evt}}^{\mathrm{LB}}}\left(N_{\mathrm{trk}}^{i}\right)^{2}-\left(\frac{1}{N_{\mathrm{evt}}^{\mathrm{LB}}}\sum_{i=1}^{N_{\mathrm{evt}}^{\mathrm{LB}}}N_{\mathrm{trk}}^{i}\right)^{2}}. \tag{8.3}$$

8.1.3 Calibration of the absolute scale

The relative measurement, given by $\langle\mu\rangle_{\mathrm{vis}}^{\mathrm{TC}}$, must be calibrated to obtain the absolute luminosity scale. While the vdM method, described in Section 7.5, in theory could be used for the track-counting measurement, it is in practice limited by the read-out rate of the tracking data. In the tails of the vdM scans, where $\mu \ll 1$, very few bunch crossings will contain inelastic collisions. Calibrating track counting with the vdM method would therefore require a random trigger with sufficient rate to capture enough tracks for a statistically meaningful measurement. This has not been attempted in Run 2. Instead, the absolute scale of the track-counting luminosity is determined by normalising the measurement to the calibrated LUCID value. Adopting the same notation as in Eq. 7.5, the track-counting luminosity can be expressed as

$$\mathcal{L}^{\mathrm{TC}}=\frac{\langle\mu\rangle_{\mathrm{vis}}^{\mathrm{TC}}\cdot f_{\mathrm{r}}\cdot n_{\mathrm{b}}}{\sigma_{\mathrm{vis}}^{\mathrm{TC}}}=\frac{\langle\mu\rangle_{\mathrm{vis}}^{\mathrm{TC}}\cdot f_{\mathrm{r}}\cdot n_{\mathrm{b}}}{n^{\mathrm{TC}}\cdot\sigma_{\mathrm{inel}}}=\frac{\mathcal{L}_{\mathrm{uncalib}}^{\mathrm{TC}}}{n^{\mathrm{TC}}}, \tag{8.4}$$

where the *anchor factor* n^{TC} is given by

$$n^{\mathrm{TC}}=\frac{\langle\mu\rangle_{\mathrm{vis}}^{\mathrm{TC}}}{\langle\mu\rangle}=\frac{\sigma_{\mathrm{vis}}^{\mathrm{TC}}}{\sigma_{\mathrm{inel}}}=\frac{\mathcal{L}_{\mathrm{uncalib}}^{\mathrm{TC}}}{\mathcal{L}^{\mathrm{TC}}}, \tag{8.5}$$

and the uncalibrated track-counting luminosity by

$$\mathcal{L}_{\mathrm{uncalib}}^{\mathrm{TC}}=\frac{\langle\mu\rangle_{\mathrm{vis}}^{\mathrm{TC}}\cdot f_{\mathrm{r}}\cdot n_{\mathrm{b}}}{\sigma_{\mathrm{inel}}}. \tag{8.6}$$

The anchor factor is the parameter needed to determine the absolute scale of the track-counting measurement and must be measured with suitable data. In the vdM runs, the scanning periods are typically interleaved with periods where the beams collide head-on at a low and constant μ. These so-called *quiescent periods* provide an ideal dataset to transfer the absolute calibration of LUCID to track counting, since it can be assumed that the data-taking conditions between the scanning and head-on periods remain the same. In this dataset, the denominator in Eq. 8.5 can thus be set to the calibrated LUCID luminosity. To increase the statistical precision, the anchor factor is determined based on the track-counting measurement and the

Table 8.2: Track-counting anchor factors for 2018 data, including the statistical error.

Working point	**Anchor factor**
Tight	3.86 ± 0.14
TightLumi	3.70 ± 0.13
TightModLumi	1.69 ± 0.06
TightModSiPlusLumi	1.67 ± 0.06

calibrated LUCID luminosity integrated over all LBs in the quiescent periods of the vdM runs according to

$$n^{\mathrm{TC}} = \frac{L^{\mathrm{TC}}_{\mathrm{uncalib}}(\mathrm{vdM})}{L^{\mathrm{LUCID}}(\mathrm{vdM})}\,. \tag{8.7}$$

Once the calibration factor is known, it can be used in any other ATLAS run to convert $\langle\mu\rangle^{\mathrm{TC}}_{\mathrm{vis}}$ to a calibrated luminosity value according to Eq. 8.4, and hence provide a luminosity estimate from track counting. Similarly, the average number of *pp* interactions as measured by track counting can be computed from Eq. 8.5. One anchor factor is determined for each WP for each year. Table 8.2 lists the factors measured in the quiescent periods in the 2018 vdM run. The anchor factors measured in the other Run 2 data-taking years are very similar and agree within the statistical uncertainties.

Figure 8.1 shows the ratio of the luminosity measured by LUCID and the different track-counting WPs divided by the luminosity obtained from TightModLumi as a function of LB in the quiescent periods of the 2018 vdM run. All track WPs are corrected to LUCID by one anchor factor each, determined from the integrated luminosity ratios. Deviations in the luminosity ratios from one throughout the run could indicate changing detector conditions, and complicate the anchoring using the integrated quantities. The ratio in Figure 8.1 is however stable as a function of LB, and track counting can safely be anchored to LUCID.

8.2 Basic principle of track counting

The basic principle of the track-counting measurement is the linear relationship between the average number of tracks and the average number of inelastic *pp* interactions, μ, per bunch crossing. In what follows, a toy simulation method is used to illustrate the linearity. The model is based on the normalised distribution of the number of tracks per *pp* interaction. This can be thought of as the track-multiplicity probability density function (pdf) and is shown in Figure 8.2 for the TightLumi and TightModLumi WPs. The distributions are extracted from a MC sample containing minimum bias events with exactly one *pp* interaction per bunch

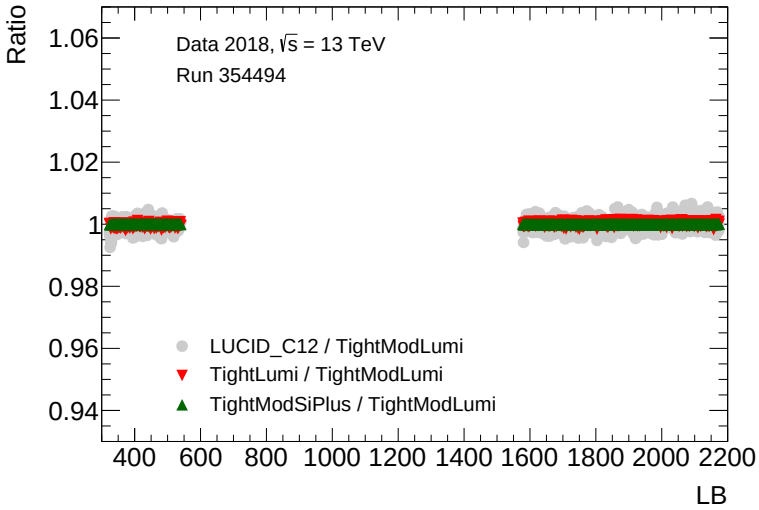

Figure 8.1: Ratios of LUCID, TightLumi, and TightModSiPlusLumi to TightMod-Lumi, as a function of LB in the quiescent periods of the 2018 vdM run.

crossing. The sample is simulated using Pythia8.186 [104] with the `NNPDF2.3lo` set of PDFs [25], and the A3 tune [118]. A total of one million events are simulated. The tracks in the MC sample are reconstructed using information from all the ID subdetectors, including the TRT, and therefore differ from the tracks in the data. This was the only suitable MC sample available in ATLAS at the time of the studies. The extra secondaries resulting from the TRT tracking will however be efficiently removed by the d_0-significance requirement.

The number of tracks measured in a given bunch crossing is a random variable described by the product of the probability distributions for the number of pp interactions per bunch crossing and the number of tracks per pp interaction. The former is described by a Poisson distribution, and the latter is described by the track-multiplicity pdf. A bunch crossing at a given value of μ is thus simulated by sampling a random number of pp interactions from a Poisson distribution with a mean equal to μ. For each pp interaction, a number of tracks is then generated randomly based on the the track-multiplicity pdf. Figure 8.3 shows the average number of tracks as a function of μ in bunch crossings generated with the toy simulation model, for the TightLumi and TightModLumi WPs. The points are well described by a straight line, with the slope being in agreement with the mean value of the track-multiplicity pdf. This is the average number of tracks per pp interaction measured with a specific track selection, and corresponds to the proportionality constant between the average number of tracks and μ. The equivalent quantity in data is the anchor factor used to calibrate the track-counting measurement.

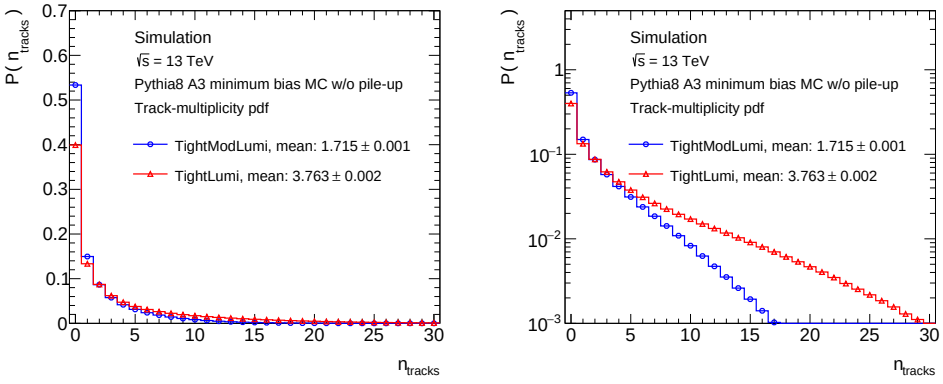

(a) (b)

Figure 8.2: Normalised track-multiplicity distributions in the minimum bias MC with one *pp* interaction per bunch crossing, in linear (a) and logarithmic (b) scale, for the TightLumi and TightModLumi WPs.

Table 8.3: Mean values of the track-multiplicity pdfs measured in minimum bias MC, together with the anchor factors repeated from Table 8.2.

Working point	**Track pdf mean**	**Anchor factor**
Tight	3.916 ± 0.002	3.86 ± 0.14
TightLumi	3.763 ± 0.002	3.70 ± 0.13
TightModLumi	1.715 ± 0.001	1.69 ± 0.06

Table 8.3 summarises the mean values of the track-multiplicity pdfs for the Tight, TightLumi, and TightModLumi WPs, together with the anchor factors repeated from Table 8.2. The values are very close to the anchor factors, suggesting a good modelling of the track-multiplicity.

The differences between the mean values of the different WPs can be understood by considering the details of the track selections. TightModLumi applies looser quality requirements on the reconstructed tracks but restricts the acceptance to $|\eta| < 1$. This leads to fewer reconstructed tracks in each *pp* interaction compared to the TightLumi selection, despite the looser tracking requirements.

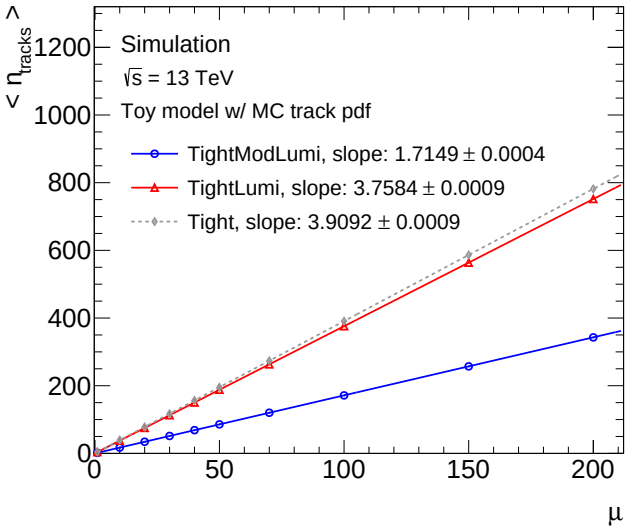

Figure 8.3: Mean number of tracks per bunch crossing as a function of μ in the toy simulation model using the Tight, TightLumi or TightModLumi track-multiplicity pdf. For each value of μ, 100,000 toy bunch crossings are simulated and used to compute the average number of tracks. A first order polynomial is fitted to the simulated points.

8.2.1 Validation of the track-multiplicity pdf in data

The procedure to derive the data track-multiplicity pdf is more involved, since it is not possible to select bunch crossings with only one pp interaction at event level. Instead, a dataset with $\mu \ll 1$ is used to ensure that the probability for an event to contain two or more pp interactions is negligible. The data is taken from an LHC μ scan performed in 2018, in which the beams are separated similarly to the vdM scans, but where beam parameters similar to those of physics data-taking are used. Three LBs with $\langle\mu\rangle = 0.0086$, as determined by LUCID, are selected for the study. Figure 8.4 shows the distribution of the number of tracks in the selected LBs for the TightModLumi WP in light blue. For $\langle\mu\rangle = 0.0086$, the Poisson probabilities to observe zero, one, and two pp interactions are 0.99144, 0.00853 and 0.00004 respectively. The high probability of zero interactions explains why the zero-track bin dominates the track distribution. Furthermore, the ratio of the probabilities of two and one interactions is $0.00004/0.00853 = 0.0047$. In the track distribution, the bins containing one or more tracks will therefore almost exclusively be populated by events with one pp interaction, with a contamination from events with two interactions of approximately 0.5%.

The zero-track bin is populated by events containing zero pp interactions and events

containing one or more interactions with zero reconstructed tracks. In order to derive a track-multiplicity pdf, only the events containing one *pp* interaction should be kept. The contribution from events with zero interactions must therefore be removed from the zero-track bin. This is done using the $\langle\mu\rangle$ value measured by LUCID, computing the Poisson probability for zero interactions and then subtracting the corresponding fraction of events from the zero track bin according to

$$N_0^{\text{corr}} = N_0 - P \cdot N_{\text{evt}} = N_0 - P \cdot (N_0 + N_{>0}) = (1 - P) \cdot N_0 - P \cdot N_{>0}, \quad (8.8)$$

where $P = P(N_{\text{int}} = 0|\langle\mu\rangle)$ is the Poisson probability for zero interactions given the $\langle\mu\rangle$ of the data, N_0 is number of events containing zero tracks, $N_{>0}$ is the number of events containing one or more tracks and N_{evt} is the total number of events in the LB.

The uncertainty on the LUCID $\langle\mu\rangle$ measurement is also accounted for by repeating the procedure above but varying the $\langle\mu\rangle$ value up and down by its uncertainty to compute an upper and lower estimate. The uncertainty used in computing the correction error is 1.4 % and considers the statistical LUCID uncertainty as well as the uncertainty on the luminosity calibration in the vdM scan. Figure 8.4 shows the corrected zero-track bin content overlaid on the original track distribution, for the TightModLumi WP.

In order to achieve the final track-multiplicity pdf, the corrected track distribution must be normalised. The systematic error assigned to the corrected zero-track bin is propagated to the normalised distribution by computing a maximum and minimum normalisation factor. Figure 8.5 shows a comparison between the resulting data track-multiplicity pdfs and the MC pdfs for the TightModLumi WP. The overall shape of the pdf is in agreement between data and MC, meaning that the track-multiplicity per *pp* interaction is well modelled. This motivates the use of the minimum bias events simulated by Pythia8 for further track-counting studies.

8.3 Study of the statistical error

In the track-counting measurement, the average number of tracks is computed over all events in one LB according to Eq. 8.1 and the statistical uncertainty is taken as the standard error on the mean according to Eq. 8.2. This is a symmetric error and assumes that the average number of tracks is distributed according to a Gaussian distribution. The statistical error is mainly important for the calibration transfer and long-term stability studies of the offline luminosity measurement, where track counting is compared to other luminometers. The measured values of $\mu_{\text{vis}}^{\text{TC}}$ in all LBs at the same μ or in the same run are then combined using either an unweighted average or a weighted average which takes the uncertainties into account. Provided that the uncertainties are properly estimated, the two methods should give consistent results. In what follows, the toy simulation model described in Section 8.2

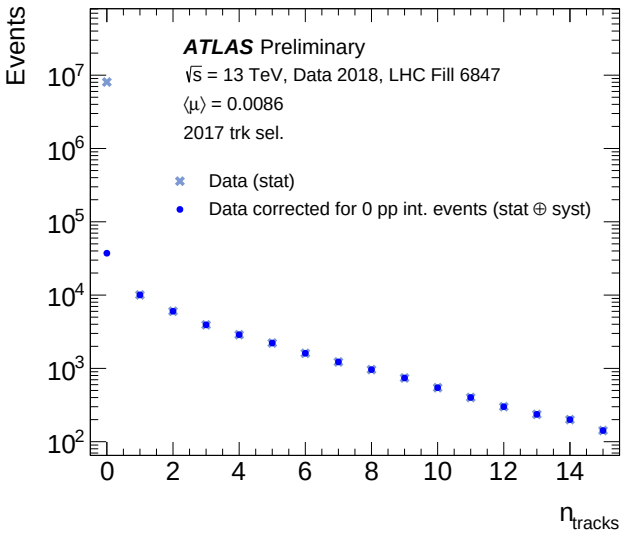

Figure 8.4: Number of tracks per event in 2018 data with $\langle\mu\rangle = 0.0086$, using the TightModLumi WP [2]. The raw data distribution (light blue cross) is shown together with the corrected distribution (filled blue circles). All points include statistical errors, but the error bars are smaller than the symbol size. The corrected distribution is equal to the raw data in all bins except for the first, where the fraction of events with zero *pp* interactions is subtracted.

is used to assess the validity of the symmetric error, and to test the weighted and unweighted average approaches.

The model is first used to generate toy LBs containing a fixed number of events N_{evt} at a given value of μ, and computing the average number of tracks in each LB. Example distributions of the average number of tracks for varying values of μ and N_{evt} in the toy LBs are shown in Figure 8.6 for the TightModLumi WP. While the mean value is kept constant within the statistical uncertainty for a fixed value of μ, the shape of the distribution varies with both N_{evt} and μ. The product of μ and N_{evt} corresponds on average to the total number of *pp* interactions N_{int} in the LB. This in turn dictates the number of tracks used to compute the average and steers the shape of the distribution. The track distribution is highly asymmetric for small values of μ and N_{evt}, corresponding to a small value of N_{int}, and approaches a Gaussian distribution for higher values.

In order to quantify the statistics needed for the symmetric errors to be applicable, the weighted and unweighted averages are studied as a function of the number of *pp* interactions. For each value of N_{int}, a large ensemble of N_{LB} toy LBs containing exactly N_{int} interactions each is generated. For each interaction in each toy LB, the

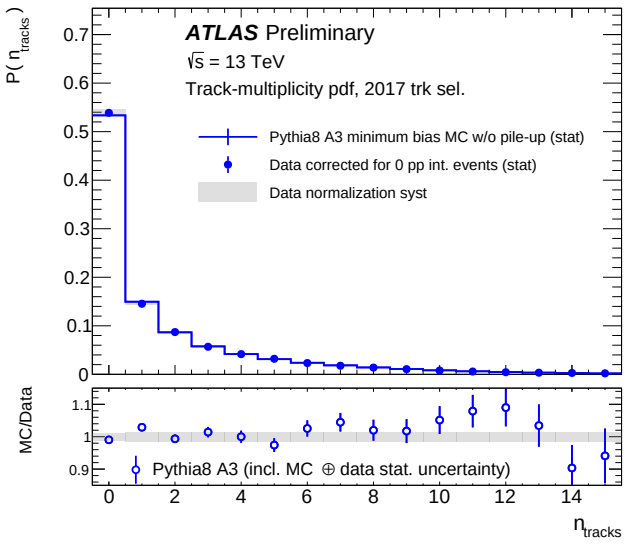

Figure 8.5: Normalised track distributions in minimum bias MC and data with one *pp* interaction per bunch crossing for the TightModLumi WP [2]. The shaded error band on the data is computed from the difference in the normalisation factor resulting from varying the content in the zero track bin by its systematic uncertainty. The lower panel shows the ratio between MC and data, with the combined MC and data statistical uncertainty shown by the error bars, and the normalisation factor uncertainty by the shaded band.

number of tracks is sampled from the TightModLumi track-multiplicity pdf and the average is then computed from the total number of tracks in each LB. Finally, all LBs in the ensemble are combined using either the unweighted or weighted average, and the result is divided by the number of interactions. Figure 8.7 shows the weighted and unweighted LB averages as a function of the number of *pp* interactions in each LB. Each point corresponds to one ensemble of LBs described above. The unweighted average is centred at 1.71 which corresponds to the expected average number of tracks per event for the TightModLumi WP according to Table 8.3. It is thus an unbiased estimate. The weighted average is biased for low values of N_{int}, but approaches the unweighted value for higher values.

The tools for the luminosity measurement in ATLAS are not built to handle an asymmetric error. Switching to a new error definition for track counting is therefore not feasible. The unweighted mean is always an unbiased estimate and any potential error in the definition of the statistical uncertainty therefore does not pose a direct problem to the offline luminosity analysis. In addition, the results above can be used to estimate the number of tracks necessary for the symmetric uncertainty to be representative of the underlying distribution of the number of tracks. In

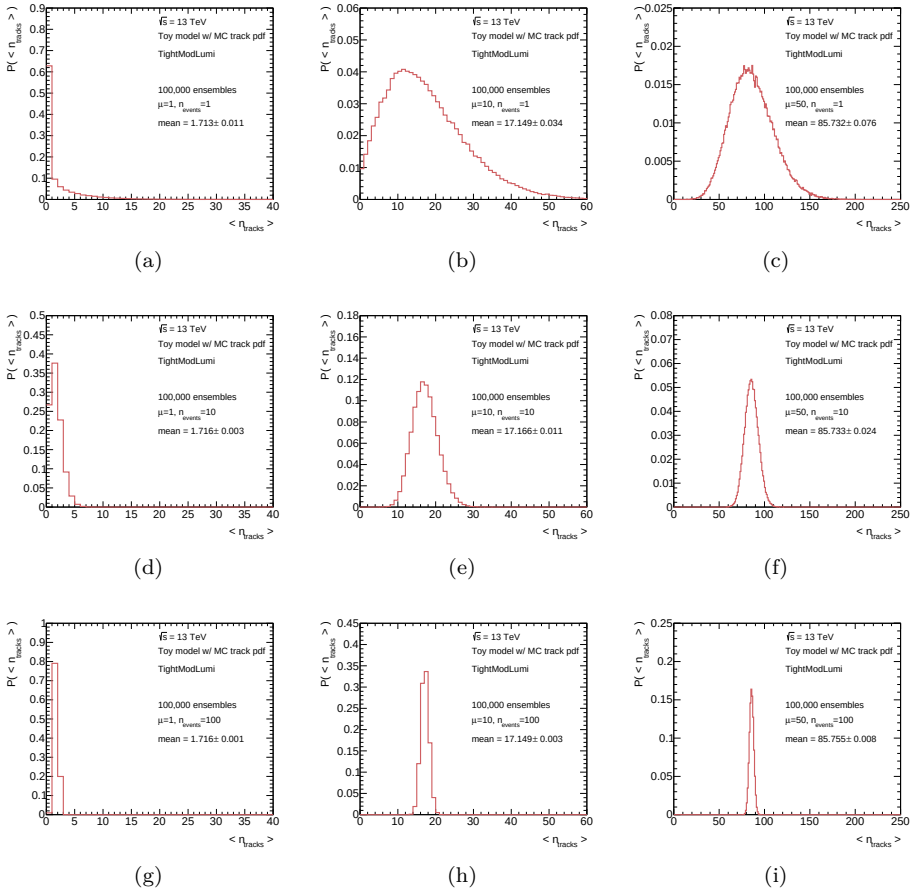

Figure 8.6: Toy simulation distributions of the average number of tracks, for $\mu = 1$, 10 and 50, for $N_{\text{evt}} = 1$ (a - c), $N_{\text{evt}} = 10$ (d - f) and $N_{\text{evt}} = 100$ (g - i). Each figure contains 100,000 toy LBs. Here, the TightModLumi WP track-multiplicity pdf is used as input to the toy simulation model.

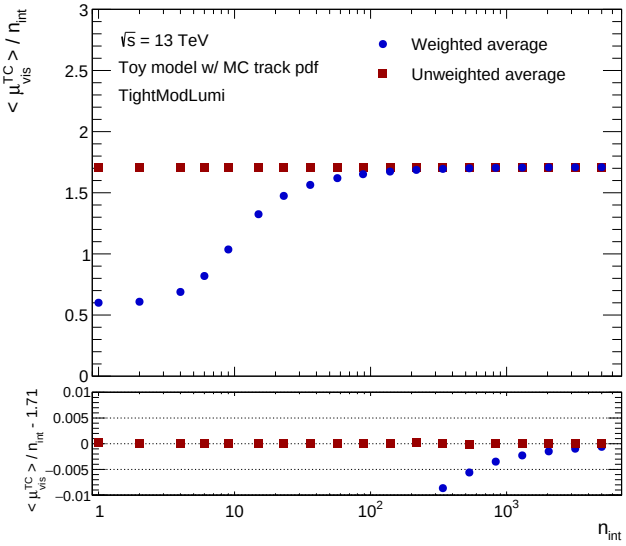

Figure 8.7: Weighted and unweighted average of $\mu_{\text{vis}}^{\text{TC}}/N_{\text{int}}$ as a function of the number of *pp* interactions in the LB. Here, the TightModLumi track-multiplicity pdf is used to sample the number of tracks per *pp* interaction. The lower panel shows the difference between the averages and the mean of the track-multiplicity pdf.

Figure 8.7, the bias in the weighted mean is less than one per mille for $N_{\text{int}} > 4500$. Compared to the offline luminosity uncertainty which is at the percent level, this is an acceptable bias. Using the expected average number of tracks, the requirement on the number of *pp* interactions per LB can be converted into a requirement of a minimum of $1.7 \cdot 4500 \approx 8000$ tracks per LB. In principle all LBs in standard physics data-taking are well above this requirement, where the uncertainty on the track-counting measurement, computed according to Eq. 8.2, gives an accurate representation of the confidence interval for the average number of tracks.

8.4 Study of non-linear effects

The toy simulation studies illustrate the proportionality between the average number of tracks and μ in the absence of any non-linear effects introduced by pileup-dependent tracking effects. In reality, the efficiency of track reconstruction, the selection, and the probability for fake tracks to be reconstructed from random combinations of hits, may vary with μ. The number of tracks measured in an event can be expressed as

$$N_{\text{trk}}(\mu) = N_{\text{true}}(\mu) \cdot \epsilon_{\text{r}\times\text{s}}(\mu) + N_{\text{trk}}^{\text{fake}}(\mu), \tag{8.9}$$

where N_{true} is the true number of charged primary particles produced in the event, $\epsilon_{\text{r}\times\text{s}}$ is the combined track reconstruction and selection efficiency and N_{fake} is the number of fake tracks. In this expression, the number of true charged particles per pp interaction is linear with μ. The tracking efficiency and the number of fake tracks vary slightly with μ however, and introduce non-linearities into the track-counting measurement. A fully simulated sample, which takes the correlations between various pp interactions into account, is used to study these non-linear effects.

8.4.1 MC sample

The non-linearity studies are based on a MC sample, produced specifically for the track-counting studies, where the hard-scatter process is overlaid with minimum bias events with a flat μ distribution covering the range from 0 to 100. The sample is reconstructed with the special tracking configuration without TRT hits and is produced and stored keeping the full truth record as described in Section 6.5. A total of 60 million events are simulated.

Technically for the MC generation, a hard-scatter event must be present in every event and for this sample it is chosen to be $Z \to \mu\mu$. However, the studies presented in this thesis only make use of the overlaid minimum bias events and ignores the hard scatter. These inelastic pp events are generated with Pythia8.186 [104] using the `NNPDF2.3lo` set of parton distribution functions (PDFs) [25] and the A3 tune [118]. The section below provides further details on the truth record in the sample and explains how it can be used to estimate the tracking efficiency and the rate of fake tracks.

8.4.2 Classification of truth particles and tracks

The truth record is used to categorise truth particles and reconstructed tracks in the $Z \to \mu\mu$ MC sample. The studies rely on charged primary particles, selected by requiring the barcode to be below 200,000. The particles are furthermore subject to truth-level requirements corresponding to the selection of reconstructed tracks in the WPs, according to Table 8.4. All reconstructed tracks in the sample are classified as primary or fake by means of the truth match probabilities and truth links. Figure 8.8 shows the truth match probability of tracks in the $Z \to \mu\mu$ sample for tracks selected with the Tight WP, as well as the fractions of tracks with and without a valid truth link. Primary tracks are those with a truth match probability above 0.5 and a valid truth link to a charged particle. Fake tracks are those with a truth match probability below 0.5, without any requirement on the validity of the truth link. It should be noted that the sum of the two categories does not correspond to all tracks in the event. Table 8.5 summarises the track categories.

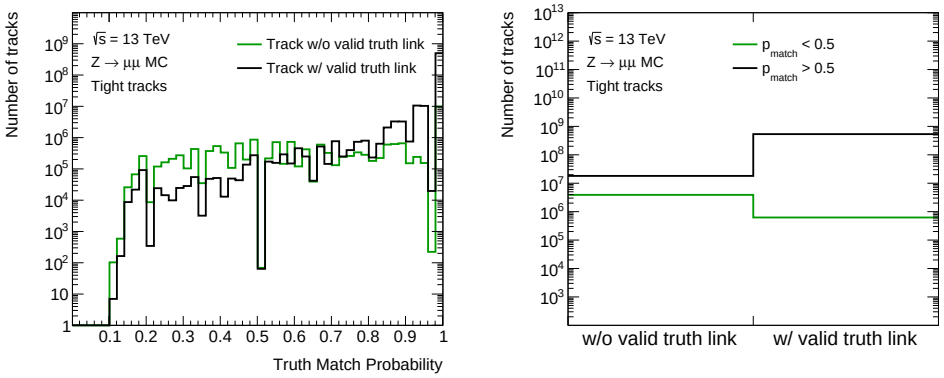

(a) (b)

Figure 8.8: Distribution of the truth match probability (a) and fraction of tracks without and with a valid truth link (b), for tracks selected with the Tight WP, in the $Z \to \mu\mu$ MC sample.

Table 8.4: Charged truth primary particle selection, based on information in the truth record of the $Z \to \mu\mu$ MC sample.

Property	**Requirement**
Status	1, i.e. final state particle
Barcode	≤ 200000, i.e. reject GEANT4 particles
p_{T}	≥ 900 MeV
$\lvert\eta\rvert$	≤ 2.5 (≤ 1.0 for TightModLumi and TightModSiPlusLumi)
Charge	≥ 1

The truth information in the sample contains information of the type of process that generated a given particle. For each track with a truth link, it is therefore possible to check if it is linked to the hard-scatter process or to a pileup interaction. The number of tracks, efficiencies, and fake rates are computed excluding the tracks from the hard scatter and are studied as functions of μ_{truth} which is the mean number of overlaid pileup interactions per event, excluding the hard-scatter interaction.

In the fake-track category, tracks without a valid truth link are included. For these tracks it is not possible to check the origin and thus reject the ones linked to the hard scatter. As a result, there will be in a small contamination of tracks originating from the hard scatter in the fake category. The size of the contamination is estimated by measuring the average number of fake-classified tracks in events with $\mu_{\mathrm{truth}} = 0$. The obtained values are 0.14×10^{-3}, 0.08×10^{-3}, 0.02×10^{-3}, and 0.005×10^{-3} for Tight, TightLumi, TightModLumi, and TightModSiPlusLumi respectively. This

Table 8.5: Definition of track categories, based on the truth match probability p_{match}, and the truth links of tracks in the $Z \to \mu\mu$ MC sample.

Category	$\boldsymbol{p_{\text{match}}}$	**Truth Link**
All tracks	-	-
Primary track	≥ 0.5	Valid link to charged truth particle
Fake track	< 0.5	-

correction factor is subtracted from the measured average number of tracks for each μ value.

8.4.3 Non-linear effects

Figure 8.9 shows the average number of tracks per event, selected with the various WPs, as a function of μ_{truth} in the MC sample. The distributions are fitted with a straight line in the low-μ range. For all WPs, good linearity is observed over the fitted μ range and the slopes are in agreement with the slopes in the toy simulation studies. For larger μ values, deviations from linearity are visible for all WPs. These non-linearities will be assessed in the following sections.

Tracking efficiency

The track reconstruction times selection efficiency is measured in MC according to

$$\epsilon_{\text{r}\times\text{s}} = \frac{N_{\text{trk}}^{\text{primary}}}{N_{\text{true}}}, \tag{8.10}$$

where N_{true} is the number of true charged primary particles and $N_{\text{trk}}^{\text{primary}}$ is the number of primary tracks selected with a given WP.

Figure 8.10 shows the track reconstruction times selection efficiency for the various WPs, together with a fitted constant in the range $0 \leq \mu_{\text{truth}} \leq 20$. TightModLumi and TightModSiPlusLumi exhibit a larger overall efficiency which is the result of the restriction to $|\eta|$ below 1.0 where only nine Silicon hits are required. TightModLumi also shows a reduced μ-dependence which can be understood by studying the efficiencies as a function of the position in the bunch train. TightModLumi exhibits the highest and flattest efficiency as a function of μ, which is the main motivation for using it as the nominal track selection for the track-counting measurement in Run 2.

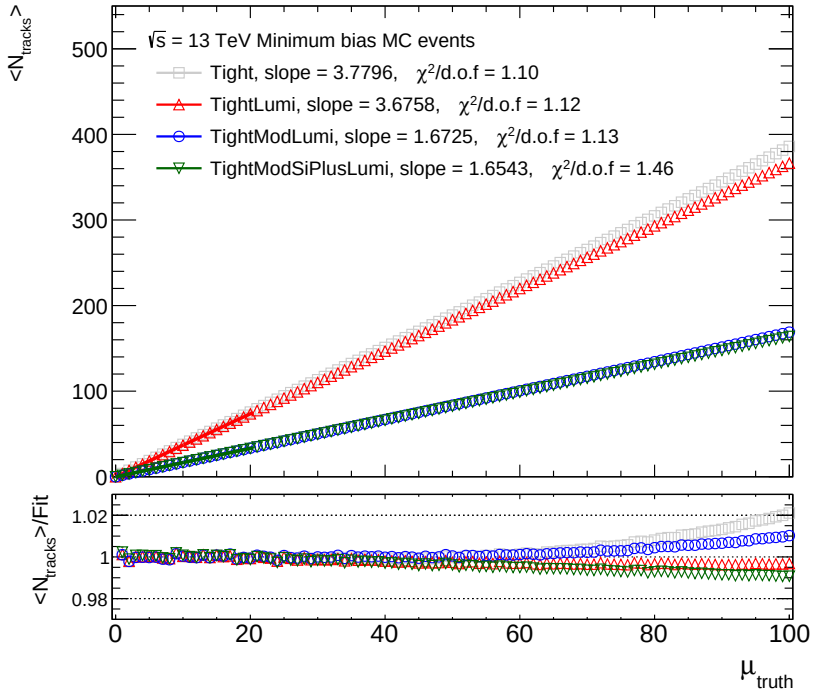

Figure 8.9: Average number of tracks per event as a function of μ_{truth}, for the various WPs, measured for the minimum bias events in the $Z \to \mu\mu$ MC sample. The points are fitted with a linear function over the range $0 \leq \mu_{\text{truth}} \leq 20$ and the slope of the fit and the χ^2/n_{dof} are indicated in the legend. The lower panel shows the ratio of the MC with respect to the fit.

Fake-track probability

The fake-track probability, or fake fraction, is defined as

$$f_{\text{fake}} = \frac{N_{\text{trk}}^{\text{fake}}}{N_{\text{trk}}}, \tag{8.11}$$

where N_{trk} is all tracks selected with a given WP and $N_{\text{trk}}^{\text{fake}}$ is the number of fake tracks passing the same WP selection.

Figure 8.11 shows the fake-track probability for the various WPs. All WPs show a reduced fake-track probability compared to the Tight selection, as a result of the additional requirement on the d_0-significance. The nominal WP, TightModLumi, shows a higher fake probability compared to TightLumi, due to the relaxed Pixel hole requirement. As a result of the requirement of an additional silicon hit, TightModSiPlusLumi exhibits the smallest fake probability. The fake-track contribution

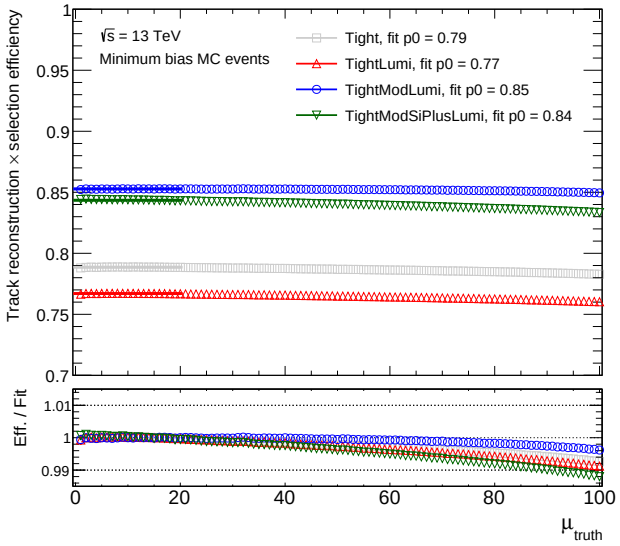

Figure 8.10: Track reconstruction times selection efficiency for the various WPs as a function of μ_{truth}, measured for the minimum bias events in the $Z \to \mu\mu$ MC sample. The points are fitted to a constant in the range $0 \leq \mu_{\text{truth}} \leq 20$, and the fitted constant is quoted in the legend. The lower panel shows the ratio of the MC to the fit.

to the TightModLumi measurement may become an important factor at μ values above 60 and require a dedicated treatment. However, in Run 2 most of the data was recorded at $\mu \approx 40$, where the fake probabilities only differ by 10-20 % between TightModLumi and TightLumi. The benefits of a higher and flatter efficiency for TightModLumi outweigh the drawbacks of the marginally higher fake-track contribution.

8.5 Performance of track counting

An essential asset of a good luminometer is the stability of the measurement over a long time during the data-taking year. One way to monitor the stability is to investigate luminosity ratios, either internally between several algorithms of the same luminometer, or with respect to external luminometers, that rely on completely different detector components. The sections below first discuss tools to monitor the detector conditions affecting the track reconstruction and then present the internal track-counting stability comparisons.

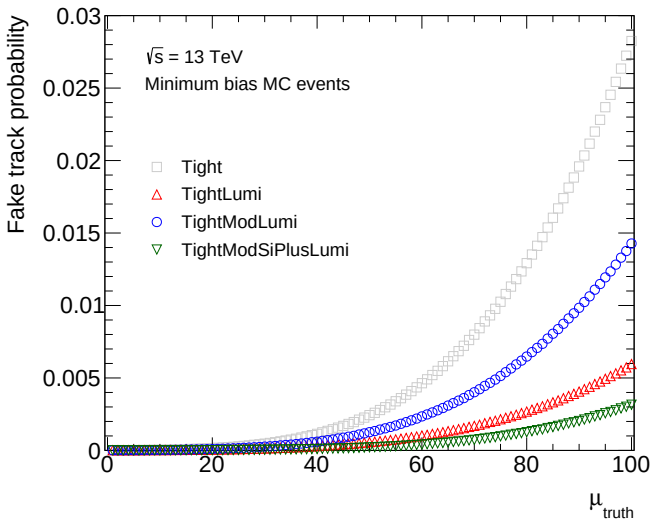

Figure 8.11: Fake-track probability for the various WPs, as a function of μ_{truth}, measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

8.5.1 Inner Detector monitoring

For each run analysed for track counting, a number of control distributions are produced to assess the quality of the track reconstruction. These monitor the number of tracks as a function of η and ϕ and provide a tool to discover changes in the ID detector conditions leading to tracking inefficiencies. As an example, Figure 8.12 shows the η-ϕ maps for tracks selected with the TightLumi WP for two runs taken in April and September 2018. The map from the beginning of the data-taking period shows several regions with reduced number of tracks. This was found to be a problem with inactive modules not being properly marked in the reconstruction software, meaning that more holes were recorded than should have been present. The issue was corrected in May, and the map from the end of the year shows a much smoother behaviour.

8.5.2 Tracking efficiency monitoring

Another tool used to monitor the performance of the track-counting measurement is data containing $Z \to \mu\mu$ decays. Muons are reconstructed with a high efficiency and high purity and are therefore suitable to isolate the effect of the track selection only. The muons used for the measurement are reconstructed as combined muons, with matched tracks in the ID and Muon Spectrometer, and are required to have p_{T} above 20 GeV. Muons are kept they can be paired with another muon in the event

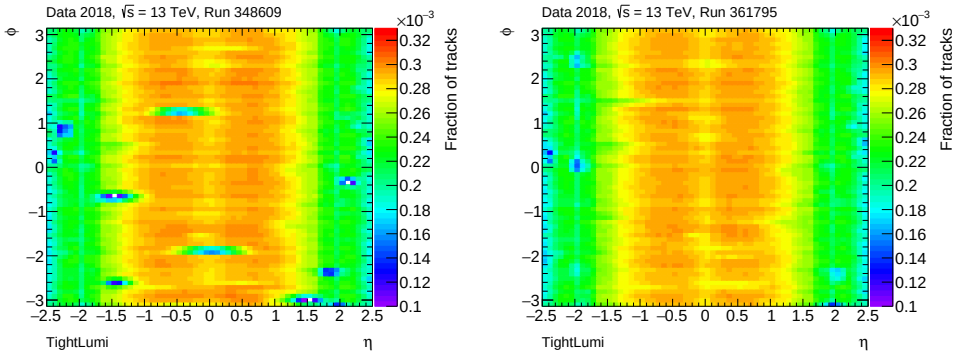

(a) (b)

Figure 8.12: Distribution in η and ϕ of the number of tracks selected with the TightLumi WP, normalised by the total number of tracks, for the 2018 runs 348609 (a) taken in April and 361795 (b) taken in September.

to form an invariant mass within 20 GeV of the nominal Z-boson mass. In addition, the tracks are required to pass the Loose selection, as was listed in Table 5.1. The fraction of ID tracks associated with these muons that also pass the track-counting WPs gives a measure of the track selection efficiency.

Figures 8.13 and 8.14 show the efficiency measured in $Z \to \mu\mu$ events for each Run 2 data-taking year, as a function of pileup and data-taking date respectively. TightModLumi and TightModSiPlusLumi show a relatively flat behaviour as a function of pileup while the TightLumi WP shows a reduction in the efficiency with increasing $\langle\mu\rangle$. The same observations are made in the MC simulation studies and the reason for the drop in efficiency is discussed in Section 8.6. The $Z \to \mu\mu$ efficiencies measured for TightModLumi show excellent stability over time in all years, another motivation for selecting it as the nominal WP for the Run 2 luminosity measurement. For TightLumi, larger variations from run to run can be observed. This is expected since the pileup changes between runs, and the efficiency for TightLumi has the largest dependency on $\langle\mu\rangle$. The TightLumi 2016 efficiency as a function of date shows a steady decrease. This is correlated with the increase in pileup throughout the data-taking year and the pileup dependence of the efficiency. In 2017, the TightLumi WP shows a decrease in efficiency at the end of the data-taking year. This coincides with the introduction of the 8b4e filling scheme which involved higher pileup values, as discussed in Section 7.9. In 2018, the TightLumi WP exhibits a clear trend, with lower efficiency in the first part of the data-taking year. This is due to the issue with inactive Pixel modules. TightModLumi and TightModSiPlusLumi are insensitive to the issue thanks to the relaxed Pixel hole

veto, and hence show no effect in the efficiency.

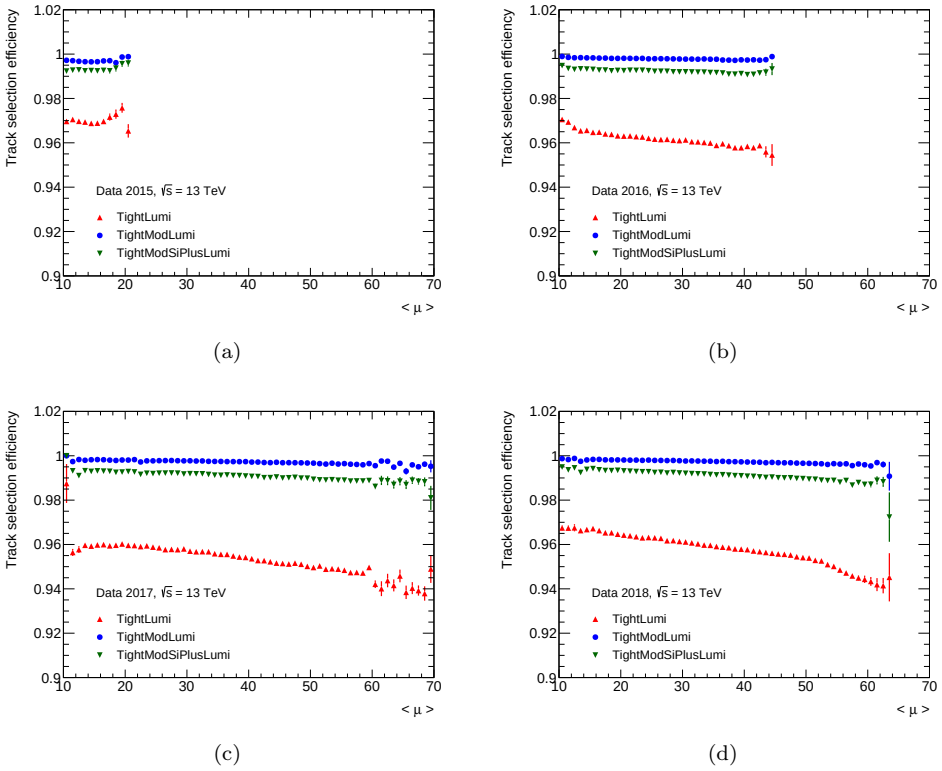

Figure 8.13: Track selection efficiency for TightLumi, TightModLumi, and TightModSiPlusLumi, measured in $Z \to \mu\mu$ events, as a function of $\langle\mu\rangle$ in 2015 (a), 2016 (b), 2017 (c), and 2018 (d). The value of $\langle\mu\rangle$ is determined from the offline LUCID measurement.

8.5.3 Internal long-term stability

The track-counting stability is monitored internally from ratios of the various WPs as functions of the data-taking date. These ratios will be sensitive to changes in the detector conditions that affect the WPs differently over time. The internal track-counting stability is shown in Figure 8.15 with TightModLumi as reference. For 2015 and 2016, the ratios are rather stable throughout the year with the same decrease in the TightLumi ratio for 2016 as observed in the Z efficiencies. Similarly, the 2017 and 2018 ratios exhibit drops in the TightLumi ratio in the end and beginning of the years, for the same reasons as the drops in the Z efficiency.

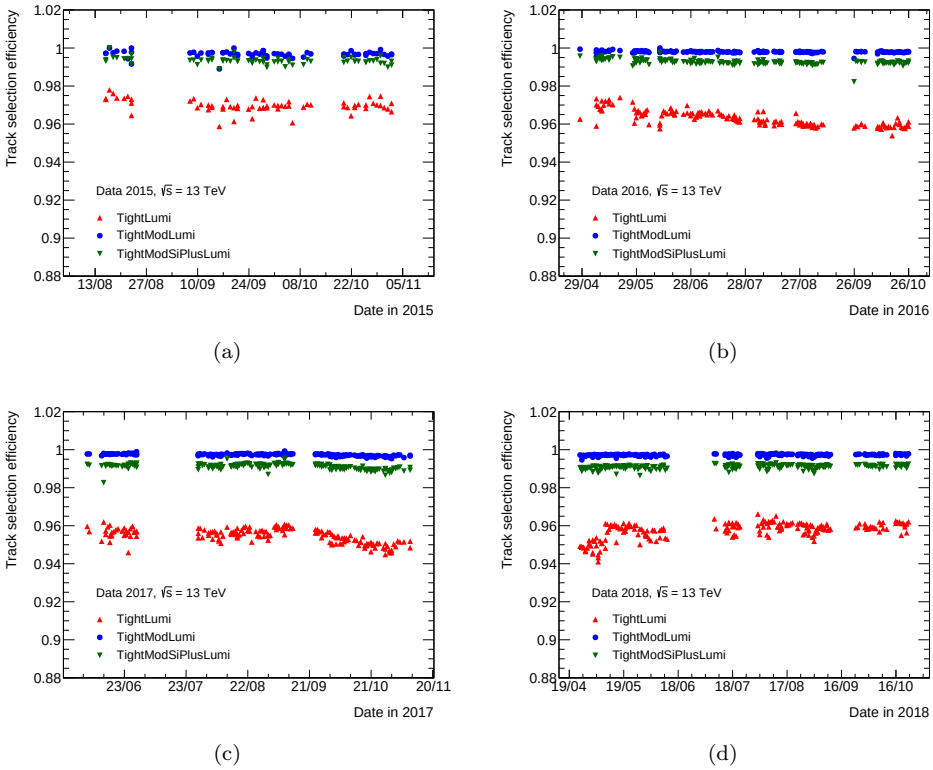

Figure 8.14: Track selection efficiency for TightLumi, TightModLumi, and TightModSiPlusLumi, measured in $Z \to \mu\mu$ events, as a function of data-taking date in 2015 (a), 2016 (b), 2017 (c), and 2018 (d). The value of $\langle\mu\rangle$ is determined from the offline LUCID measurement.

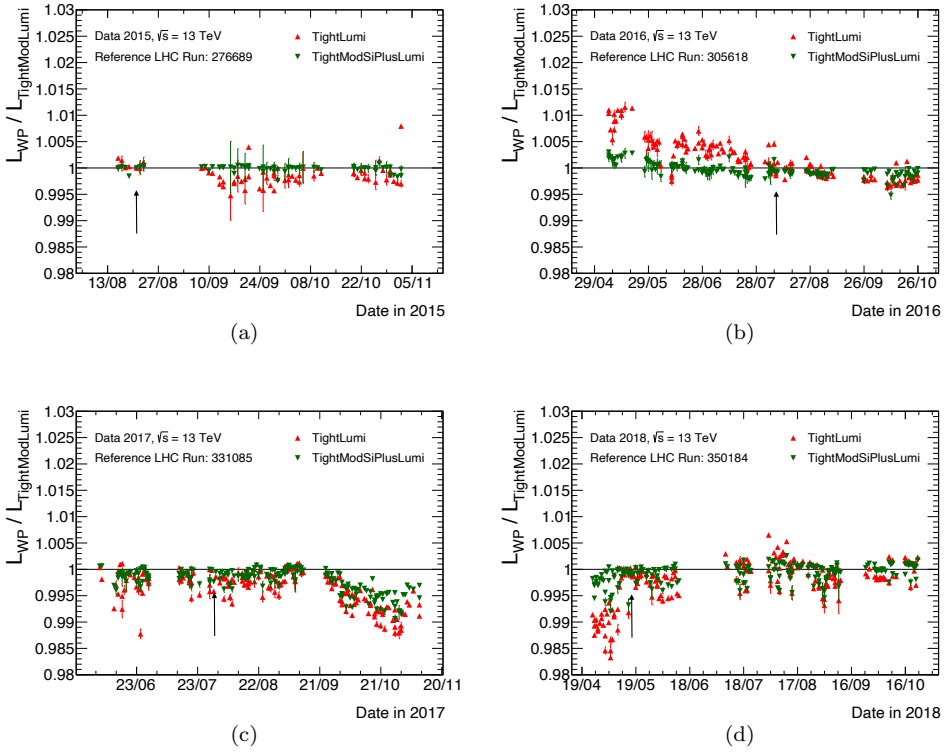

Figure 8.15: Ratio of TightLumi and TightModSiPlusLumi to TightModLumi for track-counting data as a function of data-taking date throughout 2015 (a), 2016 (b), 2017 (c), and 2018 (d). The WPs are normalised to TightModLumi in one selected run which is indicated by the arrow in each figure.

8.6 Bunch-train effects

Several of the effects seen in the tracking efficiencies and fake rates are related to the performance of track counting as a function of the position in the bunch train. The $Z \to \mu\mu$ MC sample simulates 25 ns bunch trains with 48 consecutive bunches. Figure 8.16 displays the average tracking efficiencies for the track-counting WPs as a function of the position in the train. A clear drop can be seen over the first two bunches, before the efficiency stabilises at a constant value. In order to understand this behaviour, the tracking efficiency is studied as a function of μ_{truth} for different positions in the bunch train. The result is shown in Figure 8.17 for TightLumi and TightModLumi. For both WPs, the efficiency for the first bunch increases with μ_{truth}. This is the result of a combination of reconstruction and selection

effects. In the reconstruction, the likelihood that a random hit fills up a hole in a track increases with pileup. If the resulting track still has a high truth match probability it will be counted as a primary track and lead to an increase in the measured efficiency. For the track selection, various effects may affect the efficiency as a function of pileup. Figure 8.18 shows the number of hits and inactive sensors per track in the Pixel detector and SCT as a function of μ_{truth}. The number of inactive sensors is seen to increase with μ_{truth}, meaning that the the track is more likely to pass the requirement on the number of Silicon hits which includes the number of inactive sensors.

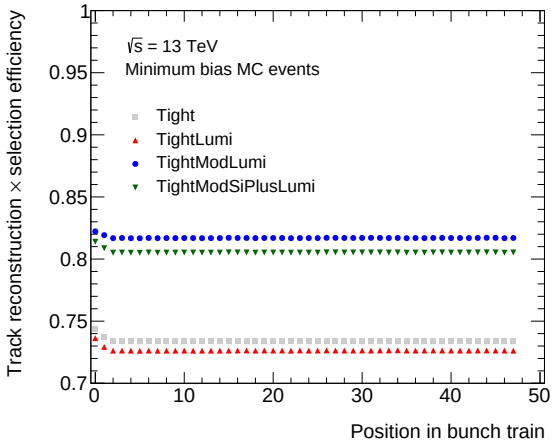

Figure 8.16: Track reconstruction times selection efficiency for tracks passing the track-counting WPs, integrated over μ_{truth}, as a function of the position in the bunch train. The efficiency is measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

The second bunch in the train shows a flatter behaviour while later bunches show a clear drop in the efficiency with increasing μ_{truth}, and are all in agreement with each other. This reduction can be partly explained by the read-out mode of the SCT in the Run 2 data taking [119], where channels which exhibit a signal above threshold in the current bunch crossing are only read out if they have no hit in the preceding bunch crossing. For increasing pileup, the probability to have SCT hits in the same channel in consecutive bunch crossings increases. The read-out mode can then lead to a rejection of SCT hits, resulting in an increase of the number of SCT holes on a track. In the case of bunch trains, the first bunch has no probability for a preceding hit, so the read-out efficiency is not reduced. The second bunch has a non-negligible probability for a preceding hit, as do the following bunches. Starting from the second bunch, the SCT read-out effect therefore counteracts the increase with μ_{truth} seen in the first bunch. Figure 8.19 shows the average number of Pixel and SCT hits per track in the MC sample, as a function of the bunch-train

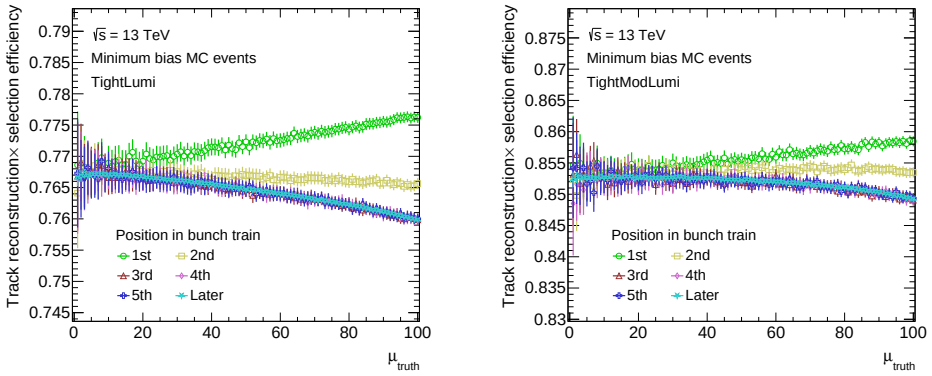

(a) (b)

Figure 8.17: Track reconstruction times selection efficiency for tracks passing the TightLumi (a) or TightModLumi (b) WPs, as a function of μ_{truth}, for different positions in the bunch train. The efficiency is measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

position, for tracks selected with the TightModLumi WP. The number of Pixel hits is approximately flat as a function of μ_{truth}, while the number of SCT hits decreases, clearly showing the SCT read-out effect.

Figure 8.20 shows the average fake-track probability as a function of the bunch-train position and Figure 8.21 displays the fake rate as a function of pileup for different positions in a bunch train. For all WPs, the first bunch in the train shows a smaller fake probability compared to other positions in the train due to the SCT read-out mode.

The bunch-train dependence in data is studied with the same Z-boson events as were used for the efficiency measurement described in Section 8.5.2. In the 2018 physics data taking, the LHC operated with 25 ns bunch trains with 48 bunches in each train, similar to the simulated structure in the $Z \to \mu\mu$ MC sample. Figure 8.22a shows the muon track selection efficiency for the TightLumi, TightModLumi and TightModSiPlusLumi WPs, as a function of the bunch-train position in three different trains for a subset of the 2018 dataset. For TightModLumi and TightModSiPlusLumi, the trends are in agreement with the MC efficiencies in Figure 8.16, with a drop over the first couple of bunches. Figure 8.22b shows the Z-boson efficiencies measured for a subset of the 2017 dataset, corresponding to the 8b4e filling scheme, which exhibits the same behaviour.

For TightLumi, the reduction in the efficiency continues throughout the train in data, while it is flat after the first two bunches in MC. This difference is understood

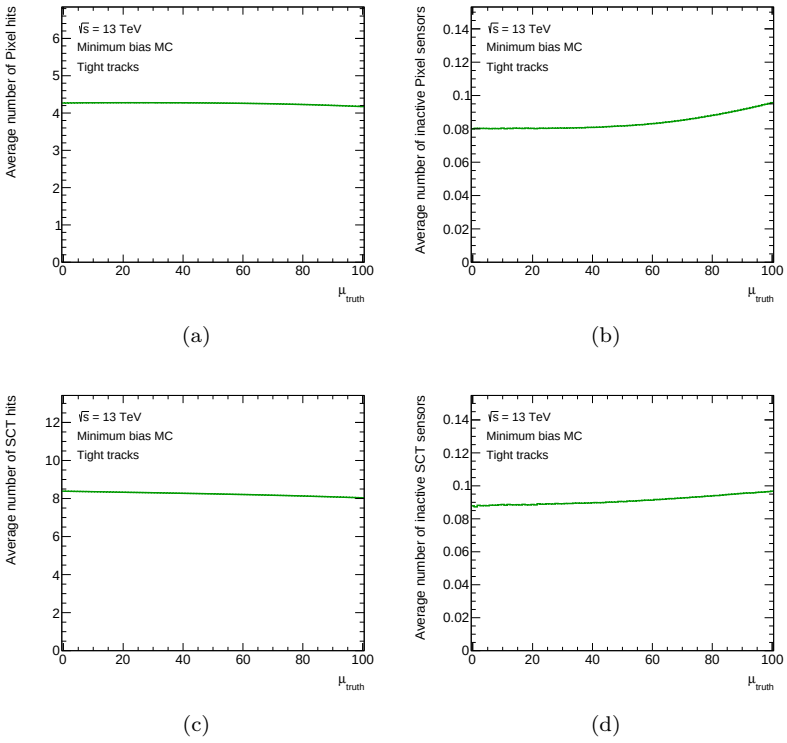

Figure 8.18: Number of hits and inactive sensors per track in the Pixel detector (a,b) and in the SCT (c,d), for tracks selected with the Tight WP, as a function of μ_{truth}. The numbers are measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

to be related to the pixel hole veto in combination with the so-called time-over-threshold (ToT) effect in the Pixel detector. When a pixel is hit twice in a row, with the second hit arriving when the signal is still over threshold, the second hit will not be registered. Instead, a hole will be attributed to the corresponding track. Within a bunch train of consecutive filled bunches, the probability that a pixel has been hit in the preceding bunch crossings, and is over threshold, increases. Saturation of this effect is reached when the time between the start of the train and a considered bunch crossing is larger than the ToT and some pixels have recovered. The resulting effect on the track selection can be seen in the efficiency measured for the TightLumi WP, which contains a veto against holes in the Pixel detector. Based on the comparison between the MC and the data efficiencies, it is clear that the SCT read-out effect is modelled in MC, while the ToT effect is missing. This is an important observation for the comparison of the non-linearities between MC and data, presented in Section 8.7.

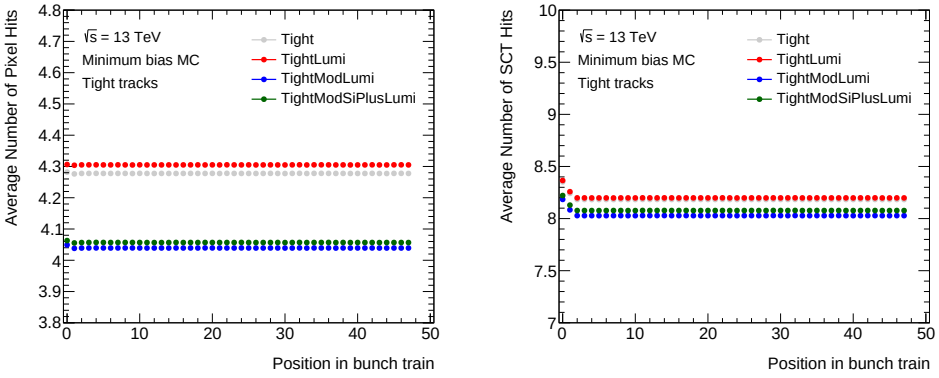

(a) (b)

Figure 8.19: Average number of hits in the Pixel (a) and SCT (b) detectors as a function of the bunch-train position, for the various WPs. The numbers are measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

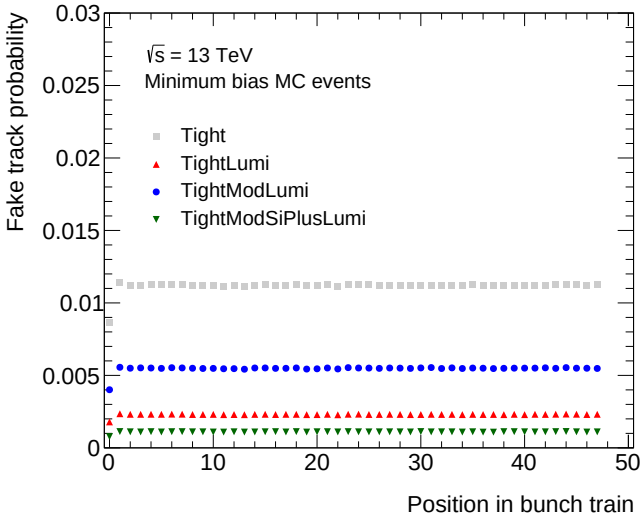

Figure 8.20: Fake-track probability for tracks passing the various WPs, integrated over μ_{truth}, as a function of the position in the bunch train, measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

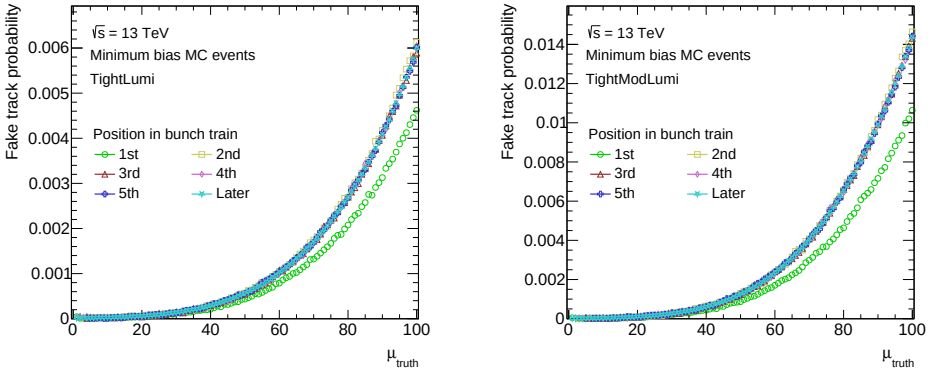

(a) (b)

Figure 8.21: Fake probability for tracks passing the TightLumi (a) or TightModLumi (b) WPs, as a function of μ_{truth}, for different positions in the bunch train. The fake-track probability is measured for the minimum bias events in the $Z \to \mu\mu$ MC sample.

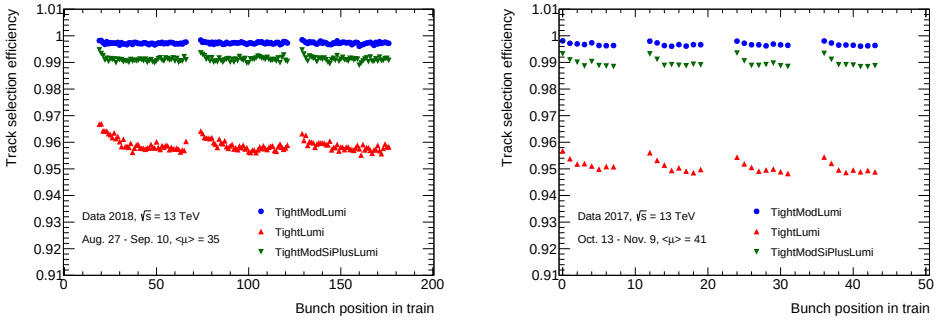

(a) (b)

Figure 8.22: Track selection efficiency measured in $Z \to \mu\mu$ events as a function of position in the bunch train for a subset of the 2018 (a) and 2017 (b) datasets. The 2018 data corresponds to the standard 25 ns running while the 2017 data corresponds to the 8b4e running.

8.7 Studies of fake tracks in data

All monitoring studies show that the track-counting measurement, using the TightModLumi WP, is stable and highly linear with μ in Run 2. The systematic uncertainty assigned to the measurement is to date estimated entirely based on external comparisons to other luminometers in the offline luminosity analysis. A goal for the future is to derive a stand-alone uncertainty for the track-counting measurement, to help reduce the overall uncertainty on the luminosity. Depending on the pileup conditions in Run 3 it may also be necessary to correct for fake tracks in the measurement. A stand-alone uncertainty estimate and a potential fake-track correction require a detailed understanding of the various effects that impact the track-counting measurement and its pileup dependence. Some uncertainties are already determined independently by the tracking conditions monitoring and the $Z \to \mu\mu$ track selection efficiencies. A dedicated measurement of the fake-track probability has however not yet been performed on data.

An approximate estimate of the fake-track probability can be achieved by studying the non-linearity in the average number of tracks as a function μ, similar to the MC study shown in Figure 8.9. Based on the findings of the MC studies, this must however be done in a μ range well above 50 in order for the fake tracks to dominate the non-linearity in the measurement. In addition, the data must correspond to a wide range of μ values in order to fully capture the pileup dependence. Most efficiency non-linearities are caused by various bunch-train effects. In order to limit such dependencies, only runs with individual bunches or alternatively runs with trains, where only the first bunch in the train is used, are suitable for the studies. A few runs from the Run 2 dataset fulfil these requirements. The first selected run corresponds to the 2018 μ scan, described in Section 8.2.1, which used a filling scheme with individual bunches and in which the $\langle\mu\rangle$ varied between 0.001-50. The second run was taken in 2016 and corresponds to a special LHC filling scheme with three individual bunches and two 25 ns trains with 48 bunches each. For the individual bunches, a pileup of 130 was achieved, while a pileup of around 50 was achieved for the trains.

In MC, the measured average number of tracks can be studied as a function of the true μ. In data, the situation is more complicated and $\langle\mu\rangle_{\text{vis}}^{\text{TC}}$ must be studied as a function of a measured μ value, which in turn also incorporates any non-linearity. A possible solution is to study ratios of the various WPs with respect to TightModLumi, as a function of μ measured with TightModLumi. The WP ratios for the 2018 and 2016 special runs are shown in Figures 8.23a and 8.23b. The ratios are normalised to one at low μ. Figure 8.24 shows the corresponding ratios of WPs measured in the $Z \to \mu\mu$ sample, put on the same scale as the data, to allow for a comparison. The MC ratios are also expressed with respect to the $\langle\mu\rangle$ measured with the TightModLumi WP in MC. The comparison between the data and the MC suggests that the fake-track probability is underestimated in MC. From these ratios it is however not possible to deduce the exact normalisation and shape differences,

as the non-linearity depends on the interplay between the efficiencies and fake-track probabilities between the various WPs.

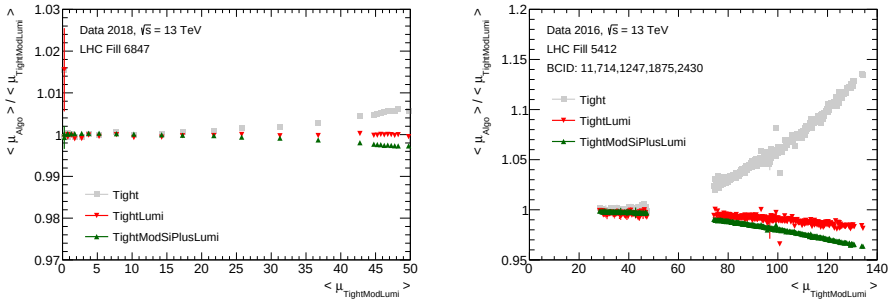

(a) (b)

Figure 8.23: Ratio of $\langle\mu\rangle$ measured in data with various WPs to TightModLumi, as a function of $\langle\mu\rangle$ measured with TightModLumi. The data corresponds to the LHC fill 6847 in which a μ scan took place (a) and to the LHC fill 5412 using only BCIDs corresponding to individual bunches and the first bunch in each train (b).

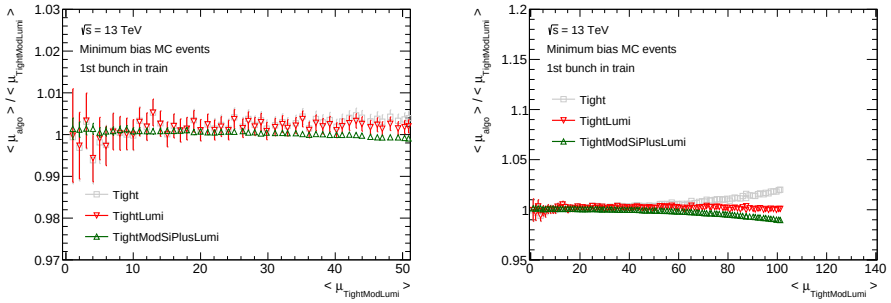

(a) (b)

Figure 8.24: Ratio of $\langle\mu\rangle$ measured in MC with various WPs to TightModLumi, as a function of $\langle\mu\rangle$ measured with TightModLumi. The x axis is restricted to 50 (a) and 140 (b) to allow for a comparison to data.

Part IV

Searches for new exotic particles

Chapter 9

Search for squarks and gluinos in a dilepton final state

Searches for SUSY particles at the LHC are motivated by the naturalness limit, which constrains the masses of the sparticles to be at the TeV scale. This chapter focuses on the search for sparticles predicted by R-parity conserving SUSY. In such models, the sparticles are produced in pairs in the pp collisions and the LSP is stable and escapes detection, resulting in final states with SM particles and large missing transverse momentum.

The production mechanisms for SUSY particles at the LHC were discussed in Section 2.2 and the predicted cross sections for pair-production of various SUSY particles were shown in Figure 2.5. Production via the strong force has the largest cross section and gluino and squark pairs are therefore expected to be the most abundantly produced SUSY particles at the LHC. Final states for such processes contain energetic jets, $E_{\mathrm{T}}^{\mathrm{miss}}$ and sometimes leptons. Electroweak SUSY production has a much smaller cross section and is typically characterised by multi-lepton and $E_{\mathrm{T}}^{\mathrm{miss}}$ final states. In ATLAS, the SUSY searches are grouped around the different production modes and decay topologies. The Strong-2L analysis searches for signs of squark or gluino production in events containing two same-flavour opposite-charge leptons, jets, and $E_{\mathrm{T}}^{\mathrm{miss}}$, and uses the $36.1\,\mathrm{fb}^{-1}$ pp dataset collected by ATLAS in 2015 and 2016[1].

[1]The Strong-2L analysis was conducted in 2017 and used preliminary luminosity values. The final Run 2 offline luminosity analysis presented in Chapter 7 estimates the 2015+2016 dataset to be $36.2\,\mathrm{fb}^{-1}$.

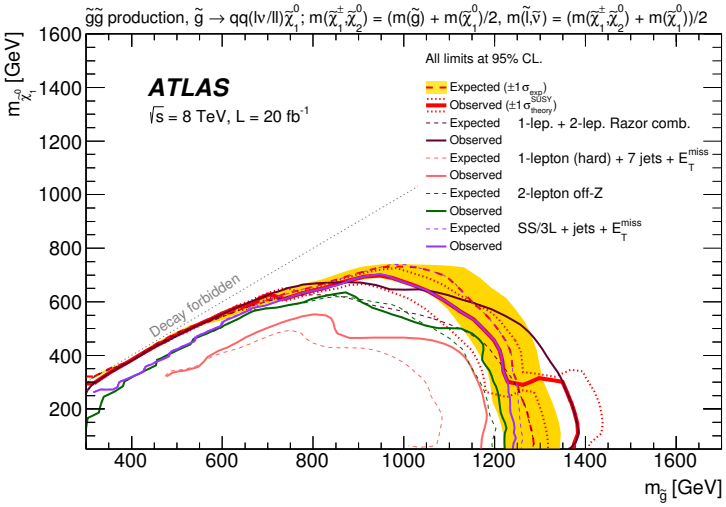

Figure 9.1: The results at the beginning of Run 2, showing 95 % CL exclusion limits in the $(m_{\tilde{g}}, m_{\tilde{\chi}_1^0})$ plane for a simplified model of gluino pair production with a two-step decay into jets, leptons, and missing transverse momentum via sleptons [120]. The solid red line and the dashed red line show respectively the combined observed and expected exclusion limits. Observed and expected limits from the individual searches are also shown for comparison.

The sensitivity to a given SUSY model depends on the production cross section as well as the integrated luminosity of the analysed dataset. Figure 9.1 shows exclusion limits on gluino and neutralino masses in a simplified SUSY model, set by several 8 TeV ATLAS Run 1 searches. This illustrates the knowledge at the start of the LHC Run 2, with gluino masses up to almost 1.4 TeV excluded by the combined result. When moving to the Run 2 centre-of-mass energy of 13 TeV, the gluino production cross section is predicted to increase by over an order of magnitude. Even just a small part of the Run 2 dataset is therefore enough to reach sensitivities well above those of Run 1 with similar search strategies and background rates. This is an important motivation for the Strong-2L analysis which was one of the first Run 2 searches for strongly produced SUSY.

9.1 Analysis overview

The SUSY scenarios targeted by the Strong-2L analysis are shown in Figures 9.2 and 9.3. All models involve pair-produced squarks or gluinos decaying into the second lightest neutralino which further decays into two same-flavour opposite-charge leptons and the lightest neutralino. The kinematics of the leptons in the

final state depend on the mass difference between the two neutralinos, and their invariant mass spectrum is therefore used to characterise the various signal models. Two complementary search strategies are defined; the *On-Z* search and the *Edge* search.

The On-Z search targets the signal scenarios shown in Figure 9.2 where the neutralino mass difference is larger than the mass of the Z boson. In this case, the same-flavour opposite-charge lepton pair can be produced from an on-shell Z boson in the decay

$$\tilde{\chi}_2^0 \to Z\tilde{\chi}_1^0 \to \ell^+\ell^-\tilde{\chi}_1^0. \tag{9.1}$$

The process is characterised by a peak in the dilepton invariant mass distribution at the Z-boson mass. The Edge search considers the two signal models shown in Figures 9.3a and 9.3b, which are referred to as the Z^* and *Slepton* models respectively. In the Z* model, the neutralino mass difference is smaller than the Z-boson mass. The same-flavour opposite-charge lepton pair is then produced via an off-shell Z boson in the decay

$$\tilde{\chi}_2^0 \to Z^*\tilde{\chi}_1^0 \to \ell^+\ell^-\tilde{\chi}_1^0. \tag{9.2}$$

This process leads to a dilepton invariant mass spectrum that is rising up to a kinematic endpoint, given by

$$m_{\ell\ell}^{\max} = m_{\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0}, \tag{9.3}$$

that occurs below the Z-boson mass. Finally, the Slepton model considers the production of sleptons through the decay chain

$$\tilde{\chi}_2^0 \to \tilde{\ell}^{\pm}\ell^{\mp} \to \ell^+\ell^-\tilde{\chi}_1^0. \tag{9.4}$$

Such decays are allowed if the masses of the sleptons are lower than the $\tilde{\chi}_2^0$ mass and lead to an invariant mass spectrum with a kinematic endpoint given by

$$m_{\ell\ell}^{\max} = \sqrt{(m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\ell}}^2)(m_{\tilde{\ell}}^2 - m_{\tilde{\chi}_1^0}^2)/m_{\tilde{\ell}}^2}, \tag{9.5}$$

that may occur at any mass. Typical shapes of the dilepton invariant mass spectra for the On-Z and two Edge signal scenarios are shown in Figure 9.4.

All considered models are expected to contain at least two jets initiated by the quarks in the two $\tilde{\chi}_2^0$ decays. The $\tilde{\chi}_1^0$ is the LSP and escapes the ATLAS detector leaving $E_{\mathrm{T}}^{\mathrm{miss}}$ in the final state. Since the SUSY particles are expected to be heavy, both the jet energies and the $E_{\mathrm{T}}^{\mathrm{miss}}$ are expected to be large compared to SM processes. Events selected in this analysis are therefore required to contain an opposite-charge electron or muon pair, two or more highly energetic jets and large $E_{\mathrm{T}}^{\mathrm{miss}}$ in the final state. The On-Z search is performed as a counting analysis in a window around the Z-boson mass whereas the Edge search is carried out across the full invariant dilepton mass spectrum.

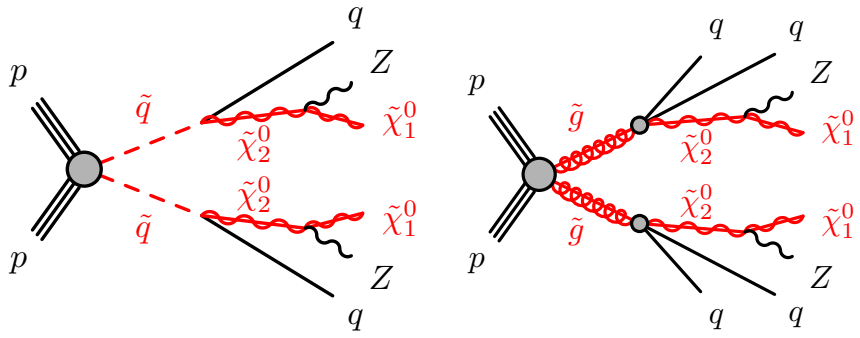

(a) (b)

Figure 9.2: Decay topologies for the signal scenarios involving squark (a) and gluino (b) pair production considered for the On-Z search.

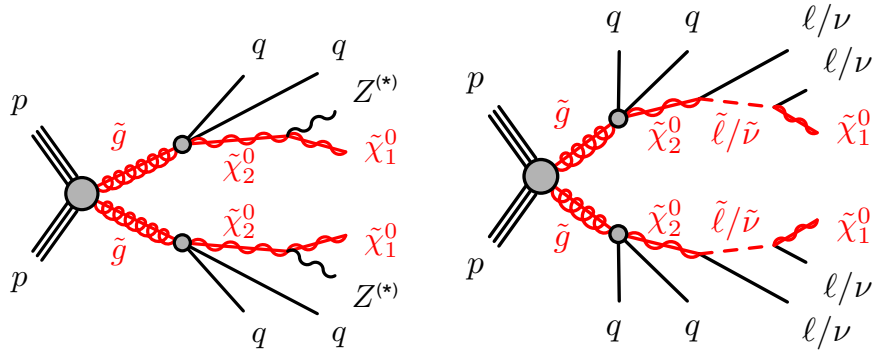

(a) (b)

Figure 9.3: Decay topologies for the Z* (a) and Slepton (b) signal scenarios considered for the Edge search.

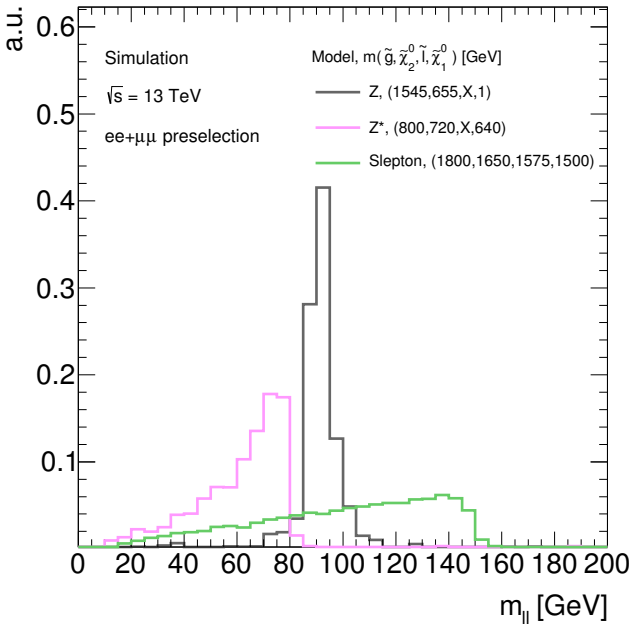

Figure 9.4: Dilepton invariant mass spectrum for three example signal models for the On-Z, Z*, and Slepton signal scenarios considered in the analysis. Events are selected with the preselection requirements detailed in Section 9.4.2, and are required to contain an electron or muon pair.

There are several SM processes that can mimic the signals targeted by the analysis. The largest backgrounds result from production of top-antitop quark pairs ($t\bar{t}$), diboson production ($WZ/ZZ/WW$), and Drell-Yan processes in which a quark and an antiquark annihilate to produce a virtual photon or Z boson (Z/γ^*). Other processes contributing to the SM background at lower rates are single top production (Wt), and the associated productions $t\bar{t}+W$, $t\bar{t}+Z$ and $t\bar{t}+WW$. The analysis event selection is optimised to reject as much as possible of these backgrounds while keeping a high acceptance for the considered signals.

Searches for SUSY in final states with a same-flavour opposite-charge lepton pair have previously been performed by ATLAS [5,121] and CMS [122–124] at 8 TeV and with the very first part of the 13 TeV dataset. The ATLAS Edge search performed at 8 TeV showed no significant excess above the predicted SM background, and exclusion limits were set on the gluino and neutralino masses in the Slepton model. These limits are included in the combined result in Figure 9.1. The On-Z search at 8 TeV however observed an excess of events above the SM background with a significance of 3σ [121]. In order to further assess the excess, the event selection criteria were kept almost identical when the search was repeated with the first $14.7\,\text{fb}^{-1}$ of 13 TeV Run 2 data [5]. In CMS, the searches have a similar history, with a slight excess in the 8 TeV analysis [122], and another iteration with the very first 13 TeV data using the same strategy [123]. The excess seen by CMS was however in the Edge search, and neither of the excesses were confirmed in the initial 13 TeV dataset. The analysis presented in this chapter is performed on a larger 13 TeV dataset and has been reoptimised to increase the sensitivity to the signal models at the Run 2 centre-of-mass energy.

9.2 Data and MC samples

The analysis is performed on $36.1\,\text{fb}^{-1}$ of 13 TeV pp collision data recorded during 2015 and 2016 by the ATLAS detector. Events are collected using a combination of lepton triggers which are further discussed in Section 9.4.1. Simulated samples are used to define the analysis strategy, for the interpretation of the analysis results, and for the estimation of some of the SM backgrounds. Table 9.1 lists the signal and background samples. All samples are overlaid with pileup simulated with Pythia 8.186 [104] with the A2 underlying event tune [125] and the MSTW 2008 PDF set [126].

The signal scenarios shown in Figures 9.2 and 9.3 represent SUSY-inspired simplified models. Signal grids are generated by varying the gluino or squark mass and either of the $\tilde{\chi}_1^0$ or $\tilde{\chi}_2^0$ masses. For the On-Z search, three different signal grids are considered. Two of the grids, one with gluino-pair and one with squark-pair production, have the $\tilde{\chi}_1^0$ mass fixed to 1 GeV and vary the gluino or squark mass and the $\tilde{\chi}_2^0$ mass. In the third grid, the $\tilde{\chi}_2^0$ mass is fixed to 100 GeV above the $\tilde{\chi}_1^0$ mass in order to target scenarios with higher-mass LSPs while still allowing decays

Table 9.1: Simulated samples used in the Strong-2L analysis, with the corresponding matrix element and parton shower generators, cross-section calculation order, set of tuned parameters, and PDF set. Table adapted from Ref. [4].

Process	**Generator**	**Parton shower**	**Cross section**	**Tune**	**PDF set**
Signal	MG5_aMC@NLO [99]	Pythia 8.186 [104]	NLO+NLL [127–131]	A14 [132]	NNPDF2.3LO [25]
$t\bar{t}+W$, $t\bar{t}+Z$	MG5_aMC@NLO	Pythia 8.186	NLO [133, 134]	A14	NNPDF2.3LO
$t\bar{t}+WW$	MG5_aMC@NLO	Pythia 8.186	LO [99]	A14	NNPDF2.3LO
$t\bar{t}$	Powheg Box v2 [101–103]	Pythia 6.428 [94]	NNLO+NNLL [135, 136]	Perugia2012 [137]	NLO CT10 [138]
Wt	Powheg Box v2	Pythia 6.428	Approx. NNLO [139]	Perugia2012	NLO CT10
WW, WZ, ZZ	Sherpa 2.2.1 [107]	Sherpa 2.2.1	NLO [140]	Sherpa default	NNPDF3.0nnlo [141]
Z/γ^*	Sherpa 2.2.1	Sherpa 2.2.1	NNLO [142]	Sherpa default	NNPDF3.0nnlo

into on-shell Z bosons. The grid is generated by varying the gluino and $\tilde{\chi}_1^0$ masses. For the two Edge models, the mass of the $\tilde{\chi}_2^0$ is fixed to the average of the gluino and $\tilde{\chi}_1^0$ masses which vary in the grid. This mass splitting is chosen to enhance the topological differences between the considered models and other simplified models with only one intermediate particle between the gluino and the LSP [143]. The Z^* model can result in either on-shell or off-shell Z bosons depending on the mass splitting between the neutralinos in the decay. Table 9.2 summarises the various signal grids used in the analysis.

Table 9.2: Summary of the signal grids used in the analysis. The symbols x and y denote the x-y plane across which the sparticle masses are varied to construct the signal grid. For the Slepton model, the masses of the superpartners of the left-handed leptons are given by $[m(\tilde{\chi}_2^0) + m(\tilde{\chi}_1^0)]/2$, while the superpartners of the right-handed leptons are decoupled. The three slepton flavours are taken to be mass-degenerate. For the gluino or squark decays, the models have equal branching fractions for the quark flavours listed. Table adapted from Ref. [4].

Model	**Production**	**Quark flavours**	$m(\tilde{g}/\tilde{q})$	$m(\tilde{\chi}_2^0)$	$m(\tilde{\chi}_1^0)$
On-shell $\tilde{g}$-$\tilde{\chi}_1^0$	$\tilde{g}\tilde{g}$	u, d, c, s	x	$m(\tilde{\chi}_1^0) + 100$ GeV	y
On-shell $\tilde{g}$-$\tilde{\chi}_2^0$	$\tilde{g}\tilde{g}$	u, d, c, s	x	y	1 GeV
On-shell $\tilde{q}$-$\tilde{\chi}_2^0$	$\tilde{q}\tilde{q}$	u, d, c, s	x	y	1 GeV
Z^*	$\tilde{g}\tilde{g}$	u, d, c, s, b	x	$[m(\tilde{g}) + m(\tilde{\chi}_1^0)]/2$	y
Slepton	$\tilde{g}\tilde{g}$	u, d, c, s, b	x	$[m(\tilde{g}) + m(\tilde{\chi}_1^0)]/2$	y

9.3 Identification and selection of physics objects

The physics objects used in the analysis are electrons, muons, jets, and photons, for which the identification and reconstruction is described in Section 5.2. All

analysis objects are first required to pass a *baseline* selection which is used for the calculation of $E_{\mathrm{T}}^{\mathrm{miss}}$ and for the overlap removal procedure. The requirements are then tightened for the final *signal* objects used in the analysis.

All baseline electrons are required to reside within $|\eta| < 2.47$, have $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ and satisfy the Loose likelihood quality criteria described in Ref. [88]. The requirement on pseudorapidity ensures that all tracks are within the acceptance of the ID and that the associated clusters are within the high-granularity portion of the calorimeter. Signal electrons must have $p_{\mathrm{T}} > 25\,\mathrm{GeV}$ and pass the Medium likelihood quality criteria, as well as the 95 % efficient isolation requirements described in Ref. [88]. The associated tracks are further required to point to the PV by imposing $|d_0|/\sigma_{d_0} < 5$ and $|z_0 \sin\theta| < 0.5\,\mathrm{mm}$.

All baseline muons are required to have $p_{\mathrm{T}} > 10\,\mathrm{GeV}$, pass the Medium likelihood quality requirements described in Ref. [89], and reside within $|\eta| < 2.5$, which corresponds to the acceptance region of the ID. Signal muons must have $p_{\mathrm{T}} > 25\,\mathrm{GeV}$, and the tracks should satisfy $|d_0|/\sigma_{d_0} < 3$ and $|z_0 \sin\theta| < 0.5\,\mathrm{mm}$. The 95 % efficient isolation criteria described in Ref. [89] are also applied to the muons.

All baseline-level jets are required to have $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and reside within $|\eta| < 4.5$. Signal jets are further required to reside within $|\eta| < 2.5$ and have $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. The tightened requirement on the pseudorapidity allows for the use of tracking information to select jets associated to the PV by means of the JVT [93].

Photon candidates are required to satisfy the Tight selection criteria described in Ref. [144], have $p_{\mathrm{T}} > 25\,\mathrm{GeV}$ and reside within the region $|\eta| < 2.37$. Signal photons are further required to have $p_{\mathrm{T}} > 50\,\mathrm{GeV}$ and to be isolated from other objects in the event, according to p_{T}-dependent requirements on both track-based and calorimeter-based isolation. Photons are only used in the analysis to estimate the Drell-Yan background and for the calculation of the $E_{\mathrm{T}}^{\mathrm{miss}}$.

All MC samples have MC-to-data corrections applied to take into account small differences between data and MC simulation in identification, reconstruction and trigger efficiencies. The p_{T} values of leptons in the MC samples are additionally smeared to match the momentum resolution in data. All MC correction factors are derived centrally in ATLAS by the Combined Performance groups.

9.4 Event selection

The searches are performed in signal regions (SRs) designed to maximise the selection efficiency of signal events and minimise the SM backgrounds. In addition, several control regions (CRs) are used for the estimation of the SM backgrounds. These regions are designed to be enriched in the SM process of interest and to have a low signal contamination. Dedicated validation regions (VRs) are then used to validate the background estimates. These are designed to be kinematically close

to the SRs and CRs. The following sections present the trigger strategy, the preselection of events, and the discriminating variables that are used to suppress the background and divide preselected events into the various analysis regions.

9.4.1 Trigger strategy

Data events are collected using a combination of triggers chosen to provide a good efficiency in each of the SRs, VRs, and CRs. The analysis primarily depends on triggers on leptons. In order to simplify the application of the trigger scale factors, which correct the MC to match the data efficiencies for a given trigger, the events are split into a series of kinematic ranges, with designated triggers. The regions are listed in Table 9.3 together with the details of the corresponding triggers. A combination of unprescaled single- and dilepton triggers is employed. The dilepton triggers select ee, $\mu\mu$, or $e\mu$ events and have lower lepton-p_{T} thresholds than the single electron and muon triggers. In kinematic regions where single-lepton triggers are fully efficient, they are preferentially used in order to increase the efficiency for high-p_{T} leptons. The dilepton triggers are used in lower-p_{T} ranges. For the selection of $e\mu$ events at high p_{T}, the electron triggers are used preferentially since their efficiency is higher than the muon triggers. The trigger menu in ATLAS changed between 2015 and 2016 in order to cope with the increase in instantaneous luminosity and pileup. Events are therefore treated differently depending on the year they were recorded.

An additional control sample of events containing photons is collected using a set of single-photon triggers with p_{T} thresholds in the range 45-140 GeV. Most of the employed photon triggers are prescaled, meaning that only a subset of events satisfying the trigger requirements are retained.

9.4.2 Preselection of events

Triggered events are preselected for analysis if they pass the trigger strategy and contain at least two signal leptons and at least two signal jets. The two leptons with the highest p_{T}, referred to as the *leading* and *subleading* leptons, are required to be oppositely charged and to geometrically match the triggered objects in ΔR. In events selected with a dilepton trigger both leptons must be matched to the trigger objects, and in events selected with a single-lepton trigger one of them must be matched. The leading lepton in the event must fulfil the requirement $p_{\mathrm{T}} > 50\,\mathrm{GeV}$ which is seen to efficiently remove background from fake leptons. The subleading lepton must have $p_{\mathrm{T}} > 25\,\mathrm{GeV}$, to ensure that it is well above the thresholds of the dilepton triggers. Events where the invariant mass of the two leading leptons is below 12 GeV are vetoed in order to reject low-mass Drell-Yan processes and low-mass hadron decays. In the photon-triggered sample, the events are required to contain a signal photon with $p_{\mathrm{T}} > 50\,\mathrm{GeV}$.

Table 9.3: Lepton offline p_{T} requirements and HLT triggers [145, 146] used for the analysis in different regions of lepton-p_{T} phase space. The trigger names `1e`, `1μ` encode single-electron and single-muon triggers, and `2e`, `2μ`, `2L` encode dielectron, dimuon, and dilepton triggers. The numbers quoted for the triggers correspond to the trigger lepton-p_{T} thresholds. All triggers are unprescaled.

Offline p_{T} requirement	**2015 HLT trigger**	**2016 HLT trigger**
Dielectron channel		
$p_{\mathrm{T}}(e) > 65\,\mathrm{GeV}$	`1e`, 60 GeV	`1e`, 60 GeV
$p_{\mathrm{T}}(e) \leq 65\,\mathrm{GeV}$	`2e`, 12 GeV, 12 GeV	`2e`, 17 GeV, 17 GeV
Dimuon channel		
$p_{\mathrm{T}}(\mu) > 52.5\,\mathrm{GeV}$	`1μ`, 50 GeV	`1μ`, 50 GeV
$p_{\mathrm{T}}(\mu) \leq 52.5\,\mathrm{GeV}$	`2μ`, 18 GeV, 8 GeV	`2μ`, 22 GeV, 8 GeV
Electron-muon channel		
$p_{\mathrm{T}}(e) > 65\,\mathrm{GeV}$	`1e`, 60 GeV	`1e`, 60 GeV
$p_{\mathrm{T}}(e) \leq 65\,\mathrm{GeV}$, $p_{\mathrm{T}}(\mu) > 52.5\,\mathrm{GeV}$	`1μ`, 50 GeV	`1μ`, 50 GeV
$p_{\mathrm{T}}(e) \leq 65\,\mathrm{GeV}$, $p_{\mathrm{T}}(\mu) \leq 52.5\,\mathrm{GeV}$, $p_{\mathrm{T}}(e) < p_{\mathrm{T}}(\mu)$	`2L`, 7 GeV(e), 24 GeV(μ)	`2L`, 7 GeV(e), 24 GeV(μ)
$p_{\mathrm{T}}(e) \leq 65\,\mathrm{GeV}$, $p_{\mathrm{T}}(\mu) \leq 52.5\,\mathrm{GeV}$, $p_{\mathrm{T}}(e) > p_{\mathrm{T}}(\mu)$	`2L`, 17 GeV(e), 14 GeV(μ)	`2L`, 17 GeV(e), 14 GeV(μ)

9.4.3 Discriminating variables

Three variables are used in the analysis to discriminate between the signal and the SM backgrounds:

- The **missing transverse momentum**, $E_{\mathrm{T}}^{\mathrm{miss}}$, is used as proxy for the energy of the LSPs in the final state. It is computed using all baseline electrons, muons, photons, and jets and includes a track-based soft term. Signal events are expected to contain a significantly higher $E_{\mathrm{T}}^{\mathrm{miss}}$ compared to the SM backgrounds, due to the high-mass nature of the sparticles. For a fixed initial squark or gluino mass, the magnitude of the $E_{\mathrm{T}}^{\mathrm{miss}}$ in the signal depends on the mass splitting between the two neutralinos in the decay chain.

- The **scalar jet-p_{T} sum**, H_{T}, is a measure of the total jet energy in the event, and is measured from all signal jets. Just like the $E_{\mathrm{T}}^{\mathrm{miss}}$, signal events are expected to exhibit a higher H_{T} than the SM backgrounds. The magnitude of the H_{T} in the signal depends on the mass splitting between the gluino or squark and the $\tilde{\chi}_2^0$, which steers the energy available for the jets.

- The **stransverse mass**, m_{T2} is an extension of the transverse mass given by

$$m_{\mathrm{T}}^2\left(\mathbf{p}_{\mathrm{T}}^{\mathrm{vis}}, \mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}\right) = 2 \cdot \left(p_{\mathrm{T}}^{\mathrm{vis}} \cdot E_{\mathrm{T}}^{\mathrm{miss}} - \mathbf{p}_{\mathrm{T}}^{\mathrm{vis}} \cdot \mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}\right), \tag{9.6}$$

  which is often used as a proxy for the invariant mass in a two-body decay with one invisible final-state particle. Here, $\mathbf{p}_{\mathrm{T}}^{\mathrm{vis}}$ is the two-vector transverse momentum of the visible particle and $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ and $E_{\mathrm{T}}^{\mathrm{miss}}$ are the missing transverse

momentum vector and magnitude associated to the invisible particle. The quantity has a maximum value corresponding to the mass of the mother particle. In the stransverse mass this is generalised to the case where two invisible particles are present in the final state. The variable is defined according to

$$m_{\mathrm{T2}} = \min_{\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}(1)}+\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}(2)}=\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}} \left[\max\left(m_{\mathrm{T}}^{(1)}, m_{\mathrm{T}}^{(2)}\right)\right], \tag{9.7}$$

where $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}(i)}$ are trial vectors obtained by decomposing $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ [147, 148]. The transverse masses $m_{\mathrm{T}}^{(i)}$ are obtained by pairing either of these trial vectors with one of the two visible objects, and the minimisation is performed over all trial momenta satisfying the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ constraint. The quantity can be interpreted as the minimum mother mass compatible with all p_{T} and invariant mass constraints in the decay chain. When building m_{T2} from the two selected leptons its distribution exhibits a sharp decline around the mass of the W boson for $t\bar{t}$ events and is therefore well suited to suppress this background.

The distributions of the dilepton invariant mass and the three discriminating variables described above are shown in Figure 9.5 for data and background MC, using preselected events with an electron or muon pair. Figure 9.6 shows the expected $E_{\mathrm{T}}^{\mathrm{miss}}$, H_{T}, and m_{T2} distributions for three points in the Slepton signal grid. A comparison between the background and signal distributions illustrates the discriminating power of the three variables. A minimum $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement at around 300 GeV and a minimum H_{T} requirement at around 500 GeV removes a large portion of the background while still retaining a high efficiency for all the grid points. Similarly, a requirement on the minimum m_{T2} at around 30 GeV removes the peak of the background distribution, while still maintaining most of the signal.

The Slepton grid points in Figure 9.6 are chosen to represent different mass splittings between the gluino and LSP masses. Since the $\tilde{\chi}_2^0$ mass is fixed to the average of the gluino and LSP masses, the $\tilde{g}$-$\tilde{\chi}_1^0$ mass difference also steers the $\tilde{\chi}_2^0$-$\tilde{\chi}_1^0$ and $\tilde{g}$-$\tilde{\chi}_2^0$ mass differences. This can be seen in Figure 9.6a where the $E_{\mathrm{T}}^{\mathrm{miss}}$ is shifted to larger values for the point with the greatest mass splitting, resulting in the largest energy available for the LSP. The effect is also visible in Figure 9.6b, with a later onset in H_{T} for the same point, which provides the largest energy available for the jets. In all figures the yield of MC events is scaled to 36.1 fb^{-1}. The difference in overall yield between the three signal points is a manifestation of the cross section which decreases with gluino mass.

Figure 9.7 shows the same set of distributions for three points in the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ model grid. This model has the $\tilde{\chi}_1^0$ mass fixed to 1 GeV and a large $\tilde{\chi}_2^0$ mass thereby entails a large mass splitting between the two neutralinos. This is evident in Figure 9.7a where the signal point with the largest $\tilde{\chi}_2^0$ mass exhibits the most pronounced tail in $E_{\mathrm{T}}^{\mathrm{miss}}$. At the same time, this point corresponds to the smallest splitting between the gluino and $\tilde{\chi}_2^0$ masses which explains the softer H_{T} distribution compared to the other models in Figure 9.7b.

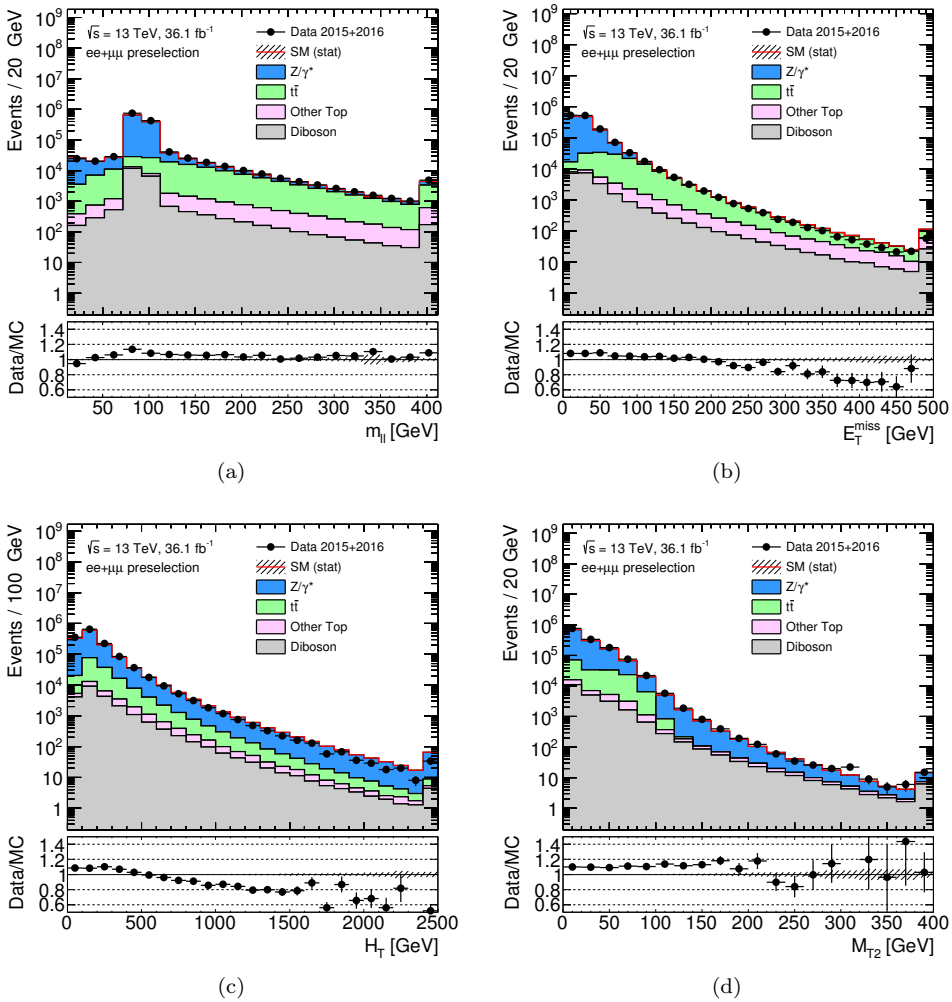

Figure 9.5: Distributions of $m_{\ell\ell}$ (a), $E_{\mathrm{T}}^{\mathrm{miss}}$ (b), H_{T} (c), and m_{T2} (d) for events that pass the preselection and contain an electron or muon pair. The error band includes only the statistical uncertainty on the MC samples.

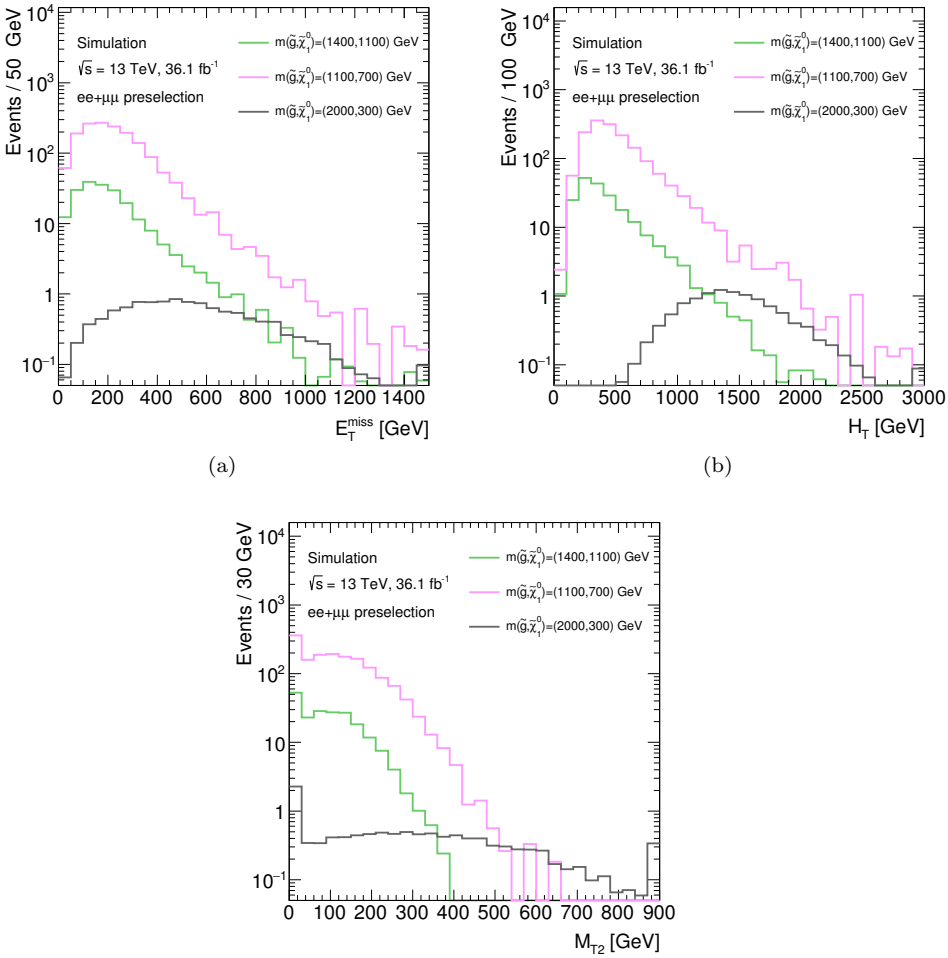

(a) (b)

(c)

Figure 9.6: Distributions of $E_{\mathrm{T}}^{\mathrm{miss}}$ (a), H_{T} (b), and m_{T2} (c) in events that pass the preselection and contain an electron or muon pair, for three points in the Slepton model signal grid.

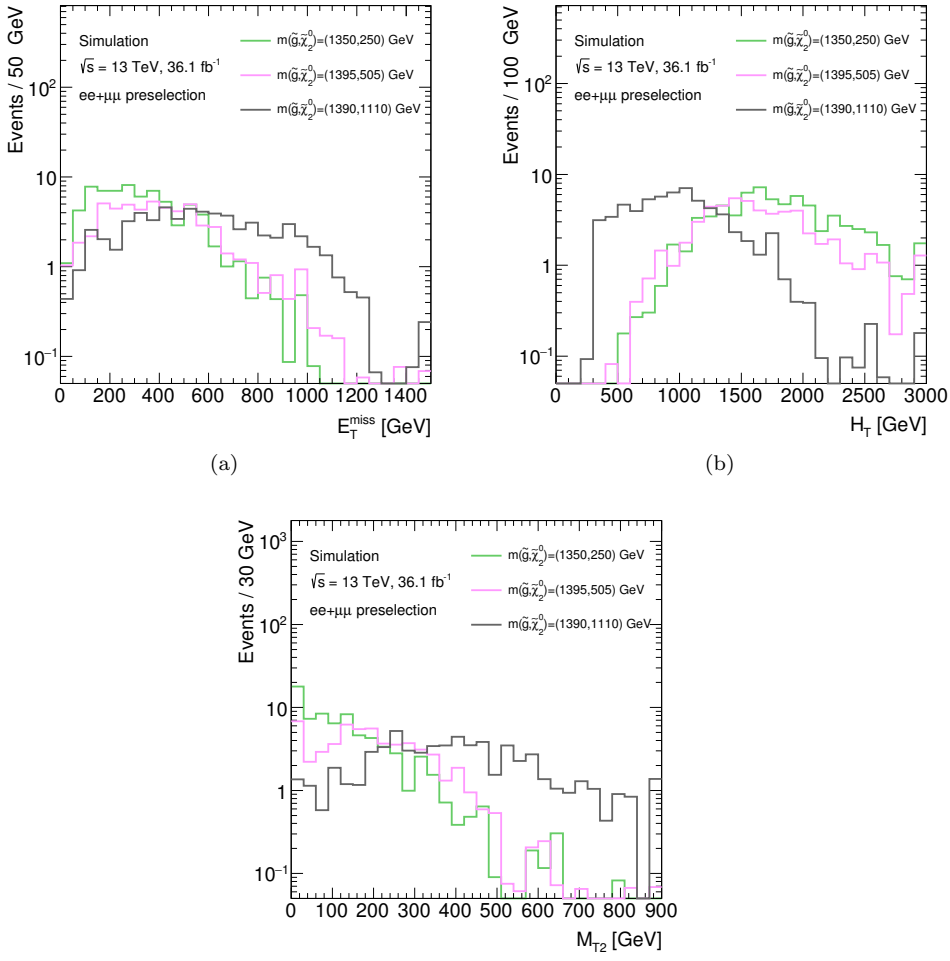

Figure 9.7: Distributions of $E_{\mathrm{T}}^{\mathrm{miss}}$ (a), H_{T} (b), and m_{T2} (c) in events that pass the preselection and contain an electron or muon pair, for three points in the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ model signal grid.

9.4.4 Signal, validation, and control regions

Three signal regions are defined for the Edge search; SR-Low, SR-Medium, and SR-High, targeting low, medium, and high mass splittings between the gluino and the $\tilde{\chi}_1^0$ in the Z* and Slepton models. Figure 9.8 shows a schematic drawing of the regions including their $E_{\mathrm{T}}^{\mathrm{miss}}$, H_{T}, and m_{T2} criteria.

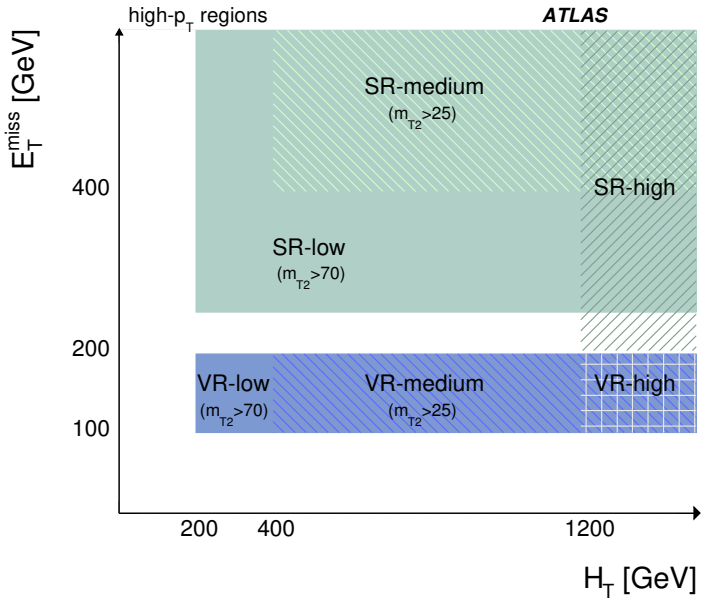

Figure 9.8: Schematic diagram of the main analysis SRs and VRs [4]. Regions where hashed markings overlap indicate the overlap between various regions.

All the regions are binned in the dilepton invariant mass according to Figure 9.9. In order to harmonise between the two searches, the On-Z SRs are designed to correspond to the $81 < m_{\ell\ell} < 101\,\mathrm{GeV}$ bin of the Edge regions. The $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement in SR-High is lowered compared to SR-Low and SR-Medium to increase sensitivity to On-Z models with small $\tilde{g}$-$\tilde{\chi}_2^0$ mass splitting. The $m_{\ell\ell}$ bins are used to make model-dependent interpretations using the invariant mass shape in each SR. The binning is chosen to ensure a sufficient number of events for a robust background estimate in each bin and a high sensitivity to the analysis signal models. In addition, the bins are used to form 29 non-orthogonal $m_{\ell\ell}$ windows which are exploited for a model-independent interpretation of the analysis results. These windows are chosen to be sensitive to a broad range of potential kinematic edge positions.

Each SR is associated with a VR with the same H_{T} and m_{T2} requirements but a lower $E_{\mathrm{T}}^{\mathrm{miss}}$ range. Several CRs for the background estimations are also defined. The m_{T2} requirement is optimised to remove a significant fraction of the $t\bar{t}$

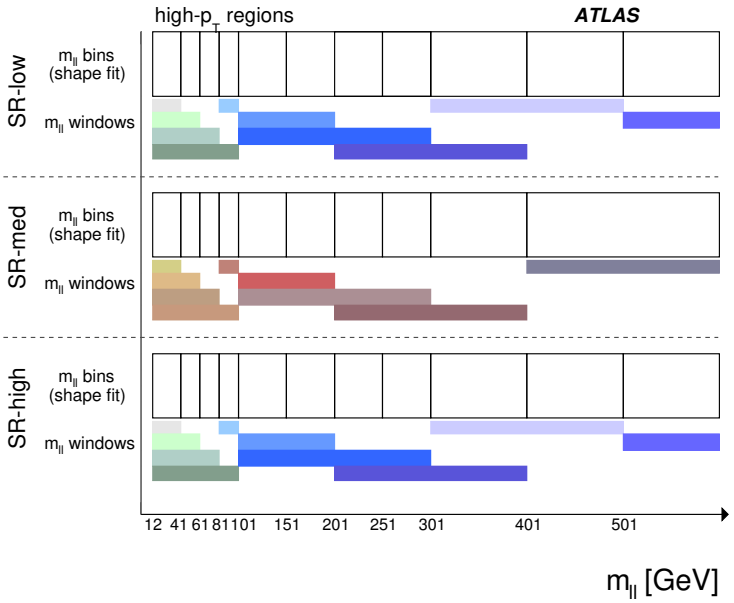

Figure 9.9: Schematic diagram showing the $m_{\ell\ell}$ binning used in the various SRs alongside the overlapping $m_{\ell\ell}$ windows used for model-independent interpretations [4]. The unfilled boxes indicate the $m_{\ell\ell}$ bin edges for the shape fits used in the model-dependent interpretations. Each filled region underneath indicates one of the $m_{\ell\ell}$ windows, formed of one or more $m_{\ell\ell}$ bins, used to derive model-independent results for the given SR. In each case, the last $m_{\ell\ell}$ bin includes the overflow.

background in the SRs while retaining sufficient statistics to perform data-driven estimates in the CRs. Table 9.4 lists the selection criteria for the analysis SRs, CRs, and VRs. The regions denoted CR-FS are used for the data-driven estimate of the $t\bar{t}$ background and the regions denoted CRγ and CRZ are used for the data-driven estimate of the Drell-Yann background. The regions denoted VR-WZ and VR-ZZ are used to validate the diboson background estimate. For most regions, the opening angle $\Delta\phi(j_{1,2}, E_{\mathrm{T}}^{\mathrm{miss}})$ between the two leading jets and the $E_{\mathrm{T}}^{\mathrm{miss}}$ is required to be above 0.4. This requirement reduces the number of events with $E_{\mathrm{T}}^{\mathrm{miss}}$ from jet mis-measurements and also suppresses $t\bar{t}$ events with a large Lorentz boost factor.

Table 9.4: Overview of the main SRs, CRs and VRs used in the Strong-2L analysis. All regions require at least two leptons, unless otherwise indicated in the *Objects* column. The flavour combination of the dilepton pair is denoted as either *SF* for same-flavour or *DF* for different-flavour. For CRγ, VR-WZ, and VR-ZZ the number of leptons, rather than a specific flavour configuration, is indicated. The regions VR-WZ, and VR-ZZ contain an additional veto against b-jets. The main requirements that distinguish the CRs and VRs from the SRs are indicated in bold. More details are given in the text. Table adapted from Ref. [4].

Regions	$E_{\mathrm{T}}^{\mathrm{miss}}$ [GeV]	H_{T} [GeV]	n_{jets}	$m_{\ell\ell}$ [GeV]	m_{T2} [GeV]	Objects	$\Delta\phi(\mathrm{jet}, \mathrm{p}_{\mathrm{T}}^{\mathrm{miss}})$
Signal regions							
SR-Low	> 250	> 200	≥ 2	> 12	> 70	SF	> 0.4
SR-Medium	> 400	> 400	≥ 2	> 12	> 25	SF	> 0.4
SR-High	> 200	> 1200	≥ 2	> 12	–	SF	> 0.4
Control regions							
CR-FS-Low	> 250	> 200	≥ 2	> 12	> 70	**DF**	> 0.4
CR-FS-Medium	> 400	> 400	≥ 2	> 12	> 25	**DF**	> 0.4
CR-FS-High	> 100	> 1100	≥ 2	> 12	–	**DF**	> 0.4
CRγ-Low	–	> 200	≥ 2	–	–	$\mathbf{0\ell}$, $\mathbf{1\gamma}$	–
CRγ-Medium	–	> 400	≥ 2	–	–	$\mathbf{0\ell}$, $\mathbf{1\gamma}$	–
CRγ-High	–	> 1200	≥ 2	–	–	$\mathbf{0\ell}$, $\mathbf{1\gamma}$	–
CRZ-Low	$\mathbf{< 100}$	> 200	≥ 2	> 12	> 70	SF	–
CRZ-Medium	$\mathbf{< 100}$	> 400	≥ 2	> 12	> 25	SF	–
CRZ-High	$\mathbf{< 100}$	> 1200	≥ 2	> 12	–	SF	–
Validation regions							
VR-Low	$\mathbf{100 - 200}$	> 200	≥ 2	> 12	> 70	SF	> 0.4
VR-Medium	$\mathbf{100 - 200}$	> 400	≥ 2	> 12	> 25	SF	> 0.4
VR-High	$\mathbf{100 - 200}$	> 1200	≥ 2	> 12	–	SF	> 0.4
VR-WZ	$\mathbf{100 - 200}$	$\mathbf{> 200}$	≥ 2	> 12	–	$\mathbf{3\ell}$	> 0.4
VR-ZZ	$\mathbf{< 50}$	$\mathbf{> 100}$	$\mathbf{\geq 1}$	> 12	–	$\mathbf{4\ell}$	> 0.4

9.5 Signal acceptance and efficiency

The acceptances and efficiencies for the signal are assessed for each analysis SR. The acceptance is defined as the fraction of signal events that kinematically match the SR. It is calculated by applying the SR kinematic requirements to truth-level objects in the MC simulation, which do not suffer from identification inefficiencies or measurement resolutions. The efficiency is defined as the fraction of accepted events that are correctly reconstructed and identified in ATLAS, and is computed using the truth links of the reconstructed objects. Figure 9.10 shows the acceptances and efficiencies for the Slepton signal grid in the three SRs. Typical values of the acceptance times efficiency are 1-25 % inclusively in $m_{\ell\ell}$, with a drop towards the diagonal where $m_{\tilde{g}} = m_{\tilde{\chi}_1^0}$. Figure 9.11 shows the acceptance and efficiency for the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ model in the On-Z bins of SR-Medium and SR-High. The acceptance is small because the simplified model includes all possible Z-boson decays, and only about 7 % of these result in an electron or muon pair. In these regions, the acceptance times efficiency varies between 1-4 %. The On-Z bin in SR-Low is not used in the analysis since it is found to not increase the sensitivity of the search.

All CRs and VRs are also checked for signal contamination using the same signal models. At gluino masses above 1200 GeV and neutralino masses above 600 GeV the contamination is negligible. For lower masses the contamination can be significant in some regions but such SUSY models are already excluded by the previous analyses in the same final state.

9.6 Standard Model backgrounds

The main background processes for this analysis can be grouped into the *flavour-symmetric* (FS), *diboson*, and $Z/\gamma^* + jets$ backgrounds. All SRs are dominated by the FS background which mainly consists of $t\bar{t}$ with smaller contributions from Wt and WW, as well as production of Z bosons which decay as $Z \to \tau\tau$. An example Feynman diagram for $t\bar{t}$ production and decay is shown in Figure 9.12. The final state contains two opposite-charge leptons, jets from the two b-quarks, and $E_{\mathrm{T}}^{\mathrm{miss}}$ from the neutrino. In all FS backgrounds the leptons in the final state result from two independent $W \to \ell\nu$ decays and their flavours are uncorrelated. The rate of produced ee, $\mu\mu$, and $e\mu$ events should therefore follow a $1:1:2$ ratio.

The diboson background consists of WZ and ZZ processes, where the bosons decay leptonically. These processes have smaller cross sections than $t\bar{t}$ production and constitute the second largest background in most analysis SRs. Two example Feynman diagrams are shown in Figure 9.13. In both cases, the final state contains two same-flavour opposite-charge leptons and at least one neutrino which provides the $E_{\mathrm{T}}^{\mathrm{miss}}$. The diboson processes can mimic the signal if they are produced together with jets from ISR where one of the partons radiates a gluon before interacting.

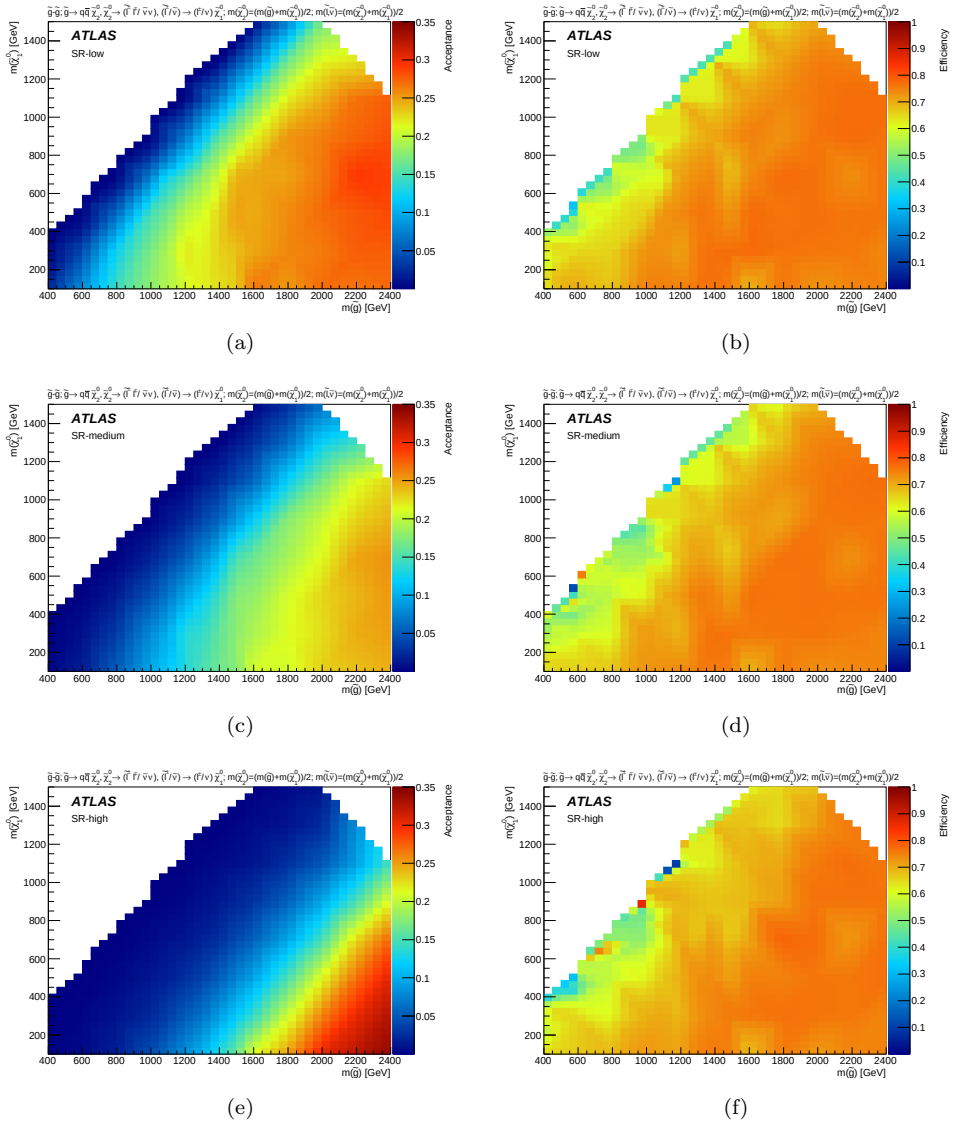

Figure 9.10: Signal region acceptance and efficiency across the Slepton signal model grid in SR-Low (a,b), SR-Medium (c,d) and SR-High (e,f) [4].

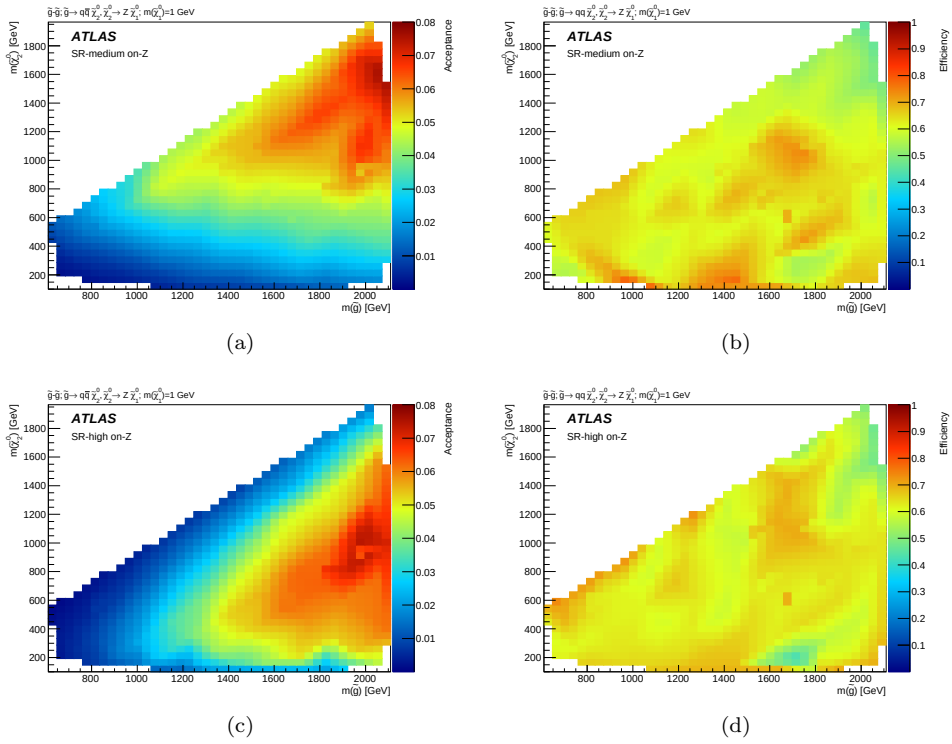

Figure 9.11: Signal region acceptance and efficiency across the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ signal model grid in the On-Z SR bin of SR-Medium (a,b) and SR-High (c,d) [4].

The $Z/\gamma^* +$ jets background results from Drell-Yan processes with jets from ISR as shown in Figure 9.14. Drell-Yann processes have a relatively large cross section, but the final state does not contain any natural $E_{\mathrm{T}}^{\mathrm{miss}}$. Events where the Z boson or photon decays into two electrons or muons can however mimic the signal if they contain $E_{\mathrm{T}}^{\mathrm{miss}}$ resulting from instrumental effects or from neutrinos in jets. This background is particularly important for the On-Z search as it peaks in the Z-boson mass window.

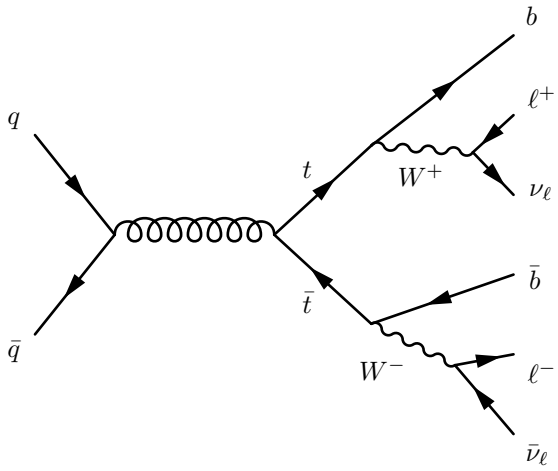

Figure 9.12: Example Feynman diagram for $t\bar{t}$ production and decay.

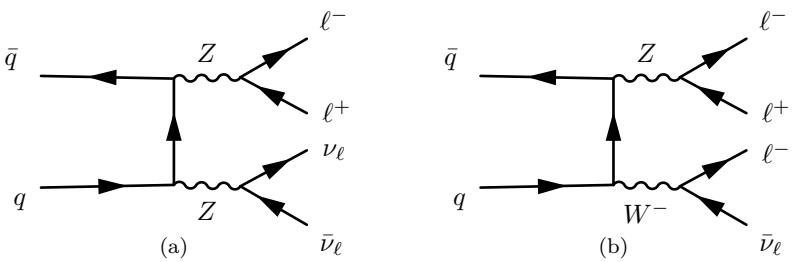

Figure 9.13: Example Feynman diagrams for ZZ (a) and WZ (b) production and decay.

Other processes can contribute at lower rates. The associated production of $t\bar{t}+W$, $t\bar{t}+Z$ and $t\bar{t}+WW$ give signatures that mimic the signal and are referred to as *rare-top* backgrounds. These are estimated directly from MC simulation. Production of $t\bar{t}$ where only one of the top quarks decays leptonically can pass the dilepton selection through lepton fakes, primarily from the decays of b-hadrons. The fake background contribution to the SRs is studied in MC and is found to be negligible.

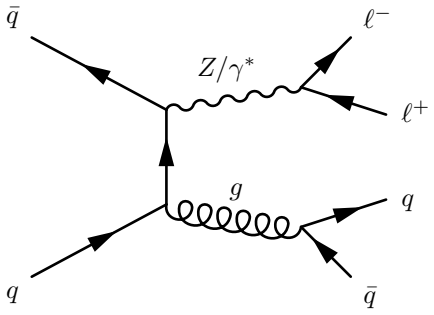

Figure 9.14: Example Feynman diagram for $Z/\gamma^* +$ jets production and decay.

The distributions shown in Figure 9.5 can be used to assess the modelling of the FS, diboson, and $Z/\gamma^* +$ jets backgrounds. A 10 % normalisation difference between data and MC can be noted around the Z-boson mass in the $m_{\ell\ell}$ distribution, suggesting an imperfect modelling of the Z/γ^* background. Furthermore, the $E_{\mathrm{T}}^{\mathrm{miss}}$ and H_{T} distributions exhibit data-MC differences in the shape, with a slight underestimation at low values and an overestimation at high values in MC. These issues motivate the choice to primarily use data-driven techniques to estimate the main backgrounds. While the FS and $Z/\gamma^* +$ jets backgrounds are estimated entirely from data, the diboson estimate is taken from MC but is carefully validated in dedicated VRs. The background estimates are described in the following sections.

9.6.1 Flavour-symmetric background

The FS background estimation method relies on data in CR-FS-Low, CR-FS-Medium and CR-FS-High introduced in Table 9.4. These regions have the same kinematic requirements and $m_{\ell\ell}$ binning as the corresponding SRs, but use a different-flavour selection which makes them orthogonal to the SRs. The CRs have a purity in the FS processes $t\bar{t}$, Wt, WW, and $Z \to \tau\tau$ greater than 95 %, as measured using the background MC samples. Other non-FS processes that contribute to the regions are WZ, ZZ, and rare-top production. The event yields in data in the CRs are 87 for CR-FS-Low, 30 for CR-FS-Medium, and 64 for CR-FS-High.

The FS background is estimated by exploiting the 1:1:2 ratio between produced ee, $\mu\mu$, and $e\mu$ final states. This ratio can however not be applied directly to the events in CR-FS to estimate the SR yield, since the measured rates may be affected by difference in the trigger and identification efficiencies between electrons and muons. Each $e\mu$ event in CR-FS is therefore corrected with two factors accounting for such efficiency differences. The small contribution from non-FS processes is also accounted for by estimating the contamination in MC and subtracting from the data. For a given SR, the combined $ee + \mu\mu$ yield is estimated according to

$$N_{ee+\mu\mu}^{\text{est}} = \frac{f_{\text{SR}}}{2} \cdot \Big[\sum_{i}^{N_{e\mu}^{\text{data}}} \Big(k_e(p_{\text{T}}^{i,\mu}, \eta^{i,\mu}) + k_\mu(p_{\text{T}}^{i,e}, \eta^{i,e}) \Big) \cdot \alpha(p_{\text{T}}^{i,l_1}, \eta^{i,l_1}) - \sum_{i}^{N_{e\mu}^{\text{MC}}} \Big(k_e(p_{\text{T}}^{i,\mu}, \eta^{i,\mu}) + k_\mu(p_{\text{T}}^{i,e}, \eta^{i,e}) \Big) \cdot \alpha(p_{\text{T}}^{i,l_1}, \eta^{i,l_1}) \Big], \tag{9.8}$$

where the first sum spans over all data events and the second sum spans over all non-FS MC events in the corresponding CR-FS. Each component is this equation is described in detail below.

Trigger and identification efficiency correction factors

The k_e, k_μ, and α factors in Eq. 9.8 are combinations of the trigger efficiencies $\epsilon_{ee}^{\text{trig}}$, $\epsilon_{\mu\mu}^{\text{trig}}$, and $\epsilon_{e\mu}^{\text{trig}}$ for ee, $\mu\mu$, and $e\mu$ events, and the identification efficiencies ϵ_e^{id} and ϵ_μ^{id} for electrons and muons. They are measured using data in a loose region similar to the analysis preselection. Events are required to pass the analysis trigger requirements, contain two opposite-charge leptons and two or more jets, have leading and subleading lepton p_{T} above 25 GeV, and have a dilepton invariant mass inside the Z-boson mass window $81 < m_{\ell\ell} < 101$ GeV. The requirement on the invariant mass lowers the statistics available to measure the efficiencies but ensures a pure sample with a minimal contribution from fake leptons.

The definitions of the k_e and k_μ factors are given in Eqs. 9.9 and 9.10, and the definition of the α factor is given in Eq. 9.11.

$$k_e = \sqrt{\frac{N_{ee}^{\text{id+trig}}}{N_{\mu\mu}^{\text{id+trig}}}} = \sqrt{\frac{\epsilon_{ee}^{\text{id}} \cdot \epsilon_{ee}^{\text{trig}}}{\epsilon_{\mu\mu}^{\text{id}} \cdot \epsilon_{\mu\mu}^{\text{trig}}}} \tag{9.9}$$

$$k_\mu = \frac{1}{k_e} = \sqrt{\frac{N_{\mu\mu}^{\text{id+trig}}}{N_{ee}^{\text{id+trig}}}} = \sqrt{\frac{\epsilon_{\mu\mu}^{\text{id}} \cdot \epsilon_{\mu\mu}^{\text{trig}}}{\epsilon_{ee}^{\text{id}} \cdot \epsilon_{ee}^{\text{trig}}}} \tag{9.10}$$

$$\alpha = \frac{\sqrt{\frac{N_{ee}^{\text{id+trig}}}{N_{ee}^{\text{id}}} \cdot \frac{N_{\mu\mu}^{\text{id+trig}}}{N_{\mu\mu}^{\text{id}}}}}{\frac{N_{e\mu}^{\text{id+trig}}}{N_{e\mu}^{\text{id}}}} = \frac{\sqrt{\epsilon_{ee}^{\text{trig}} \cdot \epsilon_{\mu\mu}^{\text{trig}}}}{\epsilon_{e\mu}^{\text{trig}}} \tag{9.11}$$

The first expression in these equations explain how the factors are measured using the events in the inclusive On-Z selection. Here, N_{ee}^{id} ($N_{ee}^{\text{id+trig}}$), $N_{\mu\mu}^{\text{id}}$ ($N_{\mu\mu}^{\text{id+trig}}$), and $N_{e\mu}^{\text{id}}$ ($N_{e\mu}^{\text{id+trig}}$) are the numbers of ee, $\mu\mu$ and $e\mu$ events before (after) the trigger

requirement. When $k_{e\mu}$ and α are combined, each event is corrected for the trigger and efficiency differences between ee, $\mu\mu$ and $e\mu$ events according to

$$\alpha \cdot k_e = \frac{\epsilon_{ee}^{\text{trig}}}{\epsilon_{e\mu}^{\text{trig}}} \cdot \sqrt{\frac{\epsilon_{ee}^{\text{id}}}{\epsilon_{\mu\mu}^{\text{id}}}}, \tag{9.12}$$

$$\alpha \cdot k_\mu = \frac{\epsilon_{\mu\mu}^{\text{trig}}}{\epsilon_{e\mu}^{\text{trig}}} \cdot \sqrt{\frac{\epsilon_{\mu\mu}^{\text{id}}}{\epsilon_{ee}^{\text{id}}}}. \tag{9.13}$$

Here, the first combination is used to transform an FS $e\mu$ event into an ee event and the second to transform an $e\mu$ event into a $\mu\mu$ event. Assuming that the two leptons have similar efficiencies, the square root gives the efficiency ϵ_e^{id} or ϵ_μ^{id} for a single lepton.

The k_e and k_μ factors are calculated separately for each year for the leading and subleading lepton and are binned in the lepton η and p_{T}. When the factors are applied to the $e\mu$ event in CR-FS the appropriate k_e or k_μ factor is selected depending on the year, whether the lepton to be replaced is leading or subleading and based on its p_{T} and η. Figure 9.15 shows the measured k_e factor for the leading and subleading lepton as a function of p_{T} and η for data collected in 2015 and 2016 respectively. The Muon Spectrometer lacks coverage for $|\eta| < 0.1$ which results in a lower efficiency for detecting muons and a large k_e factor in this region. Furthermore, electrons are accepted up to $|\eta| = 2.47$ while muons are accepted up to $|\eta| = 2.5$. This difference explains the low values of k_e in the region $|\eta| > 2.4$. The increase with p_{T} can be explained by the single muon trigger, used at high p_{T}, which only covers the region up to $|\eta| = 2.4$. In the other regions the correction factor is relatively flat and within 10 % from unity.

The main purpose of the α factor is to correct for the trigger efficiency difference, and to cancel the trigger dependence in the k_e and k_μ factors. The calculation of the α factor is complicated because the trigger used to select the events depends on the p_{T} of the leptons. Each of the efficiencies $\epsilon_{ee}^{\text{trig}}$, $\epsilon_{\mu\mu}^{\text{trig}}$ and $\epsilon_{e\mu}^{\text{trig}}$ in the α factor are therefore computed separately for the different triggers in Table 9.4 and are binned in the p_{T} and η of the leading lepton. When correcting an $e\mu$ event in CR-FS the $\epsilon_{e\mu}^{\text{trig}}$ is selected based on the year and the p_{T} and η of the leading lepton in the event. Similarly, the $\epsilon_{ee}^{\text{trig}}$ and $\epsilon_{\mu\mu}^{\text{trig}}$ are selected based on the leading lepton properties and the trigger that would have been used to select a corresponding ee or $\mu\mu$ event. Figure 9.16 shows the α factor for data collected in 2015 and 2016 as a function of p_{T} and η. Here, the trigger efficiencies $\epsilon_{ee}^{\text{trig}}$, $\epsilon_{\mu\mu}^{\text{trig}}$ and $\epsilon_{e\mu}^{\text{trig}}$ that enter α are averaged over the triggers in the analysis. For 2015 there are no triggered $e\mu$ events above $p_{\text{T}} = 300\,\text{GeV}$ and α can therefore not be computed.

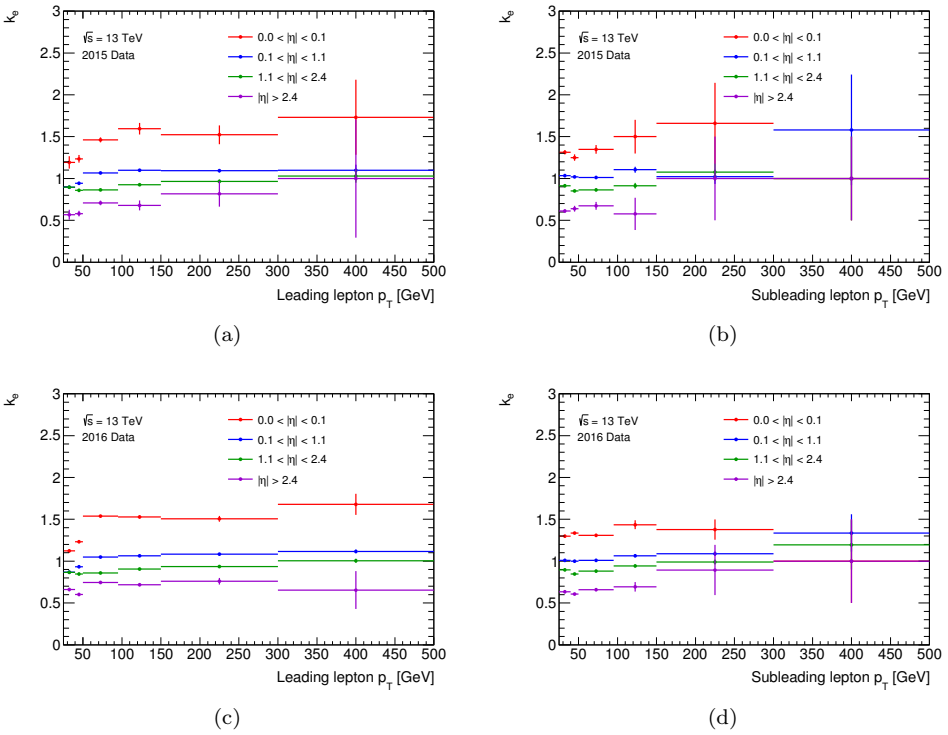

Figure 9.15: The k_e correction factor as a function of η and p_{T} derived for the leading and subleading leptons with data from 2015 (a,b) and 2016 (c,d). The uncertainties on the points correspond to the statistical uncertainty.

Extrapolation

For SR-High, the corresponding CR-FS-High is expanded by relaxing the $E_{\mathrm{T}}^{\mathrm{miss}}$ and H_{T} requirements to 100 GeV and 1100 GeV respectively. Figure 9.17 shows the invariant mass distribution of $e\mu$ events in the nominal selection compared to the relaxed selection. In addition to the increase in statistics, the relaxed selection also removes the downward fluctuation seen in the distribution around the Z-boson mass for the nominal selection. This fluctuation would otherwise lead to an underestimated FS contribution in the On-Z bin in SR-High.

The extrapolation factor f_{SR} is derived from a MC sample including the FS processes $t\bar{t}$, Wt, WW, and $Z \to \tau\tau$. It is computed as the fraction of observed $e\mu$ events in CR-FS-High that fall into the nominal $E_{\mathrm{T}}^{\mathrm{miss}}$ and H_{T} selection according to:

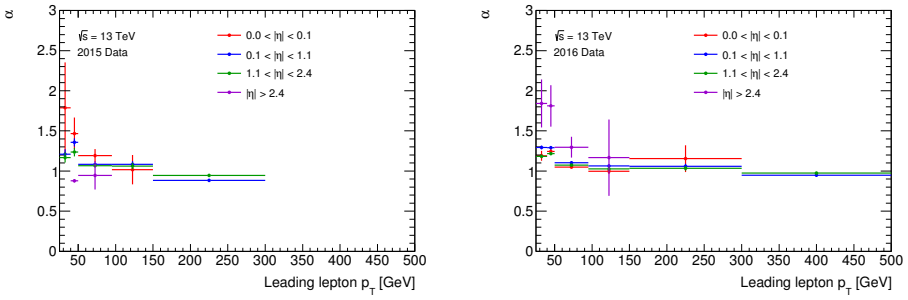

(a) (b)

Figure 9.16: The α correction factor as a function of η and p_{T} derived from 2015 (a) and 2016 (b) data. The uncertainties on the points correspond to the statistical uncertainty.

$$f_{\mathrm{SR}} = \frac{N_{e\mu}^{\text{CR-FS-High}}(E_{\mathrm{T}}^{\mathrm{miss}} > 200\ \mathrm{GeV}, H_{\mathrm{T}} > 1200\ \mathrm{GeV})}{N_{e\mu}^{\text{CR-FS-High}}(E_{\mathrm{T}}^{\mathrm{miss}} > 100\ \mathrm{GeV}, H_{\mathrm{T}} > 1100\ \mathrm{GeV})}. \tag{9.14}$$

Figure 9.18 shows the extrapolation factor as a function of $m_{\ell\ell}$. The factor is approximately flat in $m_{\ell\ell}$, and therefore an overall factor of $f_{\mathrm{SR}} = 0.40 \pm 0.01$ is used for all bins. For SR-Medium and SR-Low the same selections are used for the CR-FS and the SR, and the extrapolation factor in Eq. 9.8 is $f_{\mathrm{SR}} = 1$.

Subtraction of non-flavour-symmetric events

The contamination in CR-FS from WZ, ZZ, and rare top is estimated using MC simulation. All MC events in these samples entering CR-FS are corrected using the k_e, k_μ, and α factors and are then subtracted from the number of corrected data events. This corresponds to the last term in Eq. 9.8 and represents around 5 % of the total estimate in each region.

MC validation

The k_e, k_μ, and α factors are validated in a MC closure test where the $ee + \mu\mu$ yields predicted by the FS method are compared to the actual number of MC events in each SR. Figure 9.19 shows the results of the test for each SR. The $m_{\ell\ell}$ distribution is seen to be well modelled after applying the FS method to the $e\mu$ MC events. This is true in particular in SR-High, where the $E_{\mathrm{T}}^{\mathrm{miss}}$- and H_{T}-based extrapolation is applied. The small differences between the predictions and the observed distributions are used to assign a MC non-closure uncertainty to the estimate.

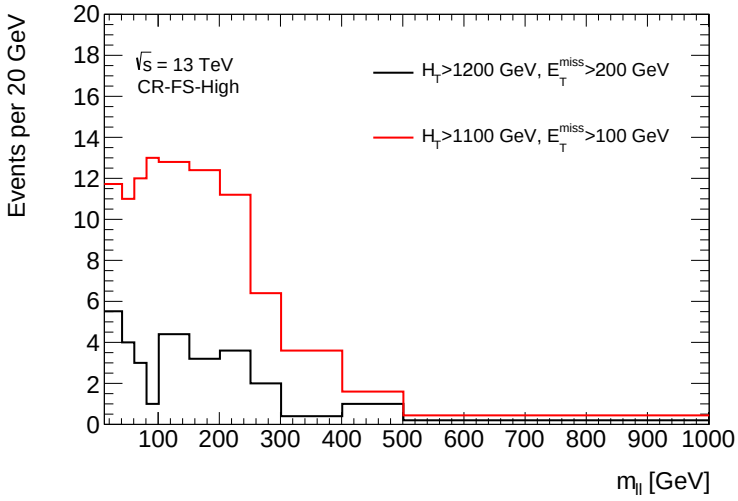

Figure 9.17: Invariant mass shape in CR-FS-High with the nominal SR-High $E_{\mathrm{T}}^{\mathrm{miss}}$ and H_{T} selections (black) and the relaxed selection (red).

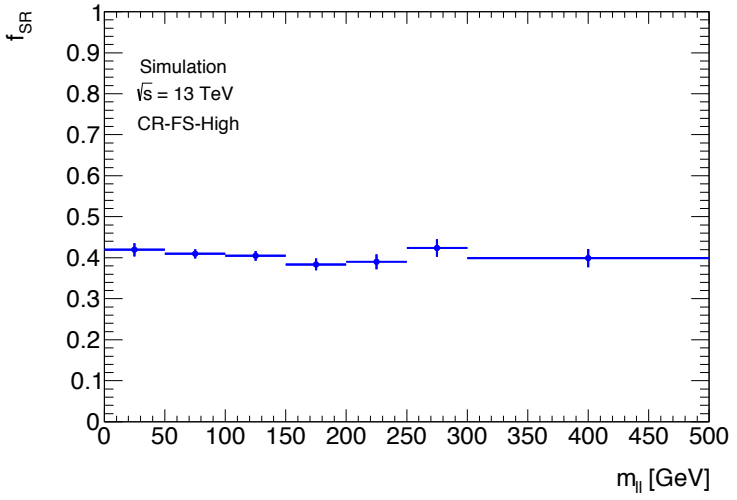

Figure 9.18: Extrapolation factor from the relaxed CR-FS-High selection to the nominal selection, computed in MC as a function of $m_{\ell\ell}$.

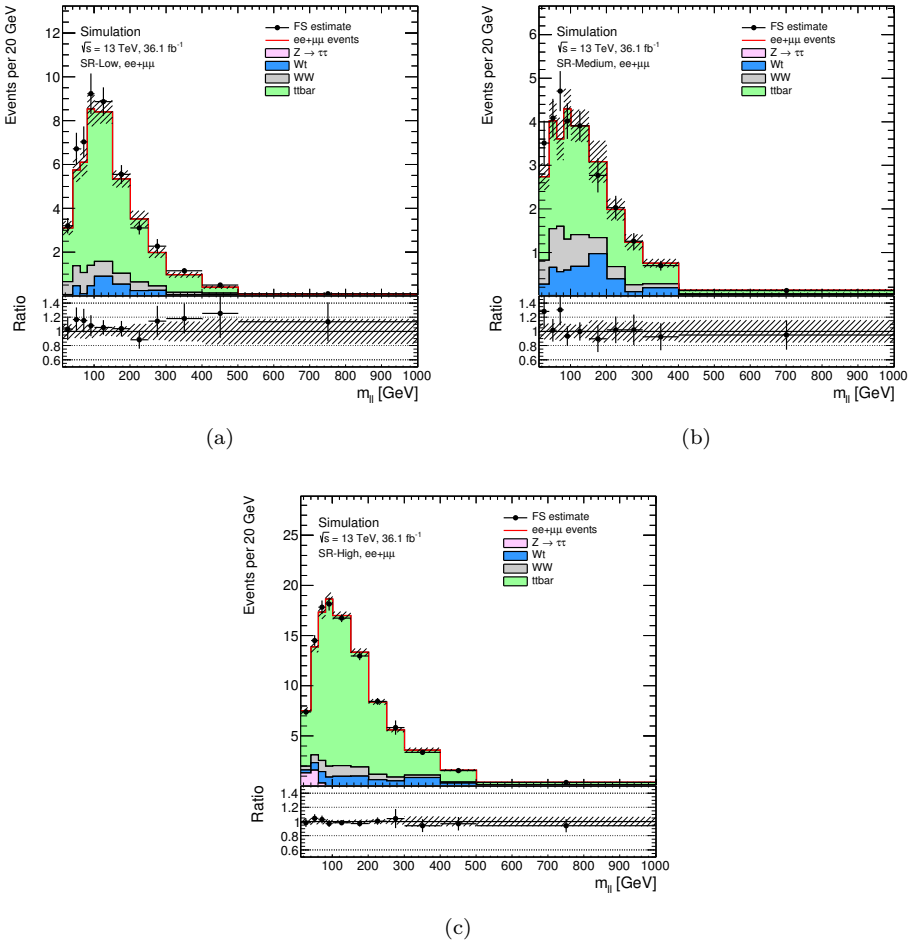

Figure 9.19: MC closure test for SR-Low (a), SR-Medium (b) and SR-High (c). The FS estimate is shown with its statistical uncertainty as black points and the actual $ee + \mu\mu$ yield for each FS background is shown by the stacked histograms. The hashed uncertainty band corresponds to the statistical uncertainty on the raw $ee + \mu\mu$ yield. The lower panel shows the ratio of the actual and predicted yields.

Systematic uncertainties on the FS estimate

The systematic uncertainties assigned to the FS estimate are associated with the statistics in the CRs, the MC non-closure, and the efficiency correction factors. Each uncertainty is described below and the relative contributions to the total uncertainty on the FS estimate are shown in Figure 9.20.

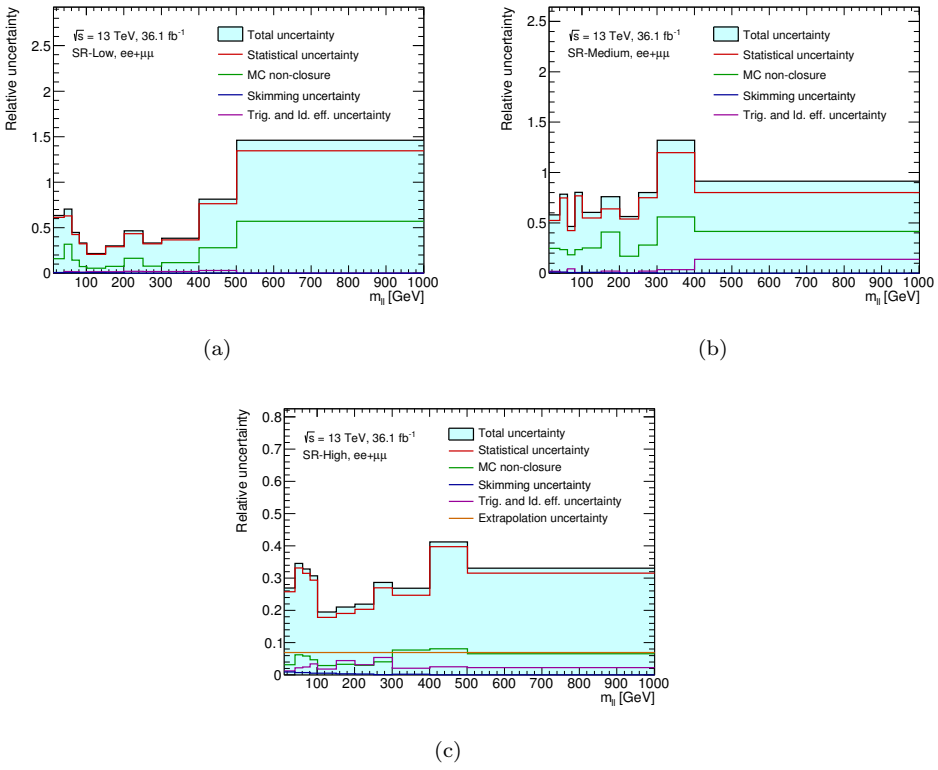

Figure 9.20: Relative systematic uncertainties on the estimated FS yields in SR-Low (a), SR-Medium (b) and SR-High (c). The total uncertainty is broken down into statistical uncertainty from the number of CR-FS events, MC non-closure uncertainty, uncertainty from the statistical error on the $k_{e/\mu}$ and α factors, uncertainty from the comparison of skimmed and unskimmed data samples, and the extrapolation uncertainty for SR-High

The statistical uncertainty from the limited number of data events in CR-FS makes up the dominant uncertainty on the estimate. Also included is a smaller contribution from the statistical uncertainty on the number of non-FS MC events subtracted from the data in each CR-FS. The second-to-largest uncertainty comes from the MC non-closure. To make a conservative estimate, the difference between the estimated and actual $ee + \mu\mu$ yields in the MC is compared to the sum in quadrature of the statistical uncertainties, and the largest value is retained. The MC statistical uncertainties tend to be larger than the non-closure, and dominate the final uncertainty.

The first uncertainty from the efficiency correction factors considers their precision.

Here, the factors are varied by their statistical uncertainties to derive a maximum and minimum FS prediction. The symmetrised difference from the nominal prediction in each $m_{\ell\ell}$ bin is taken as a systematic uncertainty on the FS estimate.

The second uncertainty from the efficiency correction factors considers a potential bias in the way the α factor is calculated. The trigger efficiencies $\epsilon_{ee}^{\text{trig}}$, $\epsilon_{\mu\mu}^{\text{trig}}$ and $\epsilon_{e\mu}^{\text{trig}}$ that enter α in Eq. 9.11 are taken as the ratio of events passing the analysis specific trigger requirements to the total number of events in each channel. Since the factors are taken from data, the events that enter the denominator in these efficiencies do not actually correspond to all possible events but rather to all events passing the triggers used in the data collection. As a consequence the measured trigger efficiency might be artificially high. In order to assess this bias, the trigger efficiencies that enter α are computed with an additional trigger requirement applied to both the numerator and denominator. The extra trigger is based on the $E_{\text{T}}^{\text{miss}}$ in the event and is completely orthogonal to the analysis specific triggers, and thereby isolates their effect on the measured efficiencies. In this test, a relative difference to the nominal α below 1 % is observed.

The data files used in the analysis are filtered based on a set of triggers required by the Strong-2L and various other ATLAS SUSY analyses. This process is called skimming and also affects the denominator in the trigger efficiencies. The bias resulting from the data format is assessed in MC using the nominal analysis $t\bar{t}$ sample stored in a different format where no skimming is applied. FS estimates are then derived in $t\bar{t}$ MC using both the skimmed and the unskimmed sample. Figure 9.21 shows the α factor as a function of $E_{\text{T}}^{\text{miss}}$ for data, the $t\bar{t}$ MC sample, and the unskimmed $t\bar{t}$ MC sample for 2015 and 2016. Here, the factor is averaged over all η and p_{T} bins. There is a relative difference below 1 % between the skimmed and unskimmed values, which is assigned as a systematic uncertainty on the measurement in data to account for the effect of skimming.

For the estimate in SR-High, an additional uncertainty is assigned for the extrapolation from the relaxed CR-FS-High selection. It is taken from the difference between the nominal extrapolation factor computed in MC and an extrapolation factor computed in data. Also included is the statistical uncertainty on the extrapolation factor in MC. Since an overall scale factor is used for all $m_{\ell\ell}$ bins in CR-FS-High, the uncertainty is flat across the $m_{\ell\ell}$ spectrum.

9.6.2 Z/γ^* + jets background

The Z/γ^* + jets background is challenging to model in simulation since the $E_{\text{T}}^{\text{miss}}$ in the events results mainly from jet mis-measurements. A data-driven method is therefore used to estimate the background, based on the $E_{\text{T}}^{\text{miss}}$ spectrum is γ + jets events. This is as suitable control sample as the two processes have similar topologies with one well-measured boson recoiling against jets. Both processes also have $E_{\text{T}}^{\text{miss}}$ originating mostly from instrumental effects. The estimate uses the

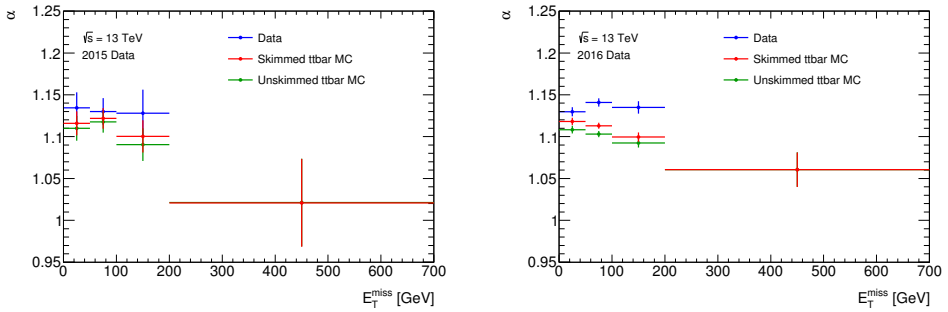

Figure 9.21: The α correction factor as a function of $E_{\mathrm{T}}^{\mathrm{miss}}$ derived from 2015 (a) and 2016 (b) data and skimmed and unskimmed $t\bar{t}$ MC. The factor is averaged over η and p_{T} and the uncertainties on the points correspond to the statistical uncertainty. The data is blinded above $E_{\mathrm{T}}^{\mathrm{miss}} = 200\,\mathrm{GeV}$.

photon-triggered data sample, and selects events in $CR\gamma$-Low, $CR\gamma$-Medium and $CR\gamma$-High, detailed in Table 9.4. These regions veto leptons and have the same kinematic selection as the SRs, but impose no $E_{\mathrm{T}}^{\mathrm{miss}}$ or $\Delta\phi(\mathrm{jet}, \mathrm{p}_{\mathrm{T}}^{\mathrm{miss}})$ requirements.

The photon events must be corrected for kinematic differences between the $Z/\gamma^* +$ jets and $\gamma +$ jets processes. In contrast to the Z boson the photon is massless and must be produced in association with jets in order to conserve energy and momentum. The production of photons also requires less energy. As a consequence, several kinematic variables including jet energies and jet multiplicities differ between the two processes. To mitigate the differences, the boson p_{T} distribution in the $\gamma +$jets data sample is reweighted to match the distribution in $Z/\gamma^* +$ jets in CRZ-Low, CRZ-Medium and CRZ-High. These CRs, detailed in Table 9.4, are similar to the SRs but require the $E_{\mathrm{T}}^{\mathrm{miss}}$ to be below 100 GeV, making them orthogonal. The boson p_{T} is used as a proxy for the energy scale of the event but the reweighting can also be performed in other representative variables. A systematic uncertainty is assigned to the procedure by comparing the p_{T} reweighting to a reweighting in the H_{T} distribution.

All $\gamma +$ jets events must also be corrected for experimental differences between photons and Z bosons. The photon is stable while the Z boson rapidly decays in the detector and is measured from its decay products. Electrons and photons are both measured in the electromagnetic calorimeter with the same energy resolution but small differences can result from the overlap removal procedure. Muons are measured using tracks for which the momentum resolution degrades at high p_{T} making it significantly different from the photon resolution. The differences in momentum resolution are corrected for by smearing the photon p_{T} according to a

$Z \to ee$ or $Z \to \mu\mu$ resolution function. This smearing function is derived from data in a 1-jet CR by comparing the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$-projection $E_{T,\parallel}^{\mathrm{miss}}$ along the boson momentum in $Z/\gamma^* + \mathrm{jets}$ and $\gamma + \mathrm{jets}$ MC events. A deconvolution procedure is used to avoid including the photon resolution in the Z-boson p_{T} resolution function. For each event, a photon p_{T} smearing Δp_{T} is obtained by sampling the smearing function. The photon p_{T} is shifted by Δp_{T}, with the parallel component of the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ vector being correspondingly adjusted by $-\Delta p_{\mathrm{T}}$.

After the smearing and reweighting procedures the $E_{\mathrm{T}}^{\mathrm{miss}}$ in each $\gamma + \mathrm{jets}$ event is recalculated, and the final $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution is obtained after applying the $\Delta\phi(\mathrm{jet}, \mathrm{p}_{\mathrm{T}}^{\mathrm{miss}}) > 0.4$ requirement. The resulting $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution is normalised to data in CRZ-Low, CRZ-Medium and CRZ-High, before the $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement of the SR is applied. In order to estimate the final SR yields, the requirement on m_{T2} and the binning in $m_{\ell\ell}$ must also be applied. Since the $\gamma + \mathrm{jets}$ events do not naturally contain any leptons, the final steps in the background estimation are dedicated to the modelling of these two distributions. The dilepton mass distribution is modelled from MC by parametrising the $m_{\ell\ell}$ spectrum in $E_{T,\parallel}^{\mathrm{miss}}$. Each $\gamma + \mathrm{jets}$ event is assigned an $m_{\ell\ell}$ value by randomly sampling from the MC distribution. The parametrisation in $E_{T,\parallel}^{\mathrm{miss}}$ takes the correlation between the momentum measurement and the invariant mass into account. For the modelling of the m_{T2} variable, each photon is split into two imaginary leptons. The system is boosted to the rest frame of the boson where the two imaginary leptons are assigned randomly selected θ and ϕ directions and half of the $m_{\ell\ell}$ in the event as momentum. After boosting back to the lab frame the final four-momenta of the two leptons is computed and the m_{T2} variable is calculated.

Finally, the production of photons together with a W or Z boson results in events with real $E_{\mathrm{T}}^{\mathrm{miss}}$. Such $V\gamma$ events can contaminate the photon data sample used to estimate the $Z/\gamma^* + \mathrm{jets}$ background. The full procedure is therefore applied to $V\gamma$ events in MC and the estimate is subtracted from the data estimate. A systematic uncertainty on the normalisation of the $V\gamma$ MC sample is derived from a data-to-MC comparison in a $V\gamma$ enriched region and is applied to the final estimate.

In addition to the uncertainties from the reweighting, smearing, and $V\gamma$ subtraction, the statistical uncertainty on the $\gamma + \mathrm{jets}$ data enters as a systematic uncertainty on the final estimate. A validation of the estimation technique is also performed by applying the full method to MC. The estimated $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution is compared to the raw $Z/\gamma^* + \mathrm{jets}$ distribution in MC and the difference is taken as an additional systematic uncertainty. Results from this closure test are shown in Figure 9.22. The total uncertainty associated to the method varies from 10-100 % depending on the $m_{\ell\ell}$ range, and is dominated by the smearing and reweighting uncertainties.

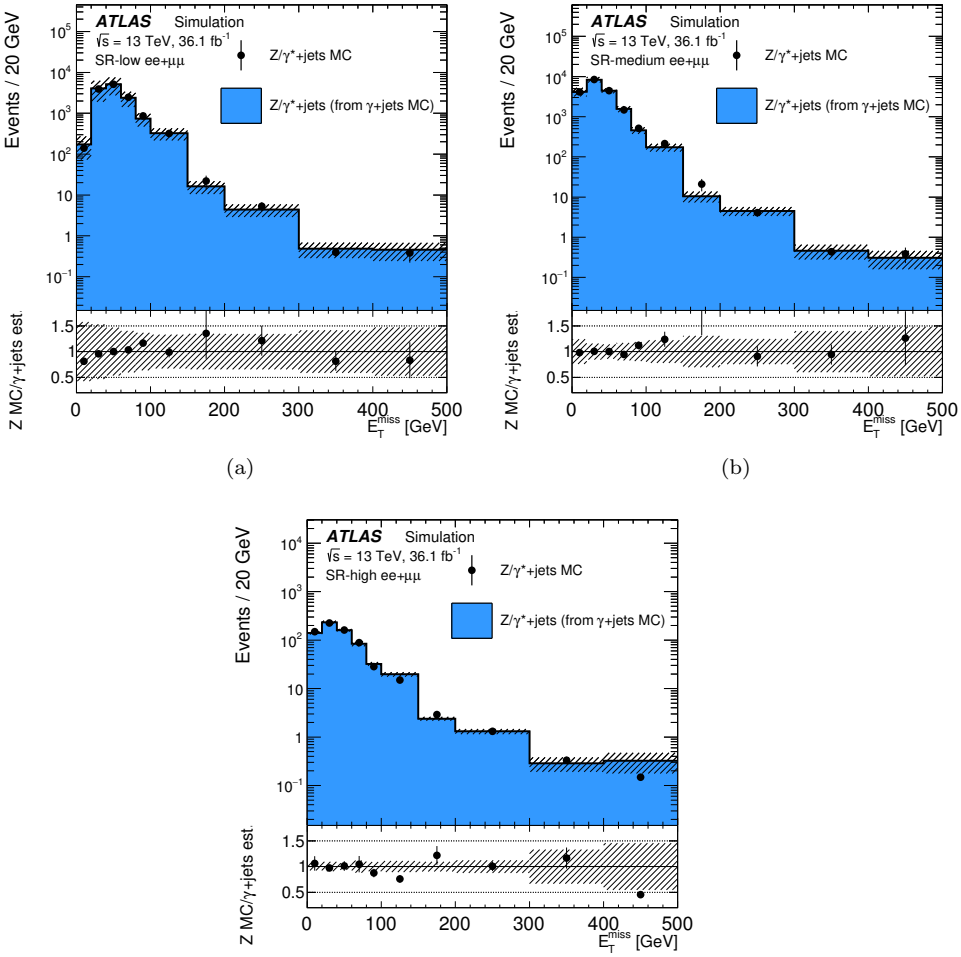

Figure 9.22: The $E_{\mathrm{T}}^{\mathrm{miss}}$ spectrum in $Z/\gamma^* + \text{jets}$ MC simulation compared to that of the $\gamma + \text{jets}$ method applied to $\gamma + \text{jets}$ MC simulation in SR-Low (a), SR-medium (b) and SR-high (c) [4]. The error bars on the points indicate the statistical uncertainty of the $Z/\gamma^* + \text{jets}$ MC simulation, and the hashed uncertainty bands indicate the statistical and reweighting systematic uncertainties of the $\gamma + \text{jets}$ background method. The lower panel shows the ratio of the observed and predicted yields.

9.6.3 Diboson background

The diboson background is taken directly from MC simulation and the accuracy of the WZ and ZZ modelling is assessed in the two regions VR-WZ and VR-ZZ detailed in Table 9.4. VR-WZ uses a three-lepton selection to create a region dominated by the process $WZ \to \ell\ell\ell\nu$. VR-ZZ applies a four-lepton selection to isolate the process $ZZ \to \ell\ell\ell\ell$, which is very similar to the $ZZ \to \ell\ell\nu\nu$ in the MC simulation. Both VRs use a b-jet veto to reduce the $t\bar{t}$ background. Additional kinematic selections are applied to take the regions closer to the SRs while still retaining sufficient statistics to assess the modelling. Figure 9.23 shows the invariant mass distributions in VR-WZ and VR-ZZ in data and simulation. Both the yields and the shapes of the distributions are observed to be well-modelled by the simulation in the VRs.

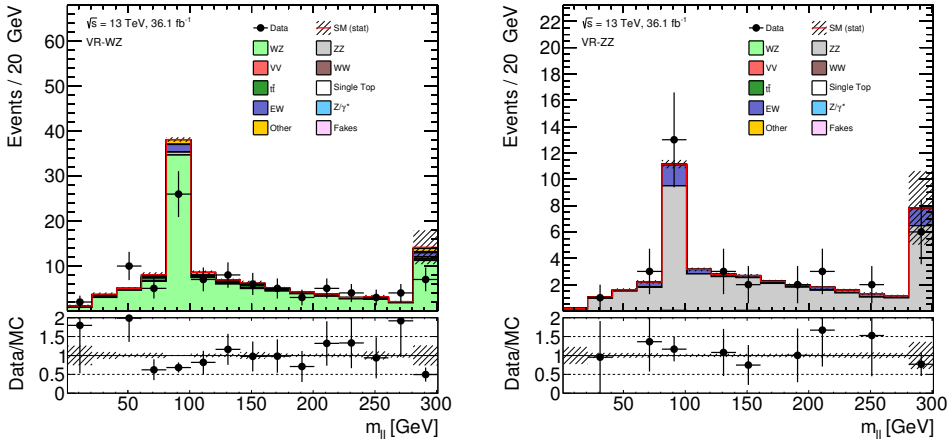

(a) (b)

Figure 9.23: The invariant mass distribution in VR-WZ (a) and VR-ZZ (b) in data and simulation. Only statistical uncertainties are applied to the simulation. The bottom panel shows the ratio between the yields observed in data and MC.

9.7 Systematic uncertainties

For the final interpretation of the analysis results, systematic uncertainties must be taken into account on the estimated background yields as well as the signal model yields in the SRs. The uncertainties associated with the data-driven FS and $Z/\gamma^* + \text{jets}$ estimates were discussed in Sections 9.6.1 and 9.6.2 respectively, and are the dominant uncertainties on the total background estimate.

The MC-based estimates and the signal are assigned experimental and theoretical uncertainties on the simulation. All samples are normalised to the integrated luminosity of the dataset used for the analysis, for which an uncertainty of 2.1 % is assigned. Other experimental uncertainties include the jet energy scale and resolution [149], the $E_{\mathrm{T}}^{\mathrm{miss}}$ soft-term resolution and scale [150], lepton energy scales and resolutions, and uncertainties on lepton trigger, reconstruction and identification efficiencies [88, 89]. All these uncertainties are derived centrally in ATLAS by the Combined Performance groups from comparisons between simulation and data.

The theoretical uncertainties correspond to all assumptions made in the MC simulation. They include uncertainties on the renormalisation and factorisation scales, cross section, and PDF set. For the diboson background, the uncertainties on the scales are computed by varying the nominal scale values up and down by a factor of two in the simulation. Additional uncertainties are also derived from a comparison between the nominal SHERPA sample and a POWHEG sample, as well as a variation of parton showering schemes. The generator comparison dominates the total uncertainty, which is generally below 20 %. For rare-top processes, a total uncertainty of 26 % is assigned to the cross section [99, 151–153]. The signal models have both the central value and uncertainty on cross sections taken from an envelope of predictions using different scales and PDF sets [127].

9.8 Validation of background estimates

The combined background estimate is validated in VR-Low, VR-Medium, and VR-High, detailed in Table 9.4. These regions include all kinematic selections of the SRs, but require the $E_{\mathrm{T}}^{\mathrm{miss}}$ to be between 100-200 GeV, making them orthogonal to the SRs. The full FS procedure is applied to data in three analogous different-flavour regions, while the $Z/\gamma^{*}+$jets estimate simply changes the final requirement on the modelled $E_{\mathrm{T}}^{\mathrm{miss}}$ spectrum. Figure 9.24 shows the estimated background compared to the observed data in all $m_{\ell\ell}$ bins with all systematic uncertainties applied. The estimates are seen to model the observed data well within the uncertainties.

9.9 Results

Expected and observed numbers of events in the Edge and On-Z searches are shown in Table 9.5, integrated over all $m_{\ell\ell}$ bins and for the On-Z bins separately, for each SR. The uncertainties are partially correlated, resulting in a final uncertainty which does not correspond to a simple addition in quadrature of each background uncertainty. Figure 9.25 shows the observed and expected yields for the full $m_{\ell\ell}$ distribution. No significant excess of data above the SM expectation is observed.

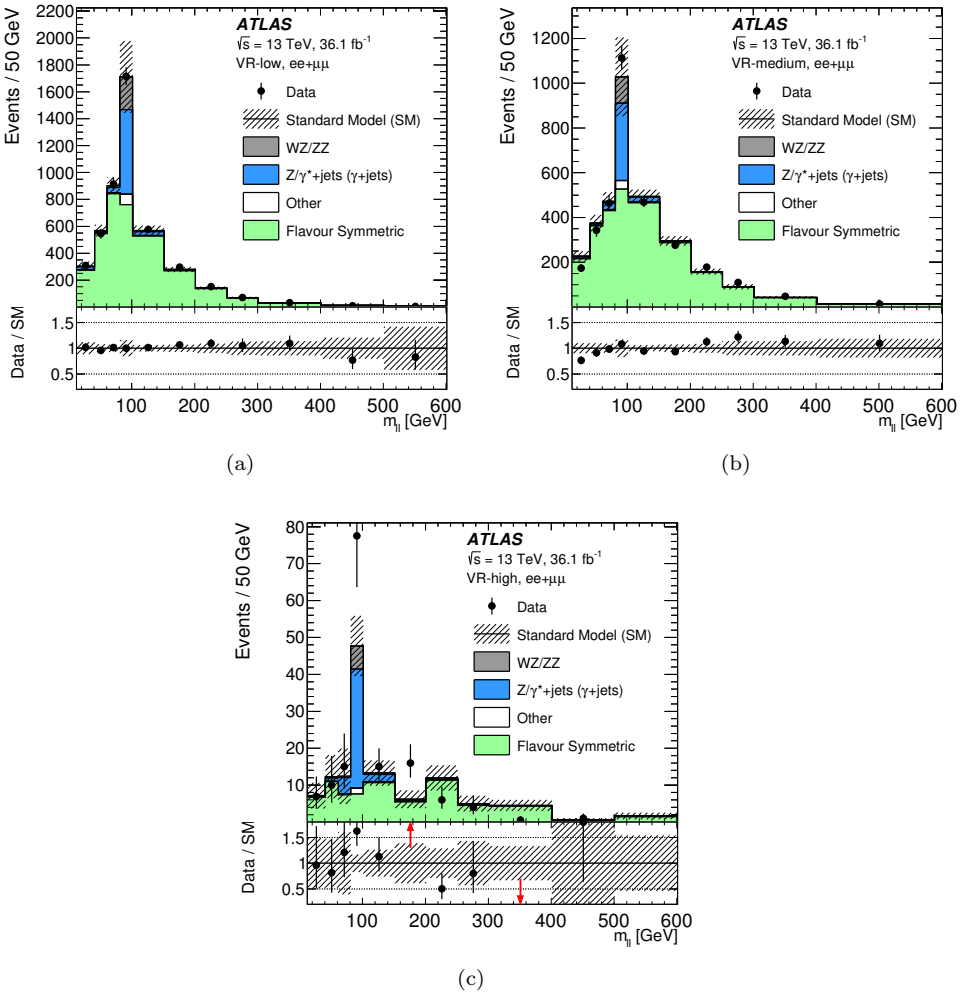

Figure 9.24: Observed and expected yields in all $m_{\ell\ell}$ bins in VR-Low (a), VR-Medium (b) and VR-High (c) [4]. The lower panels show the ratio between the data and the background prediction. The hashed bands include all statistical and systematic uncertainties.

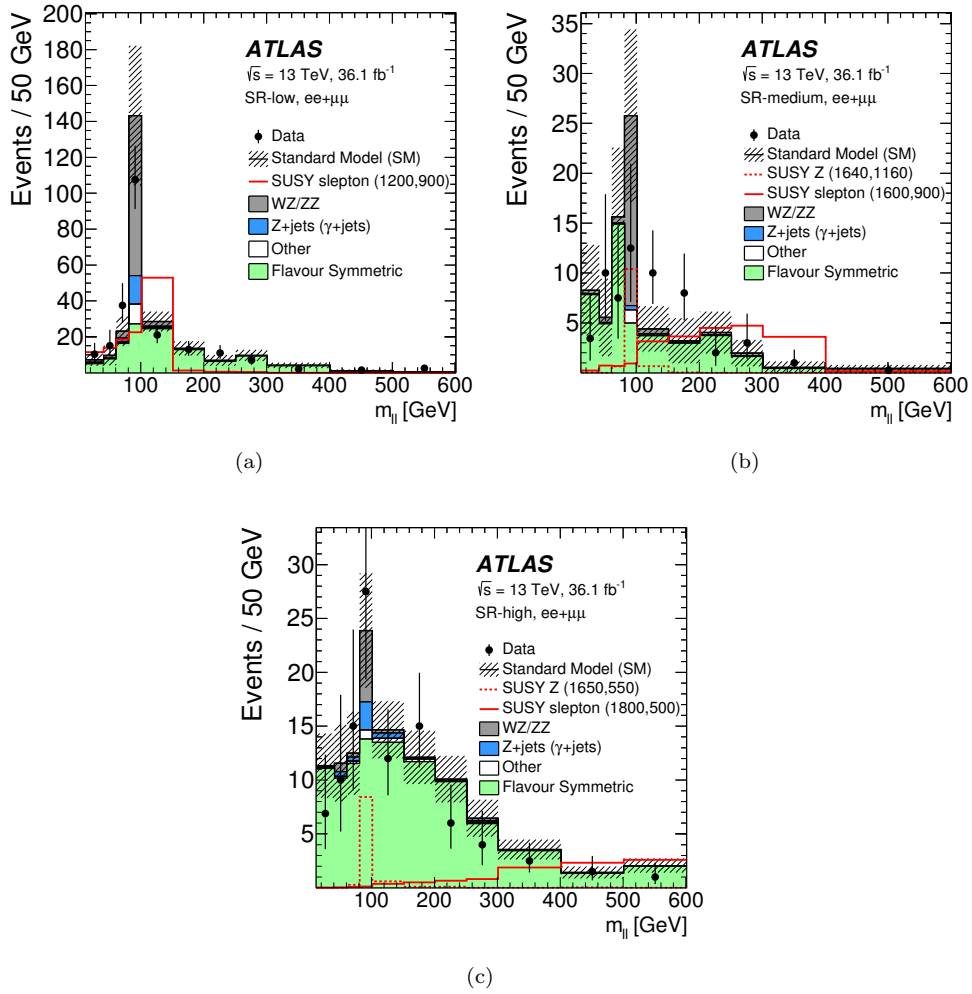

Figure 9.25: Observed and expected yields in all $m_{\ell\ell}$ bins in SR-Low (a), SR-Medium (b) and SR-High (c) [4]. The hashed bands includes all statistical and systematic uncertainties. Two example signal models from the Slepton and On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ grids are overlaid. For the Slepton model, the numbers in parentheses in the legend indicate the gluino and $\tilde{\chi}_1^0$ masses of the example model point. In the case of the On-Z model, the numbers in parentheses indicate the gluino and $\tilde{\chi}_2^0$ masses

Table 9.5: Breakdown of the expected background and observed data yields for the analysis SRs. The quoted uncertainties include statistical and systematic contributions.

	SR-Low	SR-Medium	SR-High	SR-Medium On-Z	SR-High On-Z
Observed events	134	40	72	5	11
Expected background events	144 ± 22	40 ± 10	83 ± 9	10.3 ± 3.4	9.6 ± 2.1
FS events	86 ± 12	29 ± 9	75 ± 8	2.0 ± 1.6	5.5 ± 1.7
$Z/\gamma^* +$ jets events	9^{+13}_{-9}	$0.2^{+0.8}_{-0.2}$	2.0 ± 1.2	$0.3^{+0.9}_{-0.3}$	$0.95^{+1.9}_{-1.0}$
WZ/ZZ events	43 ± 12	9.8 ± 3.2	4.1 ± 1.2	7.6 ± 2.9	2.7 ± 1.0
Rare top events	6.7 ± 1.8	1.20 ± 0.35	1.8 ± 0.5	0.54 ± 0.14	0.3 ± 0.1

9.10 Interpretation

Given that the observed data is in agreement with the SM model expectation, the results can be used to set exclusion limits on the targeted signal models. The limits are constructed with the CL_s technique [154, 155] which uses a maximum likelihood fit [156] based on the expected signal and backgrounds in the SR as well as the observed data. In the fit, the signal strength is given by the parameter μ, and subsequent fits are performed at discrete μ values to determine the relative likelihood of each value. Using the likelihood, the probability of a background-only hypothesis p_{b} can be determined by setting $\mu = 0$, and the probability of a signal plus background hypothesis $p_{\mathrm{s+b}}$ can be obtained by setting $\mu = 1$. The confidence limit is constructed as the ratio

$$\mathrm{CL}_s = \frac{p_{\mathrm{s+b}}}{1 - p_{\mathrm{b}}}, \tag{9.15}$$

which considers both the agreement with the signal plus background hypothesis and the agreement with the background-only hypothesis. An observation that agrees with neither hypothesis will have a weakened exclusion on the signal plus background hypothesis. Strong exclusions are therefore only achieved if the backgrounds are well understood and agree with the observed data. If CL_s falls below 5 %, the signal plus background hypothesis can be excluded at 95 %.

The analysis results are used to derive both expected and observed limits. All statistical and systematic uncertainties are treated as nuisance parameters in these fits. Gaussian models for the nuisance parameters are used for all signal and background uncertainties, except for the statistical uncertainties on the background estimates, which are interpreted as Poissonian. Expected exclusion limits are constructed by assuming that the observed data precisely matches the prediction, and 1σ uncertainty bands are derived from the shape of the likelihood curve around the minimum. The observed limit uses the actual observation of data in the SR to set

exclusion limits. Any excess above the expected background will result in weaker limits than expected, and any deficit will result in stronger limits. The observed exclusion is displayed with error bands that represent a 1σ variation in the cross section of the signal models.

For the Edge search, a profile likelihood shape fit [156] is performed on each of the binned $m_{\ell\ell}$ distributions in Figure 9.25 to derive the limits. The resulting exclusion contours for the Slepton model are shown in Figure 9.26a, where the limit from the SR with the best expected exclusion is selected for each grid point. In addition, the contours include limits from a complementary ATLAS low lepton-p_{T} search [4] which is optimised for a compressed scenario with a small neutralino mass difference. The region where the low lepton-p_{T} search starts to dominate the combination can be seen close to the $m_{\tilde{g}} = m_{\tilde{\chi}_1^0}$ diagonal, where there is a kink in the contour around $m_{\tilde{g}} = 1400\,\mathrm{GeV}$. For higher mass splittings, the results of the Strong-2L analysis completely dominate the contour. Around a gluino mass of 1.85 TeV, the observed limit drops below the expected limit by 200 GeV. This is caused by a slight excess in the highest $m_{\ell\ell}$ bin in SR-Low. Figure 9.26b shows the exclusion contours for the Z^* model, where similar trends can be observed. Here, the result in SR-Medium provides the strongest constraint at high gluino mass, reaching beyond 1.6 TeV. Overlaid on the figure is the observed limit from the previous Run 2 Strong-2L analysis using the $14.7\,\mathrm{fb}^{-1}$ [5] dataset. The large improvements in the sensitivity result partly from the reoptimisation of the analysis and partly from the increased dataset.

Figure 9.27 shows the limits resulting from the likelihood fit using only the On-Z bins in the SRs for the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ and $\tilde{q}$-$\tilde{\chi}_2^0$ models. The limits correspond to a combination of the results in SR-Medium and SR-High, with SR-Medium providing the strongest constraint close to the $m_{\tilde{g}} = m_{\tilde{\chi}_1^0}$ diagonal. A kink can be seen in the observed limit contour at the point at which the SR with the best-expected sensitivity changes from SR-Medium to SR-High. Gluinos with a mass below 1.65 TeV for a $\tilde{\chi}_2^0$ mass of 1.2 TeV and squarks with masses below 1.26 TeV for a $\tilde{\chi}_2^0$ mass of 900 GeV are excluded.

The limits for the On-shell $\tilde{g}$-$\tilde{\chi}_1^0$ model are shown in Figure 9.28, where a similar kink when the best limit changes from SR-Medium to SR-High can be observed. Gluino masses are excluded up to 1.6 TeV for a $\tilde{\chi}_1^0$ mass of 200 GeV. The large improvement in sensitivity at large gluino mass compared to the previous analysis is mainly due to the reoptimisation with the low $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement in SR-High.

In the two grids displayed in Figure 9.27, the exclusion extends all the way to the $m_{\tilde{g}} = m_{\tilde{\chi}_1^0}$ diagonal, where the $\tilde{\chi}_2^0$ is produced with a very small p_{T}. Since the mass of the $\tilde{\chi}_1^0$ is fixed to 1 GeV, the neutralino mass difference is large and there is enough energy in the system to produce a Z boson and a $\tilde{\chi}_1^0$ with significant p_{T}, leading to high $E_{\mathrm{T}}^{\mathrm{miss}}$ and lepton p_{T}. In the grid shown in Figure 9.28 the neutralino mass difference is fixed to 100 GeV. A $\tilde{\chi}_2^0$ produced at rest therefore has just enough energy to decay to a Z boson and a $\tilde{\chi}_1^0$, and neither decay product has

significant p_{T}. In this case, there is very little $E_{\mathrm{T}}^{\mathrm{miss}}$ and the the leptons will have lower p_{T} values. The SRs are therefore not sensitive to the signal and the limit does not reach all the way up to the $m_{\tilde{g}} = m_{\tilde{\chi}_1^0}$ diagonal.

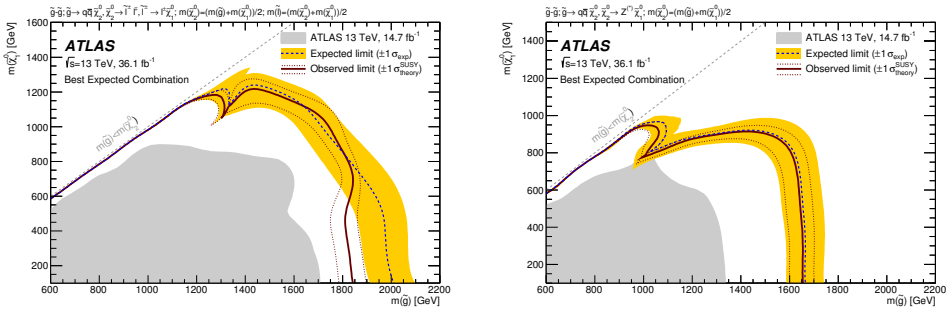

(a) (b)

Figure 9.26: Expected and observed limits for the Slepton model (a) and the Z* model (b) derived from the combination of the results in SR-Low, SR-Medium and SR-High [4].

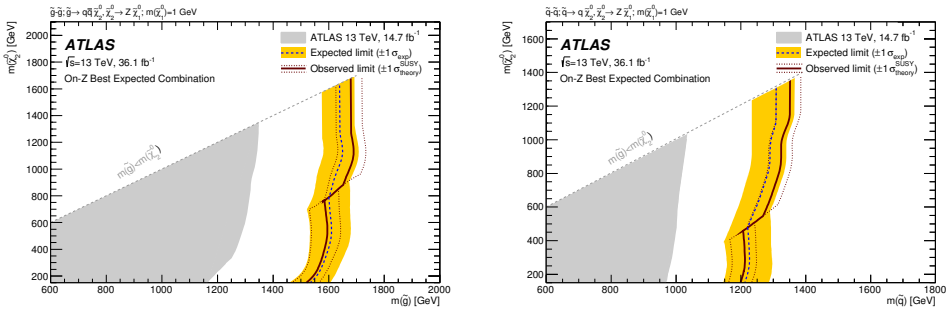

(a) (b)

Figure 9.27: Expected and observed limits for the On-shell $\tilde{g}$-$\tilde{\chi}_2^0$ (a) and $\tilde{q}$-$\tilde{\chi}_2^0$ (b) signal models derived from the combination of the results in the On-Z bins of SR-Medium and SR-High [4].

A model-independent interpretation of the results is performed by deriving upper limits at 95 % CL on the number of observed events that could be attributed to non-SM processes based on the predicted backgrounds. The overlapping $m_{\ell\ell}$ windows are used for this purpose. In cases where an edge-like signal stretches over a large $m_{\ell\ell}$ range, the exclusive bins used in the shape fit may truncate the lower-$m_{\ell\ell}$ tail, making them less sensitive. In addition, the windows with $m_{\ell\ell} < 81\,\mathrm{GeV}$ are sensitive to generic models without sleptons, with neutralino mass splittings below

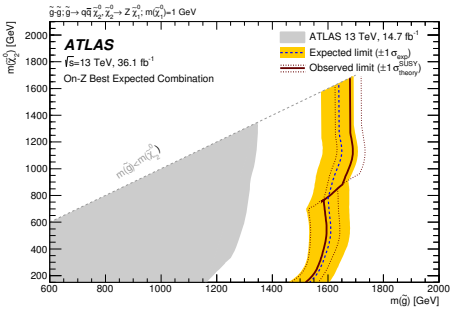

Figure 9.28: Expected and observed limits for the On-shell $\tilde{g}$-$\tilde{\chi}_1^0$ signal models derived from the combination of the results in the On-Z bins of SR-Medium and SR-High [4].

the Z-boson mass, and the window with $81 < m_{\ell\ell} < 101\,\text{GeV}$ is sensitive to models with mass splittings above the Z-boson mass. Figure 9.29 shows the observed data and expected background in all windows. The largest excess, with a local significance of 2σ, is observed in SR-Medium in the $101 < m_{\ell\ell} < 201\,\text{GeV}$ window. The various $m_{\ell\ell}$ windows exclude signals with visible cross sections between 0.12-1.56 fb, where the visible cross section includes the acceptances and efficiencies of the SRs.

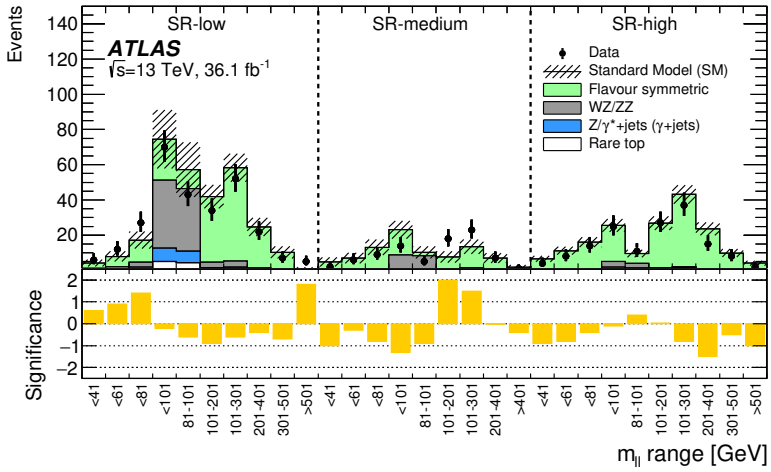

Figure 9.29: Observed and expected yields in the overlapping $m_{\ell\ell}$ windows in SR-Low, SR-Medium and SR-High [4]. The hashed band includes all statistical and systematic uncertainties. The lower panel shows the significance of the difference between the observed data and the predicted background.

Chapter 10

Search for displaced vertices in multijet events

The standard ATLAS event reconstruction techniques are designed to reconstruct particles that originate from the pp interaction point. Searches for BSM physics based on the standard physics objects, like the Strong-2L analysis, are therefore largely insensitive to models which predict LLPs. Such models were introduced in Section 2.3 and may include anything from BSM particles decaying shortly after the hard-scatter interaction, whose decay products can be used to point to an origin that is not the PV, to particles that have a sufficient lifetime to pass through the ATLAS detector without decaying. In order to regain sensitivity to these scenarios, it is important to perform dedicated searches using non-standard reconstruction techniques.

Several analyses targeting various LLP models are performed by the experiments at the LHC and correspond to either direct or indirect searches. The first category uses experimental signatures arising from direct interaction of the LLP with the detector. Examples include searches for heavy, slow-moving LLPs that give rise to delayed detector signals or a high ionisation energy loss signature [157–160], as well as searches for charged LLPs decaying partway through the tracking detectors giving rise to disappearing tracks [161–164]. The second category indirectly searches for LLPs by attempting to reconstruct their decays into SM particles. This includes searches for jets with unusual or delayed calorimeter deposits [165–167] and searches for tracks or vertices which are displaced from the pp interaction point [168–174]. The DV+jets analysis described in this chapter uses the displaced vertex signature to search for LLPs decaying within the ATLAS ID in events with multiple highly energetic jets.

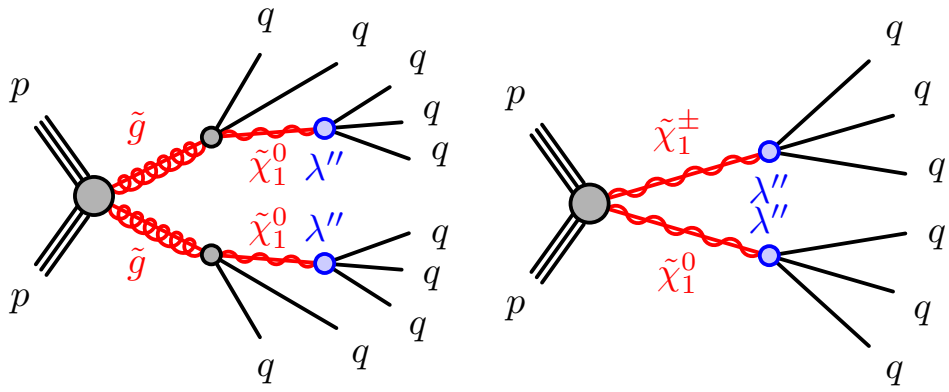

(a) (b)

Figure 10.1: The strong production (a) and electroweak production (b) signal models used for the interpretation of the DV+jets analysis.

10.1 Analysis overview

The DV+jets analysis is a signature-driven search and it is therefore sensitive to many different BSM models. Limits are set on two specific SUSY scenarios which contain the same long-lived neutralino that decays to three SM quarks via the RPV baryon-number-violating coupling λ''. The first scenario is shown in Figure 10.1a and involves gluino pair production where the gluinos decay promptly to two quarks and the long-lived neutralino. This model, which will be referred to as the *strong production* model, results in events with two displaced decays and ten jets. The second scenario, shown in Figure 10.1b, involves direct production of neutralinos and charginos that are long lived due to the same RPV coupling. This model is referred to as the *electroweak production* model and results in events with two displaced decays and six jets.

In the signal models, several charged particles are created in the jet fragmentation of the three quarks produced from the neutralino decay. Figure 10.2a shows the multiplicity of charged particles for a few different neutralino masses. On average, a 50 GeV neutralino produces around 10 charged particles while a 2050 GeV neutralino produces around 70 charged particles. Figure 10.2b shows the combined invariant mass of all charged particles with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$. This distribution peaks at around half of the neutralino mass, as expected if approximately half of the decay products are charged. The corresponding reconstructed quantity is the *visible DV mass* m_{DV}, defined as the invariant mass of all tracks associated to the DV. As the neutralino is expected to be more massive than SM particles, the DVs selected for the analysis are required to have a large track multiplicity and a large visible mass.

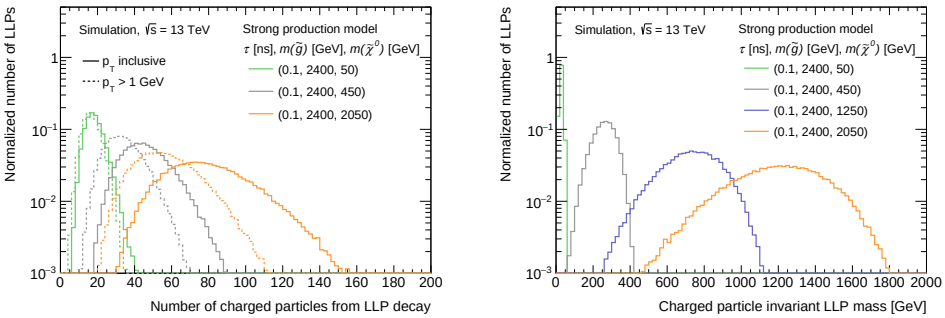

(a) (b)

Figure 10.2: Characteristics of the neutralino decay for different neutralino masses in the strong production model. (a): Number of charged particles resulting from the decay, inclusive in the p_{T} and limiting to $p_{\mathrm{T}} > 1\,\mathrm{GeV}$, (b): Invariant mass of all charged particles with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ resulting from the decay.

The reconstruction of DVs relies on event processing with LRT and secondary vertexing, described in Sections 5.1.5 and 5.1.6. The performance of these algorithms vary with the radius in the detector, and the sensitivity to a given model is therefore largely determined by the LLP decay position. For neutralino decays, the secondary vertex efficiency, shown in Figure 5.8, is above 70 % up to a radius of around 300 mm after which it quickly decreases. The sensitivity is therefore highest for decays within this volume. Figure 10.3a shows the radius of the neutralino decay for a few different mean proper lifetimes τ. For the two shortest neutralino lifetimes, 0.01 ns and 0.1 ns, all decays are contained within 300 mm. The measured decay length however also depends on the Lorentz boost factor of the neutralino. For a particle with mean proper lifetime τ, the average measured decay length is given by $\beta\gamma c\tau$ where $\beta = v/c$ and $\gamma = E/m$ is the Lorentz factor. The effect of the Lorentz boost is visible in Figure 10.3b which shows the decay radius for $\tau = 0.1$ ns neutralinos with different masses in the strong production model. The sample with the smallest neutralino mass is produced with the highest Lorentz factor, and the decay radius distribution therefore extends towards the largest values.

The main sources of background in the analysis stem from hadronic interactions where primary particles interact with the material in the detector, and from instrumental effects where unrelated DVs or tracks are merged in the reconstruction algorithms. Small contributions can result from decays of long-lived SM particles, but most such processes are efficiently removed by requiring high visible DV masses and track multiplicities.

The analysis is influenced by a number of previous ATLAS LLP searches for DVs in combination with an additional signature used to trigger and select the events. One

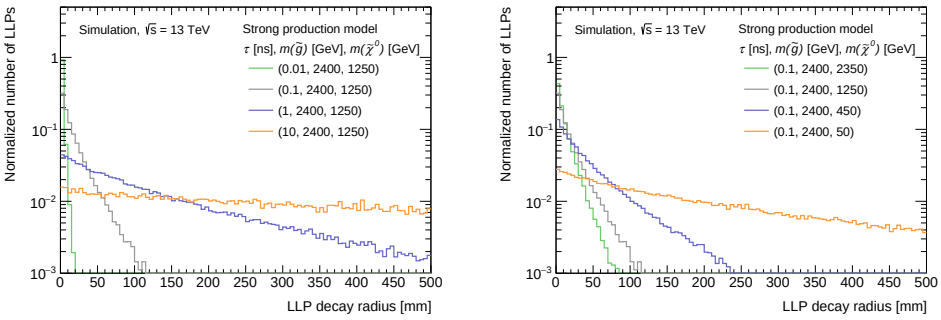

(a) (b)

Figure 10.3: Radius of the neutralino decay in the strong production model with varying mean proper neutralino lifetime (a) and varying neutralino mass (b).

example is the $DV{+}E_T^{miss}$ analysis [168] which searches for a DV in combination with large missing transverse momentum and is performed with 32.8 fb^{-1} of Run 2 pp collision data. Another example is the *DV+muon* analysis [169] which searches for events with a DV together with a displaced muon and exploits 136 fb^{-1} of the Run 2 dataset. These analyses use similar reconstruction techniques and analysis strategies and interpret the results in terms of various SUSY models. Both searches are in principle free from background, with expected event yields below one, and show no significant excess in the observed number of events. Figure 10.4 shows lower mass limits set for long-lived gluinos by several analyses performed by the ATLAS collaboration, including the Run 2 DV+$E_{\mathrm{T}}^{\mathrm{miss}}$ search. For gluinos decaying within the ID the DV search provides the best sensitivity to the signal model. This is a result of a high vertex reconstruction efficiency in combination with the low background, and demonstrates the power of the DV signature in searches for LLPs.

10.2 Data and MC samples

The DV+jets analysis is performed on the full Run 2 dataset corresponding to 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data. Events are collected using a combination of multijet triggers which are further described in Section 5.1.5. Simulated samples are used to model the signal scenarios, to study the backgrounds, and to validate the background estimate.

The signal scenarios shown in Figure 10.1 represent simplified SUSY models. Events are generated at NLO in MadGraph5_aMC@NLO 2.6.2 (strong production) or MadGraph5_aMC@NLO 2.8.1 (electroweak production) [99] and are passed to Pythia 8.240 (strong production) or Pythia 8.244 (electroweak production) [105]

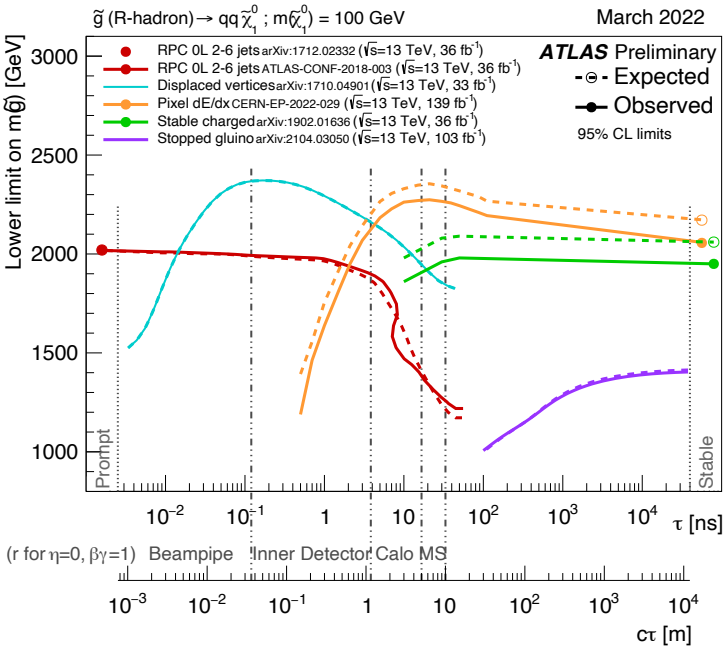

Figure 10.4: Constraints on the gluino mass-lifetime plane for a split-SUSY model with a long-lived gluino [175]. The solid lines indicate the observed limits, while the dashed lines indicate the expected limits. The dots represent results for which the particle is assumed to decay promptly or be stable on the scale of the detector.

using the A14 tune [132] and NNPDF2.3lo PDF set [25], to model showering, hadronisation, and the underlying event. EvtGen 1.6.0 [108] is used to simulate the decay of heavy-flavour particles.

In the strong production model, the pp interaction produces a pair of gluinos, which promptly decay with a 100 % branching ratio to the long-lived neutralino and a SM $q\bar{q}$ pair. The neutralino then decays to three SM quarks or antiquarks via the λ'' coupling. All other RPV couplings are assumed to be zero. The model has three free parameters; the neutralino mass $m_{\tilde{\chi}_1^0}$, the gluino mass $m_{\tilde{g}}$, and the mean proper neutralino lifetime τ. A signal grid is generated with $m_{\tilde{\chi}_1^0}$ varying between 10-2500 GeV, $m_{\tilde{g}}$ between 1600-2600 GeV, and lifetimes between 0.01-10 ns. For the electroweak production model, the pp interaction produces two neutralinos or charginos in pure higgsino states. The possible combinations are $\tilde{\chi}_2^0$ $\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm$ $\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm$ $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^\pm$ $\tilde{\chi}_2^0$. The neutralinos and charginos decay directly to three SM quarks or antiquarks via the same λ'' coupling. A nearly mass-degenerate spectrum is assumed, with a mass difference of 1 GeV between the lightest neutralino and the other neutralinos and charginos. This ensures that all the electroweakinos decay through the λ'' coupling. The model has two free parameters; the neutralino mass and the mean neutralino lifetime. A grid is generated with $m_{\tilde{\chi}_1^0}$ varying between 100-1700 GeV and lifetimes between 0.01-10 ns.

The background is studied in a simulated dijet sample where a pair of quarks is produced in the hard-scatter interaction. The process is simulated at LO with Pythia 8.235 [104] using the A14 tune [132] and the NNPDF2.3lo PDF set [25]. EvtGen 1.6.0 [108] is used to simulate the decay of heavy-flavour particles. The ATLAS dijet MC samples are split into several slices, depending on the p_{T} of the leading parton in the pp interaction. Only the slice corresponding to a partonic p_{T} between 400 GeV and 800 GeV is reconstructed with LRT for use in this analysis.

All MC samples include the effects of pileup, modeled by overlaying minimum bias pp events generated with Pythia 8.210 [104], using the NNPDF2.3lo [25] PDF set and the A3 tune [118]. Both the signal and the background samples are produced in three versions, with pileup profiles corresponding to 2015-2016, 2017 and 2018 data respectively. For the signal samples, 10k events are generated for each signal point and each pileup profile. For the dijet sample, 13M events are generated for the 2018 pileup profile, while 1M events are generated for each of the 2015-2016 and 2017 profiles.

10.3 Identification of physics objects

The main physics objects used are jets, tracks, and secondary vertices. Tracks are reconstructed using the the standard tracking and LRT, described in Sections 5.1.1 and 5.1.5, and secondary vertices are reconstructed using the algorithm described in Section 5.1.6. The jets are used to trigger on the events and are reconstructed

using the algorithms described in Section 5.2.3. No JVT selection is applied, since this would reject almost all displaced jets. The background estimate relies on so-called *track jets*, which are constructed with the anti-k_t algorithm, using standard tracks which pass the Loose track selection as input. The jet algorithm is run once for each primary vertex in the event, selecting only the tracks associated to that vertex by requiring $|d_0| < 2\,\text{mm}$ and $|z_0|/\sigma_{z_0} < 3$. This prevents tracks associated to different PVs to be clustered together in the same track jet. Only tracks with $p_\text{T} > 1\,\text{GeV}$ are used to build the jets.

10.4 Event selection

The event selection consists of several steps, with requirements applied either at event or DV level. Events are first selected using multijet triggers, and are filtered out for event processing with LRT using jet information from the HLT trigger. After reconstruction, events are subject to a preselection based on the offline jet information to ensure consistency with the trigger. These first selections are motivated by the need for LRT and the jet requirements are designed to reduce the rate of events to a manageable level.

The jet selections are followed by a set of *baseline* requirements on the DVs in the events. These are designed specifically for the analysis to ensure high-quality measurements of the DV properties and to reduce the rate of background DVs. Events with at least one DV passing the baseline selections are considered for further analysis. The baseline DVs are also subject to a track cleaning to reject tracks that are incorrectly associated to the DV in the secondary vertexing. Finally, two signal regions (SRs) are defined with requirements on the track multiplicity and the visible DV mass, to reduce the background from SM processes. Orthogonal validation regions (VRs) dedicated to the background estimates are also defined using similar requirements.

10.4.1 Trigger and event filter

Events are collected with a combination of unprescaled multijet triggers, each requiring between four to seven jets with various p_T thresholds, as summarised in Table 10.1. Figures 10.5 and 10.6 display the multiplicity of jets and the p_T spectra of the leading and subleading jets in events from the two signal models. The electroweak production model results in fewer and lower-p_T jets compared to the gluino model, as expected from the number of final-state quarks in the diagrams shown in Figure 10.1. In both cases, the trigger strategy is highly efficient to select events from the signal.

Two different filters are used to select events for processing with LRT; the *High-p_T* filter and the *Trackless* filter. The jet multiplicity and p_T requirements of the filters

Table 10.1: Jet multiplicity and p_{T} requirements for the lowest-p_{T} Run 2 unprescaled HLT multijet triggers, by year [145, 146, 176, 177].

Jet multiplicity	**Jet p_{T} threshold [GeV]**			
	2015	2016	2017	2018
4	100	100	100	120
5	60	70	70	85
6	45	60	60	70
7	-	45	45	45

Table 10.2: Jet multiplicity and p_{T} requirements for the Trackless jet and High-p_{T} parts of the event filter.

Jet multiplicity	**Jet p_{T} threshold [GeV]**	
	Trackless filter	High-p_{T} filter
4	100	220
5	75	170
6	50	100
7	45	75

are listed in Table 10.2. The High-p_{T} filter is primarily designed to target the strong production model while the Trackless filter has reduced jet-p_{T} thresholds to retain sensitivity to the electroweak production model. In order to keep the event rate down, the filter is complemented by a trackless jet requirement. A jet is considered trackless if the sum of the p_{T} of all tracks passing the Loose track selection inside the jet is less than 5 GeV. Events are required to have at least one trackless jet with $p_{\mathrm{T}} \geq 70\,\mathrm{GeV}$, or at least two trackless jets with $p_{\mathrm{T}} \geq 50\,\mathrm{GeV}$ to pass the Trackless jet selection.

10.4.2 Baseline jet selections

The minimum jet-p_{T} requirements are increased at the analysis level to ensure that the trigger is fully efficient for the selected events and to be above the minimum jet-p_{T} threshold in all data-taking years. Two baseline jet selections are defined, corresponding to the High-p_{T} and Trackless filters, with jet multiplicity and p_{T} requirements listed in Table 10.3. In order to pass the Trackless baseline jet selection, events are also required to have one trackless jet with $p_{\mathrm{T}} > 78\,\mathrm{GeV}$ or two trackless jets with $p_{\mathrm{T}} > 56\,\mathrm{GeV}$.

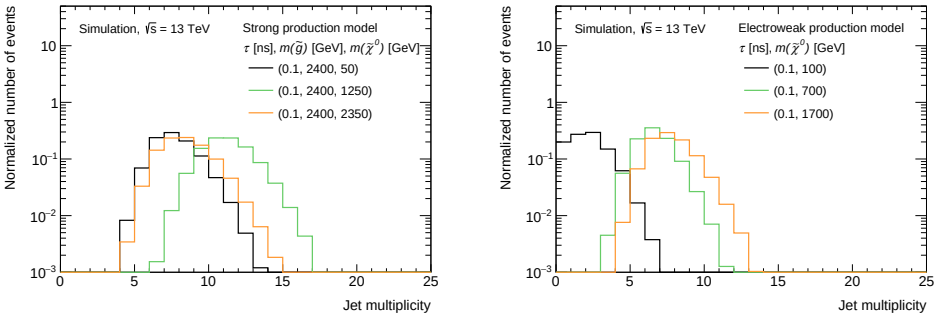

(a) (b)

Figure 10.5: Multiplicity of jets fulfilling $p_{\mathrm{T}} > 50\,\mathrm{GeV}$ in events corresponding to a selection of points in the strong production (a) and electroweak production (b) model grids. No event selection is applied.

Table 10.3: Jet multiplicity and p_{T} requirements for the Trackless and High-p_{T} baseline jet selections.

Jet multiplicity	**Jet p_{T} threshold [GeV]**	
	Trackless selection	High-p_{T} selection
4	137	250
5	101	195
6	83	116
7	55	90

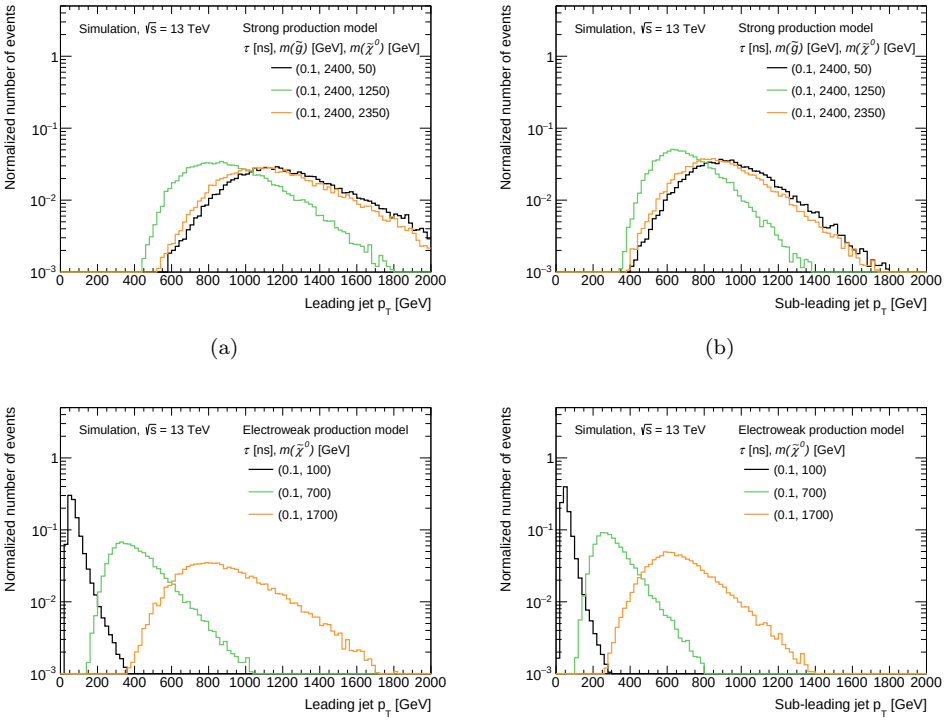

Figure 10.6: Jet p_{T} for the leading and subleading jets in events corresponding to a selection of points in the strong production (a,b) and electroweak production (c,d) model grids. No event selection is applied.

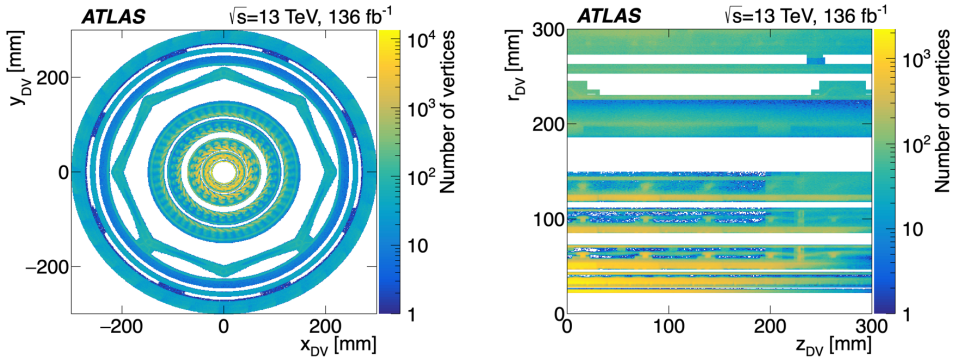

(a) (b)

Figure 10.7: The positions in the x-y plane (a) and r-z plane (b) of reconstructed DVs that fail the material map veto in the inclusive data sample used for the construction of the map [169]. The latter figure is restricted to positive z values only for presentation.

10.4.3 Baseline DV selections

Events passing the baseline jet selections are retained for analysis if they contain at least one DV satisfying the baseline DV selections in Table 10.4. The DVs are required to be within the *fiducial region*, given by $R < 300\,\text{mm}$ and $|z| < 300\,\text{mm}$, which corresponds to the detector volume where the secondary vertexing performs well. In addition, the DVs are required to be at least 4 mm from any primary vertex in the event, to reduce the background from decays of heavy-flavour particles. A goodness-of-fit criterion is also applied to ensure well-measured DV properties.

The final requirement vetoes DVs which are reconstructed within material-dense regions of the detector. This criterion is set to reduce the contribution of DVs created from hadronic interactions, which otherwise constitute a major background to the analysis. The veto is imposed via a three-dimensional map of detector material that is constructed from measured positions of low-mass vertices in an inclusive data sample. Figure 10.7 shows the material map. Regions corresponding to inside material and outside material are defined by a threshold on the minimum and maximum number of DVs in the map respectively. DVs are rejected if their position corresponds to an inside-material region, which effectively removes 48 % of the fiducial volume. Part of the inside-material region is used for some of the VRs, excluding the edges closest to the outside-material region, and corresponds to 42 % of the fiducial volume.

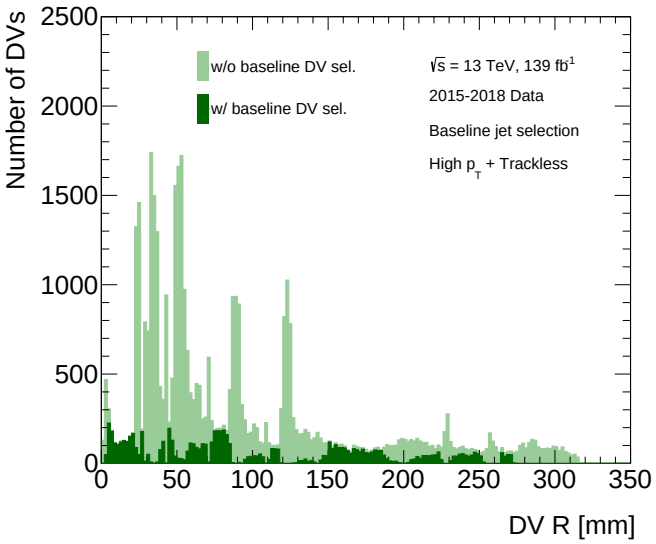

Figure 10.8: Radial positions of all DVs (light green) and of those passing the baseline DV selection (dark green) in events passing the High-p_{T} or the Trackless baseline jet selections.

Figure 10.8 shows the radial position of DVs in events passing the High-p_{T} or Trackless event selection, before and after the baseline DV selections are applied. The reduction seen close to zero radius results from the requirement on the distance to primary vertices and the lack of DVs above 300 mm from the fiducial region requirement. Before the baseline DV selection, clear peaks corresponding to material layers and support structures in the ID can be observed. These are effectively removed by the material map veto, resulting in the dips in the distribution after the baseline DV selection. The range $150 < R < 200$ mm, does not contain any material but still shows a reduction resulting from the goodness-of-fit requirement.

10.4.4 Track cleaning

All DVs passing the baseline selection are subject to a track cleaning with requirements optimised in the dijet and signal MC to reduce background DVs while still keeping a high efficiency for the two signal models. A summary of the DV-track cleaning is presented in Table 10.5 and tracks satisfying any rejection criteria are removed from the DVs. The criteria are based on the track p_{T} and the d_0-significance with respect to the PV, the *crossing angle* α_{cross}, and on the hit pattern. The crossing angle is defined as the angle in the transverse plane between the track at the point of closest approach to the DV and the vector from the PV to the DV. The

Table 10.4: Summary of the baseline DV selections.

Selection type	Criteria
Fiducial volume	$R_{\mathrm{DV}} < 300\,\mathrm{mm}$
	$\lvert z_{\mathrm{DV}}\rvert < 300\,\mathrm{mm}$
Distance to primary vertices	$r_{\mathrm{DV-PV}} > 4\,\mathrm{mm}$
Vertex fit quality	$\chi^2/N_{\mathrm{DoF}} < 5$
Outside material	Material map veto

Table 10.5: Summary of the track cleaning criteria.

Cleaning type	Track type	DV radius	Track rejection criteria
Hit pattern	All	-	Hits before the DV radius
Crossing angle	Attached	-	$\alpha_{\mathrm{cross}} > \pi/2$
	All	Outside beampipe	$\alpha_{\mathrm{cross}} < 0.2$ and $p_{\mathrm{T}} > 4\,\mathrm{GeV}$
$\boldsymbol{p}_{\mathrm{T}}$	All	-	$p_{\mathrm{T}} < 2\,\mathrm{GeV}$
	Attached	Outside beampipe	$p_{\mathrm{T}} < 3\,\mathrm{GeV}$
	Attached	Outside last pixel layer	$p_{\mathrm{T}} < 4\,\mathrm{GeV}$
$\boldsymbol{d_0}$**-significance**	All	Within beam pipe	$\lvert d_0\rvert/\sigma_{d_0} < 10$
	Attached	Within last pixel layer	$\lvert d_0\rvert/\sigma_{d_0} < 15$
	Selected	Outside last pixel layer	$\lvert d_0\rvert/\sigma_{d_0} < 10$

restriction on this angle serves to reject tracks which travel backwards from the DV and the hit pattern requirement serves to remove tracks which have hits both before and after the DV in radius.

The nature of the background is seen to change as a function of radius, and the requirements on the tracks therefore vary depending on the radial region of the DV. In addition, DVs may contain both selected and attached tracks which correspond to tighter and looser criteria on the track quality and the compatibility with the DV, as described in Section 5.1.6. Some requirements in the track cleaning are only applied to one type of tracks and some are applied to both. To be considered further in the analysis, an event must contain at least one DV with at least two selected tracks passing the DV-track cleaning.

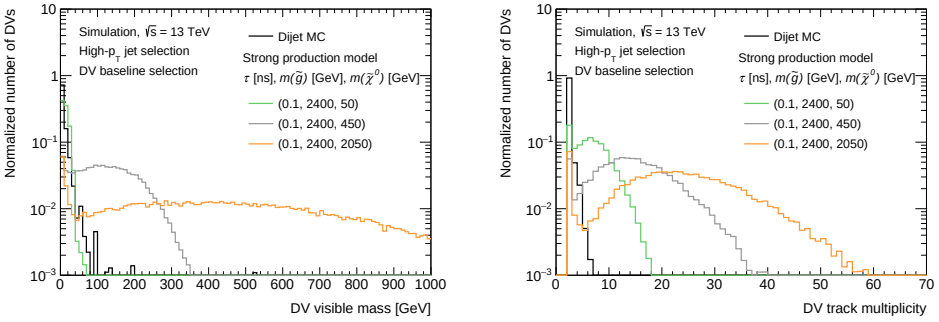

(a) (b)

Figure 10.9: Visible DV mass (a) and DV track multiplicity (b), as measured for a selection of points in the strong production signal model grid and in the dijet MC sample. The baseline High-p_{T} jet selection, baseline DV selection, and the track cleaning is applied to the events.

Table 10.6: Summary of the analysis SRs. The baseline jet selections are defined in Table 10.3. The events must contain at least one DV fulfilling the $n_{\mathrm{trk}}^{\mathrm{DV}}$ and m_{DV} requirements in the last two columns in order to be selected for a given region.

Region	Baseline jet selection	$n_{\mathrm{trk}}^{\mathrm{DV}}$	m_{DV}
High-p_{T}	Pass High-p_{T}	≥ 5	$> 10\,\mathrm{GeV}$
Trackless	Pass Trackless Fail High-p_{T}	≥ 5	$> 10\,\mathrm{GeV}$

10.4.5 Signal and validation regions

The final selection of SR events applies requirements on the track multiplicity and the visible DV mass computed with all tracks remaining after the DV-track cleaning. Figure 10.9 shows the distributions in these two variables for a few points in the strong production model grid together with the dijet background sample. For all but the lowest neutralino mass, the variables have a high discrimination power and can efficiently remove all background while keeping most of the signal. Two mutually exclusive SRs are defined as listed in Table 10.6 based on the baseline jet selections, and apply the same $n_{\mathrm{trk}}^{\mathrm{DV}} \geq 5$ and $m_{\mathrm{DV}} > 10\,\mathrm{GeV}$ requirements. The VRs utilise the sideband regions to select events with at least one DV with lower track multiplicity or lower visible mass, and also the inside-material region. The VRs are summarised in Table 10.7.

Table 10.7: Summary of the analysis VRs. The events must contain at least one DV fulfilling the material region, $n_{\text{trk}}^{\text{DV}}$, and m_{DV} requirements in the last three columns in order to be selected for a given VR. For each definition, two VRs are defined, corresponding to the two SR baseline jet selections detailed in Table 10.6. The inside material region is introduced in Section 10.4.3. The requirement that distinguishes the VR from the SR is indicated in bold.

Region	Material	$n_{\text{trk}}^{\text{DV}}$	m_{DV}
Low-$n_{\text{trk}}^{\text{DV}}$ Sideband	Outside	$\mathbf{= 4}$	$> 10\,\text{GeV}$
Low-m_{DV} Sideband	Outside	≥ 5	$\mathbf{< 10\,GeV}$
Inside Material	**Inside**	≥ 5	$> 10\,\text{GeV}$

10.5 Signal efficiency

The efficiency for the signal to pass each of the jet and DV selections is shown in Figures 10.10 and 10.11 for a few selected points in the strong and electroweak production signal grids. For the strong production model, the High-p_{T} event filter and baseline jet selection are highly efficient and retain more than 90 % of the signal. The first baseline DV selections are also highly efficient, with a maximum decrease in efficiency for the lowest neutralino mass from 90 % to 70 %. All grid points are however affected by the material map veto. The requirements on the DV-track multiplicity and visible mass have a minimal effect on the highest neutralino masses but reduce the sensitivity to the lowest neutralino masses where it is very difficult to discriminate between signal and background. Since the Trackless jet selection is orthogonal to the very efficient High-p_{T} selection, the Trackless SR has in principle no sensitivity to the strong production signal and will not be used for the interpretation. For the electroweak production signal model, the two SRs have complementary sensitivity. The High-p_{T} SR is sensitive to large neutralino masses which have enough high-momentum jets to pass the high jet-p_{T} thresholds, and the trackless SR recovers sensitivity to lower neutralino masses. Figures 10.12 and 10.13 show the final SR efficiencies after applying all jet and DV selections, for the two production modes. The highest sensitivity is achieved for lifetimes of approximately 0.1 ns, where the vast majority of neutralinos decay inside the fiducial volume with a decay length long enough to pass the displacement requirement.

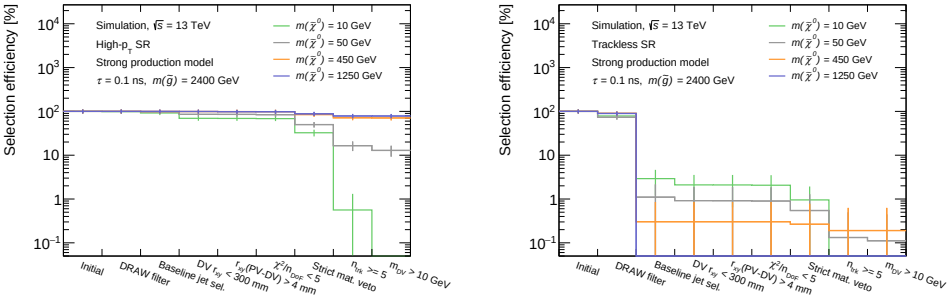

(a) (b)

Figure 10.10: Signal selection efficiency for each requirement in the jet and DV selection in the High-p_{T} SR (a) and the Trackless SR (b), for a few selected points in the strong production model grid. The gluino mass is fixed to 2400 GeV, the neutralino lifetime to 0.1 ns, and the neutralino mass is varied.

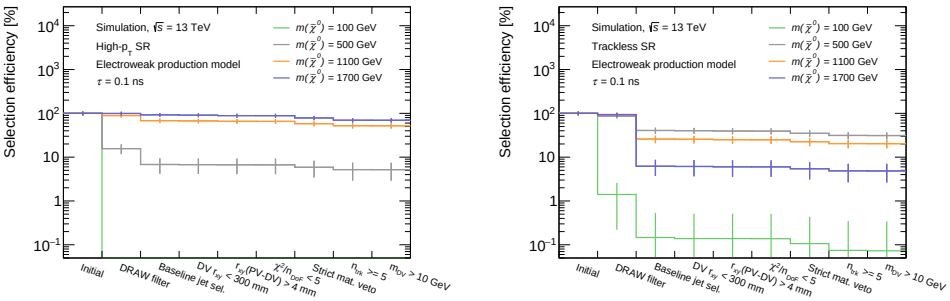

(a) (b)

Figure 10.11: Signal selection efficiency for each requirement in the jet and DV selection in the High-p_{T} SR (a) and the Trackless SR (b), for a few selected points in the electroweak production model grid. The neutralino lifetime is fixed to 0.1 ns and its mass is varied.

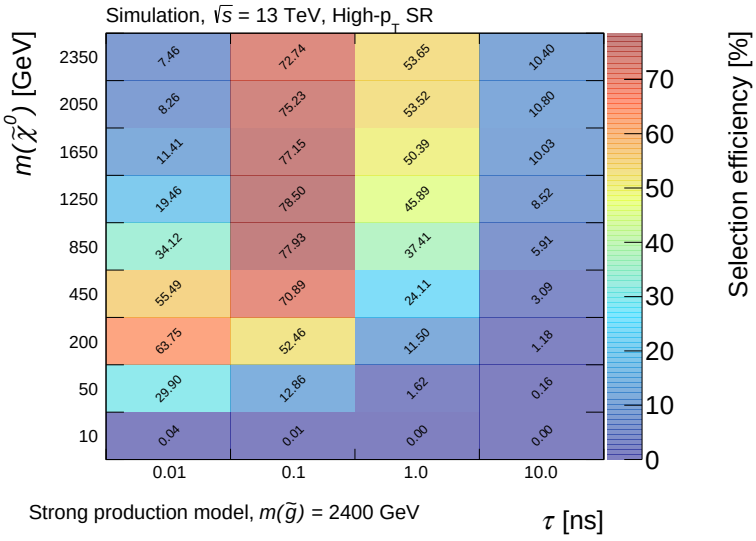

Figure 10.12: High-p_{T} SR selection efficiency for the strong production signal model as a function of neutralino mass and lifetime. The gluino mass is fixed to 2400 GeV.

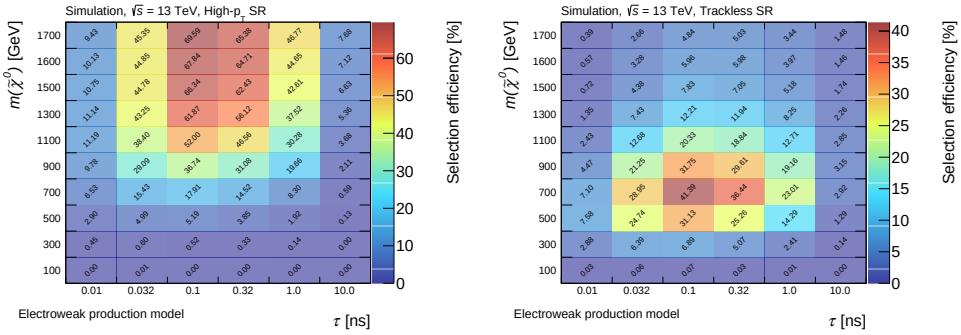

(a) (b) Trackless signal region

Figure 10.13: High-p_{T} SR (a) and Trackless SR (b) selection efficiency for the electroweak production signal model as a function of neutralino mass and lifetime.

10.6 Backgrounds

Searches for displaced, multi-track vertices benefit from a small background which is caused predominantly by material interactions or other instrumental effects. The three main sources of background-DV formation are listed below.

Hadronic interactions occur when hadrons produced in the event interact with the nuclei of detector material or gas. These interactions can produce multi-track vertices that may pass the SR selection if the mass is high. The majority of this background is removed by the material map veto. Residual background can however result from interactions with gas outside of the material, imperfections in the material map, or the resolution of the reconstructed hadronic-interaction DVs which may cause their positions to fall outside the material in the map.

Accidental crossings refer to any low-mass DV that by itself would not satisfy the SR selections, but is promoted to a higher mass and track multiplicity by an unrelated track passing close enough to be associated to the vertex.

Merged vertices occur when displaced, low-mass DVs from SM processes, such as K^0_S meson decays, are in close proximity and are merged in the secondary vertex algorithm thereby forming a higher-mass vertex.

The characteristics of the three background sources are studied in detail using the truth record in the dijet MC sample. The studies are described in Section 10.6.1 and the findings are important for the design and validation of the background estimation techniques. Two data-driven approaches are used for the background estimation; an inclusive method where all backgrounds are estimated simultaneously, and individual methods where the three sources of background are estimated separately. The inclusive method is used for the nominal background estimate and it is described in Section 10.6.2. It is based on the assumption that, in a multijet final state, the production of DVs is correlated with the presence of jets. The individual background estimates are based on the different mechanisms of background DV formation and they are discussed briefly in Section 10.6.3. Each estimation method requires independent CRs and VRs and introduces new systematic uncertainties to the total background estimate. The individual estimates are only used to cross check the nominal estimate.

10.6.1 Background studies in simulation

The dijet MC sample is used to classify DVs as hadronic interactions, accidental crossings, or merged vertices, based on the truth origin of the tracks attached to the vertices. Figure 10.14 illustrates how the origins of all tracks in a DV can be used to determine the type of process that gave rise to the vertex. In a hadronic interaction vertex, all tracks share the same origin whereas accidental crossing DVs or merged vertices contain tracks from multiple origins. A vertex from a long-lived

SM particle decaying in flight through its natural decay mode resembles a hadronic interaction, with all tracks sharing the same origin, but is as probable inside as outside detector material.

MC event selection

The simulation studies are performed using the three dijet MC samples described in Section 10.2, corresponding to the 2015-2016, 2017, and 2018 pileup profiles. All MC events are subject to the baseline DV selections and the track cleaning detailed in Sections 10.4.3 and 10.4.4, with the exception of the material veto. A map derived from MC, based on the location of detector material in the GEANT4 ATLAS model, is instead used to classify DVs as being inside or outside material. In order to increase the number of DVs available for the study, the baseline jet selection is not applied, as it otherwise would remove the majority of the available dijet events.

Classification of tracks and DVs in MC

The dijet MC is a standard sample which only contains the truth information about the hard-scatter process. A track associated to the hard scatter will therefore have a valid truth link to a charged truth particle in the sample, while a track associated to the overlaid pileup interactions will have an invalid truth link.

Tracks are therefore first classified as either *hard scatter* or *pileup*. Tracks with a truth match probability p_{match} above 0.5 and a valid truth link belong to the hard-scatter category and all remaining tracks to the pileup category. For each hard-scatter track, the truth link is used to retrieve the associated truth particle and access its barcode. Depending on the value, the track is then classified as either GEANT4 or generator. The track categories are summarised in Table 10.8.

Table 10.8: Definition of track categories, based on information in the truth record in the dijet MC sample.

Track category	p_{match}	**Truth Link**
Hard scatter	≥ 0.5	Valid truth link
Hard scatter GEANT4	≥ 0.5	Valid truth link to particle with barcode $> 200,000$
Hard scatter Generator	≥ 0.5	Valid truth link to particle with barcode $< 200,000$
Pileup	< 0.5	Valid or invalid truth link
	OR	
	Any	Invalid truth link

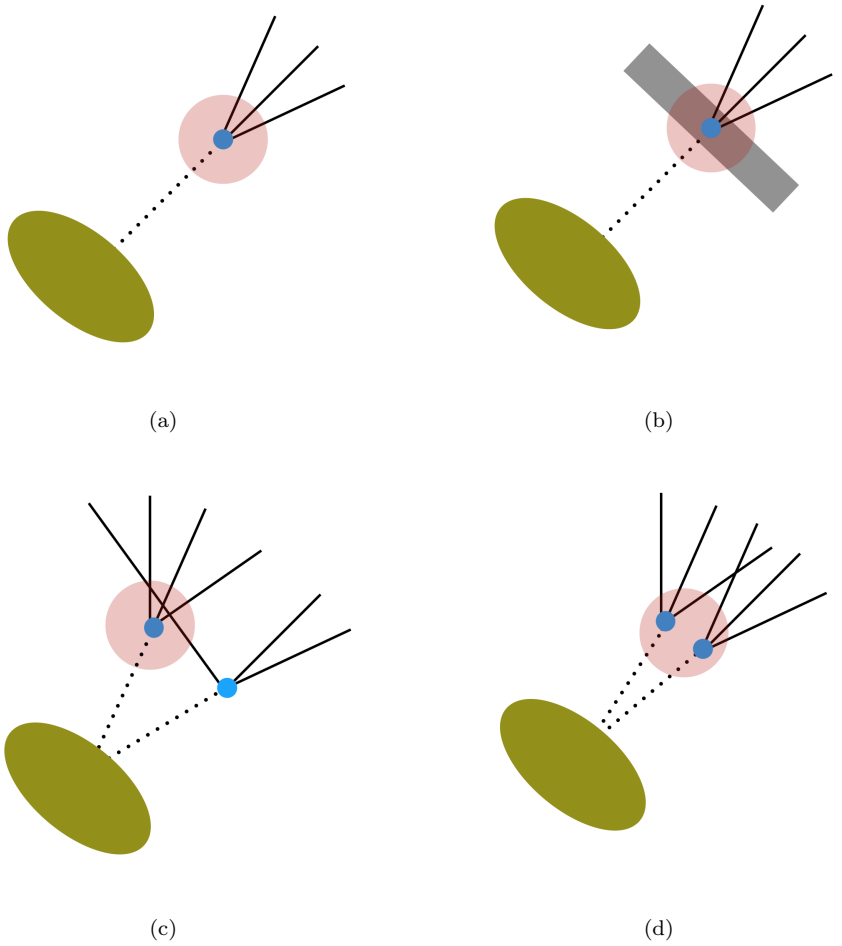

Figure 10.14: Schematic drawing of a background DV resulting from a long-lived SM particle decay (a), a hadronic interaction (b), an accidental crossing (c), and two merged vertices (d). The green oval represents the beamspot, the dashed lines and blue points the trajectories and decay positions of truth particles, the pink circle the reconstructed DV, and the solid lines the DV tracks.

The key to determine the nature of a given DV is to understand, for each of its tracks, which displaced decay it originates from. This is however complicated by the fact that the vertices, represented by the blue points in the schematic in Figure 10.14, may contain chains of intermediate particles that are created in the simulation of the decay. In order to find the long-lived parent of the track it is therefore necessary to trace its full ancestry in the truth record. Figure 10.15 shows the radial production and decay positions of all truth particles in the dijet sample. Two distinct populations can be seen with a boundary at $r \approx 1\,\text{mm}$. Based on this, the original parent is defined as the truth particle in the ancestry which crosses the 1 mm boundary, meaning that it is produced before $r = 1\,\text{mm}$ and decays after $r = 1\,\text{mm}$. A few alternative definitions, varying the boundary within the region between the two populations, are tested and show negligible differences in the final results.

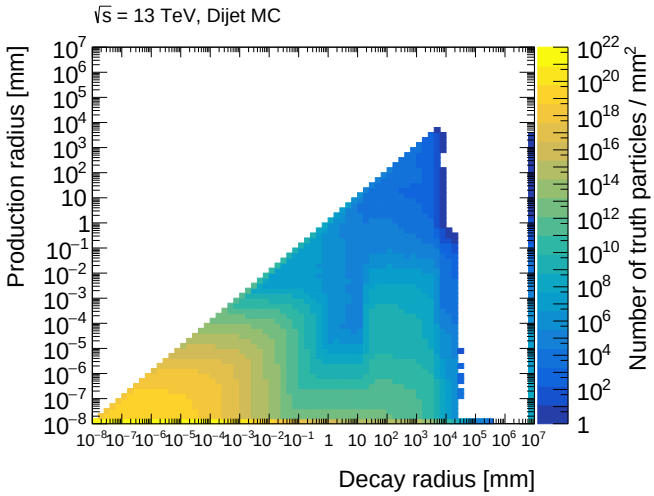

Figure 10.15: Two-dimensional distributions of the radial production and decay positions of all truth particles in the dijet MC sample.

Based on the track categories and original parents of all associated hard-scatter tracks, each DV is categorised according to the DV classes in Table 10.9. In this classification, the *Single Origin DVs* correspond to hadronic interactions and long-lived SM decays. The *Single Origin DV + Track* and *Mixed Origin DVs* correspond to accidental crossings and merged vertices, but the distinction between the two is ambiguous. A vertex that appears to be an accidental crossing can also be two merged vertices where only one track passes the DV-track selections in one of them. Similarly, two merged vertices might also be one vertex crossed by two tracks. These

cases are not possible to distinguish because the vertex merging information is not preserved in the truth record.

Table 10.9: Definition of DV categories, based on information in the truth record in the dijet MC sample.

DV category	**Track setup**
Single Origin DV	All tracks share the same origin and belong to one of the hard-scatter track categories.
Single Origin DV + Track	All tracks, except one, are hard-scatter and share the same origin. The remaining track may be either hard-scatter or pileup.
Mixed Origin DV	All tracks are hard-scatter and come from two different origins.
Pileup DV	All tracks are pileup.
Other DV	Any other DV.

DV characteristics

The visible mass distribution for vertices with different track multiplicities produced inside material is shown in Figure 10.16. The Single Origin DVs dominate the distribution at low mass while the Single Origin + Track category has a more pronounced tail towards larger mass values. This behaviour can be understood by considering the nature of the DVs. An accidentally crossing track can pass the DV with a large crossing angle, thereby promoting it to a significantly higher visible mass. The hadronic interactions are fixed-target processes without large angular spread between the decay products. Figure 10.17 shows the same distributions, but selecting DVs outside material. All categories, except the pileup DVs, decrease in magnitude outside material. This observation is expected for the hadronic interactions component and suggests that also the accidental crossings and merged vertices predominantly involve vertices from material interactions.

Figure 10.18 shows the visible mass distribution for the $n_{\text{trk}}^{\text{DV}} = 4$ Single Origin category with a breakdown into DVs where all tracks come from either GEANT4 or the generator, or are a mixture of both. The all generator component is similar in magnitude and shape inside and outside material, which is expected from natural decays of SM particles. A study of the PDG IDs shows that these DVs mainly correspond to b-hadrons. The component is localised at very low visible mass and is therefore not an important background to the analysis.

In conclusion, the study suggests that the hadronic interactions are localised at low DV mass, but that there may be some contribution in the SR above 10 GeV. The accidental crossings component takes over at higher mass, and it is the dominant

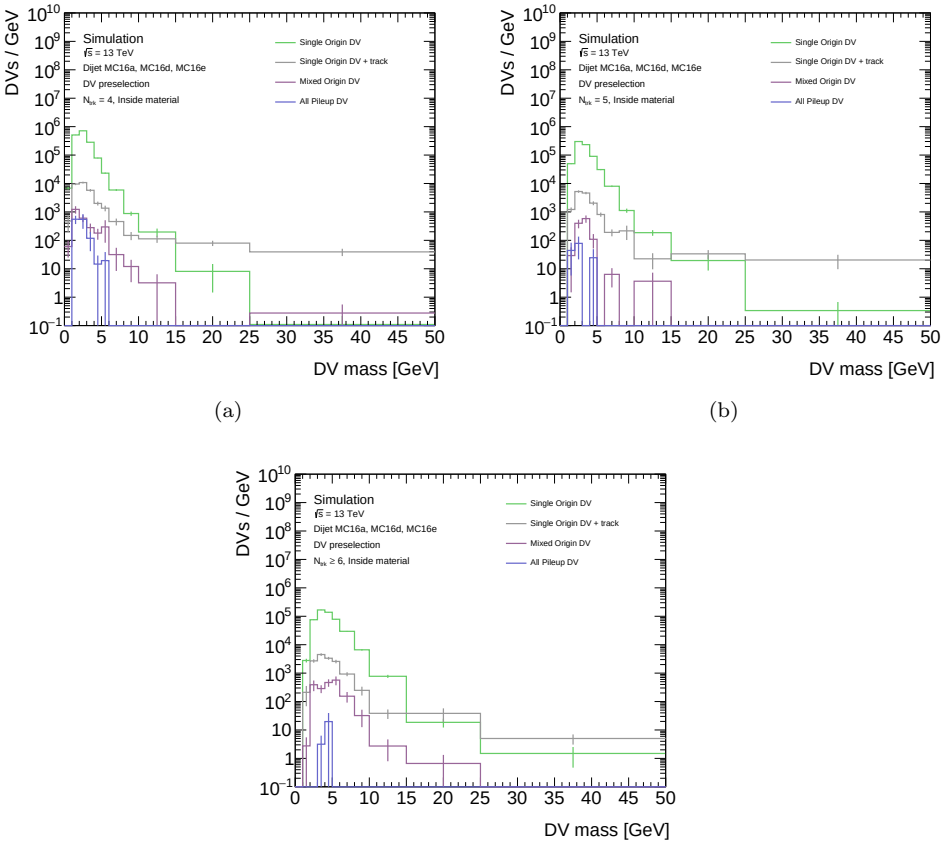

Figure 10.16: Invariant mass distributions for the various DV categories, inside material, for $n_{\text{trk}}^{\text{DV}} = 4$ (a), $n_{\text{trk}}^{\text{DV}} = 5$ (b), and $n_{\text{trk}}^{\text{DV}} \geq 6$ (c).

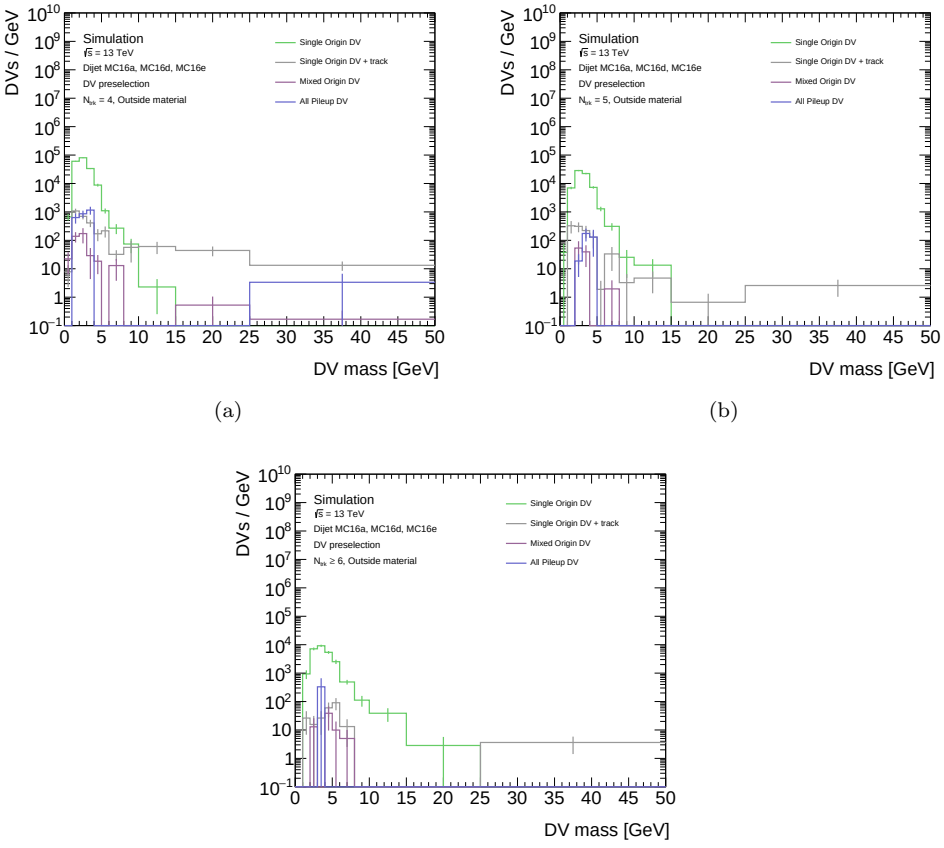

Figure 10.17: Invariant mass distributions for the various DV categories, outside material, for $n_{\text{trk}}^{\text{DV}} = 4$ (a), $n_{\text{trk}}^{\text{DV}} = 5$ (b), and $n_{\text{trk}}^{\text{DV}} \geq 6$ (c).

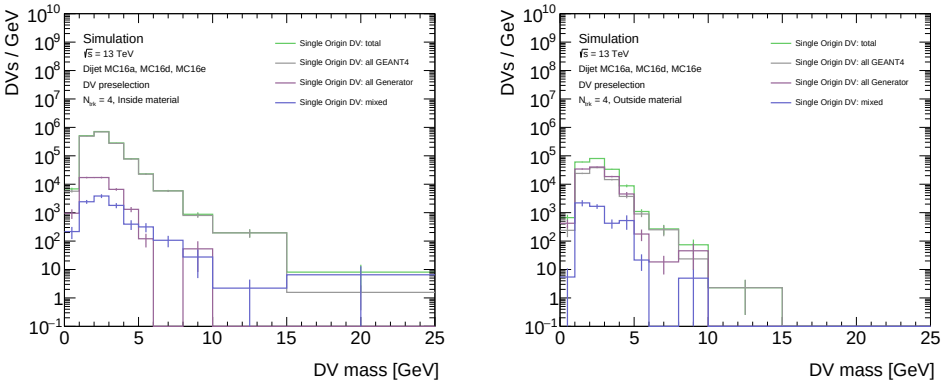

(a) (b)

Figure 10.18: Invariant mass distributions for $n_{\text{trk}}^{\text{DV}} = 4$ Single Origin DVs with all track coming from GEANT4, the generator, or from both, inside material (a) and outside material (b).

background component in the SR. The merged vertex component appears to be small and corresponding to small m_{DV}. It is however not possible to fully separate this component from the accidental crossings.

10.6.2 Inclusive background estimate

The data-driven inclusive background estimate relies on the assumption that in a multijet final state, DV production is correlated to the presence of jets. This assumption is made because all three sources of DV production are correlated with the density of tracks which in turn is correlated with jets. In order to strengthen the correlation between DVs and jets, the estimate is performed using the track jets described in Section 10.3. The estimate proceeds in two steps. First, the probability for a DV to be produced in proximity to a jet is calculated using jets and DVs in a CR of photon-triggered data. Then, this *jet-DV probability* is applied to jets in events passing the baseline jet selections to estimate the total number of expected background DVs in the two SRs.

Control region

The CR is designed to have a similar composition of jets as the SRs, while having a minimal contamination from the signal models. Events in the CR are selected using a single-photon trigger with an E_{T} threshold at 140 GeV and are further required to fail the baseline jet requirements of both SRs. The events are also

required to contain at least three track jets with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$. DVs in the CR which pass the baseline DV selections are used for the estimate, with their m_{DV} and $n_{\mathrm{trk}}^{\mathrm{DV}}$ computed after the track cleaning procedure. These DVs are referred to as *SR-like*. Figure 10.19 shows the fraction of events in the CR which contain a SR-like DV with $m_{\mathrm{DV}} > 3\,\mathrm{GeV}$ and $n_{\mathrm{trk}}^{\mathrm{DV}} \geq 3$, as a function of the track-jet multiplicity. A clear trend can be seen where events with more track jets are more likely to contain a SR-like DV, displaying the correlation between jets and DV production.

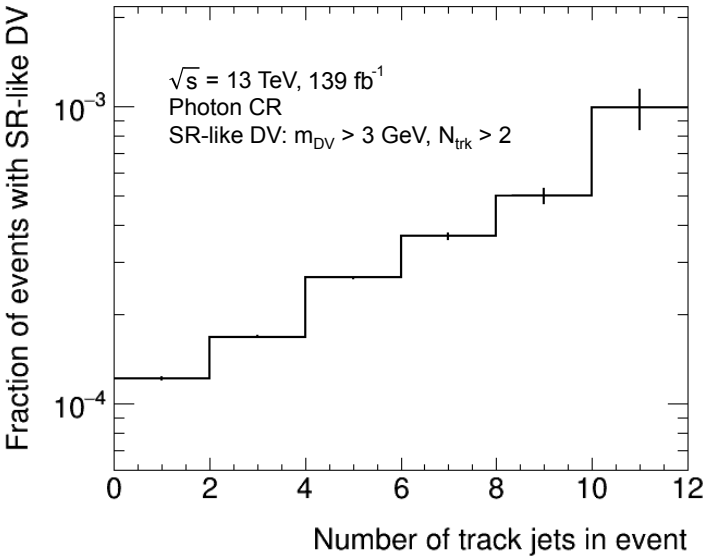

Figure 10.19: The fraction of events in the CR with an SR-like DV satisfying $m_{\mathrm{DV}} > 3\,\mathrm{GeV}$ and $n_{\mathrm{trk}}^{\mathrm{DV}} \geq 3$, as a function of the track-jet multiplicity.

As the background estimate is based on the presence of jets, it is important to assess any differences in the multijet nature of events in the CR compared to events in the SRs. Figure 10.20 shows the track-jet multiplicity distributions for events in the three regions. The average jet multiplicity is lower in the CR, and any impact on the estimate from the difference will be taken as a systematic uncertainty. Figure 10.21 shows the jet p_{T} and the number of tracks in jets in the various regions. Events passing the High-p_{T} selection have track jets with higher p_{T} and more tracks than those in the CR and events passing the Trackless selection have jets with slightly lower p_{T} and fewer tracks. In order to take these differences into account, the jet-DV probability is parametrised as a function of both variables.

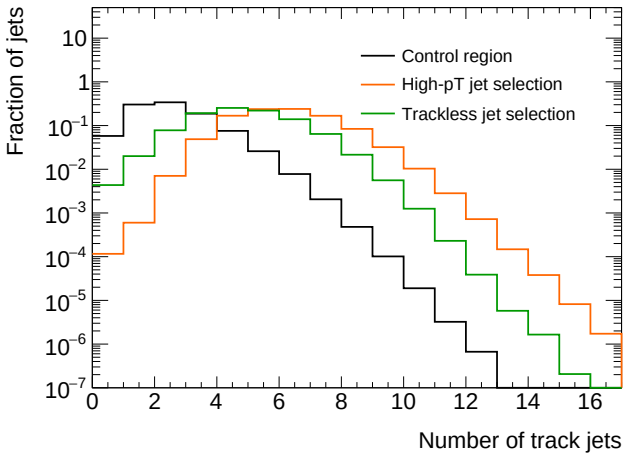

Figure 10.20: A comparison of the track-jet multiplicity in the CR and in events passing the High-p_{T} and Trackless baseline jet selections.

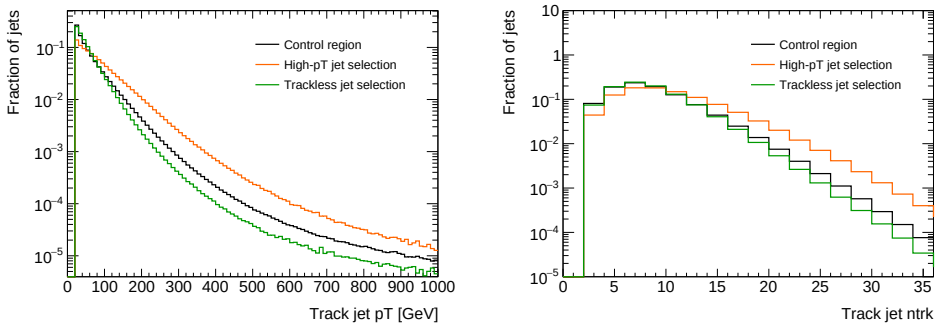

(a) (b)

Figure 10.21: A comparison of the track-jet p_{T} (a) and track multiplicity (b) in the CR and in events passing the High-p_{T} and Trackless baseline jet selection.

Jet-DV probability

The jet-DV probability is calculated in the CR by first matching all SR-like DVs to the closest track jet, and then computing the fraction of jets that have a DV matched to them. The matching is based on the angular distance parameter ΔR which is computed between all DVs and jets in the event. Each DV is matched to the jet which corresponds to the minimum ΔR. The jet DV probability is then

given by

$$P(\text{DV} \mid \text{jet}) = \frac{N_{\text{jets}}^{\text{match}}}{N_{\text{jets}}}, \tag{10.1}$$

where $N_{\text{jets}}^{\text{match}}$ is the number of jets with a matched DV and N_{jets} is the total number of jets in the CR. While this should ideally be computed with DVs restricted to the SR requirements, the selection is loosened to $n_{\text{trk}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5\,\text{GeV}$ in order to have adequate statistics for the estimate. The number of such DVs in the CR is shown in Figure 10.22a and the number of track jets is shown in Figure 10.22b. The calculated jet-DV probability is shown in Figure 10.23 and is of the order 10^{-6}. The parametrisation in track-jet p_{T} and multiplicity is developed using a looser track cleaning procedure in order to increase the number of DVs in the CR. With more statistics, clear trends are seen with increased jet-DV probability for higher jet p_{T} and jet multiplicities.

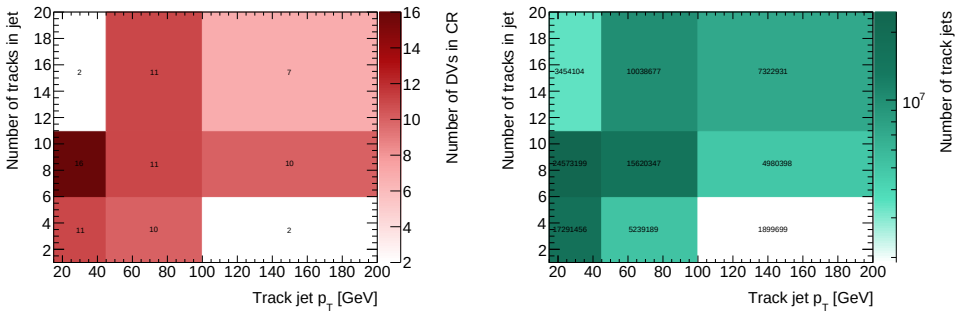

(a) (b)

Figure 10.22: Number of SR-like DVs satisfying $n_{\text{trk}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5\,\text{GeV}$ (a) and number of track jets (b) in the CR, as a function of the track-jet p_{T} and track multiplicity.

Application of jet-DV probability to SR jets

The number of background DVs expected in the SRs is calculated by weighting each jet in the regions by the jet-DV probability in the corresponding p_{T} and track-multiplicity bin. The final estimate is given by the sum

$$N_{\text{Bkg}}^{\text{SR}} = f_{\text{SR}} \cdot \sum_{\text{i}=0}^{N_{\text{jet}}^{\text{SR}}} P(\text{DV} \mid \text{jet}_i), \tag{10.2}$$

where f_{SR} is a scale factor which extrapolates the estimate from the looser CR to the m_{DV} and $n_{\text{trk}}^{\text{DV}}$ requirements of the SRs. It is calculated from the fraction of CR

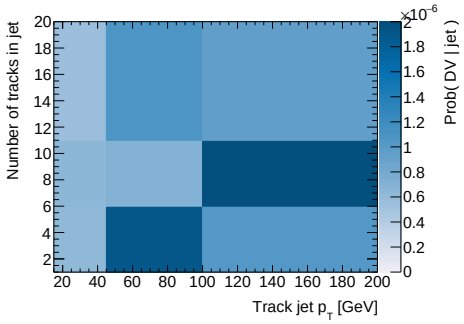

Figure 10.23: The jet-DV probability of SR-like DVs satisfying $n_{\text{trk}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5\,\text{GeV}$ in the CR, as a function of the track-jet p_{T} and track multiplicity.

events with at least a DV fulfilling the SR requirements according to

$$f_{\text{SR}} = \frac{N_{\text{Event}}^{\text{CR}}(m_{\text{DV}} > 10\,\text{GeV}, n_{\text{trk}}^{\text{DV}} \geq 5)}{N_{\text{Event}}^{\text{CR}}(m_{\text{DV}} \geq 5\,\text{GeV}, n_{\text{trk}}^{\text{DV}} \geq 4)} = 3/80 = 0.0375. \tag{10.3}$$

The total predicted background amounts to 0.49 DVs in the High-p_{T} SR and 0.69 DVs in the Trackless SR. Given the small size of the jet-DV probability, there is a negligible chance to have two DVs produced in the same event. The predicted number of DVs therefore translates directly to a predicted number of events.

Validation of the estimate in data

The performance of the inclusive background estimate is tested in the VRs listed in Table 10.7. For each VR, the full estimate is repeated, including the calculation of the jet-DV probability, and compared to the observed number of events. The results of the validation tests for the Low-$n_{\text{trk}}^{\text{DV}}$ and Low-m_{DV} Sideband VRs are shown in Figures 10.24 and 10.25 respectively. Agreement within the statistical uncertainties is observed between the estimated event counts and the observed data in all regions. The results of the validation tests in the Inside Material VRs are shown in Figure 10.26. For events passing the Trackless jet selection, closure within the statistical uncertainties is observed. In events passing the High-p_{T} jet selection, there is however a significant underestimation of the background for low visible DV mass. In order to further test the closure, the validation tests are repeated in the dijet MC sample.

Validation of the estimate in MC

The MC validation tests are performed in the same set of VRs as in data using jet-DV probabilities derived from dijet events which pass a single-photon trigger and

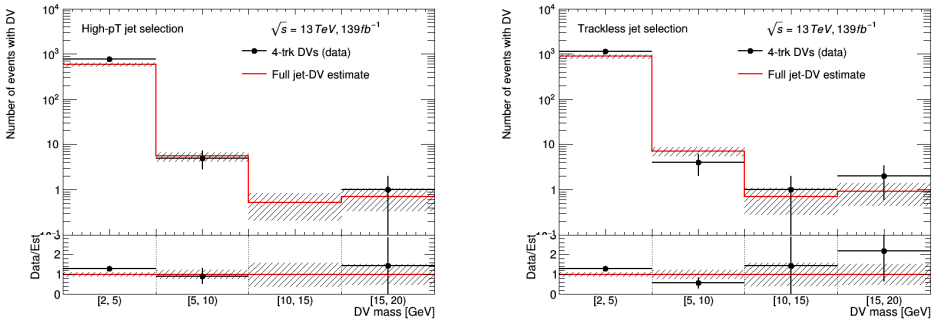

(a) (b)

Figure 10.24: Validation of the inclusive background estimate in the High-p_{T} (a) and Trackless (b) Low-$n_{\mathrm{trk}}^{\mathrm{DV}}$ Sideband VRs in data. Only the statistical uncertainties propagated from the corresponding CRs are drawn.

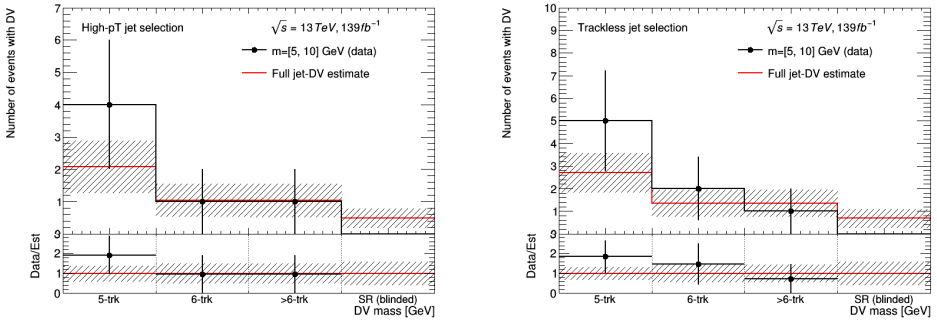

(a) (b)

Figure 10.25: Validation of the inclusive background estimate in the High-p_{T} (a) and Trackless (b) Low-m_{DV} Sideband VRs in data. Only the statistical uncertainties propagated from the corresponding CRs are drawn. The estimate in the SR is also included and is shown in the last bin, but the observed SR yield is blinded.

fail the High-p_{T} and Trackless jet selections. The MC tests in the Inside Material VRs are shown in Figure 10.27. For DVs with a high mass, there are not enough events in the MC to compute the estimate and test the closure. At low mass there is sufficient statistics, and an underestimation similar to the one observed in data, is observed for events with $m_{\mathrm{DV}} = 5 - 10\,\mathrm{GeV}$ in the High-p_{T} jet selection.

In order to further investigate the source of the non-closure inside material, the validation test in MC is repeated using the DV truth classification described in

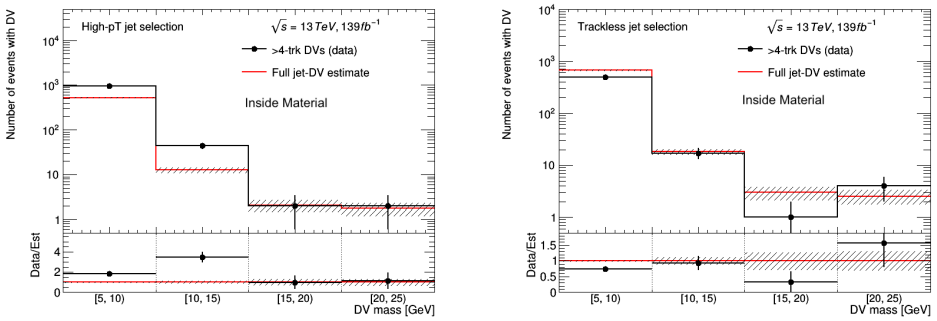

(a) (b)

Figure 10.26: Validation of the inclusive background estimate in the High-p_{T} (a) and Trackless (b) Inside Material VRs in data. Only the statistical uncertainties propagated from the corresponding CRs are drawn.

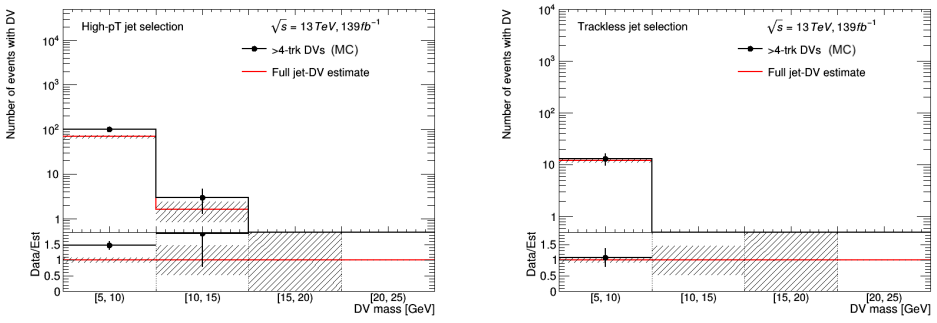

(a) (b)

Figure 10.27: Validation of the inclusive background estimate in the High-p_{T} (a) and Trackless (b) Inside Material VRs in the dijet MC sample. Only the statistical uncertainties propagated from the corresponding CRs are drawn.

Section 10.6.1. The estimate is performed separately for the Single Origin DV and Single Origin DV + Track categories. The results of the test are shown for the Inside Material VRs in Figure 10.28 with events passing the High-p_{T} jet selection, and DVs with $m_{\mathrm{DV}} = 5 - 10\,\mathrm{GeV}$. The non-closure is seen to be caused by the Single Origin class of DVs.

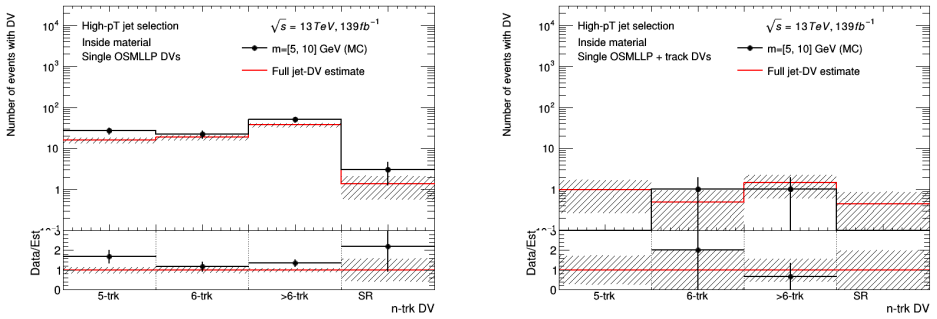

(a) (b)

Figure 10.28: Validation of the inclusive background estimate in the High-p_{T} Inside Material VRs in the dijet MC sample, for Single Origin DV (a) and Single Origin DV + Track (b) categories. Only the statistical uncertainties propagated from the corresponding CRs are drawn.

Systematic uncertainties

Three sources of uncertainty are quantified for the background estimate; a statistical uncertainty from the number of SR-like DVs in the CR, a systematic uncertainty to account for differences in the jet multiplicity between the CR and the SRs, and a non-closure uncertainty based on the validation of the method. The three uncertainties are assumed to be uncorrelated and are added in quadrature to estimate the total systematic uncertainty. The final background estimates in the SRs together with the associated uncertainties are listed in Table 10.10. The statistical uncertainty amounts to 58 % and completely dominates the total.

The systematic uncertainty takes into account any residual dependence of the jet-DV probability on the amount of jet activity in the event. It is calculated by computing the background estimate as a function of the number of jets, assuming the same shape in the jet multiplicity distribution in the SR as in the CR. This is compared to the estimate using the actual distribution and the relative difference is taken as a systematic uncertainty on the estimate. The resulting uncertainty is 0.10 % in the High-p_{T} SR and 3.8 % in the Trackless SR.

The non-closure in the Inside Material regions at low visible DV mass stems from the hadronic interactions component, which constitutes a very small fraction of the DVs in the SRs. The non-closure observed in the Inside Material VRs is therefore not translated directly into a systematic uncertainty but is scaled by the fraction of hadronic interactions in the SR. This results in a 4 % uncertainty assigned to the background estimate in the High-p_{T} SR. A similar uncertainty calculated for the Trackless SR is negligible and therefore not included.

Table 10.10: Event yields predicted by the inclusive background estimate in the High-p_{T} and Trackless SRs. The CR statistical uncertainty refers to the propagation of the statistical variations of the CR to the estimate, while the non-linearity uncertainty refers to the residual dependence of the jet-DV probability on the amount of jet activity in the event.

	High-p_{T} SR	**Trackless SR**
Estimated event yield	0.49 ± 0.28	0.7 ± 0.4
Total relative uncertainty	58 %	58 %
CR statistical uncertainty	58 %	58 %
Non-linearity	0.10 %	3.8 %
VR non-closure	4 %	-

10.6.3 Individual background estimates

In addition to the inclusive estimate, all three background sources are estimated individually using data-driven methods. The individual estimates are combined and compared to the inclusive estimate, providing an important cross check of the predicted background.

Hadronic interactions

The hadronic-interactions estimate builds on the findings from the MC background studies in Section 10.6.1. The low-mass region of the m_{DV} distribution is dominated by hadronic interactions which then exhibit an approximate exponential decrease as a function of increasing vertex mass. Therefore, to estimate this background in the SR, the m_{DV} distribution in the region $m_{\mathrm{DV}} < 10\,\mathrm{GeV}$ in data is fit to an exponential distribution and extrapolated to the SR with $m_{\mathrm{DV}} > 10\,\mathrm{GeV}$.

Merged vertices

The merged-vertex background is estimated based on the knowledge of the secondary vertexing algorithm, described in Section 5.1.6, where pairs of DVs with a distance significance $S < 10$ are merged. The estimation method proceeds in two steps. Firstly, the rate of vertex merging is computed by comparing the distance distributions for pairs of DVs taken from the same event and from different events where no merging can occur. Secondly, a mass template is produced by creating pairs of DVs from different events which satisfy $S < 10$ and computing the visible mass from all tracks associated to the DVs. The mass template is scaled by the vertex merging rate and the integral above $m_{\mathrm{DV}} = 10\,\mathrm{GeV}$ gives the estimated SR yield.

Table 10.11: Event yields in the High-p_{T} and Trackless SRs, as predicted by the inclusive background estimate and the combination of the three individual estimates.

Estimate	**High-p_{T} SR**	**Trackless SR**
Inclusive	0.49 ± 0.28	0.7 ± 0.4
Combined	1.1 ± 0.8	1.8 ± 1.3

Accidental crossings

Accidentally crossing tracks may be associated to secondary vertices in several steps of the vertexing algorithm, making it harder to estimate the background purely based on the mechanism that causes it. Instead, it is estimated using data events with K^0_{S} decays identified to have an accidentally crossing track. The method proceeds in three steps with the first consisting in identifying all 3-track DVs where two of the tracks have an invariant mass within 50 MeV of the K^0_{S} mass. Such DVs are considered to be K^0_{S} decays with an accidentally crossing track which is stored together with its kinematic properties in a track database. In the second step, DV-mass templates are generated by attaching tracks from the database to existing DVs and recomputing the visible DV mass including the extra track. The final step of the estimate calculates the probability that an accidentally crossing track is associated to a DV, to correctly normalise the mass templates. This is computed by measuring the fraction of all K^0_{S} decays reconstructed in data which have an accidentally crossing track.

Combined background estimate

The combined predictions from the three individual background estimates in the High-p_{T} and Trackless SRs are shown in Table 10.11 together with the prediction from the inclusive background estimate. The estimates from the two methods are consistent, and the individual background estimates therefore provide a successful cross check of the inclusive background estimate.

10.7 Systematic uncertainties

In addition to the systematic uncertainties assigned to the inclusive background estimate, various sources of uncertainty need to be taken into account for the signal MC used in the interpretation of the analysis results. These correspond to either theoretical uncertainties on the modelling, or experimental uncertainties that account for performance differences between data and MC.

The strong production cross sections are calculated at approximate NNLO+NNLL [127–130, 178–180], and the electroweak production cross sections are computed at NLO+NLL [181–185]. The nominal cross sections and their uncertainties are derived using the PDF4LHC15 PDF set, following the recommendations of Ref. [186]. To exemplify, the cross section for gluino production ranges from $(8.9\pm1.4)\times10^{-3}$ pb for a 1600 GeV gluino to $(4.6 \pm 1.4) \times 10^{-5}$ pb, corresponding to a 15-25 % uncertainty. For the electroweak production it ranges from 17.1 ± 0.7 pb for a 100 GeV neutralino to $(2.1 \pm 0.4) \times 10^{-5}$ pb for a 1700 GeV neutralino, corresponding to uncertainties between 4-20 %.

The experimental uncertainties include the 1.7 % uncertainty on the luminosity used to normalise the signal yield, as well as the uncertainties on the physics objects used in the analysis provided by ATLAS Combined Performance groups. Since this search uses unconventional objects, additional uncertainties need to be evaluated specifically for the analysis to assess the impact of the displacement from the *pp* interaction point. This applies to the tracks, vertices, and jets, and is described in the following sections.

10.7.1 Tracking and vertexing uncertainties

The MC uncertainty related to the LRT is estimated from the yield of reconstructed K^0_S decays in different radial regions, using the full analysis dataset and the dijet MC samples. The tracks originating from the K^0_S decays at small radii are typically reconstructed as standard tracks, which means that the reconstruction efficiency and its uncertainty are well understood. Therefore, the K^0_S yield in MC is normalised to data in a small-radius region before the yields are compared at larger radius. The MC is seen to overestimate the K^0_S yield at large radius, pointing towards inefficiencies in data which are not fully described in the simulation. The deviation from unity of the ratio between data and MC is taken as a per-track systematic uncertainty and is added in quadrature to the 1.7 % uncertainty from the standard tracking prescription [82]. The uncertainty varies between 2-20 % depending on the radius.

The tracking uncertainties are propagated to the SR selection efficiency by randomly removing tracks from the DVs in signal MC with a probability equal to the relative per-track uncertainty. Several DV samples are produced in this way, and the average efficiency of these samples is calculated. The difference with respect to the nominal efficiency is taken as the uncertainty on the signal yield. For low-mass neutralinos, the uncertainty on the signal efficiency reaches a maximum of 15 %, while it is negligible for high-mass neutralinos. This is because the high-mass neutralinos are more likely to be reconstructed with higher track multiplicities, so that the removal of one track does not exclude the DV from the SR.

10.7.2 Jet uncertainties

Standard jet uncertainties

The standard jet uncertainties are assessed by varying the jet energy scale (JES) and the jet energy resolution (JER) according to a set of systematic uncertainties derived by the Combined Performance groups in ATLAS [187]. The signal efficiency for each variation is compared to the nominal signal efficiency and the difference is taken as the corresponding systematic uncertainty. For the strong production signal grid, the JES and JER uncertainties are found to be negligible compared to the statistical and cross-section uncertainties. In the electroweak production model, the JES and JER uncertainties are comparable in size to the other uncertainties and will be included in the statistical treatment. Figures 10.29 and 10.30 show the JES and JER uncertainties for the electroweak model in the High-p_{T} and Trackless SRs respectively. The uncertainties are a few percent for most grid points, with larger values resulting mainly from statistical fluctuations for points with very low SR yields.

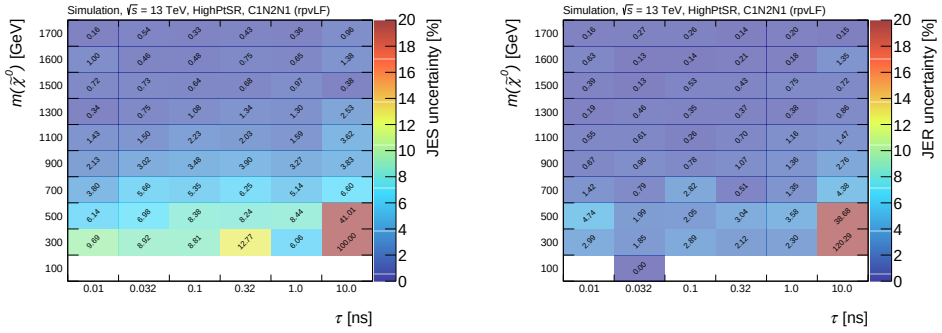

(a) (b)

Figure 10.29: Relative JES (a) and JER (b) uncertainties for the electroweak production signal grid in the High-p_{T} SR. Empty bins correspond to grid points with zero observed events.

Jet uncertainties from displacement

The effect on the jet reconstruction of the displacement from the pp interaction point is studied by measuring the so-called jet-p_{T} *response* in signal MC. This is defined as the p_{T} ratio between a reconstructed jet and the corresponding truth-level jet according to

$$R_{p_{\mathrm{T}}} = \frac{p_{\mathrm{T}}^{\mathrm{reco}}}{p_{\mathrm{T}}^{\mathrm{truth}}}. \tag{10.4}$$

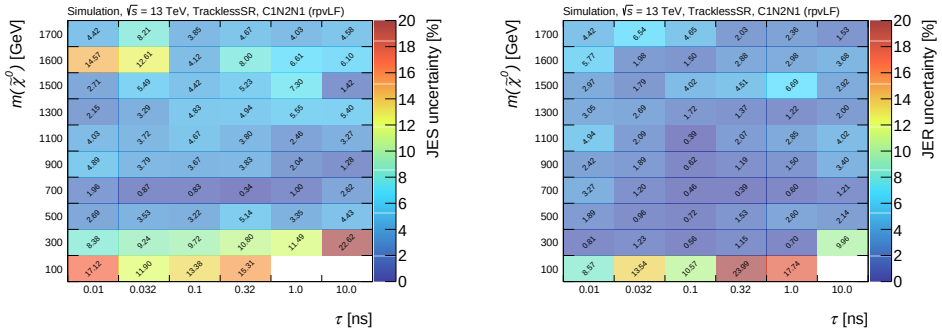

(a) (b)

Figure 10.30: Relative JES (a) and JER (b) uncertainties for the electroweak production signal grid in the Trackless SR. Empty bins correspond to grid points with zero observed events.

Well-measured jets are expected to a have response close to unity. The truth jets are constructed with an $R = 0.4$ anti-k_t algorithm using stable final-state particles and are matched to reconstructed jets by means of ΔR. For the studies presented here, the truth jets are also matched geometrically to individual truth particles in order to determine the jet origin.

The jet-p_{T} response is studied for the strong production model using an inclusive jet selection where both the reconstructed and truth jets are required to reside within $|\eta| < 2.8$. In addition, truth jets must have $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and reconstructed jets $p_{\mathrm{T}} > 50\,\mathrm{GeV}$, to ensure that their properties are well measured. Only one truth jet is allowed to be present within a cone of $\Delta R = 1.0$ and no additional reconstructed jet within $\Delta R = 0.6$.

Truth-jet origin In order to study the jet-p_{T} response as a function of displacement, the origin of each truth jet must be determined. This is accomplished by matching the truth jets to the truth particles resulting from neutralino decays. The closest truth particle within $\Delta R = 0.3$ from the jet is taken as the match, and its production vertex is used as the jet origin.

Truth-jet coordinates In MC, truth particles are stored with a four momentum expressed in a local coordinate system. This means that the η and ϕ directions are computed with respect to the origin of the truth particle instead of the beamspot, as is the case for reconstructed objects. In most cases this is a reasonable way of modeling the behaviour of particles, but for decay products of LLPs it can result in difficulties when trying to match truth particles to reconstructed objects using the angular coordinates. Figure 10.31 illustrates the difference between the local

and global angular coordinates of a jet hitting the calorimeter surface at a given position. In order to mitigate the difference between the truth and reconstructed jets, a geometric correction is applied to the angular coordinates of the truth jets to express them with respect to the beamspot. The correction, derived in Ref. [188], ignores the effects of the magnetic field, and extrapolates the jet axis linearly to the calorimeter surface from its truth origin. Figure 10.32 shows the minimum ΔR between truth jets and reconstructed jets in signal events, for grid points with varying neutralino lifetime. As the lifetime increases, the truth and reconstructed jet directions become less correlated, as seen by the flattened ΔR distribution, and the correction becomes essential for studying the performance of jets produced from decay products of LLPs.

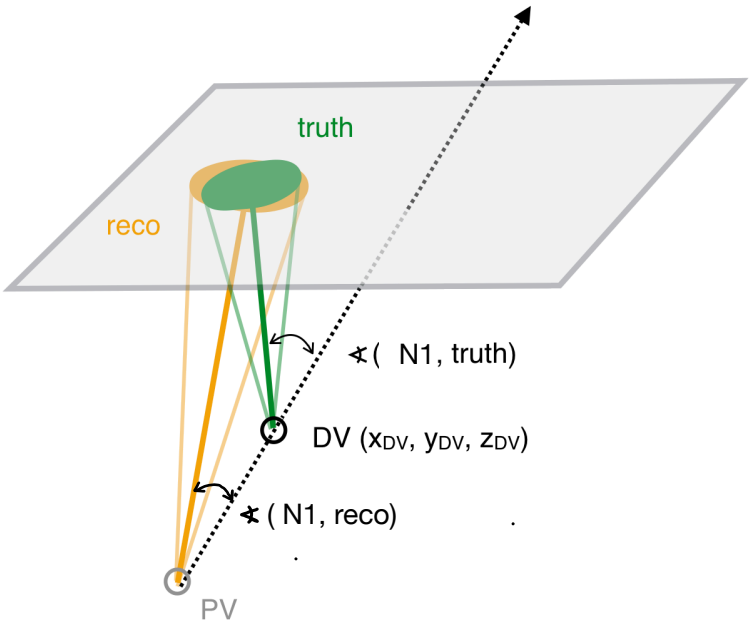

Figure 10.31: Illustration of the default axes for reconstructed and truth jets. The truth jet is reconstructed from particles with four-momenta defined relative to the DV at which the neutralino decays, whereas the reconstructed jet is initially treated as being produced from the beamspot. Figure adapted from Ref. [189].

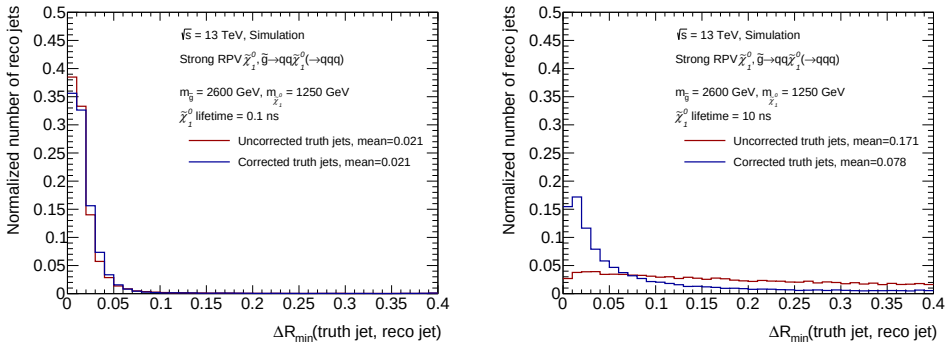

(a) (b)

Figure 10.32: The minimum ΔR between reconstructed and truth jets before and after the geometric correction is applied to truth jets, for the strong production model with $m_{\tilde{g}} = 2600\,\text{GeV}$, $m_{\tilde{\chi}_1^0} = 1250\,\text{GeV}$, and $\tau = 0.1\,\text{ns}$ (a) or $\tau = 10\,\text{ns}$ (b).

Jet response After the geometric correction is applied to all truth jets, each reconstructed jet is matched to the closest truth jet fulfilling $\Delta R \leq 0.4$ and the jet-p_T response is computed. In order to assess the dependence of the response on the decay topology, two sets of signal grid points are studied. The first set includes grid points with a small mass splitting between the neutralino and the gluino, and the second set includes grid points with a large mass splitting and a small neutralino mass. All lifetimes are included in both sets. Figure 10.33 shows the jet-p_T response as a function of the matched neutralino decay position for the small and large mass-splitting sets. For each bin in R and z, the jet-p_T response is fitted with a Gaussian to the core of the distribution. In the small mass-splitting set, both the deviation from unity and the width of the jet-p_T response increases with the displacement of the jet origin, while the effect is negligible for the set with a large mass splitting. The difference can be attributed to the Lorentz boost of the neutralino. A large boost results in more collinear decay products and more standard-like jets regardless of the displacement of the decay point. The fraction of jets in each R and z bin is shown in Figure 10.34 for the small mass-splitting set. 90 % of the jets originate from within 300 mm in radius and $|z|$. In this region the mean jet p_T response is consistent with unity also for this set. This implies that the standard jet calibration performs well for the vast majority of the jets considered in the analysis. No additional jet uncertainty is therefore considered for the displacement.

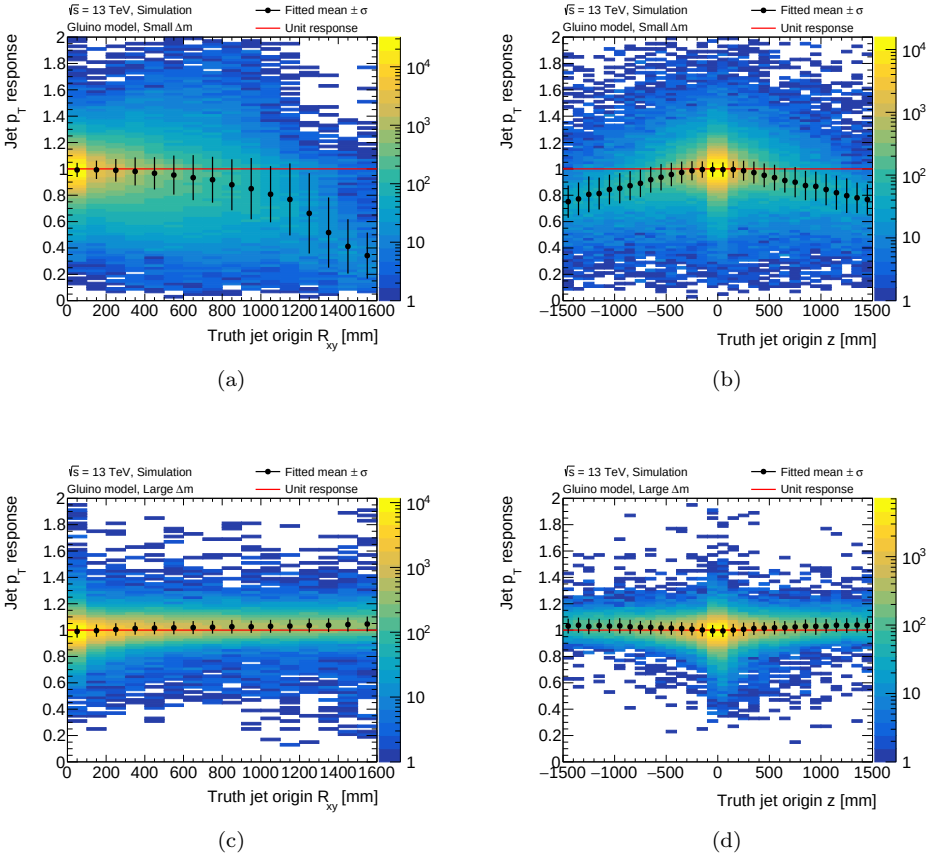

Figure 10.33: Jet-p_{T} response as a function of the decay position R and z of the neutralino matched to the truth jet, for the set of signal samples with small (a,b) and large (c,d) gluino-neutralino mass splitting in the strong production model grid. Each bin in R or z is fitted with a Gaussian to the core of the distribution and the fitted mean is shown together with the width as black points. The samples included in the small mass-splitting set have neutralino masses varying between $1250-2550\,\mathrm{GeV}$ and gluino masses varying between $1600-2600\,\mathrm{GeV}$, corresponding to a mass splitting between $50-350\,\mathrm{GeV}$. The samples included in the large mass-splitting set have neutralino masses varying between $10-50\,\mathrm{GeV}$ and gluino masses varying between $1600-2600\,\mathrm{GeV}$, corresponding to a mass splitting between $1590-2550\,\mathrm{GeV}$

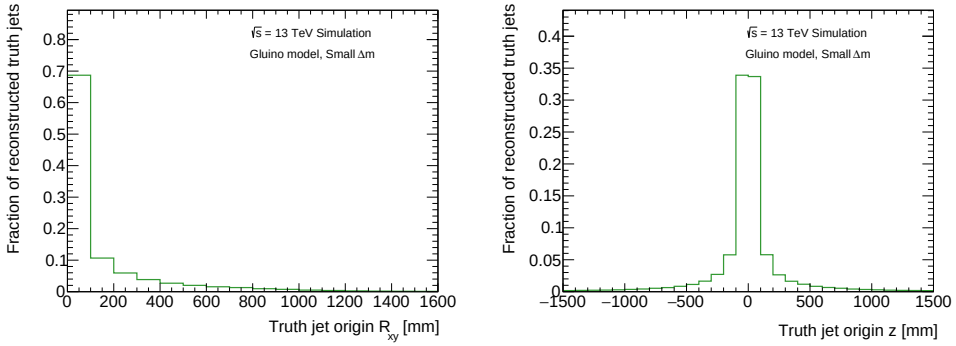

(a) (b)

Figure 10.34: Fraction of truth jets matched to reconstructed jets in each bin in R (a) and z (b), for the set of signal samples with small gluino-neutralino mass splitting in the strong production model grid.

10.8 Results

No events are observed in the Trackless SR where 0.7 ± 0.4 events are expected in the absence of signal, while a single event is observed in the High-p_{T} SR where 0.49 ± 0.28 events are expected. The observed event yields are in good agreement with the background-only expectations. The single observed event belongs to the 2017 dataset, and contains one DV with $m_{\mathrm{DV}} = 32.6\,\mathrm{GeV}$, $n_{\mathrm{trk}}^{\mathrm{DV}} = 5$, and is located at $R = 7.5\,\mathrm{mm}$ inside the beampipe. Figure 10.35 shows the DV mass and track multiplicity distributions in events passing the baseline jet selection and containing at least one baseline DV.

10.9 Interpretation

The High-p_{T} SR is designed primarily to target the strong production signal. The expected number of signal events for in the High-p_{T} SR is shown in Figure 10.36a for a fixed gluino mass of 2400 TeV as a function of the neutralino mass and lifetime. Figure 10.36b shows the expected signal yields when fixing the neutralino lifetime to 0.1 ns, corresponding to the maximum expected sensitivity, and varying the gluino and neutralino masses. For the electroweak production signal, both SRs are important and the expected signal yields are shown in Figure 10.37 as a function of the mass and lifetime of the neutralino.

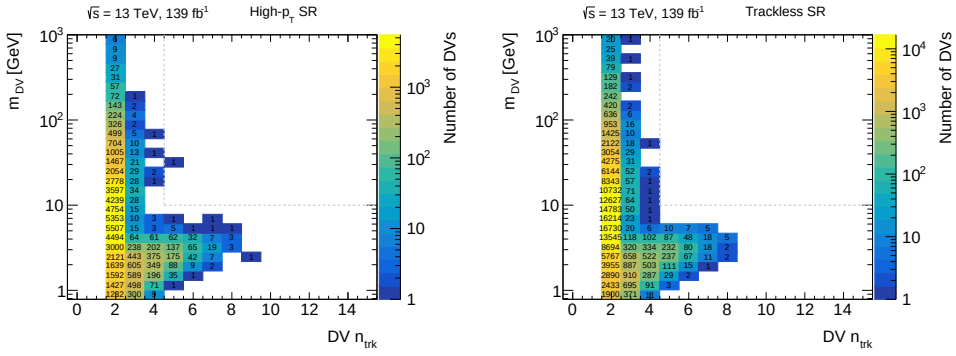

(a) (b)

Figure 10.35: Two-dimensional distributions of m_{DV} and track multiplicity for baseline DVs in events that satisfy all High-p_{T} SR (a) and Trackless SR (b) event selection criteria. The dashed line represents the boundary of the SR requirements.

In the absence of a statistically significant excess in the data, exclusion limits are placed on the two signal models following the CL_s prescription [154, 155], as described in Section 9.10. For this analysis it is performed using a two-bin fit, where the two bins correspond to the the High-p_{T} and Trackless SRs. At the time of writing, preliminary expected exclusion limits are derived including the systematic uncertainty on the background and the evaluated signal uncertainties. The resulting limits are shown in the neutralino mass-lifetime plane in Figure 10.38 for the strong production grid with the gluino mass fixed to either 2200 GeV or 2400 GeV. The electroweak limits are shown in Figure 10.39. The observed limits will be close to the expected ones, since no significant deviations from the predicted background is seen. In the strong production model, the search thus excludes gluino masses up to approximately 2400 TeV and neutralino lifetimes up to 10 ns for all but the lowest neutralino masses. For the electroweak production, the limit on the neutralino mass surpasses 1.3 TeV for all lifetimes in the range from 0.01 ns to 10 ns.

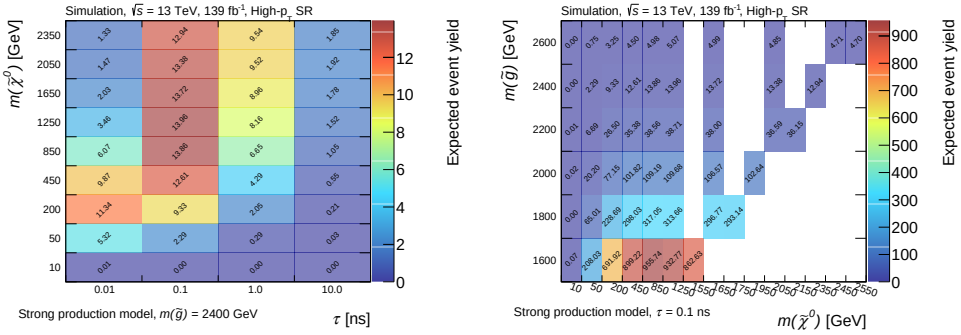

(a) (b)

Figure 10.36: Event yields in the High-p_{T} SR for the strong production signal grid as a function of neutralino mass and lifetime with the gluino mass fixed to 2400 GeV (a) and as a function of gluino and neutralino mass with the neutralino lifetime fixed to 0.1 ns (b).

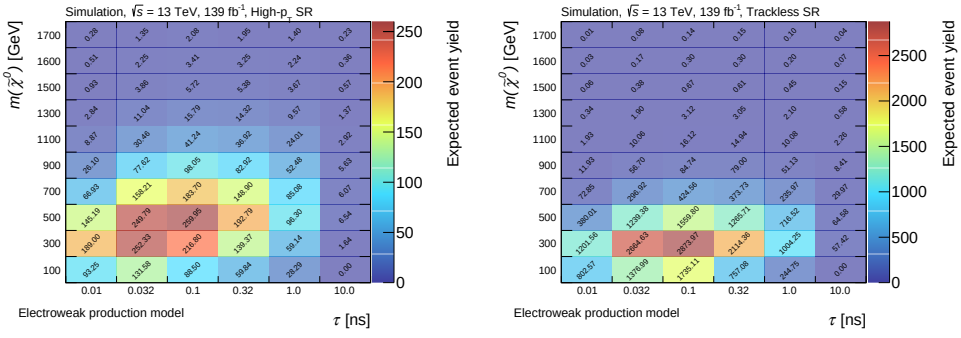

(a) (b)

Figure 10.37: Events yields in the High-p_{T} SR (a) and Trackless SR (b) for the electroweak production signal grid as a function of neutralino mass and lifetime.

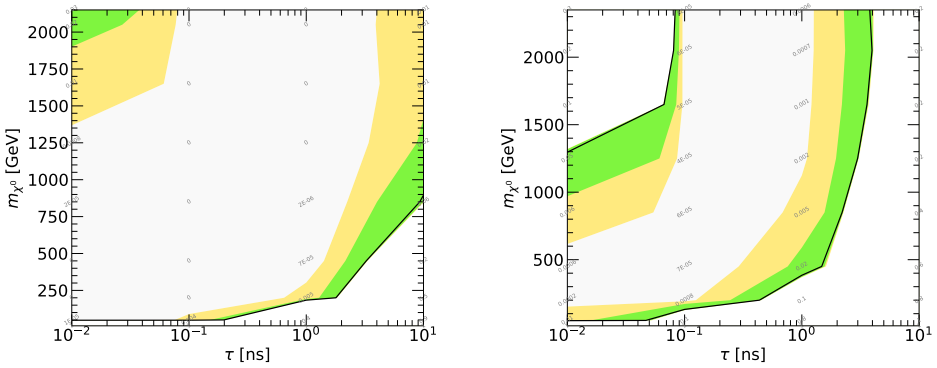

(a) (b)

Figure 10.38: 95 % CL_s exclusion limits varying $\tilde{\chi}_1^0$ lifetime and mass, for the strong production grid with the gluino mass fixed to 2200 GeV (a) and to 2400 GeV (b). The expected limits and their 1σ (green) and 2σ (yellow) bands are shown. The area between the lines is the excluded parameter space. Image credit: Filip Backman.

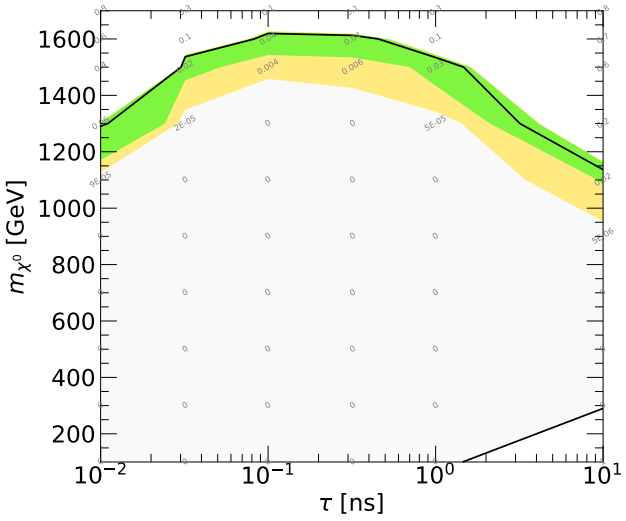

Figure 10.39: 95 % CL_s exclusion limits varying $\tilde{\chi}_1^0$ lifetime and mass, for the electroweak production grid. The expected limits and their 1σ (green) and 2σ (yellow) bands are shown. The area between the lines is the excluded parameter space. Image credit: Filip Backman.

Concluding remarks

In this thesis I have presented my work on the ATLAS luminosity measurement and on the search for new exotic particles. The work is performed with the *pp* collision dataset recorded by the ATLAS experiment in the LHC Run 2, corresponding to the years 2015-2018.

ATLAS luminosity measurement

A precise measurement of the luminosity is a key component of the ATLAS physics programme. For cross-section measurements, the uncertainty in the luminosity is often one of the main systematic uncertainties. Searches for physics beyond the Standard Model also rely on an accurate luminosity measurement to predict the sensitivity to the signal and to estimate background yields. ATLAS uses several detectors and algorithms to determine the luminosity and the absolute calibration is carried out in vdM scans with special beam conditions at low luminosity. The nominal luminosity measurement is provided by the LUCID detector and the track-counting measurement is used to determine the calibration transfer from the low-luminosity regime to the high-luminosity conditions typical of standard physics data taking. Track-counting is also used to monitor the long-term stability of the default luminosity method. The luminosity of the Run 2 dataset suitable for analysis is measured to be $139\,\text{fb}^{-1}$ with an uncertainty of 1.7 %. This is well within the precision needed to provide experimental cross-section measurements competitive with current theoretical calculations. The uncertainty is comparable in size to the uncertainties achieved for Run 1 which amounted to 1.8 % at $\sqrt{s} = 7\,\text{TeV}$ [113, 190] and 1.9 % at $\sqrt{s} = 8\,\text{TeV}$ [115]. A final Run 2 offline luminosty estimate is in preparation. It is expected to reach an even higher precision, and to be competitive with the recently published CMS measurement for the 2015 and 2016 datasets, which reported uncertainties of 1.6 % and 1.2 % respectively for the two years [191].

The track-counting method measures the luminosity in ATLAS by counting the number of Inner Detector tracks selected with specific quality criteria. At the Run 2 data-taking conditions the average number of tracks is stable as a function of time and highly linear with the average number of inelastic *pp* collisions and

therefore provides a very reliable luminosity estimate. Studies in simulation and data suggest however that the method may need to be adjusted for the future data-taking periods when the LHC will operate with higher pileup values. In particular, the contribution from fake tracks poses a major challenge for the estimate when the instantaneous luminosity increases. The work shown in this thesis has presented the first steps towards an estimate of the fake-track contribution, which can also help to reduce the uncertainty on the overall luminosity measurement.

At the pileup values of the HL-LHC, the event- and hit-counting methods exploited for the LUCID luminosity measurement will saturate and the existence of reliable alternative methods is therefore crucial. As part of the ATLAS upgrade for HL-LHC a new subdetector will be installed which is partly dedicated to the luminosity measurement using silicon hits [191]. To prepare, a pixel-hit-counting method is developed in ATLAS for Run 3, based on the tools and lessons learned from track-counting. Pixel-hit counting is already in use in CMS and it is the main luminosity method for the CMS Run 2 measurement [192].

Searches for new exotic particles

Two searches for new physics have been presented in this thesis and are both interpreted in terms of Supersymmetric signal models. The Strong-2L analysis searches for squarks and gluinos in final states with two same-flavour opposite-charge leptons, jets, and large missing transverse momentum. It is performed with the first part of the Run 2 *pp* collision dataset corresponding to $36.1\,\text{fb}^{-1}$ collected in 2015 and 2016. No significant excess is observed above the expected Standard Model background, and the analysis results are used to set exclusion limits on squark and gluino masses in terms of simplified Supersymmetry models. The analysis results exclude gluinos and squarks with masses as large as 1.85 TeV and 1.3 TeV at 95 % confidence level, respectively. The full Run 2 dataset has now been analysed with a similar ATLAS search which does not see any evidence for Supersymmetry [193], and the limits are therefore extended up to 2250 GeV and 1550 GeV. While these results appear to challenge the naturalness arguments for Supersymmetry, it should be noted that the interpretations are performed under all the hypotheses of the simplified models which correspond to very specific instantiations of the Minimal Supersymmetric Standard Model. The stringent limits are however a clear motivation for searches in less accessible regions of phase space, targeting for example long-lived particles or compressed scenarios with final state particles with very low momenta.

The DV+jets analysis searches for long-lived particles in events with displaced vertices and multiple energetic jets. It exploits the full Run 2 *pp* collision dataset and employs dedicated reconstruction techniques that significantly increase the sensitivity to long-lived particle decays inside the ATLAS Inner Detector. The search lacks irreducible Standard Model backgrounds but is sensitive to instrumental effects which merge unrelated vertices or tracks, as well as hadronic interactions in

the detector material. These background sources are predicted in the two signal regions with a data-driven method. The observed yields are one event in the first signal region where 0.49 ± 0.28 events are expected, and zero events in the second signal region where 0.7 ± 0.4 events are expected. The results are used to set limits at 95 % confidence level on scenarios with pair-production of supersymmetric particles with long-lived electroweakinos that decay via a small R-parity-violating coupling into triplets of quarks. Electroweakly produced neutralinos with masses up to 1.3 TeV are excluded for lifetimes between 0.01 ns to 10 ns. While the DV+jets result is the first to be interpreted in in terms of this model, the result is comparable to the exclusion limits on long-lived gluinos which are targeted by several ATLAS analyses, and reach 2 TeV for the same lifetime range.

All ATLAS searches for long-lived particles in final states with a displaced vertex are essentially free of backgrounds. This means that the sensitivity scales linearly with the signal yield which in turn scales with the size of the analysed dataset. This is in contrast to background-dominated searches, where the sensitivity is proportional to the square root of the integrated luminosity. The increase in accumulated luminosity during the next few years of Run 3 data taking can therefore bring large sensitivity gains for long-lived particle searches even though the the centre-of-mass energy remains approximately the same as in Run 2.

List of Figures

List of Tables

Bibliography

[1] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) S08001.

[2] ATLAS Collaboration, *Track multiplicity plots for the track-counting luminosity measurement at $\sqrt{s}$=13 TeV with the ATLAS detector*, `http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/LUMI-2021-02/`.

[3] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, ATLAS-CONF-2019-021, 2019, `https://cds.cern.ch/record/2677054`.

[4] ATLAS Collaboration, *Search for new phenomena using the invariant mass distribution of same-flavour opposite-sign dilepton pairs in events with missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, Eur. Phys. J. C **78** (2018) 625, `arXiv:1805.11381 [hep-ex]`.

[5] ATLAS Collaboration, *Search for new phenomena in events containing a same-flavour opposite-sign dilepton pair, jets, and large missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, Eur. Phys. J. C **77** (2017) 144, `arXiv:1611.05791 [hep-ex]`.

[6] G. Ripellino, *A search for Supersymmetry in final states with two same-flavor opposite-sign leptons, jets and missing transverse momentum with the ATLAS detector*. Licentiate thesis, KTH, 2018. `http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-220892/`.

[7] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264–1266.

[8] S. L. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22** (1961) 579–588.

[9] A. Salam, *Weak and Electromagnetic Interactions*, Conf. Proc. C **680519** (1968) 367–377.

[10] G. 't Hooft, *Renormalization of Massless Yang-Mills Fields*, Nucl. Phys. B **33** (1971) 173–199.

[11] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, Phys. Rev. Lett. **30** (1973) 1343–1346.

[12] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30** (1973) 1346–1349.

[13] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (1964) 321–323.

[14] P. W. Higgs, *Broken symmetries, massless particles and gauge fields*, Phys. Lett. **12** (1964) 132–133.

[15] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (1964) 508–509.

[16] D. J. Griffiths, *Introduction to elementary particles*. Physics textbook. Wiley, New York, NY, 2nd ed., 2008.

[17] M. Thomson, *Modern particle physics*. Cambridge University Press, New York, 2013.

[18] Particle Data Group Collaboration, P. Zyla et al., *Review of Particle Physics*, PTEP **2020** no. 8, (2020) 083C01.

[19] Wikipedia Contributors, *Standard model - Wikipedia, the free encyclopedia*, `https://en.wikipedia.org/wiki/Standard_Model`. Accessed: 2022-04-02.

[20] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, Phys. Rev. Lett. **10** (1963) 531–533.

[21] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, Prog. Theor. Phys. **49** (1973) 652–657.

[22] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, *Collider Searches for Long-Lived Particles Beyond the Standard Model*, Prog. Part. Nucl. Phys. **106** (2019) 210–255, `arXiv:1810.12602 [hep-ph]`.

[23] S. Höche, *Introduction to parton-shower event generators*, in *Journeys Through the Precision Frontier: Amplitudes for Colliders*. 2015. `arXiv:1411.4085 [hep-ph]`.

[24] R. Placakyte, *Parton Distribution Functions*, in *31st International Symposium on Physics In Collision*. 11, 2011. `arXiv:1111.5452 [hep-ph]`.

[25] R. D. Ball et al., *Parton distributions with LHC data*, Nucl. Phys. B **867** (2013) 244, `arXiv:1207.1303 [hep-ph]`.

[26] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, Nucl. Phys. B **126** (1977) 298–318.

[27] V. N. Gribov and L. N. Lipatov, *Deep inelastic e p scattering in perturbation theory*, Sov. J. Nucl. Phys. **15** (1972) 438–450.

[28] L. Harland-Lang, A. Martin, P. Motylinski, and R. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, Eur. Phys. J. C **75** no. 5, (2015) 204, arXiv:1412.3989 [hep-ph].

[29] UA1 Collaboration, *Experimental observation of events with large missing transverse energy accompanied by a jet or a photon(s) in $p\overline{p}$ collisions at $\sqrt{s} = 540$ GeV*, Phys. Lett. **139B** (1984) 115.

[30] UA2 Collaboration, *Evidence for Z0 → e+e- at the CERN $\overline{p}p$ collider*, Phys. Lett. **129B** (1983) 130–140.

[31] CDF Collaboration, *Observation of top quark production in $\bar{p}p$ collisions*, Phys. Rev. Lett. **74** (1995) 2626–2631, arXiv:hep-ex/9503002 [hep-ex].

[32] D0 Collaboration, *Observation of the top quark*, Phys. Rev. Lett. **74** (1995) 2632–2637, arXiv:hep-ex/9503003 [hep-ex].

[33] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B **716** (2012) 30, arXiv:1207.7235 [hep-ex].

[34] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) 1, arXiv:1207.7214 [hep-ex].

[35] ATLAS Collaboration, *Standard Model Summary Plots June 2021*, https://cds.cern.ch/record/2777014.

[36] M. Baak, J. Cúth, J. Haller, A. Hoecker, R. Kogler, K. Mönig, M. Schott, and J. Stelzer, *The global electroweak fit at NNLO and prospects for the LHC and ILC*, Eur. Phys. J. C **74** (2014) 3046, arXiv:1407.3792 [hep-ph].

[37] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, Astron. Astrophys. **641** (2020) A6, arXiv:1807.06209 [astro-ph.CO], [Erratum: Astron.Astrophys. 652, C4 (2021)].

[38] Super-Kamiokande Collaboration, *Evidence for an oscillatory signature in atmospheric neutrino oscillation*, Phys. Rev. Lett. **93** (2004) 101801, arXiv:hep-ex/0404034.

[39] B. Kayser, *Are neutrinos their own antiparticles?*, J. Phys. Conf. Ser. **173** (2009) 012013, arXiv:0903.0899 [hep-ph].

[40] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe*, Pisma Zh. Eksp. Teor. Fiz. **5** (1967) 32–35.

[41] S. P. Martin, *A Supersymmetry primer*, arXiv:hep-ph/9709356 [hep-ph], Adv. Ser. Direct. High Energy Phys. 18, 1 (1998).

[42] L. Girardello and M. Grisaru, *Soft breaking of supersymmetry*, Nuclear Physics B **194** no. 1, (1982) 65–76, https://www.sciencedirect.com/science/article/pii/0550321382905120.

[43] I. Jack and D. R. T. Jones, *Nonstandard soft supersymmetry breaking*, Phys. Lett. B **457** (1999) 101–108, arXiv:hep-ph/9903365.

[44] C. F. Kolda, *Gauge mediated supersymmetry breaking: Introduction, review and update*, Nucl. Phys. B Proc. Suppl. **62** (1998) 266–275, arXiv:hep-ph/9707450.

[45] E. Halkiadakis, G. Redlinger, and D. Shih, *Status and Implications of Beyond-the-Standard-Model Searches at the LHC*, Ann. Rev. Nucl. Part. Sci. **64** (2014) 319–342, arXiv:1411.1427 [hep-ex].

[46] J. L. Hewett, B. Lillie, M. Masip, and T. G. Rizzo, *Signatures of long-lived gluinos in split supersymmetry*, JHEP **09** (2004) 070, arXiv:hep-ph/0408248.

[47] G. F. Giudice and R. Rattazzi, *Theories with gauge mediated supersymmetry breaking*, Phys. Rept. **322** (1999) 419–499, arXiv:hep-ph/9801271.

[48] P. Meade, M. Reece, and D. Shih, *Long-Lived Neutralino NLSPs*, JHEP **10** (2010) 067, arXiv:1006.4575 [hep-ph].

[49] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[50] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002.

[51] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004.

[52] LHCb Collaboration, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005.

[53] CERN, *LHC: the guide*, Feb, 2009, http://cds.cern.ch/record/1165534.

[54] S. Baird, *Accelerators for pedestrians; rev. version*, tech. rep., CERN, Geneva, Feb, 2007. https://cds.cern.ch/record/1017689.

[55] S. Holmes, R. S. Moore, and V. Shiltsev, *Overview of the Tevatron Collider Complex: Goals, Operations and Performance*, JINST **6** (2011) T08001, arXiv:1106.0909 [physics.acc-ph].

[56] *LEP Design Report: Vol. 1: The LEP Injector Chain*. CERN, Geneva, 1983. https://cds.cern.ch/record/98881. By the LEP Injector Study Group.

[57] *LEP Design Report: Vol. 2: The LEP Main Ring*. CERN, Geneva, 1984. https://cds.cern.ch/record/102083. Copies shelved as reports in LEP, PS and SPS libraries.

[58] P. Grafström and W. Kozanecki, *Luminosity determination at proton colliders*, Prog. Part. Nucl. Phys. **81** (2015) 97–148.

[59] ATLAS Collaboration, *Public ATLAS Run-1 Luminosity Information*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults.

[60] ATLAS Collaboration, *Public ATLAS Run-2 Luminosity Information*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[61] ATLAS Collaboration, *ATLAS magnet system: Technical Design Report, 1*, https://cds.cern.ch/record/338080.

[62] ATLAS Collaboration, *ATLAS inner detector: Technical Design Report, 1*, https://cds.cern.ch/record/331063.

[63] S. Haywood, L. Rossi, R. Nickerson, and A. Romaniouk, *ATLAS inner detector: Technical Design Report, 2*, https://cds.cern.ch/record/331064.

[64] ATLAS Collaboration, *ATLAS Inner Detector Alignment Performance with February 2015 Cosmic Ray Data*, ATL-PHYS-PUB-2015-009, 2015, https://cds.cern.ch/record/2008724.

[65] ATLAS Collaboration, *Study of the material of the ATLAS inner detector for Run 2 of the LHC*, JINST **12** (2017) P12009, arXiv:1707.02826 [hep-ex].

[66] ATLAS Collaboration, *ATLAS Insertable B-Layer Technical Design Report*. Technical design report. CERN, Geneva, Sep, 2010. https://cds.cern.ch/record/1291633.

[67] ATLAS Collaboration, *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical design report. CERN, Geneva, 1996. https://cds.cern.ch/record/331061.

[68] ATLAS Collaboration, *ATLAS tile calorimeter: Technical Design Report*. Technical design report. CERN, Geneva, 1996. https://cds.cern.ch/record/331062.

[69] ATLAS Collaboration, *ATLAS muon spectrometer: Technical Design Report*. Technical design report. CERN, Geneva, 1997. https://cds.cern.ch/record/331068.

[70] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, JINST **13** no. 07, (2018) P07017.

[71] V. Cindro et al., *The ATLAS beam conditions monitor*, JINST **3** (2008) P02004, https://cds.cern.ch/record/1094819.

[72] ATLAS Collaboration, *Charged-particle distributions in $\sqrt{s} = 13$ TeV pp interactions measured with the ATLAS detector at the LHC*, Phys. Lett. B **758** (2016) 67, arXiv:1602.01633 [hep-ex].

[73] A. Salzburger, *Tracking tutorial*, in *Artemis School on Calibration and performance of ATLAS detectors*. 2008. https://indico.cern.ch/event/34087/contributions/802944/attachments/669263/919934/IDTracking_A_Salzburger.pdf.

[74] T. Cornelissen, M. Elsing, S. Fleischmann, W. Liebig, E. Moyse, and A. Salzburger, *Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)*, https://cds.cern.ch/record/1020106.

[75] R. Fruhwirth, *Application of Kalman filtering to track and vertex fitting*, Nucl. Instrum. Meth. A **262** (1987) 444–450.

[76] P. Billoir, *Progressive track recognition with a Kalman like fitting procedure*, Comput. Phys. Commun. **57** (1989) 390–394.

[77] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, Eur. Phys. J. C **77** no. 5, (2017) 332, arXiv:1611.10235 [physics.ins-det].

[78] ATLAS Collaboration, *Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-026, 2015, https://cds.cern.ch/record/2037717.

[79] R. Fruhwirth, W. Waltenberger, and P. Vanlaer, *Adaptive vertex fitting*, J. Phys. G **34** (2007) N343.

[80] ATLAS Collaboration, *Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC*, ATLAS-CONF-2010-027, 2010, https://cds.cern.ch/record/1277659.

[81] ATLAS Collaboration, *Public ATLAS Beamspot Information*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BeamSpotPublicResults.

[82] ATLAS Collaboration, *Early Inner Detector Tracking Performance in the 2015 Data at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-051, 2015, https://cds.cern.ch/record/2110140.

[83] ATLAS Collaboration, *Calibration of light-flavour b-jet mistagging rates using ATLAS proton–proton collision data at $\sqrt{s} = 13$ TeV*, ATLAS-CONF-2018-006, 2018, https://cds.cern.ch/record/2314418.

[84] ATLAS Collaboration, *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*, ATL-PHYS-PUB-2017-014, 2017, https://cds.cern.ch/record/2275635.

[85] ATLAS Collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, ATL-PHYS-PUB-2019-013, 2019, https://cds.cern.ch/record/2669425.

[86] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, JINST **14** (2019) P12006, arXiv:1908.00005 [hep-ex].

[87] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, Eur. Phys. J. C **77** (2017) 490, arXiv:1603.02934 [hep-ex].

[88] ATLAS Collaboration, *Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton–proton collision data*, ATLAS-CONF-2016-024, 2016, https://cds.cern.ch/record/2157687.

[89] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at* $\sqrt{s} = 13$ *TeV*, Eur. Phys. J. C **76** (2016) 292, arXiv:1603.05598 [hep-ex].

[90] M. Cacciari, G. P. Salam, and G. Soyez, *The Anti-k(t) jet clustering algorithm*, JHEP **04** (2008) 063, arXiv:0802.1189 [hep-ph].

[91] ATLAS Collaboration, *Determination of jet calibration and energy resolution in proton–proton collisions at* $\sqrt{s} = 8$ *TeV using the ATLAS detector*, Eur. Phys. J. C **80** (2020) 1104, arXiv:1910.04482 [hep-ex].

[92] ATLAS Collaboration, *Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*, Phys. Rev. D **96** (2017) 072002, arXiv:1703.09665 [hep-ex].

[93] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, ATLAS-CONF-2014-018, 2014, https://cds.cern.ch/record/1700870.

[94] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 physics and manual*, JHEP **05** (2006) 026, arXiv:hep-ph/0603175.

[95] A. Buckley et al., *General-purpose event generators for LHC physics*, Phys. Rept. **504** (2011) 145–233, arXiv:1101.2599 [hep-ph].

[96] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Parton Fragmentation and String Dynamics*, Phys. Rept. **97** (1983) 31–145.

[97] B. R. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*, Nucl. Phys. B **238** (1984) 492–528.

[98] ATLAS Collaboration, *Measurement of the underlying event in jet events from 7 TeV proton–proton collisions with the ATLAS detector*, Eur. Phys. J. C **74** (2014) 2965, arXiv:1406.0392 [hep-ex].

[99] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, JHEP **07** (2014) 079, arXiv:1405.0301 [hep-ph].

[100] S. Frixione and B. R. Webber, *Matching NLO QCD computations and parton shower simulations*, JHEP **06** (2002) 029, arXiv:hep-ph/0204244.

[101] S. Alioli, P. Nason, C. Oleari, and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, JHEP **06** (2010) 043, arXiv:1002.2581 [hep-ph].

[102] S. Frixione, P. Nason, and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, JHEP **11** (2007) 070, arXiv:0709.2092 [hep-ph].

[103] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, JHEP **11** (2004) 040, arXiv:hep-ph/0409146.

[104] T. Sjöstrand, S. Mrenna, and P. Skands, *A brief introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852–867, arXiv:0710.3820 [hep-ph].

[105] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159, arXiv:1410.3012 [hep-ph].

[106] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and W. J., *Event generation with SHERPA 1.1*, JHEP **02** (2009) 007, arXiv:0811.4622 [hep-ph].

[107] E. Bothmann et al., *Event generation with Sherpa 2.2*, SciPost Phys. **7** no. 3, (2019) 034, arXiv:1905.09127 [hep-ph].

[108] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. A **462** (2001) 152.

[109] GEANT4 Collaboration, S. Agostinelli, et al., GEANT4 – *a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250.

[110] A. Sidoti et al., *Minimum Bias Trigger Scintillators in ATLAS Run II*, JINST **9** no. 10, (2014) C10020.

[111] Particle Data Group Collaboration, P. A. Zyla et al., *Review of Particle Physics*, PTEP **2020** no. 8, (2020) 083C01.

[112] ATLAS Collaboration, *Precision measurement and interpretation of inclusive W^+, W^- and Z/γ^* production cross sections with the ATLAS detector*, Eur. Phys. J. C **77** (2017) 367, arXiv:1612.03016 [hep-ex].

[113] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **73** (2013) 2518, arXiv:1302.4393 [hep-ex].

[114] S. van der Meer, *Calibration of the effective beam height in the ISR*, tech. rep., CERN, Geneva, 1968. https://cds.cern.ch/record/296752.

[115] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **76** (2016) 653, arXiv:1608.03953 [hep-ex].

[116] M. Bassetti and G. A. Erskine, *Closed expression for the electrical field of a two-dimensional Gaussian charge*, tech. rep., CERN, Geneva, 1980. https://cds.cern.ch/record/122227.

[117] W. Herr, *Beam-beam effects and dynamic $\beta*$*, in *Proc. LHC Lumi Days 2012*. https://indico.cern.ch/event/162948/contributions/1417430/attachments/191879/269237/S3_WH.pdf.

[118] ATLAS Collaboration, *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model*, ATL-PHYS-PUB-2016-017, 2016, https://cds.cern.ch/record/2206965.

[119] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*, JINST **9** (2014) P08009, arXiv:1404.7473 [hep-ex].

[120] ATLAS Collaboration, *Summary of the searches for squarks and gluinos using $\sqrt{s} = 8$ TeV pp collisions with the ATLAS experiment at the LHC*, JHEP **10** (2015) 054, arXiv:1507.05525 [hep-ex].

[121] ATLAS Collaboration, *Search for supersymmetry in events containing a same-flavour opposite-sign dilepton pair, jets, and large missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, Eur. Phys. J. C **75** (2015) 318, arXiv:1503.03290 [hep-ex].

[122] CMS Collaboration, *Search for physics beyond the standard model in events with two leptons, jets, and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV*, JHEP **04** (2015) 124, arXiv:1502.06031 [hep-ex].

[123] CMS Collaboration, *Search for new physics in final states with two opposite-sign, same-flavor leptons, jets, and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV*, JHEP **12** (2016) 013, arXiv:1607.00915 [hep-ex].

[124] CMS Collaboration, *Search for new phenomena in final states with two opposite-charge, same-flavor leptons, jets, and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV*, JHEP **03** (2018) 076, arXiv:1709.08908 [hep-ex].

[125] ATLAS Collaboration, *Summary of ATLAS Pythia 8 tunes*, ATL-PHYS-PUB-2012-003, 2012, https://cds.cern.ch/record/1474107.

[126] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, *Parton distributions for the LHC*, Eur. Phys. J. C **63** (2009) 189, arXiv:0901.0002 [hep-ph].

[127] W. Beenakker, R. Hopker, M. Spira, and P. M. Zerwas, *Squark and gluino production at hadron colliders*, Nucl. Phys. B **492** (1997) 51–103, arXiv:hep-ph/9610490.

[128] A. Kulesza and L. Motyka, *Threshold resummation for squark-antisquark and gluino-pair production at the LHC*, Phys. Rev. Lett. **102** (2009) 111802, arXiv:0807.2405 [hep-ph].

[129] A. Kulesza and L. Motyka, *Soft gluon resummation for the production of gluino-gluino and squark-antisquark pairs at the LHC*, Phys. Rev. D **80** (2009) 095004, arXiv:0905.4749 [hep-ph].

[130] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen, and I. Niessen, *Soft-gluon resummation for squark and gluino hadroproduction*, JHEP **12** (2009) 041, arXiv:0909.4418 [hep-ph].

[131] W. Beenakker et al., *Squark and gluino hadroproduction*, Int. J. Mod. Phys. A **26** (2011) 2637, arXiv:1105.1110 [hep-ph].

[132] ATLAS Collaboration, *ATLAS Pythia 8 tunes to* 7 *TeV data*, ATL-PHYS-PUB-2014-021, 2014, https://cds.cern.ch/record/1966419.

[133] J. M. Campbell and R. K. Ellis, $t\bar{t}W^{\pm}$ *production and decay at NLO*, JHEP **07** (2012) 052, arXiv:1204.5678 [hep-ph].

[134] A. Lazopoulos, T. McElmurry, K. Melnikov, and F. Petriello, *Next-to-leading order QCD corrections to* $t\bar{t}Z$ *production at the LHC*, Phys. Lett. B **666** (2008) 62–65, arXiv:0804.2220 [hep-ph].

[135] M. Czakon, P. Fiedler, and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through* $O(\alpha_S^4)$, Phys. Rev. Lett. **110** (2013) 252004, arXiv:1303.6254 [hep-ph].

[136] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair cross-section at hadron colliders*, Comput. Phys. Commun. **185** (2014) 2930, arXiv:1112.5675 [hep-ph].

[137] P. Z. Skands, *Tuning Monte Carlo generators: The Perugia tunes*, Phys. Rev. D **82** (2010) 074018, arXiv:1005.3457 [hep-ph].

[138] H.-L. Lai et al., *New parton distributions for collider physics*, Phys. Rev. D **82** (2010) 074024, arXiv:1007.2241 [hep-ph].

[139] N. Kidonakis, *Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^-*, Phys. Rev. D **82** (2010) 054018, arXiv:1005.4451 [hep-ph].

[140] J. M. Campbell and R. K. Ellis, *Update on vector boson pair production at hadron colliders*, Phys. Rev. D **60** (1999) 113006, arXiv:hep-ph/9905386.

[141] NNPDF Collaboration, R. D. Ball et al., *Parton distributions for the LHC run II*, JHEP **04** (2015) 040, arXiv:1410.8849 [hep-ph].

[142] S. Catani and M. Grazzini, *An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC*, Phys. Rev. Lett. **98** (2007) 222002, arXiv:hep-ph/0703012 [hep-ph].

[143] D. S. M. Alves, E. Izaguirre, and J. G. Wacker, *Where the Sidewalk Ends: Jets and Missing Energy Search Strategies for the 7 TeV LHC*, JHEP **10** (2011) 012, arXiv:1102.5338 [hep-ph].

[144] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data*, Eur. Phys. J. C **76** (2016) 666, arXiv:1606.01813 [hep-ex].

[145] ATLAS Collaboration, *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*, ATL-DAQ-PUB-2016-001, 2016, https://cds.cern.ch/record/2136007.

[146] ATLAS Collaboration, *Trigger Menu in 2016*, ATL-DAQ-PUB-2017-001, 2017, https://cds.cern.ch/record/2242069.

[147] C. G. Lester and D. J. Summers, *Measuring masses of semiinvisibly decaying particles pair produced at hadron colliders*, Phys. Lett. B **463** (1999) 99–103, arXiv:hep-ph/9906349 [hep-ph].

[148] A. Barr, C. Lester, and P. Stephens, *A variable for measuring masses at hadron colliders when missing energy is expected; mT2: the truth behind the glamour*, J. Phys. G **29** no. 10, (2003) 2343, http://stacks.iop.org/0954-3899/29/i=10/a=304.

[149] ATLAS Collaboration, *Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-015, 2015, https://cds.cern.ch/record/2037613.

[150] ATLAS Collaboration, *Expected performance of missing transverse momentum reconstruction for the ATLAS detector at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-023, 2015, https://cds.cern.ch/record/2037700.

[151] M. V. Garzelli, A. Kardos, C. G. Papadopoulos, and Z. Trocsanyi, *$t\bar{t}\,W^{+-}$ and $t\bar{t}\,Z$ Hadroproduction at NLO accuracy in QCD with Parton Shower and Hadronization effects*, JHEP **11** (2012) 056, arXiv:1208.2665 [hep-ph].

[152] J. M. Campbell and R. K. Ellis, *$t\bar{t}W^{+-}$ production and decay at NLO*, JHEP **07** (2012) 052, arXiv:1204.5678 [hep-ph].

[153] A. Lazopoulos, T. McElmurry, K. Melnikov, and F. Petriello, *Next-to-leading order QCD corrections to $t\bar{t}Z$ production at the LHC*, Phys. Lett. B **666** (2008) 62–65, arXiv:0804.2220 [hep-ph].

[154] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C **71** (2011) 1554, arXiv:1007.1727 [physics.data-an].

[155] A. L. Read, *Presentation of search results: the CL_S technique*, J. Phys. G **28** (2002) 2693.

[156] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C **71** (2011) 1554, arXiv:1007.1727 [physics.data-an].

[157] ATLAS Collaboration, *Search for heavy charged long-lived particles in proton–proton collisions at $\sqrt{s} = 13$ TeV using an ionisation measurement with the ATLAS detector*, Phys. Lett. B **788** (2019) 96, arXiv:1808.04095 [hep-ex].

[158] ATLAS Collaboration, *Search for heavy charged long-lived particles in the ATLAS detector in 31.6 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **99** (2019) 092007, arXiv:1902.01636 [hep-ex].

[159] CMS Collaboration, *Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Lett. B **806** (2020) 135502, arXiv:2004.05153 [hep-ex].

[160] CMS Collaboration, V. Khachatryan et al., *Search for long-lived charged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **94** no. 11, (2016) 112004, arXiv:1609.08382 [hep-ex].

[161] ATLAS Collaboration, G. Aad et al., *A search for the decays of stopped long-lived particles at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **07** (2021) 173, arXiv:2104.03050 [hep-ex].

[162] 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*, arXiv:2201.02472 [hep-ex].

[163] CMS Collaboration, *Search for decays of stopped exotic long-lived particles produced in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JHEP **05** (2018) 127, arXiv:1801.00359 [hep-ex].

[164] LHCb Collaboration, *Search for long-lived heavy charged particles using a ring imaging Cherenkov technique at LHCb*, Eur. Phys. J. C **75** no. 12, (2015) 595, arXiv:1506.09173 [hep-ex].

[165] ATLAS Collaboration, *Search for nonpointing and delayed photons in the diphoton and missing transverse momentum final state in 8 TeV pp collisions at the LHC using the ATLAS detector*, Phys. Rev. D **90** (2014) 112005, arXiv:1409.5542 [hep-ex].

[166] CMS Collaboration, *Search for long-lived particles using nonprompt jets and missing transverse momentum with proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Lett. B **797** (2019) 134876, arXiv:1906.06441 [hep-ex].

[167] CMS Collaboration, *Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **104** no. 1, (2021) 012015, arXiv:2012.01581 [hep-ex].

[168] ATLAS Collaboration, *Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, Phys. Rev. D **97** (2018) 052012, arXiv:1710.04901 [hep-ex].

[169] ATLAS Collaboration, *Search for long-lived, massive particles in events with a displaced vertex and a muon with large impact parameter in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Rev. D **102** (2020) 032006, arXiv:2003.11956 [hep-ex].

[170] ATLAS Collaboration, *Search for displaced vertices of oppositely charged leptons from decays of long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Lett. B **801** (2020) 135114, arXiv:1907.10037 [hep-ex].

[171] ATLAS Collaboration, *Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector*, Phys. Rev. Lett. **127** no. 5, (2021) 051802, arXiv:2011.07812 [hep-ex].

[172] CMS Collaboration, *Search for long-lived particles decaying to jets with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **104** no. 5, (2021) 052011, arXiv:2104.13474 [hep-ex].

[173] CMS Collaboration, *Search for long-lived particles decaying to leptons with large impact parameter in proton–proton collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **82** no. 2, (2022) 153, arXiv:2110.04809 [hep-ex].

[174] LHCb Collaboration, *Search for massive long-lived particles decaying semileptonically at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **82** no. 4, (2022) 373, arXiv:2110.07293 [hep-ex].

[175] ATLAS Collaboration, *SUSY Summary Plots March 2022*, tech. rep., CERN, Geneva, Apr, 2022. https://cds.cern.ch/record/2805985.

[176] ATLAS Collaboration, *Trigger Menu in 2017*, ATL-DAQ-PUB-2018-002, 2018, https://cds.cern.ch/record/2625986.

[177] ATLAS Collaboration, *Trigger Menu in 2018*, ATL-DAQ-PUB-2019-001, 2019, https://cds.cern.ch/record/2693402.

[178] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen, and I. Niessen, *NNLL resummation for squark-antisquark pair production at the LHC*, JHEP **01** (2012) 076, arXiv:1110.2446 [hep-ph].

[179] W. Beenakker, T. Janssen, S. Lepoeter, M. Krämer, A. Kulesza, E. Laenen, I. Niessen, S. Thewes, and T. Van Daal, *Towards NNLL resummation: hard matching coefficients for squark and gluino hadroproduction*, JHEP **10** (2013) 120, arXiv:1304.6354 [hep-ph].

[180] W. Beenakker, C. Borschensky, M. Krämer, A. Kulesza, E. Laenen, V. Theeuwes, and S. Thewes, *NNLL resummation for squark and gluino production at the LHC*, JHEP **12** (2014) 023, arXiv:1404.3134 [hep-ph].

[181] J. Debove, B. Fuks, and M. Klasen, *Threshold resummation for gaugino pair production at hadron colliders*, Nucl. Phys. B **842** (2011) 51–85, arXiv:1005.2909 [hep-ph].

[182] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, *Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV*, JHEP **10** (2012) 081, arXiv:1207.2159 [hep-ph].

[183] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, *Precision predictions for electroweak superpartner production at hadron colliders with Resummino*, Eur. Phys. J. C **73** (2013) 2480, arXiv:1304.0790 [hep-ph].

[184] J. Fiaschi and M. Klasen, *Neutralino-chargino pair production at NLO+NLL with resummation-improved parton density functions for LHC Run II*, Phys. Rev. D **98** no. 5, (2018) 055014, arXiv:1805.11322 [hep-ph].

[185] W. Beenakker, M. Klasen, M. Kramer, T. Plehn, M. Spira, and P. M. Zerwas, *The Production of charginos / neutralinos and sleptons at hadron colliders*, Phys. Rev. Lett. **83** (1999) 3780–3783, arXiv:hep-ph/9906298, [Erratum: Phys.Rev.Lett. 100, 029901 (2008)].

[186] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, J. Phys. G **43** (2016) 023001, arXiv:1510.03865 [hep-ph].

[187] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, arXiv:2007.02645 [hep-ex].

[188] J. K. Roloff, *Exploring the Standard Model and beyond with jets from proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS Experiment*, May, 2019. https://cds.cern.ch/record/2677419.

[189] ATLAS Collaboration, *Transverse momentum response and reconstruction efficiency for jets from displaced decays in the ATLAS detector*, ATL-PHYS-PUB-2019-025, 2019, https://cds.cern.ch/record/2682843.

[190] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **71** (2011) 1630, arXiv:1101.2185 [hep-ex].

[191] ATLAS Collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*. Technical design report. Geneva, Jun, 2020. https://cds.cern.ch/record/2719855.

[192] CMS Collaboration, *Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS*, Eur. Phys. J. C **81** no. 9, (2021) 800, arXiv:2104.01927 [hep-ex].

[193] ATLAS Collaboration, *Searches for new phenomena in events with two leptons, jets, and missing transverse momentum in* 139 fb^{-1} *of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, arXiv:2204.13072 [hep-ex], Submitted to Eur. Phys. J. C.